

# Panoptic Segmentation of Satellite Image Time Series with Convolutional Temporal Attention Networks

Vivien Sainte Fare Garnot

Loic Landrieu

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

{vivien.sainte-fare-garnot, loic.landrieu}@ign.fr

## Abstract

Unprecedented access to multi-temporal satellite imagery has opened new perspectives for a variety of Earth observation tasks. Among them, pixel-precise panoptic segmentation of agricultural parcels has major economic and environmental implications. While researchers have explored this problem for single images, we argue that the complex temporal patterns of crop phenology are better addressed with temporal sequences of images. In this paper, we present the first end-to-end, single-stage method for panoptic segmentation of Satellite Image Time Series (SITS). This module can be combined with our novel image sequence encoding network which relies on temporal self-attention to extract rich and adaptive multi-scale spatio-temporal features. We also introduce PASTIS, the first open-access SITS dataset with panoptic annotations. We demonstrate the superiority of our encoder for semantic segmentation against multiple competing architectures, and set up the first state-of-the-art of panoptic segmentation of SITS. Our *implementation* and *PASTIS* are publicly available.

## 1. Introduction

The precision and availability of Earth observations have continuously improved thanks to sustained advances in space-based remote sensing, such as the launch of the Planet [5] and the open-access Sentinel constellations [8]. In particular, satellites with high revisit frequency contribute to a better understanding of phenomena with complex temporal dynamics. Crop mapping—the driving application of this paper—relies on exploiting such temporal patterns [37] and entails major financial and environmental stakes. Indeed, remote monitoring of the surface and nature of agricultural parcels is necessary for a fair allocation of agricultural subsidies (50 and 22 billion euros per year in Europe and in the US, respectively) and for ensuring that best crop rotation practices are respected. More generally, the automated analysis of SITS represents a significant interest for a wide

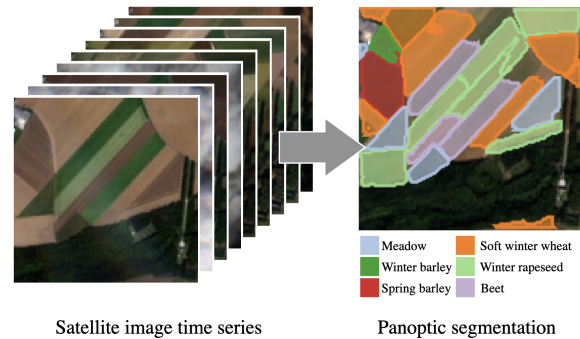


Figure 1: **Overview.** We propose an end-to-end, single-stage model for panoptic segmentation of agricultural parcels from time series of satellite images. Note the difficulty of resolving the parcels’ borders from a single image, highlighting the need for modeling temporal dynamics.

range of applications, such as surveying urban development and deforestation.

The task of monitoring both the content and extent of agricultural parcels can be framed as the panoptic segmentation of an image sequence. Panoptic segmentation consists of assigning to each pixel a class and a unique instance label, and has become a standard visual perception task in computer vision [18, 25]. However, panoptic segmentation is a fundamentally different task for SITS versus sequences of natural images or videos. Indeed, understanding videos requires tracking objects through time and space [43]. In yearly SITS, the targets are static in a geo-referenced frame, which removes the need for spatial tracking. Additionally, SITS share a common temporal frame of reference, which means that the time of acquisition itself contains information useful for modeling the underlying temporal dynamics. In contrast, the frame number in videos is often arbitrary. Finally, while objects on the Earth surface generally do not occlude one another, as is commonly the case for objects in natural images, varying cloud cover can make the analysis of SITS arduous. For the specific problem addressed in this

paper, individualizing agricultural parcels requires learning complex and specific temporal, spatial, and spectral patterns not commonly encountered in video processing, such as differences in plant phenological profiles, subpixel border information, and swift human interventions such as harvests.

While deep networks have proven efficient for learning such complex patterns for pixel classification [16, 12, 1], there is no dedicated approach for detecting individual objects in SITS. Existing work on instance segmentation has been restricted to analysing a single satellite image [32]. In summary, specialized remote sensing methods are limited to semantic segmentation or single-image instance segmentation, while computer vision’s panoptic-ready networks require significant adaptation to be applied to SITS.

In this paper, we introduce U-TAE (U-net with Temporal Attention Encoder), a novel spatio-temporal encoder combining multi-scale spatial convolutions [33] and a temporal self-attention mechanism [37] which learns to focus on the most salient acquisitions. While convolutional-recurrent methods are limited to extracting temporal features at the highest [34] or lowest [36] spatial resolutions, our proposed method can use the predicted temporal masks to extract specialized and adaptive spatio-temporal features at different resolutions simultaneously. We also propose Parcels-as-Points (PaPs), the first end-to-end deep learning method for panoptic segmentation of SITS. Our approach is built upon the efficient CenterMask network [48], which we modify to fit our problem. Lastly, we present Panoptic Agricultural Satellite Time-Series (PASTIS), the first open-access dataset for training and evaluating panoptic segmentation models on SITS, with over 2 billion annotated pixels covering over 4000km<sup>2</sup>. Evaluated on this dataset, our approach outperforms all reimplemented competing methods for semantic segmentation, and defines the first state-of-the-art of SITS panoptic segmentation.

## 2. Related Work

To the best of our knowledge, no instance or panoptic segmentation method operating on SITS has been proposed to date. However, there is a large body of work on both the encoding of satellite sequences, and the panoptic segmentation of videos and single satellite images.

**Encoding Satellite Image Sequences.** While the first automated tools for SITS analysis relied on traditional machine learning [13, 45], deep convolutional networks allow for the extraction of richer spatial descriptors [19, 12, 1, 16]. The temporal dimension was initially dealt via handcrafted temporal descriptors [2, 42, 51] or probabilistic models [3], which have been advantageously replaced by recurrent [34, 37, 27], convolutional [29, 36, 15], or differential [24] architectures. Recently, attention-based approaches have

been adapted to encode sequences of remote sensing images and have led to significant progress for pixel-wise and parcel-wise classification [38, 35, 53]. In parallel, hybrid architectures [41, 36, 28] relying on U-Net-type architectures [33] for encoding the spatial dimension and recurrent networks for the temporal dimension have shown to be well suited for the semantic segmentation of SITS. In this paper, we propose to combine this hybrid architecture with the promising temporal attention mechanism.

