

Binary Classification of Agricultural Crops Using Sentinel Satellite Data and Machine Learning Techniques

1st Paolo Bertellini
R&D - Data science
Abaco Group
Mantua, Italy
p.bertellini@abacogroup.eu

2nd Gianluca D'Addese
Department of FIM
UNIMORE
Modena, Italy
gianluca.daddese@unipr.it
0000-0001-7755-0893

3rd Giorgia Franchini
Department of FIM
UNIMORE
Modena, Italy
giorgia.franchini@unimore.it
0000-0001-9082-8087

4th Simone Parisi
R&D - Data science
Abaco Group
Mantua, Italy
s.parisi@abacogroup.eu

5th Carmelo Scribano
Department of FIM
UNIMORE
Modena, Italy
carmelo.scribano@unimore.it
0000-0003-1006-7826

6th Daniele Zanirato
Department of Statistical Sciences
University of Bologna
Bologna, Italy
daniele.zanirato@studio.unibo.it

7th Marko Bertogna
Department of FIM
UNIMORE
Modena, Italy
marko.bertogna@unimore.it
0000-0003-2115-4853

Abstract—The automated process of determining the crop type carried on plots of land, leveraging data provided by earth observation satellites, represents a highly valuable ability that can serve as a foundation for subsequent analyses or as input for calibrating models, such as Decision Support Systems. This paper presents a study on the task of crop classification starting from indices derived from imagery data provided by ESA Satellites Sentinel 1 and 2. We create a valuable tool to verify farmers' claims, especially in relation to state subsidies for specific crops of interest. To this purpose, we focus on perfecting a binary classification for each of five crops of interest (Tomatoes, Soy, Sugar Beet, Rice, and Wheat), aimed to accurately discern the target crop against any other possible crop. The paper investigates various preprocessing techniques to create a dataset suitable for traditional machine learning methods, which presumes that each land plot to classify is represented by a fixed set of features. To deal with inevitable missing observations caused by clouds or other environmental factors, we investigate different imputation strategies (linear interpolation and constant value filling). Complementary, we study the impact of imbalanced classification labels and evaluate the effectiveness of standard balancing techniques. The findings offer practical implications for monitoring and optimizing agricultural practices in the context of precision farming and sustainable agriculture.

Index Terms—Crop Classification, AgriAI, Machine Learning for Agriculture, Decision Support Systems, Sentinel, Copernicus

I. INTRODUCTION

IN LINE with the objectives for the period 2023-2027 of the European Union's Common Agricultural Policy (CAP), Italy has allocated funds for the production of protein crops (€544 million per year). The aim is to incentivize local agricultural production of crops such as Tomatoes, Soy, Sugar Beet, Rice, and Wheat in local production through

the granting of reimbursements by the state. Consequently, there is a pressing need for methods to verify the authenticity of indigenous crops. On-site inspections prove to be costly and inadequate given the scale of the problem, meanwhile satellite remote sensing is a promising technology to classify crops since it can provide periodically large-scale observations of ground objects [11], [13]. The objective of this study is to automate such verification procedures by using Machine Learning (ML) systems applied to satellite data of the relevant geographical areas. Specifically, we focus on the central area of region Emilia-Romagna (Italy). The problem can be cast as a binary classification: given the data pertaining to a particular field and the farmer's statement regarding the crop, our goal is to verify the veracity of the information given.

The usage of machine learning algorithms for similar tasks is explored in many recent contributions. Among them, Random Forest [3], [7], [12] or decision tree-based [3], [9] classifiers are among the most commonly employed methods for handling this type of data. Data for this paper has been sourced from Sentinel satellites, deployed within the Copernicus program, managed by the European Space Agency (ESA). Sentinel-2 images, have already proven to be a valuable data source for crop mapping in different regions and countries like Central Europe [7], Spain [3], [10], Lebanon [9] and China [12].

