

SELF-SUPERVISED LEARNING - A WAY TO MINIMIZE TIME AND EFFORT FOR PRECISION AGRICULTURE?

Michael L. Marszalek^{1*}, Bertrand Le Saux¹, Pierre-Philippe Mathieu¹, Artur Nowakowski², Daniel Springer³

¹ European Space Agency (ESA) Φ-lab, Frascati, Italy -
(Michael.Marszalek, Bertrand.Le.Saux, Pierre.Phillippe.Mathieu)@esa.int

² World Food Programme (WFP), Rome, Italy - artur.nowakowski@wfp.org

³ Institute of Advanced Research in Artificial Intelligence (IARAI), Vienna, Austria - daniel.springer@iarai.ac.at

KEY WORDS: Crop types, Self-Supervised Learning, Transformer, Random Forest, SimSiam, Remote Sensing

ABSTRACT:

Machine learning, satellites or local sensors are key factors for a sustainable and resource-saving optimisation of agriculture and proved its values for the management of agricultural land. Up to now, the main focus was on the enlargement of data which were evaluated by means of supervised learning methods. Nevertheless, the need for labels is also a limiting and time-consuming factor, while in contrast, ongoing technological development is already providing an ever-increasing amount of unlabeled data. Self-supervised learning (SSL) could overcome this limitation and incorporate existing unlabeled data. Therefore, a crop type data set was utilized to conduct experiments with SSL and compare it to supervised methods. A unique feature of our data set from 2016 to 2018 was a divergent climatological condition in 2018 that reduced yields and affected the spectral fingerprint of the plants. Our experiments focused on predicting 2018 using SLL without or a few labels to clarify whether new labels should be collected for an unknown year. Despite these challenging conditions, the results showed that SSL contributed to higher accuracies. We believe that the results will encourage further improvements in the field of precision farming, why the SSL framework and data will be published (Marszalek, 2021).

1. INTRODUCTION

Food sustainability is one of the grand challenges of the next decades, and a rigorous monitoring of the global food system is needed to allocate our resources (Fanzo et al., 2021). In particular, cropland use monitoring is essential, to assess the supply chain, but also evaluate the impact on natural ecosystems and keep track of the related hidden costs and subsidies (Rockström et al., 2020). Earth Observation (EO) is a global recurring proxy for land use monitoring, which is why it is widely used especially in areas where accessibility and infrastructure are a problem.

Today's opportunity is that there is much more image data available for feature extraction, even if not labeled. In recent years, new methods such as self-supervision have emerged that allow solid representations to be extracted in an unsupervised manner and could provide a more reliable representation for crop type mapping or other precision farming applications. Contrastive learning is a powerful approach to self-supervision, which aims to learn a representation where similar pairs of samples, such as time series of one crop type, are close to each other in the embedding space and different time series are far apart. In addition, only a few available labels do not impose a restriction as with supervised learning. Self-Supervised Learning (SSL) was successfully applied to detect changes (Leenstra et al., 2020), using techniques such as pretext tasks or augmentations to learn an invariant representation (Güldenring and Nalpantidis, 2021). Baevski et al. (2022) applied self-supervision for several tasks (NLP, speech, computer vision) with the objective to apply one augmentation suitable for different domains. The core idea was to mask a part of the input instead of using augmentations such as rotation or color distortion which are only suitable for certain

use cases. The increasing number of self-supervised learning methods differ mainly in terms of loss function, augmentation or architecture, with the choice of the underlying encoder playing an important role. In this work, we used a transformer (TF) as an encoder, which was verified in previous studies in a supervised manner (Rußwurm et al., 2020). SimCLR was one of the first architectures proposed where augmentation and the use of positive and negative pairs is an important property. Another similar example is MOCO which uses a memory bank in addition to negative and positive pairs. An overview as well as comparison of these siamese networks was presented in (Chen and He, 2021). We build on the recent SimSiam method, which combines a dual-stream siamese network with various data augmentations on positive pairs of input data (Chen and He, 2021). SimSiam also achieved promising results with a small batch size and a small number of epochs.