**Instance Segmentation of Satellite Images.** The first step of panoptic segmentation is to delineate all individual instances, *i.e.* instance segmentation. Most remote sensing instantiation approaches operate on a single acquisition. For example, several methods have been proposed to detect individual instances of trees [31, 54], buildings [46], or fields [32]. Several algorithms start with a delineation step (border detection) [9, 23, 47], and require postprocessing to obtain individual instances. Other methods use segmentation as a preprocessing step and compute cluster-based features [6, 7], but do not produce explicit cluster-to-object mappings. Petitjean *et al.* [30] propose a segmentation-aided classification method operating on image time series. However, their approach partitions each image separately and does not attempt to retrieve individual objects consistently across the entire sequence. In this paper, we propose the first end-to-end framework for directly performing joint semantic and instance segmentation on SITS.

**Panoptic Segmentation of Videos.** Among the vast literature on instance segmentation, Mask-RCNN [11] is the leading method for natural images. Recently, Wang *et al.* proposed CenterMask [48], a lighter and more efficient single-stage method which we use as a starting point in this paper. Several approaches propose extending instance or panoptic segmentation methods from image to video [50, 43, 17]. However, as explained in the introduction, SITS differs from natural video in several key ways which require specific algorithmic and architectural adaptations.

## 3. Method

We consider an image time sequence  $X$ , organized into a four-dimensional tensor of shape  $T \times C \times H \times W$ , with  $T$  the length of the sequence,  $C$  the number of channels, and  $H \times W$  the spatial extent.

### 3.1. Spatio-Temporal Encoding

Our model, dubbed U-TAE (U-Net with Temporal Attention Encoder), encodes a sequence  $X$  in three steps: **(a)** each image in the sequence is embedded simultaneously and independently by a shared multi-level spatial convolutional encoder, **(b)** a temporal attention encoder collapses

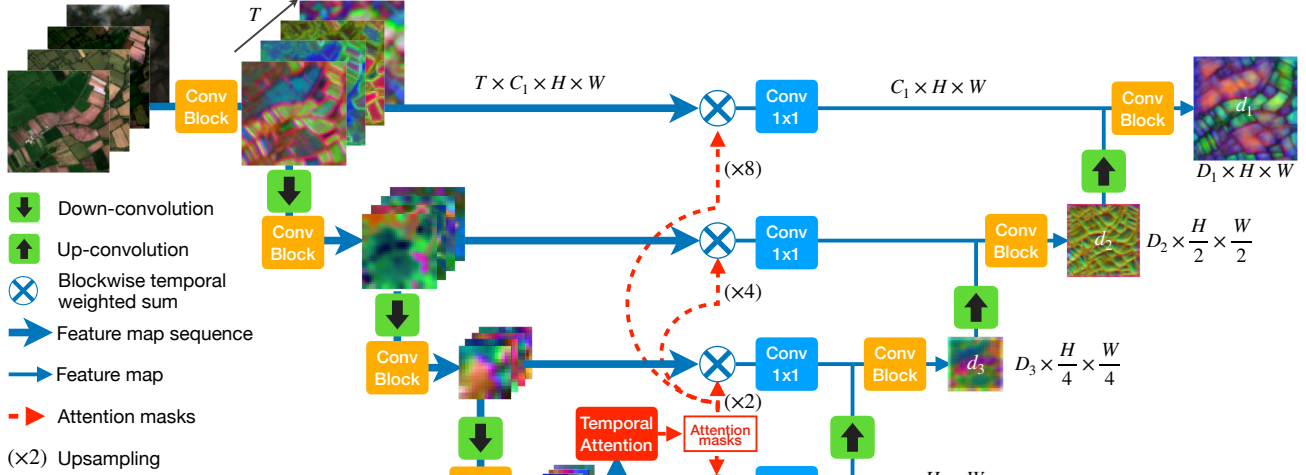


Figure 2: **Spatio-temporal Encoding.** A sequence of images is processed in parallel by a shared convolutional encoder. At the lowest resolution, an attention-based temporal encoder produces a set of temporal attention masks for each pixel, which are then spatially interpolated at all resolutions. These masks are used to collapse the temporal dimension of the feature map sequences into a single map per resolution. A convolutional decoder then computes features at all resolution levels. All convolutions operate purely on the spatial and channel dimensions, and we use strided convolutions for both spatial up and down-sampling. The feature maps are projected in RGB space to help visual interpretation.

the temporal dimension of the resulting sequence of feature maps into a single map for each level, (c) a spatial convolutional decoder produces a single feature map with the same resolution as the input images, see Figure 2.

**a) Spatial Encoding.** We consider a convolutional encoder  $\mathcal{E}$  with  $L$  levels  $1, \dots, L$ . Each level is composed of a sequence of convolutions, Rectified Linear Unit (ReLU) activations, and normalizations. Except for the first level, each block starts with a strided convolution, dividing the resolution of the feature maps by a factor 2.

For each time stamp  $t$  simultaneously, the encoder  $\mathcal{E}_l$  at level  $l$  takes as input the feature map of the previous level  $e_t^{l-1}$ , and outputs a feature map  $e_t^l$  of size  $C_l \times H_l \times W_l$  with  $H_l = H/2^{l-1}$  and  $W_l = W/2^{l-1}$ . The resulting feature maps are then temporally stacked into a feature map sequence  $e^l$  of size  $T \times C_l \times H_l \times W_l$ :

$$e^l = [\mathcal{E}_l(e_t^{l-1})]_{t=0}^T \text{ for } l \in [1, L], \quad (1)$$

with  $e^0 = X$  and  $[\cdot]$  the concatenation operator along the temporal dimension. When constituting batches, we flatten the temporal and batch dimensions. Since each sequence comprises images acquired at different times, the batches' samples are not identically distributed. To address this issue, we use Group Normalization [49] with 4 groups instead of Batch Normalization [14] in the encoder.

**b) Temporal Encoding.** In order to obtain a single representation per sequence, we need to collapse the temporal dimension of each feature map sequence  $e^l$  before using them as *skip connections*. Convolutional-recurrent U-Net networks [41, 36, 28] only process the temporal dimension of the lowest resolution feature map with a temporal encoder. The rest of the skip connections are collapsed with a simple temporal average. This prevents the extraction of spatially adaptive and parcel-specific temporal patterns at higher resolutions. Conversely, processing the highest resolution would result in small spatial receptive fields for the temporal encoder, and an increased memory requirement. Instead, we propose an attention-based scheme which only processes the temporal dimension at the lowest feature map resolution, but is able to utilize the predicted temporal attention masks at all resolutions simultaneously.

