

Early Classification for Agricultural Monitoring from Satellite Time Series

Marc Rußwurm¹ Romain Tavenard^{2,3} Sébastien Lefèvre⁴ Marco Körner¹

Abstract

In this work, we introduce a recently developed early classification mechanism to satellite-based agricultural monitoring. It augments existing classification models by an additional stopping probability based on the previously seen information. This mechanism is end-to-end trainable and derives its stopping decision solely from the observed satellite data. We show results on field parcels in central Europe where sufficient ground truth data is available for an empiric evaluation of the results with local phenological information obtained from authorities. We observe that the recurrent neural network outfitted with this early classification mechanism was able to distinguish the many of the crop types before the end of the vegetative period. Further, we associated these stopping times with evaluated ground truth information and saw that the times of classification were related to characteristic events of the observed plants' phenology.

1. Introduction

The identification of crop types from space-born imagery forms an important component of agricultural monitoring at global scale. Information obtained from satellite data is often the only source of information to assess the state of national agriculture. This is particularly important in developing countries where agricultural production is rarely monitored or controlled by central authorities. Crucial information on the effect of droughts and shortages are typically available after the end of the vegetative periods which may be too late for effective counter measures. Obtaining knowledge about the expected crop type as early as possible with space-born crop monitoring can play a significant role in providing food security, preventing famine, and determining

¹Technical University of Munich ²Littoral - Environnement - Télédétection - Géomatique-Rennes ³Institut de Recherche en Informatique et Systèmes Aléatoires ⁴Université de Bretagne Sud. Correspondence to: Marc Rußwurm <marc.russwurm@tum.de>.




Figure 1: An example of the predictions of the early classification mechanism $\hat{y}, p_t = f(\mathbf{X}_{\rightarrow t})$. A binary stopping decision t_{stop} is sampled from the evaluated stopping probability p_t . The upper illustration shows the *Sentinel 2* time series x based on which the class probabilities \hat{y} and the stopping time t_{stop} are produced.

policies for sustainable agriculture. Monitoring of vegetation has been in increasing focus of Earth observation programs. NASA's Landsat or ESA's Sentinel satellites observe the surface of the Earth at weekly intervals at spatial resolutions of 10-60m. This allows for regular observations of comparatively small objects, such as field parcels. While enormous quantities of data is collected, methods to extract information, *e.g.*, Torbick et al. (2018); Chen et al. (2018); Rußwurm & Körner (2018), are often applied at later times when the vegetative period has ended to ensure a certain accuracy of the prediction. With the proposed early classification mechanism, we enable a generic multi-temporal classification model to choose the time of classification t_{stop} sampled from a probability p_t . Our results show empirically that the classification stopped when sufficient information has been observed to justify a confident classification, as shown in an example in Fig. 1.

2. Related Work

Multi-temporal approaches commonly rely on the classification of the entire time series either using features extracted

from expert knowledge (Bailly et al., 2016) or data-driven end-to-end learning (Sharma et al., 2018; Interdonato et al., 2019; Rußwurm & Körner, 2018; Pelletier et al., 2019). Especially for seasonally variable classes, such as vegetation monitoring or crop type mapping, region-specific expert knowledge is used to restrain the period of the time series to the period where the observed classes are deemed observable, *i.e.*, the growing season (Torbick et al., 2018; Chen et al., 2018). Early classification is the task of predicting the class of an incoming time series $\mathbf{X} = (x_0, x_1, \dots, x_T)$ of observations x_t as early as possible. Common approaches formalize a mixed cost function that could take into account both earliness and accuracy, and optimizing this cost function or a surrogate of it (Dachraoui et al., 2015; Tavenard & Malinowski, 2016; Mori et al., 2017). Recently, an end-to-end fine-tuning framework has been proposed (Rußwurm et al., 2019) to jointly learn (i) an embedding for partial time series, (ii) a classifier and (iii) a stopping rule, where both the classifier and the stopping rule are computed from the embedding. We will use this formulation in this work for the usecase of agricultural monitoring with a novel loss function that allows jointly optimizing parameters for classification model and stopping decision in one training phase.