The remainder of this paper is structured as follows. Section II, provides a description of the composition of the dataset, highlighting the type of utilized data, the characteristics of the satellite observations, and the distribution of the different crop types. Section III provides a theoretical overview of the Machine Learning method employed: Random Forests



Fig. 1: Study area and locations of ground truth samples (red area). The analyzed fields are located in the central part of the region. This specific area of interest tends to prioritize the production of Wheat, Soy, and Sugar Beet over Rice and Tomatoes.

(RF) with an emphasis on the data pre-processing work, from handling missing dates due to adverse weather conditions to addressing the dataset’s imbalance. Section IV presents the performance of the proposed pipeline for crop classification. The aim of this section to assess the accuracy of our model using the complete time frame of the planting process, from seeding to product harvest. In Section V we briefly go through the main result of this work to provide possible branches of further research on the topic.

II. DATASET AND PREPROCESSING

All the data used for the remote sensing purpose in this work is acquired from the earth observation satellites deployed by ESA during the missions Sentinel-1 and Sentinel-2.¹ Each of these missions deployed in a near-polar sun-synchronous orbit a twin pair of satellites (named Sentinel-1A and Sentinel-1B, Sentinel-2A and Sentinel-2B), which provide sensor observation capabilities depending on the objective of the mission. In particular, Sentinel-1 satellites are equipped with C-band synthetic-aperture radar (C-SAR), while Sentinel-2 satellites are instead equipped with passive Multi-spectral camera operating in 13 distinct bands spanning the spectrum of visible, near-infrared and short wave infrared. Sentinel-1A has a revolution period, hence a temporal revolution, of 12 days, whereas Sentinel-2A and Sentinel 2B of 3 to 5 days, depending on the area. Spatial resolution of Sentinel-1 observations is 10 meters, hence the crop field is discretized in squares of $100m^2$. Sentinel-2 data has a different spatial resolution according to the sensor by which they are collected, either 10 or 20 meters. From a qualitative standpoint, Sentinel-1 active radar sensors imply it always collects the data independently of

atmospheric conditions, while Sentinel-2 satellites, relying on passive optical sensors, can produce missing or fragmentary observations due to clouds presence. Starting from the raw observations obtained from the Sentinel satellites, we leverage a total of 16 numerical indices, 12 of them are obtained by Sentinel-2A and Sentinel-2B, and 4 came from Sentinel-1A and Sentinel-1B. Sentinel-1B stopped working in December 2021 and is currently unavailable. Overall, the time frame of the whole dataset spans across 14 months.

A. Data Preparation

Each of the 2 indexes coming from Sentinel-1, named *backscatter* and *coherence*, is further divided in the two polarization VV (co-polarized) and VH (cross-polarized). Backscatter defines the portion of the radar signal that get reflected from the earth’s surface straight to the radar antenna, while coherence is defined as the normalized value of the complex cross-correlation between a pair of SAR observation spaced by a period of 12 days. Intuitively, a very low value of the coherence, might indicate a big change in how the field presents itself in 12 days time-difference. To prioritize the number of available observations over the homogeneity of those observations, we also leverage the multiple observations of fields that are visible from partially overlapping orbits This is acceptable since we use pixel statistics (as detailed below), not raw observations, to analyze our data. The differences between observations from different orbits are therefore considered negligible.

The most popular index obtainable from Sentinel-2 observations is the NDVI (Normalized Difference Vegetation index), eq. (1), defined as the ratio of the difference and the sum of the reflected radiation in the near infrared and red, which is a good indicator of the amount of chlorophyll in a field.

$$NDVI := \frac{NIR - RED}{NIR + RED} \quad (1)$$

In order to produce a suitable time series representation for each field, we devise a simple yet effective post-processing strategy of the raw satellite observations. For each record, using its vector geometry, the corresponding patch is cut out for each available observation. Then, each patch is reduced to a set of five statistics (mean, mode, standard deviation, maximum and minimum). The full set of observation obtained over the growing season constitute the field time series that we aim to classify, and it’s labeled with a single class identifying the crop being grown.