Based on its performance and computational efficiency, we choose the Lightweight-Temporal Attention Encoder (L-TAE) [10] to handle the temporal dimension. The L-TAE is a simplified multi-head self-attention network [44] in which the attention masks are directly applied to the input sequence of vectors instead of predicted *values*. Additionally, the L-TAE implements a channel grouping strategy similar to Group Normalization [49].

We apply a shared L-TAE with  $G$  heads independently at each pixel of  $e^L$ , the feature map sequence at the lowest level resolution  $L$ . This generates  $G$  temporal attention masks for each pixel, which can be arranged into  $G$  tensors

$a^{L,g}$  with values in  $[0, 1]$  and of shape  $T \times H_L \times W_L$ :

$$a^{L,1}, \dots, a^{L,G} = \text{LTAE}(e^L), \text{ applied pixelwise.} \quad (2)$$

In order to use these attention masks at all scale levels  $l$  of the encoder, we compute spatially-interpolated masks  $a^{l,g}$  of shape  $T \times H_l \times W_l$  for all  $l$  in  $[1, L - 1]$  and  $g$  in  $[1, G]$  with bilinear interpolation:

$$a^{l,g} = \text{resize } a^{L,g} \text{ to } H_l \times W_l. \quad (3)$$

The interpolated masks  $a^{l,g}$  at level  $l$  of the encoder are then used as if they were generated by a temporal attention module operating at this resolution. We apply the L-TAE channel-grouping strategy at all resolution levels: the channels of each feature map sequence  $e^l$  are split into  $G$  contiguous groups  $e^{l,1}, \dots, e^{l,G}$  of identical shape  $T \times C_l/G \times W_l \times H_l$ . For each group  $g$ , the feature map sequence  $e^{l,g}$  is averaged on the spatial dimension using  $a^{l,g}$  as weights. The resulting maps are concatenated along the channel dimension, and processed by a shared  $1 \times 1$  convolution layer  $\text{Conv}_{1 \times 1}^l$  of width  $C_l$ . We denote by  $f^l$  the resulting map of size  $C_l \times W_l \times H_l$  by :

$$f^l = \text{Conv}_{1 \times 1}^l \left( \left[ \sum_{t=1}^T a_t^{l,g} \odot e_t^{l,g} \right]_{g=1}^G \right), \quad (4)$$

with  $[\cdot]$  the concatenation along the channel dimension and  $\odot$  the term-wise multiplication with channel broadcasting.

**c) Spatial Decoding.** We combine the feature maps  $f^l$  learned at the previous step with a convolutional decoder to obtain spatio-temporal features at all resolutions. The decoder is composed of  $L - 1$  blocks  $\mathcal{D}_l$  for  $1 \leq l < L$ , with convolutions, ReLU activations, and BatchNorms [14]. Each decoder block uses a strided transposed convolution  $\mathcal{D}_l^{\text{up}}$  to up-sample the previous feature map. The decoder at level  $l$  produces a feature map  $d^l$  of size  $D_l \times H_l \times W_l$ . In a U-Net fashion, the encoder’s map at level  $l$  is concatenated with the output of the decoder block at level  $l - 1$ :

$$d^l = \mathcal{D}_l([\mathcal{D}_l^{\text{up}}(d^{l+1}), f^l]) \text{ for } l \in [1, L - 1], \quad (5)$$

with  $d^L = f^L$  and  $[\cdot]$  is the channelwise concatenation.

### 3.2. Panoptic Segmentation

Our goal is to use the multi-scale feature maps  $\{d^l\}_{l=1}^L$  learnt by the spatio-temporal encoder to perform panoptic segmentation of a sequence of satellite images over an area of interest. The first stage of panoptic segmentation is to produce instance proposals, which are then combined into a single panoptic instance map. Since an entire sequence of images (often over 50) must be encoded to compute  $\{d^l\}_{l=1}^L$ , we favor a simple approach for our panoptic segmentation module. Furthermore, given the relative

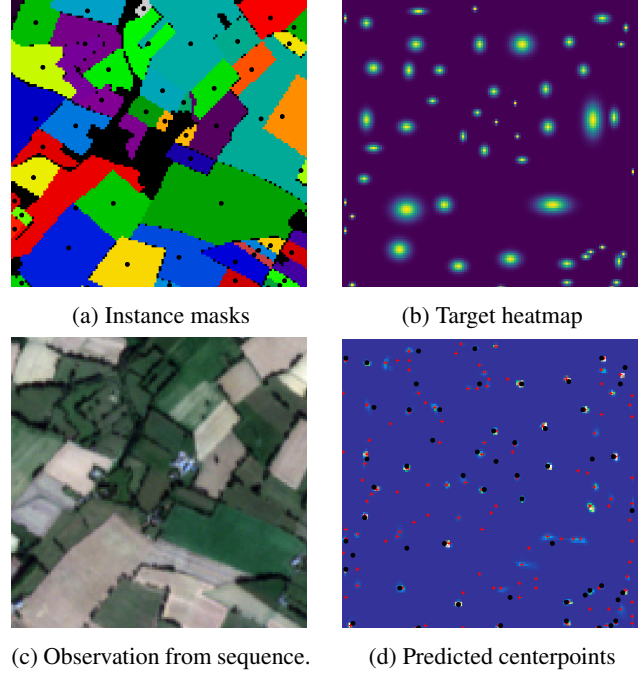


Figure 3: **Centerpoint Detection.** The ground truth instance masks (a) is used to construct a target heatmap (b). Our parcel detection module maps the raw sequence of observation (c) to a predicted heatmap (d). The predicted centerpoints (red crosses) are the local maxima of the predicted heatmap (d). The black dots are the true parcels centers.

simplicity of parcels’ borders, we avoid complex region proposal networks such as Mask-RCNN. Instead, we adapt the single-stage CenterMask instance segmentation network [48], and detail our modifications in the following paragraphs. We name our approach *Parcels-as-Points* (PaPs) to highlight our inspiration from CenterNet/Mask [55, 48].

