



Crop Rotation Modeling for Deep Learning-Based Parcel Classification from Satellite Time Series

Félix Quinton, Loic Landrieu

► To cite this version:

Félix Quinton, Loic Landrieu. Crop Rotation Modeling for Deep Learning-Based Parcel Classification from Satellite Time Series. Remote Sensing, 2021, 13 (22), pp.4599. 10.3390/rs13224599 . hal-03500526

HAL Id: hal-03500526

<https://hal.science/hal-03500526>

Submitted on 22 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Crop Rotation Modeling for Deep Learning-Based Parcel Classification from Satellite Time Series

Félix Quinton, Loic Landrieu

LASTIG, Univ. Gustave Eiffel, ENSG, IGN, F-94160 Saint-Mandé, France

Abstract

While annual crop rotations play a crucial role for agricultural optimization, they have been largely ignored for automated crop type mapping. In this paper, we take advantage of the increasing quantity of annotated satellite data to propose the first deep learning approach modeling simultaneously the inter- and intra-annual agricultural dynamics of parcel classification. Along with simple training adjustments, our model provides an improvement of over 6.6 mIoU points over the current state-of-the-art of crop classification. Furthermore, we release the first large-scale multi-year agricultural dataset with over 300 000 annotated parcels.

1 Introduction

The Common Agricultural Policy (CAP) is responsible for allocating agricultural subsidies in the European Union, which nears 50 billion euros each year [36]. Consequently, monitoring the crop types for subsidy allocation represents a significant challenge for payment agencies, which have encouraged the development of automated crop classification tools based on machine learning [24]. In particular, The Sentinels for Common Agricultural Policy (Sen4CAP) project [19] aims to provide EU member states with algorithmic solutions and best practice studies on crop monitoring based on satellite data from the Sentinel constellation [10]. Despite the inherent difficulty of differentiating between the complex growth patterns of plants, this task is made possible by the nearly limitless access to data and annotations. Indeed, Sentinel-2 offers multi-spectral observations at a high revisit time of five days on average, which are particularly appropriate for characterizing the complex spectral and temporal characteristics of crop phenology. Moreover, farmers declare the crop cultivated in each of their parcels every year. This represents over 10 million of annotations each year for France alone [29], all openly accessible in the Land-Parcel Identification System (LPIS). However, the sheer scale of the problem raises interesting computational challenges: Sentinel-2 gathers over 25Tb of data each year over Europe.

The state-of-the-art of yearly parcel-based crop type classification from Satellite Image Time Series (SITS) is particularly dynamic, especially since the adoption of deep learning methods [15, 27, 31]. However, most methods operate on a single year worth of temporal acquisitions and ignore inter-annual crop rotations. In this paper, we propose the first deep learning framework for classifying yearly crop types from information spanning several years, as represented in Figure 1. We show that we can improve their predictions by a large margin

with straightforward alterations of the top-performing models and their training protocols.

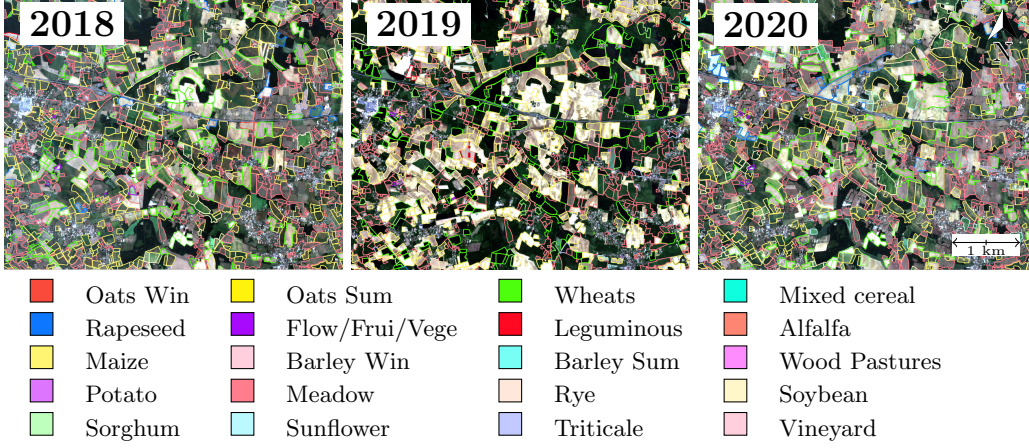


Figure 1: **Multi-Year Sentinel-2 Data.** Details of our area of interest for the three years studied in this article. The crop type of each parcel is represented by the color of a polygon following their contour according to the legend above. This color code is used throughout this article for all figures representing cultivated crops.

Single-Year Crop-Type Classification. Single-year crop-type classification involves the classification of the crop grown in a parcel from a single year worth of observation. Pre-deep learning parcel-based classification methods rely on such as support vector machines [42] or random forests [38] operating on handcrafted descriptors such as the Normalized Difference Vegetation Index. The temporal dynamics are typically handled with stacking [38], probabilistic graphical models [35], or dynamic warping method [5].

In conjunction with growing data availability, the adoption of deep learning-based methods has allowed for a large increase in performance for parcel-based crop classification. The spatial dimension of parcels is typically handled with convolutional neural networks [22], parcel-based statistics [31], or set-based encoders [12]. The temporal dynamics are modeled with temporal convolutions [27], recurrent neural networks [12], hybrid convolutional-recurrent networks [30], and temporal attention [31, 15, 41].

Multiple recent studies [21, 13, 15, 32, 14] have solidified the PSE+LTAE (Pixel Set Encoder & Lightweight Temporal Attention) as the state-of-the-art of crop type classification for these reasons. Furthermore, this network is particularly parsimonious in terms of computation and memory usage, which proves well suited for training on multi-year data. Finally, the code is available (<https://github.com/VSainteuf/lightweight-temporal-attention-pytorch>, accessed on 10/10/21). For these reasons, we choose to use this network as the basis for our analysis and design modifications.

