



**HAL**  
open science

# Time-Space Tradeoff in Deep Learning Models for Crop Classification on Satellite Multi-Spectral Image Time Series

Vivien Sainte Fare Garnot, Loic Landrieu, Sébastien Giordano, N. Chehata

► **To cite this version:**

Vivien Sainte Fare Garnot, Loic Landrieu, Sébastien Giordano, N. Chehata. Time-Space Tradeoff in Deep Learning Models for Crop Classification on Satellite Multi-Spectral Image Time Series. IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium, Jul 2019, Yokohama, Japan. pp.6247-6250, 10.1109/IGARSS.2019.8900517 . hal-02386701

**HAL Id: hal-02386701**

**<https://hal.science/hal-02386701>**

Submitted on 29 Nov 2019

**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.

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/330753608>

# TIME-SPACE TRADEOFF IN DEEP LEARNING MODELS FOR CROP CLASSIFICATION ON SATELLITE MULTI-SPECTRAL IMAGE TIME SERIES

Preprint · January 2019

CITATIONS

0

READS

130

4 authors:



**Vivien Sainte Fare Garnot**

Institut national de l'information géographique et forestière

4 PUBLICATIONS 1 CITATION

SEE PROFILE



**Loic Landrieu**

Institut national de l'information géographique et forestière

22 PUBLICATIONS 234 CITATIONS

SEE PROFILE



**Sébastien Giordano**

Institut national de l'information géographique et forestière

34 PUBLICATIONS 77 CITATIONS

SEE PROFILE



**Nesrine Chehata**

Bordeaux INP

93 PUBLICATIONS 830 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



ANR HIATUS [View project](#)



Semantic segmentation of large 3D point clouds [View project](#)

# TIME-SPACE TRADEOFF IN DEEP LEARNING MODELS FOR CROP CLASSIFICATION ON SATELLITE MULTI-SPECTRAL IMAGE TIME SERIES

V. Sainte Fare Garnot<sup>1</sup>, L. Landrieu<sup>1</sup>, S. Giordano<sup>1</sup>, N. Chehata<sup>1,2</sup>

<sup>1</sup> Univ. Paris-Est, LASTIG-STRUDEL, IGN-ENSG, F-94160 Saint-Mandé, France

<sup>2</sup> EA G&E Bordeaux INP, Université Bordeaux Montaigne, France

## ABSTRACT

In this article, we investigate several structured deep learning models for crop type classification on multi-spectral time series. In particular, our aim is to assess the respective importance of spatial and temporal structures in such data. With this objective, we consider several designs of convolutional, recurrent, and hybrid neural networks, and assess their performance on a large dataset of freely available Sentinel-2 imagery. We find that the best-performing approaches are hybrid configurations for which most of the parameters (up to 90%) are allocated to modeling the temporal structure of the data. Our results thus constitute a set of guidelines for the design of bespoke deep learning models for crop type classification.

*Index Terms*— Deep learning, Crop type classification, Multi-temporal, Sentinel-2

## 1. INTRODUCTION

Agricultural crop type mapping provides a highly valuable tool for agricultural management, environmental monitoring, or policy making. With the development of publicly available medium-high resolution remote sensing data such as the multi-spectral image time series provided by the Sentinel-2 mission, crop classification can now be implemented on a large scale and thus calls for increasingly high-performing automated machine-learning based classification tools.

In order to achieve this, discriminative models such as Random Forest (RF) classifiers have been extensively used for their good performance and the easy interpretability of their classification decisions [1, 2, 3]. Yet, such algorithms fail to take advantage of the spatial and temporal structures of the data at hand. Consequently, several works attempt to explicitly model the temporal structures: modelling multi-annual crop rotation using Conditional Random Fields [3], or modelling phenological properties with Hidden Markov Models [4] for instance. More recently, advances in deep learning provided very efficient tools for complex data analysis. Specifically, Convolutional Neural Networks (CNN) [5] and Recurrent Neural Networks (RNN) [6] are proving very efficient at exploiting spatial and temporal structure respectively. Additionally, RNNs and CNNs can be combined in hybrid

recurrent convolutional networks (RCNN) in order to exploit both spatial and temporal structures in time series of images. These methods are being successfully applied for crop type mapping in a variety of manners: using CNN for unitemporal and multi-temporal classification [7], RNN to model crop phenology [8], or RCNN for multi-temporal crop segmentation [9].