We denote by  $P$  the set of ground truth parcels in the image sequence  $X$ . Note that the position of these parcels is time-invariant and hence only defined by their spatial extent. Each parcel  $p$  is associated with (i) a centerpoint  $\hat{i}_p, \hat{j}_p$  with integer coordinates, (ii) a bounding box of size  $\hat{h}_p, \hat{w}_p$ , (iii) a binary instance mask  $\hat{s}_p \in \{0, 1\}^{H \times W}$ , (iv) a class  $\hat{k}_p \in [1, K]$  with  $K$  the total number of classes.

**Centerpoint Detection.** Following CenterMask, we perform parcel detection by predicting *centerness heatmaps* supervised by the ground truth parcels’ bounding boxes. In the original approach [55], each class has its own heatmap: detection doubles as classification. This is a sensible choice for natural images, since the tasks of detecting an object’s nature, location, and shape are intrinsically related. In our setting however, the parcels’ shapes and border characteristics are mostly independent of the cultivated crop. For this

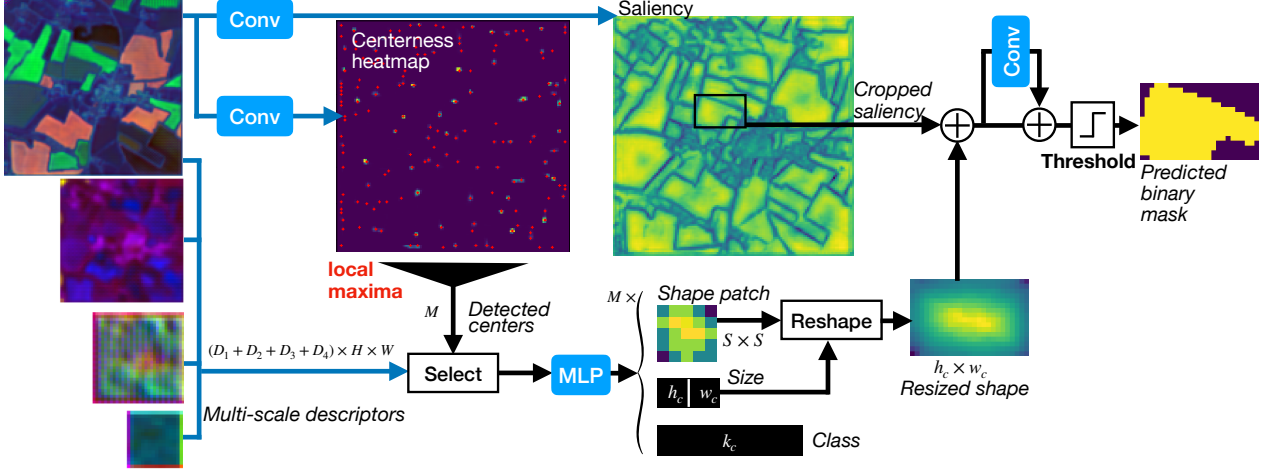


Figure 4: **Panoptic Segmentation.** The local maxima of the predicted centerness heatmap defines  $M$  tentative parcels. For each one, the pixel features at all levels are concatenated and used to predict a bounding box size, a semantic class, and an  $S \times S$  shape patch. The latter is combined with a global saliency map for predicting pixel-precise masks. The instance predictions are combined into a panoptic segmentation using the centerness as quality.

reason, we use a single centerness heatmap and postpone class identification to a subsequent specialized module. See Figure 3 for an illustration of our parcel detection method.

We associate each parcel  $p$  with a Gaussian kernel of deviations  $\sigma_p^{\text{ver}}$  and  $\sigma_p^{\text{hor}}$  taken respectively as  $1/20$  of the height and width of the parcels’ bounding box. Unlike Law and Deng [20], we use heteroschedastic kernels to reflect the potential narrowness of parcels. We then define the target centerness heatmap  $\hat{m} \in [0, 1]^{H \times W}$  as the maximum value of all parcel kernels at each pixel  $(i, j)$  in  $H \times W$ :

$$\hat{m}_{i,j} = \max_{p \in P} \exp \left( - \left[ \frac{(i - \hat{i}_p)^2}{2(\sigma_p^{\text{ver}})^2} + \frac{(j - \hat{j}_p)^2}{2(\sigma_p^{\text{hor}})^2} \right] \right) \quad (6)$$

A convolutional layer takes the highest-resolution feature map  $d^1$  as input and predicts a centerness heatmap  $m \in [0, 1]^{H \times W}$ . The predicted heatmap is supervised using the loss defined in Equation 7 with  $\beta = 4$ :

$$\mathcal{L}_{\text{center}} = \frac{-1}{|P|} \sum_{\substack{i=1 \dots H \\ j=1 \dots W}} \begin{cases} \log(m_{i,j}) & \text{if } \hat{m}_{i,j} = 1 \\ (1 - \hat{m}_{i,j})^\beta \log(1 - m_{i,j}) & \text{else.} \end{cases} \quad (7)$$

We define the predicted centerpoints as the local maxima of  $m$ , *i.e.* pixels with larger values than their 8 adjacent neighbors. This set can be efficiently computed with a single max-pooling operation. Replacing the max operator by argmax in Equation 6 defines a mapping  $H \times W \mapsto P$  between pixels and parcels. During training, we associate each true parcel  $p$  with the predicted centerpoint  $c(p)$  with highest predicted centerness  $m$  among the set of centerpoints which coordinates are mapped to  $p$ . If this set is

empty, then  $c(p)$  is undefined: the parcel  $p$  is not detected. We denote by  $P'$  the subset of detected parcels, *i.e.* for which  $c(p)$  is well defined.

**Size and Class Prediction.** We associate to a predicted centerpoint  $c$  of coordinate  $(i_c, j_c)$  the multi-scale feature vector  $\tilde{d}_c$  of size  $D_1 + \dots + D_L$  by concatenating channelwise the pixel features at location  $(i_c, j_c)$  in all maps  $d^l$ :

$$\tilde{d}_c = [d^l ([i_c/2^{l-1}], [j_c/2^{l-1}])]_{l=1}^L, \quad (8)$$

with  $[\cdot]$  the channelwise concatenation. This vector  $\tilde{d}_c$  is then processed by four different multilayer perceptrons (MLP) to obtain three vectors of sizes 2,  $K$ , and  $S^2$  representing respectively: (i) a bounding box size  $h_c, w_c$ , (ii) a vector of class probabilities  $k_c$  of size  $K$ , and (iii) a shape patch  $s_c$  of fixed size  $S \times S$ . The latter is described in the next paragraph.