Self-supervised learning has rarely been used in precision agriculture although it could make an important contribution. For instance, for field-level yield prediction, there are typically very few labels available, which is a limitation for most supervised learning methods. In addition, adapting to a new region or a year with different climatic conditions reduces the effort required to record new labels. In fact, plant morphology varies in different climates, and plant growth and spectral response may change from year to year, due to climatic variations or agricultural practices. Models trained on specific imagery and crop type may therefore not transfer optimally to new regions or future years (Belgiu and Csillik, 2018). From a machine learning perspective, this is framed as domain adaptation. In this context, transfer learning solved similar problems by means of learned knowledge. It used other data sources such as ImageNet to pre-train a neural network and then transfer it to a downstream task (Nowakowski et al., 2021; Lucas et al., 2021).

* Corresponding author

Castillo-Navarro et al. (2022) experimented with out-of distribution data and confirmed the power of self-supervised learning with very few labels. Nyborg et al. (2022) used Thermal Positional Encoding (TPE) for attention-based crop classifiers to classify crop types in different regions in Europe and, most importantly, to reduce the effects of climate on the spectral responses of plants. Orynbaikyzy et al. (2022) identified appropriate Sentinel-1 and Sentinel-2 features to fit a new target region. Reducing the features by eliminating the weather-dependent bands likewise improved the results in our experiments. Another initial application of SSL in the agricultural sector outperformed deep learning approaches with a limited number of labels (Güldenring and Nalpantidis, 2021). Agastya et al. (2021) successfully applied self-supervised learning for irrigation detection. These promising approaches will save considerable time and reduce the need for labels.

In this study, SimSiam was applied with and without augmentation. In one experiment, we omitted augmentation and aimed to bring in our existing labels to learn an invariant representation per crop type. Data from previous years already include several variations of time series for each crop type, covering not only small climatological but also soil-related differences. This provided an additional advantage because, unlike normal augmentation with 1D time series, the risk of shifting to another crop type is reduced. Dwibedi et al. (2021) followed a similar hypothesis and added nearest neighbors from the data set to find additional positive pairs, assuming that more similar variations would be found this way.

The objectives of this research were as follows:

1. A comparison of supervised learning with self-supervised learning.
2. Experiments with/without augmentation to assess their impact on SSL performance.
3. Analysis to what extent SSL is suitable for the prediction of unknown or deviating years. We assume that this also paves the way for the prediction of crop types in new regions.

2. MATERIALS AND METHODS

2.1 Study site

The crop mapping task was evaluated on a data set which included the main crop types (corn, winter wheat, winter barley, winter rapeseed, sugar beet, and potato) in Upper Bavaria (Germany), collected for the years 2016, 2017 and 2018 (Figure 1). It is part of a larger collection, assembled and partly self-created, which included crop types and yields for different regions (Marszalek, 2021). The climatological data in Table 1 provided better insights into the various climatological conditions.

	2016	2017	2018
Crop samples	558	600	600
Mean temperature (°C)	12.3	13.0	13.7
Mean precipitation (mm)	308.5	292.9	315.4

Table 1. Data overview with average temperature and accumulated precipitation amount per year



Figure 1. Experimental site in WGS84 (EPSG:4326). All crop type samples were randomly selected in the region of Upper Bavaria. Two weather stations (in orange) were used to generate the climatological overview.




Figure 2. Data on average temperature for each year and an overview of monthly precipitation. Climatological data from March to July were used for the visualisation.

The corresponding climatological overview was visualised in Figure 2. The average temperature was compared for each year and showed that 2018 was 1.4 degrees different from 2016. 2016 and 2017 were more similar in terms of temperature trend and precipitation. Figure 2 indicates that precipitation in spring 2018 was below normal, which affected crop yields and plant growth. It was visualised as a monthly mean value. Here, climatological data from several fields near two weather stations (highlighted in orange in Fig. 2) were averaged to provide an overview of the period March through July. While this does not reflect the entire region of Upper Bavaria, it is intended to provide an overview of the climatological conditions per year. The crop samples were randomly selected from all over Upper Bavaria.