Multi-Year Agricultural Optimization. Most of the literature on multi-year crop rotation focuses on agricultural optimization, *ie*, improving agricultural practices aiming to improve yields. These models generate suggested rotations according to expert knowledge

[11], handcrafted rules [9], or statistical analysis [18], while other models are based on a physical analysis of the soil composition [7] such as the nitrogen cycle [8]. Aurbacher and Dabbert also take a simple economic model into account in their study [1]. More sophisticated models combine different sources of knowledge for better suggestions, such as ROTOR [2] or CropRota [33]. The RPG Explorer software [23] uses a second-order Markov Chain for a more advanced statistical analysis of rotations.

Given the popularity of these tools, it is clear that the careful choice of cultivated crops can substantially impact agricultural yields and is the object of meticulous attention from farmers. This is reinforced by the multi-model, multi-country meta-study of Kollas *et al.* [20], showing that multi-year modeling allows for a large increase in yield prediction. Consequently, we posit that a classification model with access to multi-year data will learn inter-annual patterns to improve its accuracy.

Multi-Year Crop Type Classification. Multi-year crop type classification refers to leveraging information (satellite observations, past declarations) to improve the classification of the grown crop type in agricultural parcels. Osman *et al.* [26] propose to use probabilistic Markov models to predict the most probable crop type from the sequence of past cultivated crops of the previous 3 to 5 years. Giordano *et al.* [16] and Bailly *et al.* [4] propose to model the multi-year rotation with a second-order chain-Conditional Random Field (CRF). Finally, Yamasu *et al.* [40] are the first to propose to analyze multi-year data with a deep convolutional-recurrent model. However, they only choose one image per year and hence do not model both inter- and intra-annual dynamics. In contrast, we propose to explicitly our model operates at both the intra-annual scale by using the sequence of yearly observation and the inter-annual scale by considering past declarations.

We list here the main contributions of this paper:

- We propose a straightforward training scheme to leverage multi-year data and show its impact on yearly agricultural parcel classification.
- We introduce a modified attention-based temporal encoder able to model both inter- and intra-annual dynamics of agricultural parcels, yielding a large improvement in precision.
- We present the first open-access multi-year dataset [25] for crop classification based on Sentinel-2 images, along with the full implementation of our model.
- Our code is open-source and accessible at the following repository: <https://github.com/felixquinton1/deep-crop-rotation>.

i

2 Materials and Methods

We present our dataset and proposed method to model multi-year SITS, along with several baseline methods to assess the performance of its components. We denote by $[1, I]$ the set of years for which satellite observations are available to us and use the compact *pixel-set* format to represent the SITS. For a given parcel and a year $i \in [1, I]$, we denote the corresponding SITS by a tensor x^i of size $C \times S \times T_i$, with C the number of spectral channels, S the

number of pixels within the parcel, and T_i the number of temporal observation available for the year i . Likewise, we denote by $l^i \in \{0, 1\}^L$ the one-hot-encoded label at year i , denoting which kind of crop is cultivated in the considered parcel among a set L of crop types. Note that, in this article, we focus on the prediction of the main culture, *ie*, only one crop type per year.

2.1 Dataset

Our proposed dataset, represented in Figure 2, is based on parcels within the 31TFM Sentinel-2 tile, covering an area of $110 \times 110 \text{ km}^2$ in the South East of France (centered around 4.31N, 46.44E in WGS84). This area is in the Auvergne-Rhône-Alpes region, a major producer of cereal with over 54 000ha of corn and 30 000ha of wheat. Extensive livestock production makes meadow the most common crop type with over 60% of declared parcels in the LPIS. The most frequent crop rotations are permanent cultures (meadows, vineyards, pasture) and alternating between corn, wheat, and rapeseed.

Our satellite time series are constituted of Sentinel-2 level 2A images. We discard the bands B01, B09, and B10 and resample the remaining 10 spectral bands to a spatial resolution of 10m per pixel with bilinear interpolation. Our data spans three years of acquisition: 2018, 2019, and 2020, with respectively 36, 27 and 29 valid entries, see Figure 3. The length of sequences varies due to the automatic discarding of cloudy tiles by the data provider THEIA [3]. We do not apply further preprocessing such as cloud removal or radiometric calibration than what is already performed by the data provider THEIA.

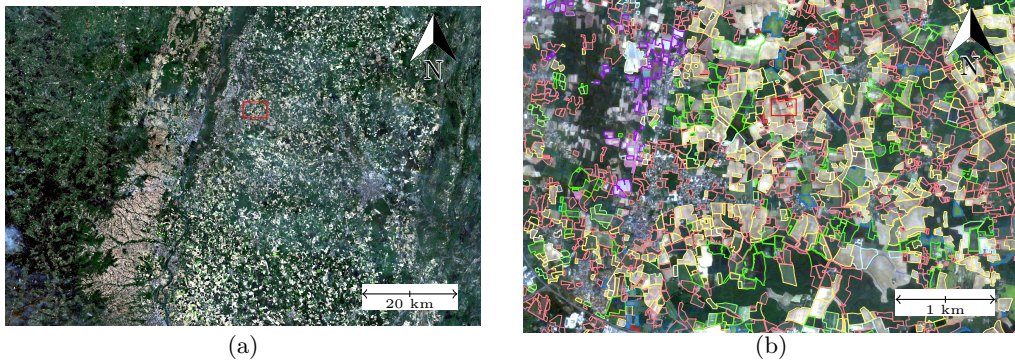


Figure 2: **Area of Interest.** The studied parcels are taken from the 31TFM Sentinel-2 tile, covering $110 \times 110 \text{ km}$ and containing over 103 602 parcels meeting our size, shape, and stability criteria. (a) Large view of the tile. (b) Detail of the area.