The class prediction  $k_{c(p)}$  associated to the true parcel  $p$  is supervised with the cross-entropy loss, and the size prediction with a normalized L1 loss. For all  $p$  in  $P'$ , we have:

$$\mathcal{L}_{\text{class}}^p = -\log(k_{c(p)}[\hat{k}_p]) \quad (9)$$

$$\mathcal{L}_{\text{size}}^p = \frac{|h_{c(p)} - \hat{h}_p|}{\hat{h}_p} + \frac{|w_{c(p)} - \hat{w}_p|}{\hat{w}_p}. \quad (10)$$

**Shape Prediction.** The idea of this step is to combine for a predicted centerpoint  $c$  a rough shape patch  $s_c$  with a full-resolution global saliency map  $z$  to obtain a pixel-precise

instance mask, see Figure 4. For a centerpoint  $c$  of coordinates  $(i_c, j_c)$ , the predicted shape patch  $s_c$  of size  $S \times S$  is resized to the predicted size  $\lceil h_c \rceil \times \lceil w_c \rceil$  with bilinear interpolation. A convolutional layer maps the outermost feature map  $d^1$  to a saliency map  $z$  of size  $H \times W$ , which is shared by all predicted parcels. This saliency map is then cropped along the predicted bounding box  $(i_c, j_c, \lceil h_c \rceil, \lceil w_c \rceil)$ . The resized shape and the cropped saliency are added (11) to obtain a first local shape  $\tilde{l}_c$ , which is then further refined with a residual convolutional network CNN (12). We denote the resulting predicted shape by  $l_c$ :

$$\tilde{l}_c = \text{resize}_c(s_c) + \text{crop}_c(z) \quad (11)$$

$$l_c = \text{sigmoid}(\tilde{l}_c + \text{CNN}(\tilde{l}_c)), \quad (12)$$

with  $\text{resize}_c$  and  $\text{crop}_c$  defined by the coordinates  $(i_c, j_c)$  and predicted bounding box size  $(\lceil h_c \rceil, \lceil w_c \rceil)$ . The shape and saliency predictions are supervised for each parcel  $p$  in  $P'$  by computing the pixelwise binary cross-entropy (BCE) between the predicted shape  $l_{c(p)}$  and the corresponding true binary instance mask  $\hat{s}_p$  cropped along the predicted bounding box  $(i_{c(p)}, j_{c(p)}, \lceil h_{c(p)} \rceil, \lceil w_{c(p)} \rceil)$ :

$$\mathcal{L}_{\text{shape}}^p = \text{BCE}(l_{c(p)}, \text{crop}_{c(p)}(\hat{s}_p)). \quad (13)$$

For inference, we associate a binary mask with a predicted centerpoint  $c$  by thresholding  $l_c$  with the value 0.4.

**Loss Function** : These four losses are combined into a single loss with no weight and optimized end-to-end:

$$\mathcal{L} = \mathcal{L}_{\text{center}} + \frac{1}{|P'|} \sum_{p \in P'} (\mathcal{L}_{\text{class}}^p + \mathcal{L}_{\text{size}}^p + \mathcal{L}_{\text{shape}}^p). \quad (14)$$

**Differences with CenterMask.** Our approach differs from CenterMask in several key ways: (i) We compute a single saliency map and heatmap instead of  $K$  different ones. This represents the absence of parcel occlusion and the similarity of their shapes. (ii) Accounting for the lower resolution of satellite images, centerpoints are computed at full resolution to detect potentially small parcels, thus dispensing us from predicting offsets. (iii) The class prediction is handled centerpoint-wise instead of pixel-wise for efficiency. (iv) Only the selected centerpoints predict shape, class, and size vectors, saving computation and memory. (v) We use simple feature concatenation to compute multi-scale descriptors instead of deep layer aggregation [52] or stacked Hourglass-Networks [26]. (vi) A convolutional network learns to combine the saliency and the mask instead of a simple term-wise product.

**Converting to Panoptic Segmentation** Panoptic segmentation consists of associating to each pixel a semantic

label and, for non-background pixels (our only *stuff* class), an instance label [18]. Our predicted binary instance masks can have overlaps, which we resolve by associating to each predicted parcel a quality measure equal to the predicted centerness  $m$  at its associated centerpoint. Masks with higher quality overtake the pixels of overlapping masks with lesser predicted quality. If a mask loses more than 50% of its pixels through this process, it is removed altogether from the predicted instances. Predicted parcels with a quality under a given threshold are dropped. This threshold can be tuned on a validation set to maximize the parcel detection F-score. All pixels not associated with a parcel mask are labelled as background.

**Implementation Details.** Our implementation of U-TAE allows for batch training on sequences of variable length thanks to a simple padding strategy. The complete configuration and training details can be found in the Appendix. A Pytorch implementation is available at <https://github.com/VSainteuf/utae-paps>.

## 4. Experiments

### 4.1. The PASTIS Dataset

We present PASTIS (Panoptic Agricultural Satellite Time Series), the first large-scale, publicly available SITS dataset with both semantic and panoptic annotations. This dataset, as well as more information about its composition, are publicly available at <https://github.com/VSainteuf/pastis-benchmark>.

**Description.** PASTIS is comprised of 2 433 sequences of multi-spectral images of shape  $10 \times 128 \times 128$ . Each sequence contains between 38 and 61 observations taken between September 2018 and November 2019, for a total of over 2 billion pixels. The time between acquisitions is uneven with a median of 5 days. This lack of regularity is due to the automatic filtering of acquisitions with extensive cloud cover by the satellite data provider **THEIA**. The 10 channels correspond to the non-atmospheric spectral bands of the Sentinel-2 satellite, after atmospheric correction and re-sampling at a spatial resolution of 10 meters per pixel. The dataset spans over 4000 km<sup>2</sup>, with images taken from four different regions of France with diverse climates and crop distributions, covering almost 1% of the French Metropolitan territory. We estimate that close to 28% of images have at least partial cloud cover.