2.2 Sentinel-2

Sentinel-2 observations were used for the experiments with the self-supervised framework. With a repeat cycle of 2-5 days and its multispectral data in e.g. near infrared (NIR) and short-wave infrared (SWIR), these data allow various applications in agriculture. One drawback is the impact of clouds because, unlike SAR-based missions, clouds limit Sentinel-2's data availability. Nevertheless, the high temporal availability allowed the collection of sufficient data over time in the test region. Initially, all 13 bands were used in this study. A more detailed band description with properties can be found in (GEE, 2012) or (Marszalek et al., 2020). For this purpose, Level-1C data were downloaded from Google Earth Engine (GEE) (Google, 2020), but not further processed to bottom-of-atmosphere (BOA) reflection information since the focus was the comparison of supervised learning with self-supervised learning. For each field a georeferenced polygon (StMELF, 2020) was applied for the collection of corresponding field images. The original time series from February to the end of August were interpolated with a resolution of 2 weeks, resulting in 14 time steps for each field. Thus, a sample was represented by 13 band time series with 14 time steps, with each band averaged over the field polygon. For our experiments clouds were partially filtered out, but only in a simplified form, which increased the noise in the time series. GEE provides the ability to filter a scene based on metadata such as "CLOUD_COVER", but this did not remove the field-level clouds because we used a cloud cover of 20% for each scene. For 2016, several time series were downloaded incorrectly, which is reflected in the original data set in non-changing values over time. These data were also removed in the pre-processing, which is why slightly fewer samples are available for 2016.

2.3 Methods

Random Forest (RF) (Biau and Scornet, 2015) and a transformer (TF) (Vaswani et al., 2017) were used for the supervised crop type classification task in this study. Overall Accuracy (OA) was used as a measure to evaluate the experiments.

The semi-supervised crop type classification follows a two-fold training regime. In the initial phase (pre-training), a SimSiam architecture with a transformer as encoder (backbone) is trained in such a way that the negative cosine similarity (symmetrized loss) is minimized in order to maximize the similarity between the augmented sample pairs x_1 and x_2 . The latent representation encoded by the transformer backbone is then used in a supervised classification task (fine-tuning). Here, a linear layer was added as a classification head on the pre-trained backbone. The SimSiam implementation is based on the original work (Chen and He, 2021). We conducted experiments without and with a few labels of the target year. An important aspect was to verify that the embeddings did not collapse during pre-training. Chen and He (2021) suggested that the l_2 -normalized outputs during training should be close to $1/\sqrt{dim}$, where dim is the output dimension.

3. EXPERIMENTS

3.1 Supervised

The supervised training for RF and the TF was performed in four different scenarios, distinguishable by the data used for

training and testing. In the first run the entire data (2016-2018) was used in the train (75%) and test (25%) split (E1). The second training was performed with data from 2016 and 2017 only (E2), while data from 2018 was only used in the test set. The training data of the third (E3) and fourth experiment (E4) is again based on the years 2016 and 2017 but was enriched with 5% and 10% of samples from 2018 respectively. The first experiment was expected to outperform the other three experiments because it was trained with data from three different years and thus provided the most information about atmospheric and weather patterns. Experiment E2 is the most realistic approach, to illustrate the performance of the framework with respect to unseen data from an entire year.

TF was initialised with 4 self-attention heads and 3 layers. All experiments were performed with a learning rate of 0.0016612, 256 as batch size and 300 epochs. As the experiment progressed, bands and time steps were reduced to minimise noise.

3.2 Self/Semi-supervised

Augmentation is a key factor not only for SSL techniques to learn an invariant representation but also to increase the training data set. In the present work we used three augmentations, hereafter referred to as Aug1, Aug2, and Aug3, with Aug1 in particular following a new approach.

Aug1 draws original (non-augmented) pairs of samples from the same crop type and treats them as augmentations of each other. In this way an invariant representation for each crop type is learned. If the pairs are from different years, the model may generalise with respect to different weather patterns.

Aug2 randomly applied noise or drift to the original time series x_1 to generate x_2 . We used a Python framework for time series augmentation (tsaug, 2020) and parameterised the drift with $max_drift = 0.1$ and $n_drift_points = 2$. The noise was initialised with $scale = 0.02$.