We select *stable* parcels, meaning that their contours only undergo minor changes across the three studied years. We also discard very small parcels (under 800m^2) small or with very narrow shapes to reflect the resolution of the Sentinel-2 satellite. Each parcel has a ground truth cultivated crop type for each year corresponding to the main culture as reported by the French LPIS, whose precision is estimated at over 97% as reported by the French Payment Agency. Note that we ignore secondary cultures for parcels with multiple growth cycles. To limit class imbalance, we only keep crop types among a list of 20 of the most cultivated species in the area of interest. In sum, our dataset is composed of 103 602 parcels, each

associated with three image time sequences and three crop annotations corresponding to the farmers’ declarations for 2018, 2019, and 2020.

The Sentinel2Agri dataset [15], composed of parcels from the same area, is composed of 191 703 parcels. We can estimate that our selection criteria exclude approximately every other parcel. A more detailed analysis of the evolving parcel partitions across different plots could lead to retaining a higher proportion of the original parcels.



Figure 3: **Intra-Year Dynamics.** Evolution of two areas across three seasons of the year 2020. The top parcels contain mainly meadow parcels, while the bottom one comprises more diverse crops. The aspect of most parcel drastically changes across one year’s worth of acquisition, corresponding to different phases in the growth cycle. (a) Winter. (b) Spring. (c) Summer

As represented in Table 1, the dataset is still imbalanced: more than 60% of declarations correspond to meadows. In comparison, potato is cultivated in less than 100 parcels each year in the area of interest.

2.2 Pixel-Set and Temporal Attention Encoders

The Pixel Set Encoder (PSE) [15] is an efficient spatio-spectral encoder which learns expressive descriptors of the spectral distribution of the observations by randomly sampling pixels within a parcel. Its architecture is inspired by set-encoding deep architecture [28, 22], and dispenses us from preprocessing parcels into image patches, saving memory and computation. The Temporal Attention Encoder (TAE) [15] and its parsimonious version Lightweight-TAE (LTAE) [13] are temporal sequence encoders based on the language processing literature [37] and adapted for processing SITS. Both networks can be used sequentially to map the sequence of observations x^i at year i to a learned yearly spatio-temporal

Table 1: **Crop distribution.** We indicate the number of parcels declarations in the LPIS for each class across all 103 602 parcels and all 3 years.

Class	Count	Class	Count
Meadow	184 489	Triticale	5114
Maize	42 006	Rye	569
Wheat	27 921	Rapeseed	7624
Barley Winter	10 516	Sunflower	1886
Vineyard	15 461	Soybean	6072
Sorghum	820	Alfalfa	2682
Oat Winter	529	Leguminous	1454
Mixed cereal	1061	Flo./fru./veg.	1079
Oat Summer	330	Potato	230
Barley Summer	538	Wood pasture	425

descriptor e^i :

$$e^i = \text{TAE} \left([\text{PSE}(x_t^i)]_{t=1}^{T_i} \right) . \quad (1)$$

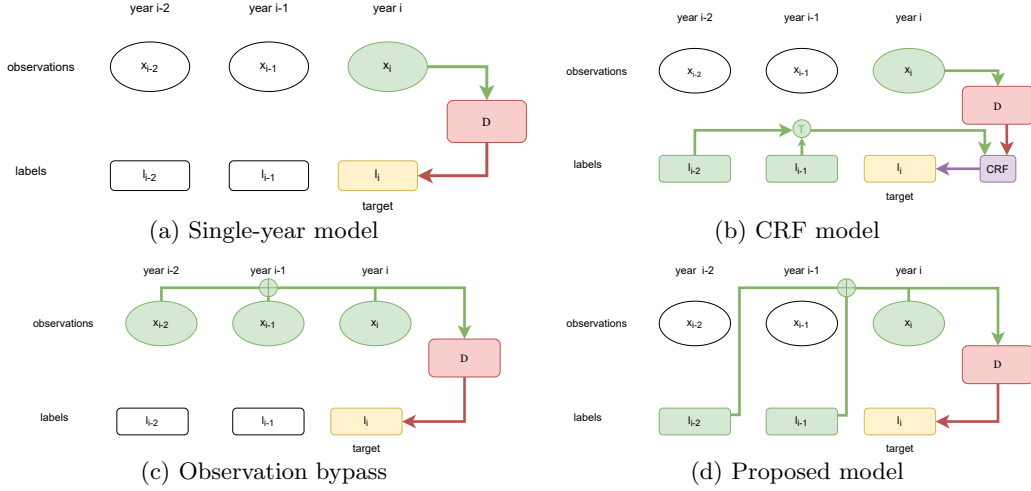


Figure 4: **Multi-Year Modeling.** Different approaches to model crop rotation dynamics: (a) the model only has access to the current year's observation; (b) a chain-CRF is used to model the influence of past cultivated crop; (c) the model has access to the observation of the past two years; (d) proposed approach: the model has access to the last two declared crops.

2.3 Multi-Year Modeling

We now present a simple modification of the PSE+LTAE network to model crop rotation. In the original PSE+LTAE approach, the descriptor e^i is directly mapped to a vector of class

scores z^i by a Multi-Layer Perceptron (MLP). In order to make the prediction z^i covariant with past cultivated crops, we augment the spatio-temporal descriptors e^i by concatenating the sum of the one-hot-encoded labels l^j for the previous two years $j = i - 1, i - 2$. Then, a classifier network \mathcal{D} , typically an MLP, maps this feature to a vector z^i of L class scores:

$$z^i = \mathcal{D}([e^i \parallel l_{i-1} + l_{i-2}]) , \quad (2)$$

with $[\cdot \parallel \cdot]$ the channelwise concatenation operator. We handle the edge effects of the first two available years by defining l^0 and l^{-1} as vectors of zero of size L (temporal zero-padding). This model can be trained end-to-end to simultaneously learn inter-annual crop rotations along with the intra-annual evolution of the parcels' spectral statistics.. Our model makes three simplifying assumptions:

- We only consider the last two previous years because of the limited span of our available data. However, it would be straightforward to extend our approach to a longer duration.
- We consider that the history of a parcel is entirely described by its past cultivated crop types, and we do not take the past satellite observations into account. In other words, the label at year i is independent from past observations conditionally to its past labels [39, Chap 2]. This design choice allows the model to stay tractable in terms of memory requirements.
- The labels of the past two years are summed and not concatenated. The order in which the crops were cultivated is then lost, but results in a more compact model.