**Annotation.** Each pixel of PASTIS is associated with a semantic label taken from a nomenclature of 18 crop types plus a background class. As is common in remote sensing applications, the dataset is highly unbalanced, with a ratio of over 50 between the most and least common classes.

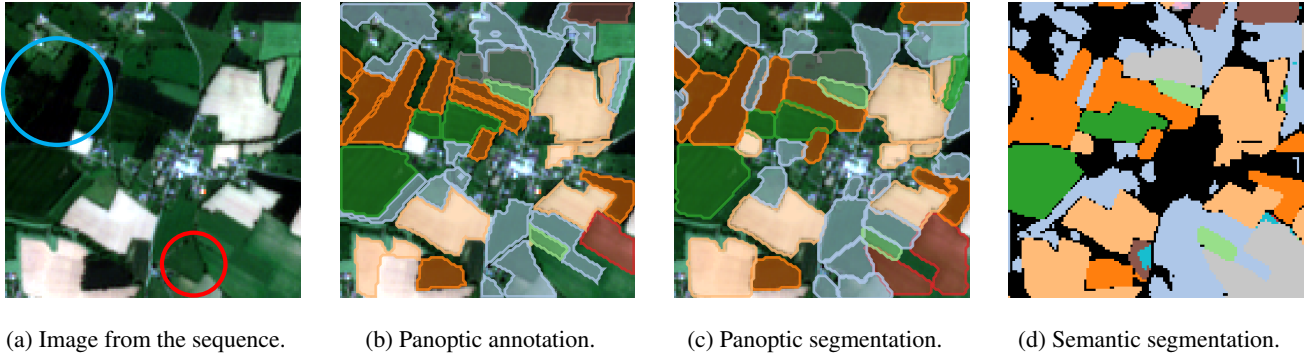


Figure 5: **Qualitative results.** We consider an image sequence (a) with panoptic annotations (b). We represent the results of our method in terms of panoptic segmentation (c) and semantic segmentation (d). The parcels’ and pixels’ color corresponds to the crop type, according to a legend given in the appendix. The predominantly correct class predictions highlight the fact that the difficulty of panoptic segmentation lies in the precise delineation of each individual parcel. We observe cases where the temporal structure of the SITS was successfully leveraged to resolve boundary ambiguities that could not be seen from a single image (cyan circle  $\circ$ ). Conversely, some visually fragmented parcels are annotated as a single instance (red circle  $\circ$ ).

Each non-background pixel also has a unique instance label corresponding to its parcel index. In total, 124 422 parcels are individualized, each with their bounding box, pixel-precise mask, and crop type. All annotations are taken from the publicly available French Land Parcel Identification System. The French Payment Agency estimates the accuracy of the crop annotations via in situ control over 98% and the relative error in terms of surfaces under 0.3%. To allow for cross-validation, the dataset is split into 5 folds, chosen with a 1km buffer between images to avoid cross-fold contamination.

## 4.2. Semantic Segmentation

Our U-TAE has  $L = 4$  resolution levels and a LTAE with  $G = 16$  heads, see appendix for an exact configuration. For the semantic segmentation task, the feature map  $d_1$  with highest resolution is set to have  $K$  channels, with  $K$  the number of classes. We can then interpret  $d_1$  as pixel-wise predictions to be supervised with the cross-entropy loss. In this setting, we do not use the PaPs module.

**Competing Methods.** We reimplemented six of the top-performing SITS encoders proposed in the literature:

- *ConvLSTM* [34, 39] and *ConvGRU*[4]. These approaches are recurrent neural networks in which all linear layers are replaced by spatial convolutions.
- *U-ConvLSTM* [36] and *U-BiConvLSTM* [22]. To reproduce these UNet-Based architectures, we replaced the L-TAE in our architecture by either a convLSTM [40] or a bidirectional convLSTM. Skip connections are temporally averaged. In contrast to the original methods, we replaced the batch normalization in the encoders with

Table 1: **Semantic Segmentation.** We report for our method and six competing methods the model size in trainable parameters, Overall Accuracy (OA), mean Intersection over Union (mIoU), and Inference Time for one fold of  $\sim 490$  sequences (IT). The second part of the table report results from our ablation study.

Model	# param $\times 1000$	OA	mIoU	IT (s)
U-TAE (ours)	1 087	<b>83.2</b>	<b>63.1</b>	<b>25.7</b>
3D-Unet [36]	1 554	81.3	58.4	29.5
U-ConvLSTM [36]	1 508	82.1	57.8	28.3
FPN-ConvLSTM [22]	1 261	81.6	57.1	103.6
U-BiConvLSTM [22]	1 434	81.8	55.9	32.7
ConvGRU [4]	1 040	79.8	54.2	49.0
ConvLSTM [34, 39]	1 010	77.9	49.1	49.1
Mean Attention	1 087	82.8	60.1	24.8
Skip Mean + Conv	1 087	82.4	58.9	24.5
Skip Mean	1 074	82.0	58.3	24.5
BatchNorm	1 087	71.9	36.0	22.3
Single Date (August)	1 004	65.6	28.3	1.3
Single Date (May)	1 004	58.1	20.6	1.3

group normalization which significantly improved the results across-the-board.

- *3D-Unet* [36]. A U-Net in which the convolutions of the encoding branch are three-dimensional to handle simultaneously the spatial and temporal dimensions.
- *FPN-ConvLSTM* [22]. This model combines a feature pyramid network [21] to extract spatial features and a bidirectional ConvLSTM for the temporal dimension.

**Analysis.** In Table 1, we detail the performance obtained with 5-fold cross validation of our approach and the six reimplemented baselines. We report the Overall Accuracy (OA) as the ratio between correct and total predictions, and (mIoU) the class-averaged classification IoU. We observe that the convolutional-recurrent methods *ConvGRU* and *ConvLSTM* perform worse. Recurrent networks embedded in an U-Net or a FPN share similar performance, with a much longer inference time for FPN. Our approach significantly outperforms all other methods in terms of precision. In Figure 5, we present a qualitative illustration of the semantic segmentation results.