In this article, we consider the problem of crop classification on multi-spectral time series, assuming that parcel segmentation is already known. We attempt to answer the following question: Given a fixed budget of parameters, should one prioritize modeling the spatial structure (with CNN), temporal structure (with RNN), or both (with RCNN)? We compare the crop classification performance of several architectures with the same number of parameters on a Sentinel-2 dataset of agricultural parcels.

The key contributions of our work are as follows:

- We provide an empirical analysis of the features extracted by convolutional and recurrent models.
- We show that recurrent architectures are acting as a memory combining multiple observations, as well as a model for temporal evolution.
- We provide empirical evidence that the temporal structure of Sentinel-2 data is richer than the spatial structure for crop type classification, and, consequently, that most of a model's complexity should be allocated to it.

## 2. METHODS

### 2.1. Neural Network Architectures

In order to assess the influence of the temporal and spatial structure, we implement four neural network architectures. All of them share the same scheme: first extract an embedding — spatial, temporal or both — from the input image sequence and then perform the classification based on this embedding. The same design of classifier unit is used across all architectures: a Multi Layer Perceptron (MLP) with two hidden layers of dimension 128 and 64.

### 2.1.1. CNN

We first implement a convolution-based neural network whose goal is to leverage the spatial structure of image time series. The images corresponding to each date are embedded independently through the same three layers composed of the following units: convolution with  $3 \times 3$  kernel size and no padding, batch normalization [10], ReLu activation [11] and Max-Pool with  $2 \times 2$  kernel size. We then compute a global embedding for the whole sequence by concatenating all the image embeddings and taking the maximum for each channel, in the manner of PointNet [12]. Finally, the global embedding is fed to the MLP for classification.

### 2.1.2. RNN

Unlike the CNN network, the RNN architecture focuses purely on the temporal dimension of image sequences. For each image, we compute a vector of parcel-level handcrafted features: the mean and standard deviation of each spectral band. These vectors are then fed in chronological order to a standard Gated Recurrent Unit (GRU) [13]. The last hidden state of the GRU is used as the embedding for classification with the MLP.

### 2.1.3. RCNN

We explore two implementations of hybrid recurrent convolutional network to simultaneously model the spatial and temporal structures.

#### 2.1.3.1. CNN+GRU

Our first implementation successively extracts spatial and temporal embeddings. Each image of the sequence is first independently embedded with a CNN-like network. The resulting sequence of spatial embeddings is then fed to a GRU unit in chronological order. The last hidden state of the GRU is used as a spatio-temporal embedding for classification.

We implement three such models with varying ratios of parameters allocated to the temporal structure. Indeed, we can reduce the number of convolutional kernels and chose a larger hidden state size to increase this ratio (see Table 1).

#### 2.1.3.2. Convolutional Long Short Term Memory(ConvLSTM)

Our second implementation follows the ConvLSTM architecture introduced by [14] and implements directly the spatial structure in the recurrent network. It uses image-shaped hidden and cell states, as well as convolutions instead of MLP layers in an LSTM structure [6]. We refer the reader to [9] for more details on this architecture.

**Table 1:** Summary of the models’ parameters.

Model	Number of Kernels	Hidden State Size	Temporal Parameter Ratio	Total Number of Parameters
CNN	16 : 32 : 96	-	0	92 899
CNN+GRU <sub>7</sub>	16 : 32 : 64	64	0.7	92 035
CNN+GRU <sub>8</sub>	16 : 32 : 36	96	0.8	93 807
CNN+GRU <sub>9</sub>	16 : 16 : 16	128	0.9	90 179
GRU	-	156	1	94 587
ConvLSTM	30 : 64	64	-	95 353

Table 1 summarizes the parameters used for the models: the number of kernels for each of the convolutions and the size of the hidden state of the recurrent unit. We do not consider three dimensional convolutional architectures to cover the spatial and temporal dimension of the data, contrarily to previous studies such as [7]. Indeed, convolutions are local computations, and hence are not as well-suited for modeling long term dependencies as recurrent architectures. Additionally, applying convolutions along the temporal axis assumes that the images of the sequence are regularly sampled in time which is not necessarily the case in practice.