Our dataset consists of 49 different types of crops and 16,684 sample fields. Among those 49 crops only five are subject to crop-specific subsidies and only for them it is necessary to verify the truthfulness of the farmer’s declarations. Getting more specific, the dataset consists of 12,496 samples with specific target crop, divided in: 743 Tomatoes, 63 Rice, 5,214 Wheat, 2,974 Sugar Beet and 3,502 Soy, and 4,188 samples with crops not subject to specific subsidies and therefore considered as ”OTHER” class.

¹<https://sentinel.esa.int/web/sentinel/missions>

TABLE I: Amount of crop types for the available fields

<i>Crop</i>	Amount
Tomatos	743
Soy	3502
Sugar Beet	2974
Rice	63
Wheat	5214
Others	4188
Total	16684

B. Analysis of the used Dataset

Therefore, we propose a qualitative analysis of the available data, this is helpful in order to build an intuition for the possibility to accurately classify the observed crop based on the data gathered. The bar plots (Figure 2a-Figure 2e) depict the variation of the NDVI value across all the fields during the season. The black line represents the behavior of the NDVI average across the fields, while the green vertical bands represent one standard deviation above and one below the mean. The sowing and harvesting phases of the different crops (See Table II) are delimited by the dashed blue and red vertical lines, respectively. The behavior of the NDVI appears to be highly indicative of the growth trend: the blue area (sowing) is usually followed by an increase in the value of the index, whereas the red bands (harvesting) correspond to a steep decrease of the NDVI value. This has a straightforward interpretation given the meaning of the NDVI (indicative of the amount of chlorophyll in the field): after the sowing period, the amount of chlorophyll increases during the vegetative growth, and decreases rapidly with the harvesting. Looking at the width of the green error bars, we can also appreciate the variability changes during the different phases and among different crop types.

III. METHOD

Since ultimately the objective is to verify if the crop is the one declared by the farmer (among the 5 ones of interest) versus any other different crop type, we propose to train a binary classifier for each of the crop of interest. After the common data preparation steps detailed above, a binary classification dataset is prepared for each of the target crops, by assigning to the samples labeled as it the label positive (1) and to any other sample the label negative (0). At deployment stage, the binary classifier is fed a new unseen crop that should belong to a specific class and assess the truthfulness of the declared crop.

It is known in the literature that Random Forest has great performances in remote sensing classification tasks [1]. A Random Forest (RF) [2] is an ensemble method whose base estimators are decision trees. This method reduces bias and variance thanks to the introduction of multiple uncorrelated voters, because each of the trees has the chance to learn a different pattern in the data and then this knowledge is combined.

A. Missing data management

Due to variable atmospheric conditions and the nature of Sentinel-2 passive optical sensors, we have a lot of images that were totally or partially covered, hence we opted to disregard covered units up to a certain threshold during pre-processing. In order to represent each field as a fixed set of features, we define a common set of observation dates. While other strategies are possible, we prefer to preserve all the remaining observations by considering all the timestamps that corresponds to at least a single observation in the whole dataset, resulting in 110 valid timestamps. Complete observations of all the 16 indices are usually not available for all the 110 timestamps for each field, Figure 3 illustrates the approximate distribution of the amount of observations per field. We experiment with two standard techniques to fill the missing values. The first is to insert an out-of-scale value in all empty dates. The other technique we evaluate is linear interpolation [8]. For each field, we linearly interpolated all the statistics (mean, mode, minimum, maximum, standard deviation) in each of the empty dates.

B. Data imbalance management

The creation of the five binary classification dataset inherently causes large data imbalance that is of potential harm when training a classifier, as it could learn a bias towards the dominant class. To tackle this problem, we investigate 3 popular ways to obtain a balanced dataset before training the classifier: undersampling, random oversampling, SMOTE (Synthetic Minority Oversampling Technique) [5].