2.4 Baseline Models

In order to meaningfully evaluate the performance of our proposed approach, we implement different baselines for classifying parcels from multi-year data . In Figure 4, we represent schematically the main idea behind these baselines and our proposed approach. Note that the choice of backbone network to handle single-year data is out of the scope of this paper.

Single-Year: M_{single} . We do not provide the previous years' labels, and directly map the current year's observations to a vector of class scores [13].

Conditional Random Fields: M_{CRF} . Based on the work of [4] and [16], we implement a simple chain-CRF probabilistic model. We use the prediction of the previous PSE+LTAE, calibrated with the method of Guo *et al.* [17] to approximate the posterior probability $p \in [0, 1]^L$ of a parcel having the label k for year i : $p_k = P(l^i = k \mid x^i)$ (see Section 3.3 for more details). We then model the second order transition probability $p(l^i = k \mid l^{i-1}, l^{i-2})$ with a three-dimensional tensor $T \in [0, 1]^{L \times L \times L}$ which can be approximated based on the observed transitions in the training set. As suggested by Bailly *et al.* , we use a Laplace regularization [34, Chap. 13] to increase robustness. The resulting probability for a given year i is given by:

$$z_{\text{CRF}}^i[k] = p \odot T[l^{i-2}, l^{i-1}, :] , \quad (3)$$

with \odot the Hadamard term-wise multiplication. This method is restricted to $i > 2$ as edge effects are not straightforwardly fixed with padding.

Observation Bypass: M_{obs} . Instead of concatenating the labels of previous years to the embedding e^i , we concatenate the average of the descriptors of the last two years e^{i-1} for e^{i-2} :

$$z_{\text{obs}}^i = \mathcal{D}_{\text{obs}} \left(\left[e^i \parallel \begin{cases} \frac{1}{2}[e^{i-1} + e^{i-2}] & \text{if } i > 1 \\ e^0 & \text{if } i = 1 \\ 0 & \text{if } i = 0 \end{cases} \right] \right). \quad (4)$$

Edge effects are handled with mirror and zero temporal padding.

Label Concatenation: $M_{\text{dec-concat}}$. Instead of concatenating the sum of the last two previous years, we propose to concatenate each one-hot-encoded vector l^{i-1} and l^{i-2} with the learned descriptor z^i . This approach is similar to Equation 2, but leads to a larger descriptor and a higher parameter count.

Single-Year Label Bypass: $M_{\text{dec-one-year}}$. In order to evaluate the impact of describing the history of parcels as the past two cultivated crops, we only concatenate the label of the previous year to the learned descriptor e^i .

2.5 Training Protocol

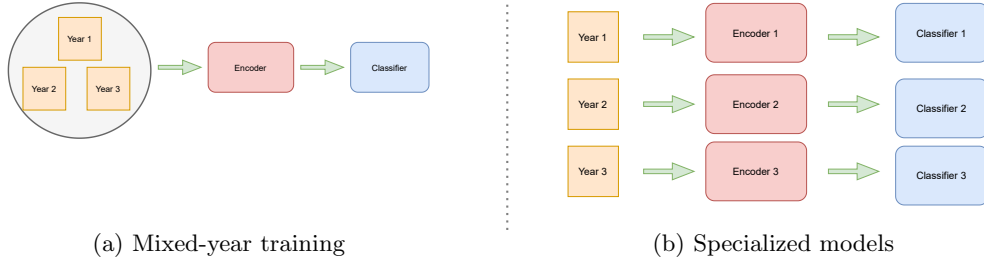


Figure 5: **Training Protocol.** A single model is trained with parcels taken from all three years (a), and three specialized models whose training set only comprises observation for a given year (b).

We propose a simple training protocol to leverage the availability of observations and farmers’ declarations from multiple years.

Mixed-year Training: We train a single model with parcels from all available years. Our rationale is that exposing the model to data from several years will contribute to learning richer and more resilient descriptors. Indeed, each year has different meteorological conditions influencing the growth profiles of crops. Moreover, by increasing the size of the dataset, mixed-year training mitigates the negative impact of rare classes on the performance.

We assess the impact of mixed-year training by considering $I = 3$ *specialized* models whose training set is restricted to a given year: M_{2018} , M_{2019} , and M_{2020} . In contrast, the model M_{mixed} is trained with all parcels across all years with no information regarding of the year of acquisition. All models share the same PSE+LTAE configuration [13]. We visualize the training protocols in Figure 5, and report the results in Table 2.

Cross-validation: We split our data into 5 folds for cross-validation. For each fold, we train on 3 folds and use the last fold for calibration and model selection, corresponding to a train/validation/test ratio of 60%, 20% and 20% in each fold. In order to avoid data contamination and self-correlation, our folds are all spatially separated: the fold separation is done parcel-wise and not for yearly observations. A parcel cannot appear in multiple folds for different years.

2.6 Evaluation Metrics

In order to assess the performance of the different approaches evaluated, we report the Overall Accuracy (OA), corresponding to the rate of correct prediction. If we denote by N_c the number of accurate predictions and N the total number of parcels, the overall accuracy writes:

$$\text{OA} = \frac{N_c}{N} . \quad (5)$$

To address the high class imbalance, we also report the mean Intersection over Union (mIoU), defined as the unweighted class-wise average of the intersection over Union (or Jaccard distance) between the prediction and the ground truth for each class. For a given class i , IoU_i is defined as the ratio between the number of elements that are both predicted and labeled by class i (the intersection, or true positives, and the number of elements that are either predicted or labeled as belonging to class i (the union). In terms of binary classification (class i *vs.* not class i), this translates into the following formula:

$$\text{IoU}_i = \frac{\text{TP}_i}{\text{TP}_i + \text{FP}_i + \text{FN}_i} , \quad (6)$$

with TP_i , FP_i and FN_i the number of true positives, false positives, and false negatives respectively. The mIoU represents the average of the IoU calculated over the K studied classes:

$$\text{mIoU} = \frac{1}{K} \sum_{i=1}^K \text{IoU}_i . \quad (7)$$

3 Results

This section presents the quantitative and qualitative impact of our design choice in terms of training protocol and architecture.