## 2.2. Implementation details

All the models are trained for 50 epochs with Adam optimizer [15] set with a batch size of 32 and a learning rate of  $10^{-3}$ . We use 5-fold cross-validation: for each fold the dataset is split into train, validation and test set with a 3:1:1 ratio. The epoch achieving the best results on the validation set is used for performance evaluation on the test set.



(a) Fragment of tile T31TFM (b) Example of an image patch for one observation of a parcel.

**Fig. 1:** Example of input image and dataset sample.

## 3. EXPERIMENTAL RESULTS

### 3.1. Dataset

We test our models using Sentinel-2 multi-spectral image sequences in top-of-canopy (TOC) reflectance (see Figure 1(a)). We leave out the atmospheric bands (bands 1, 9, and 10), keeping 10 spectral bands. The area of interest (AOI) corresponds to tile T31TFM in southern France. This tile provides a challenging use case as it presents a high diversity of crops and different terrain conditions. It covers a surface of 12 100 km<sup>2</sup> and contains 199 464 parcels, each observed at 24 dates from January to October 2017.

The values of cloudy pixels are linearly interpolated from the first previous and next available pixel using Orfeo Toolbox [16]. To each parcel we associate a rectangular image patch obtained from the bounding box of its geo-referenced polygon. We set to zero all the pixels outside the parcel as shown on Figure 1(b). The patches are resized to  $32 \times 32$  pixels which implies downsampling for only a minority of parcels. The dataset is thus comprised of 199 464 tensors of size  $24 \times 10 \times 32 \times 32$ . The images are normalized channel-wise for each date individually.

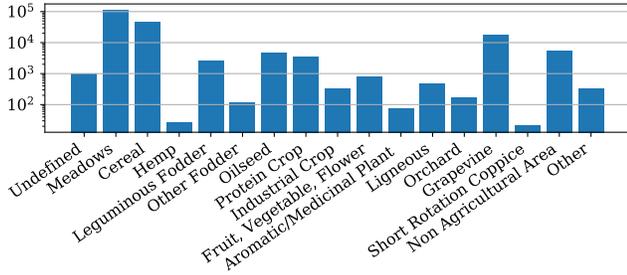


Fig. 2: Class breakdown in the AOI.

We use the French Land Parcel Identification System records to associate a label to each parcel. The labels are drawn from a comprehensive terminology that covers all observable parcel types and regroups them in 18 high-level classes. Figure 2 shows the breakdown of the different classes in the AOI, where meadows, cereals and grapevine are dominant and only the rice class is not represented.

### 3.2. Results

We present the performance of each model in Table 2. Given the high imbalance of the dataset under consideration (see Figure 2), we chose the unweighted average class-wise F-score rather than Overall Accuracy (OA) to rank the models’ performances. We insist on the fact that all models are designed with approximately the same number of parameters (see Table 1), so that the differences in performance can only be attributed to the architecture itself, and not to differences in model complexity.

Table 2: Performance metrics of the different models.

Model	OA	F-score	Precision	Recall
CNN	89.5	34.9	53.0	31.4
CNN+GRU <sub>7</sub>	93.4	53.2	59.7	49.9
CNN+GRU <sub>8</sub>	93.7	55.1	63.3	51.6
CNN+GRU <sub>9</sub>	<b>93.7</b>	<b>55.1</b>	<b>65.1</b>	<b>51.8</b>
GRU	93.5	53.8	63.0	50.0
ConvLSTM	93.1	49.2	64.0	44.5

Surprisingly, the purely recurrent GRU model outperforms the ConvLSTM and the CNN+GRU<sub>7</sub> hybrid models. Only the CNN+GRU<sub>8</sub> and CNN+GRU<sub>9</sub> models achieve higher overall performance. This shows that extracting both spatial and temporal structures allows for a better classification, only on condition that most of the parameters are allocated to the temporal structure. Only 10% of the parameters seem sufficient for spatial feature extraction with convolutions.