**Ablation Study.** We first study the impact of using spatially interpolated attention masks to collapse the temporal dimension of the spatio-temporal feature maps at different levels of the encoder simultaneously. Simply computing the temporal average of skip connections for levels without temporal encoding as proposed by [41, 36], we observe a drop of 4.8 mIoU points (Skip Mean). This puts our method performance on par with its competing approaches. Adding a  $1 \times 1$  convolutional layer after the temporal average reduces this drop to 4.2 points (Skip Mean + Conv). Lastly, using interpolated masks but foregoing the channel grouping strategy by averaging the masks group-wise into a single attention mask per level results in a drop of 3.1 points (Mean Attention). This implies that our network is able to use the grouping scheme at different resolutions simultaneously. In conclusion, the main advantage of our proposed attention scheme is that the temporal collapse is controlled at all resolutions, in contrast to recurrent methods.

Using batch normalization in the encoder leads to a severe degradation of the performance of 27.1 points (Batch-Norm). We conclude that the temporal diversity of the acquisitions requires special considerations. This was observed for all U-Net models alike. We also train our model on a single acquisition date (with a classic U-Net and no temporal encoding) for two different cloudless dates in August and May (Single Date). We observe a drop of 24.8 and 42.5 points respectively, highlighting the crucial importance of the temporal dimension for crop classification. We also observed that images with at least partial cloud cover received on average 58% less attention than their cloud-free counterparts. This suggests that our model is able to use the attention module to automatically filter out corrupted data.

### 4.3. Panoptic Segmentation

We use the same U-TAE configuration for panoptic segmentation, and select a PaPs module with 190k parameters and a shape patch size of  $16 \times 16$ . In Table 2, we report the class-averaged Segmentation Quality (SQ), Recognition Quality (RQ), and Panoptic Quality (PQ) [18]. We observe that while the network is able to correctly detect and clas-

Table 2: **Panoptic Segmentation Experiment.** We report class-averaged panoptic metrics: SQ, RQ, PQ.

	SQ	RQ	PQ
<b>U-TAE + PaPs</b>	81.3	49.2	40.4
U-ConvLSTM + Paps	80.9	40.8	33.4
$S = 24$	81.3	48.5	39.9
$S = 8$	81.0	48.6	39.8
Multiplicative Saliency	74.5	47.2	35.5
Single-image	72.3	16.9	12.4

sify most parcels, the task remains difficult. In particular, the combination of ambiguous borders and hard-to-classify parcel content makes for a challenging panoptic segmentation problem. We illustrate these difficulties in Figure 5, along with qualitative results.

Replacing the temporal encoder by a U-BiConvLSTM as described in Section 4.2 (U-BiConvLSTM+PaPs), we observe a noticeable performance drop of 8.4 RQ, which is consistent with the results of Table 1. As expected, our model’s performance is not sensitive to changes in the size  $S$  of the shape patch. Indeed, the shape patches only determine the rough outline of parcels while the pixel-precise instance masks are derived from the saliency map. Performing shape prediction with a simple element-wise multiplication as in [48] (Multiplicative Saliency) instead of our residual CNN results in a drop of over  $-6.8$  SQ. Using a single image (August) leads to a low panoptic quality. Indeed, identifying crop types and parcel borders from a single image at the resolution of Sentinel-2 is particularly difficult.

Inference on 490 sequences takes 129s: 26s to generate U-TAE embeddings, 1s for the heatmap and saliency, 90s for instance proposals, and 12s to merge them into a panoptic segmentation. Note that the training time is also doubled compared to simple semantic segmentation.

## 5. Conclusion

We introduced U-TAE, a novel spatio-temporal encoder using a combination of spatial convolution and temporal attention. This model can be easily combined with PaPs, the first panoptic segmentation framework operating on SITS. Lastly, we presented PASTIS, the first large-scale panoptic-ready SITS dataset. Evaluated on this dataset, our approach significantly outperformed all other approaches for semantic segmentation, and set up the first state-of-the-art for panoptic segmentation of satellite image sequences.

We hope that the combination of our open-access dataset and promising results will encourage both remote sensing and computer vision communities to consider the challenging problem of panoptic SITS segmentation, whose economic and environmental stakes can not be understated.