3.1 Training Protocol

Predictably, the specialized models have good performance when evaluated on a test set composed of parcels from the year they were trained, and poor results for other years, making this training procedure ill-fitted for the application at hand. In contrast, the model M_{mixed} vastly outperformed specialized models on average over the three considered years: over 15 points of mIoU. More surprisingly, the model M_{mixed} also outperforms all specialized models even when evaluated on the year of their training set. This implies that the increased diversity of the mixed-year training set allows the model to learn more robust and expressive representations.

In Figure 6, we illustrate the representations learned by the mixed model M_{mixed} and the specialized model M_{2020} . We remark that the parcel embeddings of the specialized model are inconsistent from one year to another, resulting in a higher overlap between classes. In contrast, the mixed year model M_{mixed} learns year-consistent representations. This results in embedding clusters with large margins between classes, illustrating the ability of the model to learn robust and discriminative SITS embeddings.

Table 2: **Quantitative evaluation.** Performance (mIoU and OA) of the different specialized models M_{2018} , M_{2019} , M_{2020} and of the mixed-years model M_{mixed} evaluated on each year individually and all available years simultaneously with 5-fold cross-validation. The best performances are shown in bold. Boxed values correspond to evaluations where the training and evaluation sets are drawn from the same year. The mixed-year model performs better for all years, even compared to specialized models.

Model	In the rest of the paper, we use mixed year training for all models.							
	2018		2019		2020		3 years	
	OA	mIoU	OA	mIoU	OA	mIoU	OA	mIoU
M_{2018}	97.0	64.7	90.3	45.5	90.8	43.4	92.7	49.1
M_{2019}	88.9	39.5	97.2	70.1	88.7	40.1	91.6	48.0
M_{2020}	91.4	44.2	93.7	51.8	96.7	67.3	93.9	54.0
M_{mixed}	97.3	69.2	97.4	72.2	96.8	68.7	97.2	70.4

3.2 Influence of Crop Rotation Modeling

We evaluate all models presented in Section 2.3 and Section 2.4, and provide qualitative illustration in Figure 7. All models are trained with the mixed-year training protocol and only tested on parcels from 2020 to avoid edge effects affecting the evaluation. We give quantitative cross-validated results in Table 3. Training our model on one fold takes 4 hours, and inference on all parcels takes under 3 minutes (over 500-parcels per second).

Table 3: **Performance by model.** Performances (mIoU and OA) of the models M_{single} , M_{obs} , M_{CRF} , and M_{dec} tested for the year 2020. Our proposed model M_{dec} achieve higher performance than M_{single} with a 6.3% mIoU gap.

Model	Description	OA	mIoU
M_{single}	single-year observation	96.8	68.7
M_{obs}	bypassing 2 years of observation	96.8	69.3
M_{CRF}	using past 2 declarations in a CRF	96.8	72.3
$M_{\text{dec-one-year}}$	concatenating last declaration only	97.5	74.3
$M_{\text{dec-concat}}$	concatenating past 2 declarations	97.5	74.4
M_{dec}	proposed method	97.5	75.0

We observe that our model appreciably improved on the single-year model, with over 6 points gained in mIoU. The CRF models also increase the results to a lesser margin. We attribute this inferior performance to an oversmoothing phenomenon already pointed out by Bailly *et al.*: CRFs tend to resolve ambiguities with the most frequent transition



(a) M_{mixed}



(b) M_{2020}

Figure 6: **Learned Representations.** Illustrations of the learned SITS representations of the mixed-year model M_{mixed} (a) and the specialized M_{2020} (b). T-SNE algorithm is used to plot in 2D the representation for 100 parcels over 10 classes and 3 years. We observe that M_{mixed} produced cluster of embeddings that are consistent from one year to another, and with clearer demarcation between classes.

regardless of the specificity of the observation. In contrast, our approach simultaneously models the current year’s observations and the influence of past cultivated crops. M_{obs} barely improves the quality of the single-year model. While this model has indeed access to more information than M_{single} , the same model is used to extract SITS descriptors for all three years. This means that the model’s ambiguities and errors will be the same for all three representations, which prevent M_{obs} from largely improving its prediction. Our approach injects new information into the model by concatenating the labels of previous years, which is independent of the model’s limitations. Our method is more susceptible to the propagation of mistakes in the farmers’ declarations but provides the most significant increase in performance in practice.

Lastly, we concatenate both past label vectors to keep information about the order in which past crops were cultivated and observe a slight decrease in performance. The increase in model size can explain this. We conclude that this order is not crucial information for our model conditionally to the observation of the target year. Lastly, the model’s performance with only the declaration of the last year performs almost as well as our model with two years worth of crop declarations. This suggests that yearly transition rules are sufficient to capture most inter-year dynamics, such as permanent culture. Alternatively, our two-year scheme may suffer from sharp edge effects with only three years’ worth of data. Only a quantitative analysis over a longer period may resolve this ambiguity. On average, our M_{dec} model obtains an mIoU of 84.7% and overall accuracy of 98.1% on the training set.