Undersampling consists in balancing the training set by randomly eliminating units from the majority class. The main drawback of this is the information loss, as the classifier may lose the chance to see some different and valuable examples of the adversary class. On the other hand, this approach leads to having a smaller dataset, which makes the training phase less time-consuming and energy-demanding, also reducing the environmental impact in the perspective of GreenAI. On a different note, random oversampling consists of balancing the training set by randomly sampling more units of the minority class. While this approach does not lose any information, it presents other drawbacks. The dataset becomes larger without actually adding new information: this translates in an increased computational cost and training time, without a corresponding increase in the algorithm performance. Furthermore, we run the risk to induce overfitting, reducing the classifier's ability to generalize with respect to undersampling, likely leading to many false positives [6]. Finally, SMOTE is a balancing technique that focus on generating synthetic samples for the minority class, this is achieved by interpolation in feature space between two neighbor samples from the minority class. This approach mitigates the risk of overfitting, but it is often unclear whether the newly generated units are actually realistic, therefore the risk is to create noise or introducing a bias in the training sample. For this reason, we refrain from leveraging SMOTE in the following experimental section.

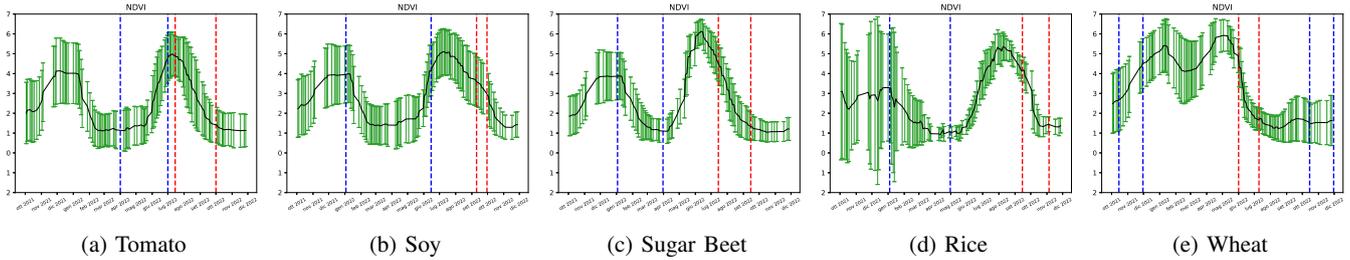


Fig. 2: Bar plots showing NDVI index across time for different crops. Vertical green error bars represent one standard deviation above and one below the average, computed among the fields. Vertical dashed blue lines enclose the sowing periods, red ones enclose the harvesting ones (See Table II)

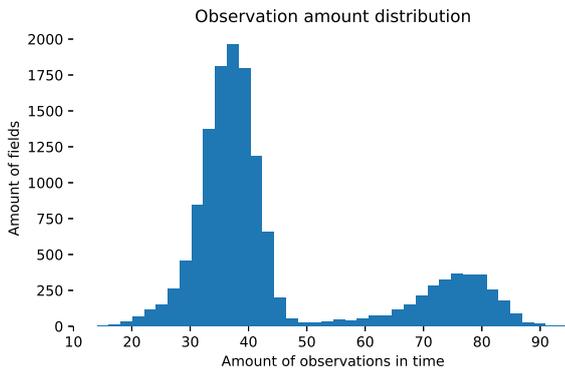


Fig. 3: Distribution of the amount of observations.

Intuitively, we could prospect that undersampling could be the best option for our case, given the dataset to be large enough for the model to see enough variability in the adversary class. Moreover, since false declarations are a rare occurrence (<5%), we could foresee that among the units flagged as suspicious (classified as negative for the class of interest), there would be a predominance of false positive. False positive could trigger and unnecessary verification, while false negatives could mean a false declaration going undetected, hence the trade-off should be carefully evaluated in the deployment scenario.

IV. EXPERIMENTS

In this section, we present and discuss the results of the methodology detailed in the first part of this manuscript. To restate, the goal of our research is to leverage the observations provided by the ESA Copernicus satellites in order to provide the regulatory agency with feedback on the truthfulness of the stated crop for a given plot of land. A negative outcome of the automatic classification might trigger an on-site inspection, hence it is crucial to be able to provide a reliable classification.

A. Experimental Setup

For the purpose of this work, five crops have been considered, which are reported in Table II along with the most relevant information of the production's life cycle. Our training data include all the available observation for the 2022 season

TABLE II: Start and end of sowing and harvesting operations for the crops of interest, along with the minimum and maximum observed days of duration of the crop cycle. The information is obtained from the data for the 2022 growth season. Dates are expressed in (mm/dd) format.