This suggests that the features extracted by RNNs are more discriminative than those extracted by CNNs. Indeed, the purely convolutional model performs significantly worse than its purely recurrent counterpart (by 19 points of F-score).

This performance gap can be explained by the fact that convolutional features are not completely relevant for the prob-

Table 3: Per-class F-score on the test set.

	CNN	GRU	CNN+GRU <sub>9</sub>	ConvLSTM
Undefined	0.0	<b>45.1</b>	35.4	38.2
Meadows	92.9	95.8	<b>96.0</b>	96.0
Cereal	93.3	97.1	<b>97.5</b>	97.5
Hemp	21.6	60.5	64.4	<b>72.7</b>
Leguminous Fodder	15.6	42.9	<b>43.8</b>	34.8
Other Fodder	0.0	25.5	<b>32.9</b>	14.8
Oilseed	93.1	95.9	<b>96.0</b>	95.1
Protein Crop	79.2	87.9	<b>89.1</b>	87.6
Industrial Crop	19.6	40.3	<b>47.3</b>	23.9
Fruit, Vegetable, Flower	42.5	60.3	<b>63.1</b>	55.1
Aromatic/Medicinal Plant	0.0	<b>32.8</b>	28.9	8.3
Ligneous	5.5	<b>39.8</b>	38.6	32.1
Orchard	0.0	<b>8.0</b>	6.0	2.5
Grapevine	81.8	<b>95.1</b>	94.3	92.4
Short Rotation Coppice	0.0	8.3	<b>18.2</b>	17.4
Non Agricultural Area	46.7	54.9	<b>59.1</b>	55.2
Other	1.2	24.0	<b>27.0</b>	14.1

lem at hand. CNNs are well-suited for extracting shape and texture information, and it appears that parcels’ shape does not strongly correlate with crop type. Furthermore, the resolution of Sentinel-2 images may not allow to capture rich texture information (see Figure 1(b)). This would also explain why the ConvLSTM model — which relies more on convolutions — performs slightly worse than the CNN+GRU ones.

Additionally, the results stress the importance of choosing a large enough hidden state to fully leverage the temporal structure of the data. Comparing the GRU and CNN+GRU<sub>7</sub> models indicates that allocating too many parameters to extract convolutional features reduces the performance compared to a model learning on simple hand crafted features with a larger hidden state.

We show the per-class performances of the four types of architectures in Table 3. For most classes, the use of convolutional features in the CNN+GRU<sub>9</sub> improves the performance compared to the GRU model. Yet some classes such as Grapevine or Ligneous present the opposite behaviour, showing that the relative importance of recurrent and convolutional features is class-dependant.

Table 4: F-score while trained on the regular image sequences, and with randomly shuffled sequences.

	Ordered	Shuffled	$\Delta$
ConvLSTM	49.2	47.9	-1.3
CNN+GRU <sub>7</sub>	53.2	<b>52.0</b>	-1.2
CNN+GRU <sub>9</sub>	<b>55.1</b>	48.8	-6.3
GRU	53.8	45.3	<b>-8.5</b>

Finally, in order to assess the importance of the temporal structure for the features extracted by the recurrent networks, we re-train several models with randomly shuffled input sequences, such that the temporal structure is lost. Table 4 summarizes the F-scores obtained.

Time-shuffling impacts all models negatively. This impact is all the more important as the ratio of temporal pa-

rameters is high. Yet, all models still outperform the purely convolutional model. These results suggest that the hidden states of the recurrent units act in two ways: first, as a memory storing information regardless of the order, and second, as a model for the temporal evolution of the crops. As our dataset only covers a single year, this chronological evolution is most probably related to the phenology of the crops.