We report the confusion matrix of M_{dec} in Figure 8, and its performance for each crop in Table 4. We also compute $\Delta = \text{IoU}(M_{\text{dec}}) - \text{IoU}(M_{\text{single}})$ the gain compared to the single-year model $\text{IoU}(M_{\text{single}})$, as well as the ratio of improvement $\rho = \Delta / (1 - \text{mIoU}(M_{\text{single}}))$. This last number indicates the proportion of IoU that have been gained by modeling crop rotations. We observe that our model provides a large performance increase across all classes but four. The improvement is particularly stark for classes with strong temporal stability such as vineyards.




Figure 7: **Qualitative Illustration.** Detail of the area of interest with the ground truth in (a) and the qualification of the prediction in (b) with correct prediction in blue and errors in red.

To further this analysis, we arrange the crop types into three groups according to the crop grown in 2018 and the number of observed class successions over the 2018 – 2019 – 2020 period:

- **Permanent Culture.** Classes within this group are such that at least 90% of the




Figure 8: **Confusion Matrix.** Confusion matrix of the prediction of \mathcal{M}_{dec} for the year 2020. The area of each entry corresponds to the square root of the number of predictions.

observed successions are constant over three years. It Contains Meadow, Vineyard, and Wood Pasture.

- **Structured Culture.** A crop is said to be structured if, when grown in 2018, over 75% of the observed three-year successions fall into 10 different rotations or less and is not permanent. It contains Rapeseed, Sunflower, Soybean, Alfalfa, Leguminous, Flowers/Fruits/vegetables, and Potato.
- **Other.** All other classes.

We report the unweighted class average for these three groups in Table 5. Predictably, our approach considerably improves the results for permanent cultures. Our model is also able to learn non-trivial rotations as the improvement for structured classes is also noticeable. On average, our method also improves the performance for other nonstructured classes, albeit to a lesser degree. This indicates that our model can learn multi-year patterns not easily captured by simple rotation statistics.

3.3 Model Calibration

Crop mapping can be used for various downstream applications, such as environmental monitoring, subsidy allocation, and price prediction. These applications carry crucial economic

Table 4: **Performance by class.** IoU per class of our model M_{dec} for the year 2020, as well as the improvement Δ compared to the single-year model M_{single} , and the ratio of improvement ρ . All values are given in %, and we sort the classes according to decreasing ratios ρ .

Class	IoU	Δ	ρ	Class	IoU	Δ	ρ
Wood Pasture	92.4	+48.2	86.3	Oat Summer	52.8	+3.6	7.0
Vineyard	99.3	+1.4	68.7	Rapeseed	98.3	+0.1	6.6
Alfalfa	68.7	+23.9	49.9	Maize	95.7	+0.2	6.3
Flo./Fru./Veg.	83.4	+14.5	46.5	Wheat	91.9	+0.3	3.9
Meadow	98.4	+0.9	36.9	Barley Summer	64.3	+1.1	3.1
Leguminous	45.2	+14.6	21.1	Potato	57.1	+0.5	1.2
Rye	54.7	+6.4	12.4	Sunflower	92.2	-0.1	-0.3
Oat Winter	57.7	+4.5	9.7	Sorghum	56.6	-0.2	-0.4
Triticale	68.7	+2.6	7.8	Soybean	91.8	-0.2	-3.1
Mix. Cereals	31.0	+5.1	6.8	Barley Winter	92.8	-0.6	-8.5

Table 5: **Improvement Relative to Structure.** Classwise IoU and mean improvement of our model compared to the single-year model according to the rotation structure of the cultivated crops.

Category	mIoU	mean Δ
Permanent	97.3	16.9
Structured	77.7	7.6
Other	66.6	2.3

and ecological stakes and hence benefit from properly *calibrated* prediction. A prediction is said to be calibrated when the confidence (*ie*, the probability associated to a given class) of the prediction corresponds to the empirical rate of correct prediction: we want 90% of the prediction with a 90% confidence to be accurate. This allows for more precise risk estimation and improves control on the rate of false positives / negatives.

Deep learning methods such as ours are notoriously badly calibrated. However, this can be corrected with the simple technique proposed by Guo *et al.* [17]. This method, called temperature scaling, consists in minimizing the discrepancy between the predicted confidence and the rate of errors binned into quantiles (we chose here 15 bins) on the validation set by adjusting the temperature parameters in the last softmax layer [6, Chap. 2.4]. As represented in Figure 9, we can improve the calibration and observe a 43% decrease of the Expected Calibration Error (ECE) at a small computation cost.

4 Discussion

In this paper, we set out to develop a deep learning method to leverage both the inter- and intra-annual dynamics of crop growth for crop mapping. We propose to enrich the learned spatio-temporal features with the last two declared cultures. Our experiments show that this simple method leads to an appreciable increase in performance compared to models operating on data drawn from a single year. Our method outperforms other approaches




Figure 9: **Model calibration.** Empirical rate of correct prediction by predicted confidence. We quantize the predicted confidence into 100 bins for visualization purposes. For a perfectly calibrated prediction, the blue histogram would exactly follow the orange line. We observe that a simple post-processing step can considerably improve calibration.

such as CRF smoothing or observation stacking. This improvement can be observed for most crop types, including those with rotation patterns beyond permanent culture. We now discuss the limitations of our method and of our analysis.

Choice of Backbone Network. Our method can be adapted to any network with a distinct classifier module mapping a spatio-temporal learned feature vector to a predicted vector of class scores. However, the choice of spatio-temporal encoder (*backbone*) is out of the scope of this article. While it may be relevant to explore the effect of our modification on other architectures, we limited our investigation to the PSE+LTAE as it is the current state-of-the-art network for crop type mapping by a large margin.

Operational Setting. We showed that training our model with samples from all available years leads to considerably improved results. However, this scenario is not compatible with the operational setting of crop monitoring, in which payment agencies may want to detect erroneous declarations before all farmers’ declarations have been received. Instead, we use the same setting as the vast majority of work in parcel classification and whose task is to classify parcels after the year is over [42, 38, 35, 5, 22, 31, 27, 12, 30, 41, 21, 13, 15, 32, 14].