Crop	Sowing		Harvesting		Cycle duration (d)	
	start	end	start	end	min	max.
Tomato	04/01	07/01	07/15	10/01	95	110
Soy	01/03	06/15	09/10	09/30	120	150
Sugar Beet	01/03	03/31	07/15	09/15	60	90
Rice	01/04	04/30	09/15	11/05	160	180
Wheat	10/15	11/30	06/01	07/10	210	230

(completed), we train a binary classifier for each of the five crops of interest, leveraging the datapoints labeled with the target class as positive examples (class 1), and the remaining datapoints as negative examples (class 0), those include the remaining 4 classes along with other crop classes. In this evaluation, we use as training data all the data available between the beginning of the sowing phase to the end of the harvesting for the target crop. This analysis is relevant for two aspects: (a) It allows for a controlled setup to experiment with multiple options for dealing with the problems of extremely unbalance between positive and negative examples (Section III-B) and to evaluate the two options for dealing with missing observations (Section III-A). (b) It allows defining a baseline for the classifier's performance in a best-case scenario.

To assess the effectiveness and reliability of the classification models we use the standard metrics in the literature for binary classification: Precision, Recall and F1-Score, defined below.

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}} \quad (2)$$

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}} \quad (3)$$

$$\text{F1-score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (4)$$

Given the extreme unbalance of the dataset uses for testing, it is best to avoid relying on the Accuracy, since a high accuracy may be achieved by simply predicting the majority class.

For each binary classifier, we isolate a uniformly sampled 25% of the available data for evaluation, while the remaining 75% is used for training. The very same train-test subset have been used for all the experiments, to ensure a fair comparison.

Before training each classifier, we further refine the training data by narrowing the observation window to only include observations between the start of the sowing season to the end of the harvesting season for the target crop. This is especially important since the same field might be grown with different crops through the year, which might induce a bias in the classifiers.

B. Results

TABLE III: Comparison of Interpolation with out-of-range placing for filling statistics of missing observation dates.

<i>Crop</i>	<i>Strategy</i>	Precision	Recall	F1-Score
Tomato	<i>Interp.</i>	0.89	0.59	0.71
	<i>Fill</i>	0.89	0.62	0.73
Soy	<i>Interp</i>	0.915	0.86	0.89
	<i>Fill</i>	0.92	0.86	0.89
Sugar Beet	<i>Interp</i>	0.99	0.93	0.96
	<i>Fill</i>	0.99	0.92	0.96
Rice	<i>Interp</i>	1.0	0.75	0.85
	<i>Fill</i>	1.0	0.375	0.54
Wheat	<i>Interp</i>	0.93	0.93	0.93
	<i>Fill</i>	0.93	0.93	0.93

a) *Management of missing observations*: The first fundamental experiment involves comparing different strategies for filling missing observations in the statistics, as introduced in Section III-A. The results obtained with two strategies, linear interpolation and filling missing values with an out-of-scale value (-1000), are presented in Table III. The choice of these strategies appears to have minimal impact on the evaluated metrics, except for the *Rice* crop, which exhibits a significantly low recall when using the constant value strategy. However, it is challenging to precisely attribute this result solely to the interpolation strategy, due to the crop's extreme under-representation across the dataset. For the above reasons, in the remainder of this analysis, we opted to use the Interpolation strategy.

b) *Dataset balancing*: The second part of the experimental evaluation focuses on determining the optimal strategy for handling the highly unbalanced dataset used to train the binary classifier. We compare the performance of the baseline classifier trained on the original imbalanced training set with the same model trained using two different resampling techniques: undersampling and oversampling. Undersampling involves reducing the number of instances from the majority class to achieve a more balanced representation of the classes, conversely oversampling involves increasing the number of instances in the minority class to address the class imbalance. These and other resampling strategies are further described in Section III-B.

By analyzing the results presented in Table IV, noteworthy observations can be made regarding the two resampling techniques. The application of the undersampling strategy resulted

TABLE IV: Effects of training data balancing on the unbalanced test data.