3. Methodology

Our model estimates classification scores \hat{y}_t and a probability of stopping p_t at each time t . It is composed of a feature extraction network whose output is $\mathbf{h}_t = f_{\text{feat}}(\mathbf{X}_{\rightarrow t})$ (Section 3.3) and two light-weight mappings for classification scores $\hat{y}_t = f_c(\mathbf{h}_t)$ and stopping probability $p_t = f_\delta(\mathbf{h}_t)$ (Section 3.1). We denote vectors and matrices by bold face and differentiate them by case. A dataset $\mathcal{D} = \{(\mathbf{X}^i, \mathbf{y}^i)\}_{i=1}^N$ with single example $\mathbf{X}^i \in \mathbb{R}^{\{T \times D\}}$ is composed of T D -dimensional observations where the subscript is used to indicate time t . Temporal sequences up to time t are denoted by an arrow, *e.g.*, $\mathbf{X}_{\rightarrow t}$, while prediction scores are represented as vector $\hat{\mathbf{y}} \in \mathbb{R}^M$ of M classes with the prediction score, *i.e.*, the probability output by the model. The probability of the correct class is denoted as \hat{y}^+ .

3.1. Early Classification Mechanism

First, we describe the temporal classification model that evaluates a classification score $\hat{y}_t = f_c(f_{\text{feat}}(\mathbf{X}_{\rightarrow t}))$ at each time t . This hidden feature vector $\mathbf{h}_t = f_{\text{feat}}(\mathbf{X}_{\rightarrow t})$ is typically obtained as a latent representation after a series of cascaded layers in deep classification models. Such models produce classification scores by applying a non-linear mapping $\hat{y}_t = f_c(\mathbf{h}_t) = \text{softmax}(\Theta_c \mathbf{h}_t + \mathbf{b}_c)$. Since this hidden state contains information of the predicted class it encodes an internal representation of the classification confidence. So, it can be potentially used as an input for the decision of stopping the classification if enough data has

been observed. A loss function $\mathcal{L}_c(\mathbf{X}_{\rightarrow t}, \mathbf{y})$ evaluates the quality of the prediction compared to ground truth labels \mathbf{y} . Gradients of this loss are then back-propagated to adjust the classification model parameters $\Theta_c \in \mathbb{R}^{H \times M}$, $\mathbf{b}_c \in \mathbb{R}^M$. The number of hidden states H is a hyperparameter of the model while M denotes the number of classes.

The early classification mechanism augments this classification model by an additional non-linear mapping $p_t = f_\delta(\mathbf{h}_t) = \sigma(\Theta_\delta \mathbf{h}_t + b_\delta)$ with sigmoidal activation function $\sigma(\cdot)$ that produces a one-dimensional probability of stopping p_t . The parameterization $\Theta_\delta \in \mathbb{R}^{H \times 1}$, b_δ of this mapping, however, can not be adjusted from gradients of the classification loss \mathcal{L}_c directly. Hence, a new loss function $\mathcal{L}_t(\mathbf{X}_{\rightarrow t}, \mathbf{y})$ is required to adjust both Θ_δ , b_δ and Θ_c , \mathbf{b}_c .

This earliness-aware loss function is then weighted by the probability of stopping

$$P(t; \delta_{\rightarrow t}) = p_t \cdot \prod_{\tau=0}^{t-1} (1 - p_\tau). \quad (1)$$

at corresponding time t . To ensure that this probability will sum to one, $\delta_{t=T}$ is set to 1 regardless of the last hidden state \mathbf{h}_T . This mechanism utilizes information about the full sequence at training time to restrain $P(t; \delta_{\rightarrow t})$ and calculate the training loss

$$\mathcal{L}(\mathbf{x}, \mathbf{y}) = \sum_{t=0}^T P(t; \delta_{\rightarrow t}) \mathcal{L}_t(\mathbf{X}_{\rightarrow t}, \mathbf{y}). \quad (2)$$

At inference time when $P(t; \delta_{\rightarrow t})$ can not be normalized to a probability, p_t is used to sample a stopping time t_{stop} . The prediction $\mathbf{y}_{t_{\text{stop}}}$ at this time is used to derive a classification label for the time series and to assess the accuracy. Note that we initialize the weights so that $p_t \approx 0$ to favor late observations in the beginning of the training period.