## References

- [1] Nicolas Audebert, Bertrand Le Saux, and Sébastien Lefèvre. Semantic segmentation of earth observation data using multi-modal and multi-scale deep networks. In *ACCV*, 2016.
- [2] Adeline Bailly, Simon Malinowski, Romain Tavenard, Laetitia Chapel, and Thomas Guyet. Dense bag-of-temporal-sift-words for time series classification. In *International Workshop on Advanced Analysis and Learning on Temporal Data*. Springer, 2015.
- [3] 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.
- [4] Nicolas Ballas, Li Yao, Chris Pal, and Aaron Courville. Delving deeper into convolutional networks for learning video representations. *ICLR*, 2016.
- [5] Christopher Boshuizen, James Mason, Pete Klupar, and Shannon Spanhake. Results from the planet labs flock constellation. *AIAA/USU Conference on Small Satellites*, 2014.
- [6] Alessandro Michele Censi, Dino Ienco, Yawogan Jean Eudes Gbodjo, Ruggero Gaetano Pensa, Roberto Interdonato, and Raffaele Gaetano. Spatial-temporal GraphCNN for land cover mapping. *IEEE Access*, 2021.
- [7] Dawa Derksen, Jordi Inglada, and Julien Michel. Spatially precise contextual features based on superpixel neighborhoods for land cover mapping with high resolution satellite image time series. In *IGARSS*, 2018.
- [8] 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. *Remote sensing of Environment*, 2012.
- [9] Angel Garcia-Pedrero, Consuelo Gonzalo-Martin, and M Lillo-Saavedra. A machine learning approach for agricultural parcel delineation through agglomerative segmentation. *International journal of remote sensing*, 2017.
- [10] 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.
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017.
- [12] Dino Ienco, Raffaele Gaetano, Claire Dupaquier, and Pierre Maurel. Land cover classification via multitemporal spatial data by deep recurrent neural networks. *Geoscience and Remote Sensing Letters*, 2017.
- [13] Jordi Inglada, Marcela Arias, Benjamin Tardy, Olivier Hagolle, Silvia Valero, David Morin, Gérard Dedieu, Guadalupe Sepulcre, Sophie Bontemps, Pierre Defourny, et al. Assessment of an operational system for crop type map production using high temporal and spatial resolution satellite optical imagery. *Remote Sensing*, 2015.
- [14] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.
- [15] Shunping Ji, Chi Zhang, Anjian Xu, Yun Shi, and Yulin Duan. 3d convolutional neural networks for crop classification with multi-temporal remote sensing images. *Remote Sensing*, 2018.
- [16] Andreas Kamilaris and Francesc X Prenafeta-Boldú. Deep learning in agriculture: A survey. *Computers and electronics in agriculture*, 2018.
- [17] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Video panoptic segmentation. In *CVPR*, 2020.
- [18] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *CVPR*, 2019.
- [19] Nataliia Kussul, Mykola Lavreniuk, Sergii Skakun, and Andrii Shelestov. Deep learning classification of land cover and crop types using remote sensing data. *Geoscience and Remote Sensing Letters*, 2017.
- [20] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, 2018.
- [21] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017.
- [22] Jorge Andres Chamorro Martinez, Laura Elena Cué La Rosa, Raul Queiroz Feitosa, Ieda Del’Arco Sanches, and Patrick Nigri Happ. Fully convolutional recurrent networks for multivariate crop recognition from multitemporal image sequences. *ISPRS*, 2021.
- [23] Khairiya Mudrik Masoud, Claudio Persello, and Valentyn A Tolpekin. Delineation of agricultural field boundaries from Sentinel-2 images using a novel super-resolution contour detector based on fully convolutional networks. *Remote sensing*, 2020.
- [24] Nando Metzger, Mehmet Ozgur Turkoglu, Stefano D’Aronco, Jan Dirk Wegner, and Konrad Schindler. Crop classification under varying cloud cover with neural ordinary differential equations. *arXiv preprint arXiv:2012.02542*, 2020.
- [25] Rohit Mohan and Abhinav Valada. Efficientps: Efficient panoptic segmentation. *International Journal of Computer Vision*, 2021.
- [26] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [27] Mehmet Ozgur Turkoglu, Stefano D’Aronco, Gregor Perich, Frank Liebisch, Constantin Streit, Konrad Schindler, and Jan Dirk Wegner. Crop mapping from image time series: deep learning with multi-scale label hierarchies. *arXiv preprint arXiv:2102.08820*, 2021.
- [28] Maria Papadomanolaki, Maria Vakalopoulou, and Konstantinos Karantzas. A deep multi-task learning framework coupling semantic segmentation and fully convolutional LSTM networks for urban change detection. *Transactions on Geoscience and Remote Sensing*, 2021.
- [29] Charlotte Pelletier, Geoffrey I Webb, and François Petitjean. Temporal convolutional neural network for the classification of satellite image time series. *Remote Sensing*, 2019.
- [30] François Petitjean, Camille Kurtz, Nicolas Passat, and Pierre Gançarski. Spatio-temporal reasoning for the classification of satellite image time series. *Pattern Recognition Letters*, 2012.

- [31] Yuchu Qin, Antonio Ferraz, Clément Mallet, and Corina Iovan. Individual tree segmentation over large areas using airborne lidar point cloud and very high resolution optical imagery. In *IGARSS*, 2014.
- [32] Christoph Rieke. Deep learning for instance segmentation of agricultural fields. [https://github.com/chrieke/InstanceSegmentation\\_Sentinel2](https://github.com/chrieke/InstanceSegmentation_Sentinel2), 2017.
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [34] Marc Rußwurm and Marco Körner. Convolutional LSTMs for cloud-robust segmentation of remote sensing imagery. *NeurIPS Workshops*, 2018.
- [35] Marc Rußwurm and Marco Körner. Self-attention for raw optical satellite time series classification. *ISPRS*, 2020.
- [36] Rose Rustowicz, Robin Cheong, Lijing Wang, Stefano Ermon, Marshall Burke, and David Lobell. Semantic segmentation of crop type in africa: A novel dataset and analysis of deep learning methods. In *CVPR Workshops*, 2019.
- [37] Vivien 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.
- [38] 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.
- [39] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *NeurIPS*, 2015.
- [40] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. *arXiv preprint arXiv:1506.04214*, 2015.
- [41] Andrei Stoian, Vincent Poulain, Jordi Inglada, Victor Poughon, and Dawa Derksen. Land cover maps production with high resolution satellite image time series and convolutional neural networks: Adaptations and limits for operational systems. *Remote Sensing*, 2019.
- [42] Romain Tavenard, Simon Malinowski, Laetitia Chapel, Adeline Bailly, Heider Sanchez, and Benjamin Bustos. Efficient temporal kernels between feature sets for time series classification. In *ECML-KDD*. Springer, 2017.
- [43] Pavel Tokmakov, Cordelia Schmid, and Karteek Alahari. Learning to segment moving objects. *International Journal of Computer Vision*, 2019.
- [44] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017.
- [45] Francesco Vuolo, Martin Neuwirth, Markus Immitzer, Clement Atzberger, and Wai-Tim Ng. How much does multi-temporal Sentinel-2 data improve crop type classification? *International journal of applied earth observation and geoinformation*, 2018.
- [46] Fabien H Wagner, Ricardo Dalagnol, Yuliya Tarabalka, Tassiana YF Segantine, Rogério Thomé, and Mayumi Hirye. U-net-id, an instance segmentation model for building extraction from satellite images—case study in the joanópolis city, brazil. *Remote Sensing*, 2020.
- [47] François Waldner and Foivos I Diakogiannis. Deep learning on edge: extracting field boundaries from satellite images with a convolutional neural network. *Remote Sensing of Environment*, 2020.
- [48] Yuqing Wang, Zhaoliang Xu, Hao Shen, Baoshan Cheng, and Lirong Yang. Centermask: single shot instance segmentation with point representation. In *CVPR*, 2020.
- [49] Yuxin Wu and Kaiming He. Group normalization. In *ECCV*, 2018.
- [50] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *CVPR*, 2019.
- [51] Lexiang Ye and Eamonn Keogh. Time series shapelets: a new primitive for data mining. In *ACM SIGKDD*, 2009.
- [52] Fisher Yu, Dequan Wang, Evan Shelhamer, and Trevor Darrell. Deep layer aggregation. In *CVPR*, 2018.
- [53] Yuan Yuan and Lei Lin. Self-supervised pre-training of transformers for satellite image time series classification. *Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2020.
- [54] Tiebiao Zhao, Haoyu Niu, Erick de la Rosa, David Doll, Dong Wang, and YangQuan Chen. Tree canopy differentiation using instance-aware semantic segmentation. In *ASABE Annual International Meeting*, 2018.
- [55] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. *arXiv preprint arXiv:1904.07850*, 2019.