<i>Crop</i>	<i>Train Data</i>	Precision	Recall	F1-Score
Tomato	Unbalanced	0.89	0.59	0.71
	Undersampling	0.39	0.92	0.54
	Oversampling	0.83	0.72	0.77
Soy	Unbalanced	0.915	0.86	0.89
	Undersampling	0.832	0.93	0.88
	Oversampling	0.9	0.91	0.91
Sugar Beet	Unbalanced	0.9	0.93	0.96
	Undersampling	0.95	0.96	0.95
	Oversampling	0.99	0.99	0.99
Rice	Unbalanced	1.0	0.75	0.85
	Undersampling	0.12	0.87	0.22
	Oversampling	1.0	0.75	0.86
Wheat	Unbalanced	0.93	0.93	0.93
	Undersampling	0.88	0.97	0.92
	Oversampling	0.92	0.95	0.93

in an increase in the number of false positives, where crops are mistakenly classified as the target crop. Consequently, this led to a notable decrease in the Precision metric, reaching drastic levels for certain cases as highlighted in the results table. This issue can be attributed to the significant reduction in the training data caused by the undersampling procedure. As a result, the negative effect is particularly pronounced for the most underrepresented crops, such as Rice and Tomato. It's worth noticing that False positive are possibly the least desirable outcome in our application scenario, while a false negative leads to further investigations on the effective crop being carried out, a false negative means that bogus declaration are more likely to be undetected.

On the other end the oversampling strategy shown promising results, with a noticeable reduction in false negatives with respect to the baseline with only a manageable increase in false positive, the overall superiority of this approach is hence validated by an increase in F1 score for all the crops. To conclude, while we can't recommend relying on an undersampling strategy, we are confident in suggesting the oversampling approach in order to reduce the number of occurrences of an investigation being triggered by mistake.

c) *Choice of the classifier*: The previous results were all obtained with a Random Forest classifier, to demonstrate the validity of our choice, therefore hereafter we evaluate the alternative usage of the very powerful and popular binary classifier SVM [4] This classifier builds a separation hyperplane by choosing support vectors, those are the defined as the harder to classify points. For a simple yet meaningful comparison, we compare the results obtained without train dataset balancing, using the interpolation strategy for the management of the missing observation.

In Table V we report the results. It's clear that RF slightly outperforms SVM for all the crops, of particular interest it is the Rice crop, that never gets correctly classified, hence scoring zero in all the metrics. The reason lies in the very small number of fields labeled as Rice (Table I), with SVM clearly requiring a larger training set. Another advantage of RF is its superior interpretability compared to SVM. With each decision

TABLE V: Comparison of the results obtained with Random Forest and SVM.

<i>Crop</i>	<i>Classifier</i>	Precision	Recall	F1-Score
Tomato	RF	0.89	0.59	0.71
	SVM	0.876	0.55	0.68
Soy	RF	0.915	0.86	0.89
	SVM	0.87	0.87	0.87
Sugar Beet	RF	0.99	0.93	0.96
	SVM	0.97	0.93	0.95
Rice	RF	1.0	0.75	0.85
	SVM	0.0	0.0	0.0
Wheat	RF	0.93	0.93	0.93
	SVM	0.9	0.94	0.92

tree acting as a transparent flowchart, it becomes easier to comprehend how the model classifies data points. The consensus-based decision-making further enhances stability and improves generalization abilities. In contrast, SVM's optimal hyperplane may lack clear interpretability.

V. CONCLUSION AND FUTURE WORKS

In conclusion, this work provides a significant contribution to agricultural research by demonstrating the applicability of machine learning techniques and the utility of satellite data for crop classification. The proposed methodology can be applied to verify the authenticity of farmers' claims, especially regarding state subsidies for specific crops of interest.

Our work lays the foundation for further research in the field of agricultural field classification using satellite imagery and ML techniques. Several avenues for future work can be explored to enhance and extend the findings presented in this paper.