3.2. Earliness Reward Loss Function

Inspired by existing works (Mori et al., 2017; Tavenard & Malinowski, 2016; Rußwurm et al., 2019), we employ a loss

$$\mathcal{L}_t(\mathbf{X}_{\rightarrow t}, \mathbf{y}; \alpha) = \alpha \mathcal{L}_c(\mathbf{X}_{\rightarrow t}, \mathbf{y}) - (1 - \alpha) \mathcal{R}_e(t, \hat{y}_t^+) \quad (3)$$

α -weights a classification loss $\mathcal{L}_c(\mathbf{X}_{\rightarrow t}, \mathbf{y})$ with a *earliness reward* term

$$\mathcal{R}_e(t, \hat{y}_t^+) = \hat{y}_t^+ \left(1 - \frac{t}{T}\right). \quad (4)$$

This term linearly scales the potential reward for earlier classifications $(1 - \frac{t}{T})$. We remind that \hat{y}_t is the prediction of the network at time t and that \hat{y}_t^+ is the probability of the annotated class. This rewards the classifier only for an




Figure 2: Schematic illustration of the multi-layer long short-term memory classification model equipped with the linear mappings for the class probabilities \hat{y} and the stopping probability p_t .

early classification if the correct class has been predicted at reasonable accuracy. By combining Eqs. (2) to (4) we see that three terms influence the earliness reward. These are the probability $P(t; \delta_{\rightarrow t})$, the evaluated classification score of the correct class \hat{y}_t^+ and a linear scaling term $(1 - \frac{t}{T})$ that rewards earlier classifications.

3.3. The Early Classification LSTM Network

To test this early classification mechanism on a real-world remote sensing vegetation dataset, we implemented a light-weight multi-layer *long short-term memory* (LSTM) (Hochreiter & Schmidhuber, 1997) recurrent network, as this architecture has shown good results in similar use-cases in recent studies (Rußwurm & Körner, 2018; Interdonato et al., 2019; Sharma et al., 2018). Figure 2 shows a detailed representation of the employed gated recurrent neural network architecture with the previously described early classification mechanism. The input vector $x_t \in \mathbb{R}^d$ at time t is first z -normalized by mean and variance parameters using *layer normalization* (Ba et al., 2016), which are computed from the training set. The normalized representation is then provided to L cascaded LSTM layers $h_t^l = \text{Lstm}^l(h_t^{l-1}, h_{t-1}^l)$. We denote the l -th layer in the network as raised index. Each of these gated recurrent layers evaluate a feature representation h_t^l based on two hidden representations: one at the current time of the previous layer h_t^{l-1} and one of the same layer at the previous time step h_{t-1}^l . Long short-term memory cells additionally utilize an internal cell state c_t to store long-term information, as indicated by the circular connection in the illustration. The last layer-normalized hidden representation h_t then computes class prediction scores \hat{y} for each class and the stopping probability p_t , as described in Section 3.1.

4. Results and Experiments

To test the early classification mechanism, we obtained the label and geometry information of 40k field parcels from local authorities of the 2018 season in central Europe. Over this region we acquired *all* available *Sentinel 2* images from *Google Earth Engine* (Gorelick et al., 2017). All $D = 13$ spectral bands located within one field parcel were mean-aggregated to a feature vector x_t , as shown in Fig. 1. From the available satellite observations we randomly choose $T = 70$ times to introduce a variance to the data samples and prevent overfitting. Geographic parcels are known to be spatially autocorrelated due to their proximity to each other (Tobler, 1970). Hence, we avoid a field wise random partitioning to training, validation and evaluation partitions. Instead we group the area of interest first in spatially separate regions and then assign all fields within one region to the respective partition, as suggested by Jean et al. (2018); Rußwurm & Körner (2018).

We determine a set of optimal hyperparameters for the classification model by training without earliness mechanism on standard cross-entropy loss $\mathcal{L}(x, y) = \frac{1}{T} \sum_{t=0}^T -\log(\hat{y}^+)$ with equal weight over all observations. For the weight updates, we used the *Adam optimizer* (Kingma & Ba, 2014) and dropped-out connections between recurrent layers at 50% probability and a batch size of 1024. We searched over the number of recurrent layers $L \in \{1, 2, 3\}$, the size of the hidden dimension $H \in \{64, 128, 256, 512\}$, and the learning rate $\nu \in \{0.1, 0.01, 0.001\}$ for 30 epochs. The parameters $L = 4, H = 64, \nu = 0.01$ that achieved best accuracy on the validation were used throughout this work. When training the model from scratch, we favor initially late classifications with $p_t \approx 0$. To achieve this, we initialize the bias term $b_\delta \sim \mathcal{N}(\mu = -0.2, \sigma = 0.1)$. The training of one model took approximately ten minutes on a NVIDIA GeForce GTX 1060.