#### 4. CONCLUDING REMARKS

In this article, we compared the performance of four deep learning architectures extracting spatial, temporal, or spatio-temporal features for crop type mapping with a constraint on the total number of parameters. Our results showed that the best performance is achieved with spatio-temporal features when up to 90% of the parameters are allocated to the temporal structure of the data.

This suggests that simple convolutional architectures are sufficient to extract expressive features from Sentinel-2 images. Moreover, this emphasizes the importance of the temporal dimension of Sentinel-2 data to classify crop types. We showed that RNNs can successfully leverage this structure by acting as a memory combining multiple observations and foremost by taking into account the temporal evolution of the different observations over a year.

In future works, it would be interesting to explore the time-modelling capacity of RNNs in the case of multi-annual observations. Such use cases can present temporal structures of higher order such as the crop rotation patterns studied by [3]. Additionally, the potential of RNNs should be even more valuable with the fully operational revisit time of five days achieved by the Sentinel-2 mission since late 2017. Furthermore, we focused here on the temporal and spatial dimensions and treated the spectral dimension with simple concatenation, while learned spectral features could be beneficial as well. This study should, therefore, be extended to cover the extraction of relevant spectral features with deep learning methods.

More generally, our results highlight the potential of deep learning models for agricultural parcel classification: all RNN and RCNN models outperform the RF baseline which achieves an average F-score of 36.9 on the same dataset (not shown here). This study provides a set of guidelines to help in designing bespoke models for the task at hand and contribute to the development of automated crop type mapping at a large scale.

#### 5. REFERENCES

- [1] J.Inglada, M.Arias, et al., “Assessment of an operational system for crop type map production using high temporal and spatial resolution satellite optical imagery,” *Remote Sensing*, vol. 7, no. 9, 2015.
- [2] J.Singh, U.Devi, et al., “Crop-identification using sentinel-1 and sentinel-2 data for indian region,” in *IGARSS. IEEE*, 2018.
- [3] S.Bailly, S.Giordano, et al., “Crop-rotation structured classification using multi-source sentinel images and lpis for crop type mapping,” in *IGARSS. IEEE*, 2018.
- [4] S.Siachalou, G.Mallinis, et al., “A hidden markov models approach for crop classification: Linking crop phenology to time series of multi-sensor remote sensing data,” *Remote Sensing*, vol. 7, no. 4, 2015.
- [5] Y.LeCun, Y.Bengio, et al., “Convolutional networks for images, speech, and time series,” *The handbook of brain theory and neural networks*, vol. 3361, no. 10, 1995.
- [6] S.Hochreiter and J.Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, 1997.
- [7] S.Ji, C.Zhang, et al., “3d convolutional neural networks for crop classification with multi-temporal remote sensing images,” *Remote Sensing*, vol. 10, no. 1, 2018.
- [8] M.Rußwurm and M.Körner, “Multi-temporal land cover classification with long short-term memory neural networks,” *ISPRS Archive*, vol. 42, 2017.
- [9] M.Rußwurm and M.Körner, “Multi-temporal land cover classification with sequential recurrent encoders,” *ISPRS*, vol. 7, no. 4, 2018.
- [10] S.Ioffe and C.Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *ICML*, 2015.
- [11] X.Glorot, A.Bordes, et al., “Deep sparse rectifier neural networks,” in *Proceedings of the international conference on artificial intelligence and statistics*, 2011.
- [12] C. R.Qi, H.Su, et al., “Pointnet: Deep learning on point sets for 3d classification and segmentation,” *Proc. CVPR, IEEE*, vol. 1, no. 2, 2017.
- [13] J.Chung, C.Gulcehre, et al., “Empirical evaluation of gated recurrent neural networks on sequence modeling,” *CoRR*, vol. abs/1412.3555, 2014.
- [14] S.Xingjian, Z.Chen, et al., “Convolutional lstm network: A machine learning approach for precipitation nowcasting,” in *NIPS*, 2015.
- [15] D. P.Kingma and J.Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014.
- [16] E.Christophe, J.Inglada, et al., “Orfeo toolbox: a complete solution for mapping from high resolution satellite images,” *ISPRS Archive*, vol. 37, no. PART B4, 2008.