- 1) Multi-class classification could provide a more comprehensive understanding of crop distribution and facilitate more accurate crop monitoring and yield estimation.
- 2) Incorporating additional data sources, such as weather data, soil composition, or historical crop records, could improve the accuracy and robustness of the classification models.
- 3) Investigating the development of dynamic classification models that can adapt to changing environmental conditions and crop phenology could enable real-time monitoring and detection of crop changes, disease outbreaks, or other significant events that affect agricultural fields.

By pursuing these future research directions, we can advance the field of agricultural field classification and contribute to the development of more accurate, efficient, and sustainable agricultural practices.

ACKNOWLEDGMENT

This work has been partially supported by the INdAM research group GNCS.

The publication was created with the co-financing of the European Union-FSE-REACT-EU, PON Research and Innovation 2014-2020 DM1062/2021.

This work has been partially supported by the industrial collaboration between ABACO Group SpA and UNIMORE.

REFERENCES

- [1] Belgiu, M., Drăguț, L.: Random forest in remote sensing: A review of applications and future directions. *ISPRS journal of photogrammetry and remote sensing* **114**, 24–31 (2016)
- [2] Breiman, L.: Random forests. *Machine learning* **45**, 5–32 (2001)
- [3] Campos-Taberner, M., García-Haro, F.J., Martínez, B., Sánchez-Ruíz, S., Gilabert, M.A.: A copernicus sentinel-1 and sentinel-2 classification framework for the 2020+ european common agricultural policy: A case study in valència (spain). *Agronomy* **9**(9) (2019). <https://doi.org/10.3390/agronomy9090556>, <https://www.mdpi.com/2073-4395/9/9/556>
- [4] Chapelle, O., Haffner, P., Vapnik, V.N.: Support vector machines for histogram-based image classification. *IEEE transactions on Neural Networks* **10**(5), 1055–1064 (1999)
- [5] Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research* **16**, 321–357 (2002)
- [6] Drummond, C., Holte, R.C., et al.: C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In: *Workshop on learning from imbalanced datasets II*. vol. 11, pp. 1–8 (2003)
- [7] Immitzer, M., Vuolo, F., Atzberger, C.: First experience with sentinel-2 data for crop and tree species classifications in central europe. *Remote Sensing* **8**(3) (2016). <https://doi.org/10.3390/rs8030166>
- [8] Lepot, M., Aubin, J.B., Clemens, F.H.: Interpolation in time series: An introductory overview of existing methods, their performance criteria and uncertainty assessment. *Water* **9**(10), 796 (2017)
- [9] Nasrallah, A., Baghdadi, N., Mhawej, M., Faour, G., Darwish, T., Belhouchette, H., Darwich, S.: A novel approach for mapping wheat areas using high resolution sentinel-2 images. *Sensors* **18**(7) (2018). <https://doi.org/10.3390/s18072089>
- [10] Piedadlobo, L., Hernández-López, D., Ballesteros, R., Chakhar, A., Del Pozo, S., González-Aguilera, D., Moreno, M.A.: Scalable pixel-based crop classification combining sentinel-2 and landsat-8 data time series: Case study of the duero river basin. *Agricultural Systems* **171**, 36–50 (2019). <https://doi.org/https://doi.org/10.1016/j.agry.2019.01.005>
- [11] Wardlow, B., Egbert, S., Kastens, J.: Analysis of time-series modis 250 m vegetation index data for crop classification in the u.s. central great plains. *Remote Sensing of Environment* **108**, 290–310 (06 2007). <https://doi.org/10.1016/j.rse.2006.11.021>
- [12] Yi, Z., Jia, L., Chen, Q.: Crop classification using multi-temporal sentinel-2 data in the shiyang river basin of china. *Remote Sensing* **12**(24) (2020), <https://www.mdpi.com/2072-4292/12/24/4052>
- [13] You, N., Dong, J.: Examining earliest identifiable timing of crops using all available sentinel 1/2 imagery and google earth engine. *ISPRS Journal of Photogrammetry and Remote Sensing* **161**, 109–123 (2020). <https://doi.org/https://doi.org/10.1016/j.isprsjprs.2020.01.001>