As the Sentinel-2 mission was only operational starting in 2017, we only have access to full-year coverage since 2018. This means that we only have three years’ worth of data at the time of writing this paper. In our opinion, this prevents us from a realistic setting in which the last year is withheld from the training set. Indeed, the inter-year meteorological variations between two years are typically too great to test for a third year and reasonably expect good results, as corroborated with preliminary experiments not shown in this paper. As more Sentinel-2 data become available, we will be able to evaluate our approach in a more realistic setting.

Scope of the Study. Given the large amount of data involved and the complexity of data collection, we have limited our analysis and our proposed open-access dataset to a single area of the French Metropolitan territory. While nothing in our method is specific to this area, some of our analysis may be biased by the preponderance of stable cultures such

as vineyards in this area. To confirm the generality of our conclusions, we would require a dataset with parcels taken from regions across the world with various meteorological conditions and agricultural practices. This task is complicated by the lack of harmonization between LPIS regarding open-access policy and even nomenclature. We hope that our results will encourage mapping agencies worldwide to release multi-year LPIS in open-source to help constitute a truly global dataset, allowing the community to assess the spatial generalizability of state-of-the-art methods. We also limit ourselves to predicting the main culture in each parcel while ignoring cases with multiple growth cycles within one year. This may be particularly detrimental to its application in subtropical regions.

Applicability of our Model. By requiring the last two grown crops to classify a parcel, our method cannot be applied to areas where the LPIS is not easily accessible. Furthermore, our training setting requires only selecting stable parcels. This can be easily obtained from the LPIS if it also contains information about the extent and position of each parcel, as is the case for the French LPIS. As a consequence, the effect of parcels with changing contours is out of the scope of our investigation.

5 Conclusions

We explored the impact of using multi-year data to improve the quality of the automatic classification of parcels from satellite image time series. We showed that training a deep learning model from multi-year observations improved its ability to generalize and better precision across the board. We proposed a simple modification to a state-of-the-art network to model both inter- and intra-year dynamics. This resulted in an increase of +6.3% of mIoU when compared to models operating on single-year data. The effect is most substantial for classes with strong temporal structures, but also impacts other crop types. We also showed how simple post-processing could improve the calibration of the models considered.

Finally, we release both our code and our data. We hope that our promising results will encourage the SITS community to develop methods modeling multiple time scales simultaneously and release more datasets spanning several years.

References

- [1] Joachim Aurbacher and Stephan Dabbert. “Generating crop sequences in land-use models using maximum entropy and Markov chains”. In: *Agricultural Systems* (2011).
- [2] Johann Bachinger and Peter Zander. “ROTOR, a tool for generating and evaluating crop rotations for organic farming systems”. In: *European Journal of Agronomy* (2007).
- [3] Nicolas Baghdadi, Marc Leroy, Pierre Maurel, Selma Cherchali, Magali Stoll, Jean-François Faure, Jean-Christophe Desconnets, Olivier Hagolle, Jérôme Gasperi, and Philippe Pacholczyk. “The Theia land data centre”. In: *Remote Sensing Data Infrastructures (RSDI) International Worksho.* 2015.
- [4] Simon Bailly, Sebastien Giordano, Loic Landrieu, and Nesrine Chehata. “Crop-rotation structured classification using multi-source Sentinel images and lpis for crop type mapping”. In: *IGARSS*. 2018.

- [5] Mariana Belgiu and Ovidiu Csillik. “Sentinel-2 cropland mapping using pixel-based and object-based time-weighted dynamic time warping analysis”. In: *Remote Sensing of Environment* (2018).
- [6] Christopher M Bishop. “Pattern recognition”. In: *Machine learning* (2006).
- [7] Gerhard Brankatschk and Matthias Finkbeiner. “Modeling crop rotation in agricultural LCAs—challenges and potential solutions”. In: *Agricultural Systems* (2015).
- [8] Nina K Detlefsen and Allan Leck Jensen. “Modelling optimal crop sequences using network flows”. In: *Agricultural Systems* (2007).
- [9] S Dogliotti, WAH Rossing, and MK Van Ittersum. “ROTAT, a tool for systematically generating crop rotations”. In: *European Journal of Agronomy* (2003).
- [10] Matthias Drusch, Umberto Del Bello, Sébastien Carlier, Olivier Colin, Veronica Fernandez, Ferran Gascon, Bianca Hoersch, Claudia Isola, Paolo Laberinti, Philippe Martimort, et al. “Sentinel-2: ESA’s optical high-resolution mission for GMES operational services”. In: *Remote sensing of Environment* (2012).
- [11] Jérôme Dury, Noémie Schaller, Frédérick Garcia, Arnaud Reynaud, and Jacques Eric Bergez. “Models to support cropping plan and crop rotation decisions. A review”. In: *Agronomy for sustainable development* (2012).
- [12] V Sainte Fare Garnot, Loic Landrieu, Sebastien Giordano, and Nesrine Chehata. “Time-Space tradeoff in deep learning models for crop classification on satellite multi-spectral image time series”. In: *IGARSS*. 2019.
- [13] Vivien Sainte Fare Garnot and Loic Landrieu. “Lightweight Temporal Self-Attention for Classifying Satellite Images Time Series”. In: *International Workshop on Advanced Analytics and Learning on Temporal Data*. Springer. 2020.
- [14] Vivien Sainte Fare Garnot and Loic Landrieu. “Panoptic Segmentation of Satellite Image Time Series with Convolutional Temporal Attention Networks”. In: *ICCV* (2021).
- [15] Vivien Sainte Fare Garnot, Loic Landrieu, Sebastien Giordano, and Nesrine Chehata. “Satellite image time series classification with pixel-set encoders and temporal self-attention”. In: *CVPR*. 2020.
- [16] Sébastien Giordano, Simon Bailly, Loic Landrieu, and Nesrine Chehata. “Improved Crop Classification with Rotation Knowledge using Sentinel-1 and-2 Time Series”. In: *Photogrammetric Engineering & Remote Sensing* (2018).
- [17] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. “On calibration of modern neural networks”. In: *ICML*. 2017.
- [18] Yekaterina Voskoboynik Kellen Myers Victoria Ferguson. *Modeling Crop Rotation with Discrete Mathematics*. Tech. rep.
- [19] B Koetz, P Defourny, S Bontemps, K Bajec, C Cara, L de Vendictis, L Kucera, P Malcorps, G Milcinski, L Nicola, et al. “SEN4CAP Sentinels for CAP monitoring approach”. In: *JRC IACS Workshop*. 2019.
- [20] Chris Kollas, Kurt Christian Kersebaum, Claas Nendel, Kiril Manevski, Christoph Müller, Taru Palosuo, Cecilia M Armas-Herrera, Nicolas Beaudoin, Marco Bindi, Monia Charfeddine, et al. “Crop rotation modelling—A European model intercomparison”. In: *European Journal of Agronomy* (2015).