4.1. Qualitative Example

We illustrated the early classification results qualitatively in Fig. 1 on a classification example using raw Sentinel 2 top-of-atmosphere reflectance data gathered over the year 2018. In this multi-temporal classification task, we estimate classification scores \hat{y}_t for each evaluated class based on observed data $X_{\rightarrow t}$ at each time t . In the example we can see that the model initially predicts randomly from the January to March period. Then some classification-relevant features are observed, which initially mislead the classifier to false predictions in the April period. Finally, the correct class is assigned the highest prediction score from May onwards. Early classification augments this classification scheme by estimating a stopping probability p_t out of which a stopping time t_{stop} is sampled. This stopping decision serves as indi-

α	accuracy	\bar{t}_{stop}	precision	recall	f_1	κ
.0	.25 ± .22	.10 ± .17	.19 ± .20	.25 ± .17	.16 ± .20	.12 ± .19
.2	.81 ± .03	.40 ± .02	.70 ± .01	.74 ± .01	.71 ± .01	.71 ± .04
.4	.80 ± .09	.47 ± .03	.71 ± .02	.74 ± .01	.71 ± .02	.71 ± .10
.6	.85 ± .02	.88 ± .07	.73 ± .04	.74 ± .03	.73 ± .03	.77 ± .03
.8	.84 ± .01	.93 ± .05	.72 ± .02	.75 ± .01	.73 ± .02	.76 ± .02
1.0	.83 ± .03	1.00 ± .00	.72 ± .03	.75 ± .01	.72 ± .03	.75 ± .04

Table 1: Varying the weighting factor α for *early reward* loss formulation (Section 3.2).

cation if a sufficient amount of data has been observed and no further data is required to allow for a confident classification.

4.2. Quantitative Accuracy vs Earliness Analysis

In Table 1, we quantitatively evaluated the effect of the trade-off parameter α on the classification performance on a series of accuracy metrics. We trained models with different weight initializations and batching sequences three times to assess the stability of the training process under a range of trade-off parameters $\alpha \in \{0.2, 0.4, 0.6, 0.8\}$. For each α we report a series of accuracy metrics and the averaged time of classification $\bar{t}_{\text{stop}} = \frac{1}{N} \sum_{i=0}^N t_{\text{stop}}$ normalized as fraction of the total sequence length. The *overall accuracy* describes the parcel-wise ratio of correctly positive and correctly negative examples in relation to the total number of samples. Precision, recall and their harmonic mean f_1 -score are calculated on a class-wise basis which is less sensitive to the frequency of samples per classes compared to overall accuracy or the mean time of classification \bar{t}_{stop} . We also report the kappa metric κ (Fleiss et al., 1969), as it is popular in the remote sensing community and similarly accounts for imbalanced class distributions by normalizing the score with a probability of random prediction. The early reward loss function was in terms of achieved accuracy resilient to the choice of α . This becomes apparent as the classification accuracies remained stable throughout the evaluated ranges of α .

4.3. Class-wise Analysis of Stopping Times

Next, we examine the stopping times per class and evaluate if these are related to phenologically characteristic events of the crop classes. We show the distribution of stopping times of all field parcels per individual class as boxplots in Fig. 3. The distributions over stopping times vary for each of the respective classes. For instance the quartiles of the *corn* distribution is narrow which means that 50% of the observed parcels are identified within one to two weeks, while crop types, such as *barley* or *meadows* are spread over one month. One can observe that the stopping times generally correlate with the vegetative period (shaded in blue) and the harvest times (marked by triangle \blacktriangledown). We obtained these dates from local authorities to be able to assess the performance of the early classification mechanism objectively.