- [21] Lukas Kondmann, Aysim Toker, Marc Rußwurm, Andrés Camero Unzueta, Devis Peressuti, Grega Milcinski, Pierre-Philippe Mathieu, Nicolas Longép , Timothy Davis, Giovanni Marchisio, et al. “DENETHOR: The DynamicEarthNET dataset for Harmonized, inter-Operable, analysis-Ready, daily crop monitoring from space”. In: (2021). URL: openreview.net/forum?id=uUa4jNMLjrL.
- [22] Nataliia Kussul, Mykola Lavreniuk, Sergii Skakun, and Andrii Shelestov. “Deep learning classification of land cover and crop types using remote sensing data”. In: *Geoscience and Remote Sensing Letters* (2017).
- [23] Florent Levavasseur, Philippe Martin, Clemence Bouty, Aude Barbottin, Vincent Bretagnolle, Olivier Th ron, Olivier Scheurer, and N Piskiewicz. “RPG Explorer: A new tool to ease the analysis of agricultural landscape dynamics with the Land Parcel Identification System”. In: *Computers and Electronics in Agriculture* (2016).
- [24] P Loudjani, W Devos, B Baruth, and G Lemoine. “Artificial Intelligence and EU Agriculture”. In: (2020).
- [25] *multi-year dataset*. <https://zenodo.org/record/5535882>. DOI: 10.5281/zenodo.5535882.
- [26] Julien Osman, Jordi Inglada, and Jean-Fran ois Dejoux. “Assessment of a Markov logic model of crop rotations for early crop mapping”. In: *Computers and Electronics in Agriculture* (2015).
- [27] Charlotte Pelletier, Geoffrey I Webb, and Fran ois Petitjean. “Deep learning for the classification of Sentinel-2 image time series”. In: *IGARSS*. 2019.
- [28] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. “Pointnet: Deep learning on point sets for 3d classification and segmentation”. In: *CVPR*. 2017, pp. 652–660.
- [29] *Registre parcellaire graphique (RPG) : contours des parcelles et  lots culturaux et leur groupe de cultures majoritaire*. <https://www.data.gouv.fr/en/datasets/registre-parcellaire-graphique-rpg-contours-des-parcelles-et-ilots-culturaux-et-leur-groupe-de-cultures-majoritaire/>. Accessed: 2021-09-24.
- [30] Marc Ru wurm and Marco K rner. “Multi-temporal land cover classification with sequential recurrent encoders”. In: *ISPRS International Journal of Geo-Information* (2018).
- [31] Marc Ru wurm and Marco K rner. “Self-attention for raw optical satellite time series classification”. In: *ISPRS Journal of Photogrammetry and Remote Sensing* (2020).
- [32] Maja Schneider and Marco K rner. “[Re] Satellite Image Time Series Classification with Pixel-Set Encoders and Temporal Self-Attention”. In: *ML Reproducibility Challenge 2020*. 2020.
- [33] Martin Sch nhart, Erwin Schmid, and Uwe A Schneider. “CropRota–A crop rotation model to support integrated land use assessments”. In: *European Journal of Agronomy* (2011).
- [34] Hinrich Sch tze, Christopher D Manning, and Prabhakar Raghavan. *Introduction to information retrieval*. Vol. 39. Cambridge University Press Cambridge, 2008.
- [35] Sofia Siachalou, Giorgos Mallinis, and Maria Tsakiri-Strati. “A hidden Markov models approach for crop classification: Linking crop phenology to time series of multi-sensor remote sensing data”. In: *Remote Sensing* (2015).

- [36] *The common agricultural policy at a glance*. https://ec.europa.eu/info/food-farming-fisheries/key-policies/common-agricultural-policy/cap-glance_en. Accessed: 2021-09-24.
- [37] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. “Attention is all you need”. In: *NeurIPS* (2017).
- [38] Francesco Vuolo, Martin Neuwirth, Markus Immitzer, Clement Atzberger, and Wai-Tim Ng. “How much does multi-temporal Sentinel-2 data improve crop type classification?” In: *International Journal of Applied Earth Observation and Geoinformation* (2018).
- [39] Martin J Wainwright and Michael Irwin Jordan. *Graphical models, exponential families, and variational inference*. Now Publishers Inc, 2008.
- [40] Raghu Yaramasu, Varaprasad Bandaru, and Koutilya Pnvr. “Pre-season crop type mapping using deep neural networks”. In: *Computers and Electronics in Agriculture* (2020).
- [41] Yuan Yuan and Lei Lin. “Self-Supervised Pretraining of Transformers for Satellite Image Time Series Classification”. In: *Journal of Selected Topics in Applied Earth Observations and Remote Sensing* (2020).
- [42] Baojuan Zheng, Soe W Myint, Prasad S Thenkabail, and Rimjhim M Aggarwal. “A support vector machine to identify irrigated crop types using time-series Landsat NDVI data”. In: *International Journal of Applied Earth Observation and Geoinformation* (2015).