Figure 3: Plot of distributions of the stopping times grouped by individual classes. The stopping times for each class follows a different distribution. The prediction for some classes, e.g., *barley*, are on median earlier than the others, e.g., *corn* or *wheat*, that are stopped at a later. The evaluated stopping times are correlated with actual phenological information, such as the vegetative season (shaded in blue) or the date of harvest (\blacktriangledown) provided by local authorities.

Note, however, that no region-specific expert knowledge is required for evaluating the stopping probabilities. Hence, the shown correlation originates solely from the observed satellite data which can be obtained globally. Overall, we can observe that our early classification model decided to stop the classification before the end of the vegetative period for a majority of field parcels. These results demonstrate the potential of early classification for agricultural monitoring where a confident classification decision can be determined independently from local expert knowledge and often before the end of the vegetative season.

5. Discussion

In this work we introduced a novel end-to-end trainable early classification mechanism for satellite-based agricultural monitoring. It can augment any multi-temporal classification models that evaluate a hidden representation at each time. Further, we analyzed the stopping times for the individual vegetation classes. Here, we saw that the parametrization of the stopping mechanism for each class was learned at different times and that the final times of stopping followed distinct distributions from crop type to crop type. We compared these evaluated stopping times with phenological and ontological information from local authorities and saw that many field parcels could be classified before the end of the vegetative period which is the usual end time of classification for regular crop type mapping approaches.

References

- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bailly, A., Arvor, D., Chapel, L., and Tavenard, R. Classification of modis time series with dense bag-of-temporal-sift-words: Application to cropland mapping in the brazilian amazon. In *2016 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*, pp. 2300–2303. IEEE, 2016.
- Chen, Y., Lu, D., Moran, E., Batistella, M., Dutra, L. V., Sanches, I. D., da Silva, R. F. B., Huang, J., Luiz, A. J. B., and de Oliveira, M. A. F. Mapping croplands, cropping patterns, and crop types using modis time-series data. *International journal of applied earth observation and geoinformation*, 69:133–147, 2018.
- Dachraoui, A., Bondu, A., and Cornuéjols, A. Early classification of time series as a non myopic sequential decision making problem. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 433–447. Springer, 2015.
- Fleiss, J. L., Cohen, J., and Everitt, B. S. Large sample standard errors of kappa and weighted kappa. *Psychological Bulletin*, 72(5):323, 1969.
- Gorelick, N., Hancher, M., Dixon, M., Ilyushchenko, S., Thau, D., and Moore, R. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202: 18–27, 2017.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Interdonato, R., Ienco, D., Gaetano, R., and Ose, K. Duplo: A dual view point deep learning architecture for time series classification. *ISPRS Journal of Photogrammetry and Remote Sensing*, 149:91–104, 2019.
- Jean, N., Wang, S., Azzari, G., Lobell, D., and Ermon, S. Tile2vec: Unsupervised representation learning for remote sensing data. *arXiv preprint arXiv:1805.02855*, 2018.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Mori, U., Mendiburu, A., Dasgupta, S., and Lozano, J. A. Early classification of time series by simultaneously optimizing the accuracy and earliness. *IEEE Transactions on Neural Networks and Learning Systems*, 2017.
- Pelletier, C., Webb, G. I., and Petitjean, F. Temporal convolutional neural network for the classification of satellite image time series. *Remote Sensing*, 11(5), 2019.
- Rußwurm, M. and Körner, M. Multi-temporal land cover classification with sequential recurrent encoders. *ISPRS International Journal of Geo-Information*, 7(4):129, 2018.
- Rußwurm, M., Lefèvre, S., Courty, N., Emonet, R., Körner, M., and Tavenard, R. End-to-end learning for early classification of time series. *arXiv preprint arXiv:1901.10681*, 2019.
- Sharma, A., Liu, X., and Yang, X. Land cover classification from multi-temporal, multi-spectral remotely sensed imagery using patch-based recurrent neural networks. *Neural Networks*, 105: 346–355, 2018.
- Tavenard, R. and Malinowski, S. Cost-aware early classification of time series. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 632–647. Springer, 2016.
- Tobler, W. R. A computer movie simulating urban growth in the detroit region. *Economic geography*, 46:234–240, 1970.
- Torbick, N., Huang, X., Ziniti, B., Johnson, D., Masek, J., and Reba, M. Fusion of moderate resolution earth observations for operational crop type mapping. *Remote Sensing*, 10(7):1058, 2018.