Aug3 is a combination of Aug1 with cloud noise simulation. A constant factor (e.g. 7000) was added to all band values for a randomly chosen time step to simulate clouds and learn an invariant representation. Both x_1 and x_2 were selected from the labeled data set per crop type before random noise was added.

For each augmentation we repeated the same four experiments as discussed in section 3.1. The impact of the different augmentations was analysed in terms of embedded clusters after pre-training and fine-tuning. The pre-training is based on 600 epochs and a learning rate of 0.0016612. SimSiam was initialised with 64 features (d_model), 0.0005 weight decay and 0.9 as momentum. The hidden dimension in the projection and prediction head were initialised with 6 and 14 as the output dimension respectively. The augmentation experiments were performed with 11 and 14 time steps to learn an invariant representation.

3.2.1 Classification by comparison of the feature representation

While the pre-training aimed to provide a robust feature encoder for the classification task, it is also possible to use this framework directly as a classification tool. To illustrate the performance of the pre-training, we trained SimSiam according to E1 and E2 and used 2018 data for evaluation. For the classification (\mathcal{P}) of a single sample from 2018 (x_1) we detected the minimum of the averaged loss \mathcal{C} between x_1 and all N_c samples

of each crop type c from 2016 and 2017 (x_2) according to

$$P = \min(\mathcal{L}_c) \quad \text{with} \quad \mathcal{L}_c = \frac{1}{N_c} \sum_n \mathcal{C}(x_1, x_2^n). \quad (1)$$

The incorrectly predicted crop types were recorded.

3.3 Results

3.3.1 Supervised RF and TF were trained with all 13 bands and time steps as a first test. Table 2 summarises the four experiments repeated throughout the paper. RF improved as the amount of data from the target year increased. In contrast, the transformer had difficulties with this noisy data set which is why the overall accuracy was reduced by about 40% for 2018 without labels.

Method	E1 (OA)	E2 (OA)	E3 (OA)	E4 (OA)
RF	0.92	0.80	0.84	0.87
TF	0.93	0.54	0.86	0.62

Table 2. Supervised benchmark with all bands and time steps: Overall accuracy (OA) with all data (E1); OA with 0% labels for 2018 (E2); with 5% labels (E3) and 10% for 2018 (E4).

As a next step, bands 1, 2, 3 and 10 were removed because in particular bands 1 and 10 are strongly influenced by aerosols and clouds. Table 3 provides an overview about the results with B4, B5, B6, B7, B8, B8A, B9, B11 and B12.

Method	E1 (OA)	E2 (OA)	E3 (OA)	E4 (OA)
RF	0.93	0.85	0.87	0.88
TF	0.93	0.67	0.88	0.75

Table 3. Supervised benchmark - 14 time steps and 9 bands: Overall accuracy (OA) with all data; OA with 0% labels for 2018; with 5% labels and 10% for 2018.

Removal of these bands stabilised the prediction of crop types in 2018 for E3 and E4. Tables 2 and 3 show that the results for E4 were worse compared to E3 for RF or TF. The last optimisation step visualised in Table 4 removed the first observations in the time series. An investigation of the time series showed that the signal of the first time steps in 2018 were contaminated by clouds.

Method	E1 (OA)	E2 (OA)	E3 (OA)	E4 (OA)
RF	0.93	0.85	0.88	0.88
TF	0.93	0.86	0.85	0.91

Table 4. Supervised benchmark - 11 time steps and 9 bands: Overall accuracy (OA) with all data; OA with 0% labels for 2018; with 5% labels and 10% for 2018.

Therefore, the time series were shortened by 3 time steps resulting in an observation period from the end of March to the end of August. This significantly affected the accuracy and improved the results for predictions where either 10% or no samples of 2018 were included in the training. While RF has dominated so far, transformer’s predictive capability was emerging here.

3.3.2 Self/Semi-Supervised The following results are based on 14 and 11 time steps. As shown in Section 3.3.2, the 14 time steps contained cloud noise in the first few weeks, but this is useful for learning a noise-invariant representation. Table 5 provides an overview with Aug1 and Aug2 and 14 time steps. Table 6 focused on Aug1 and Aug3 with 11 time steps and without cloud noise. Comparing the results in Table 3 and Table 5 shows that Aug3 improved the accuracies. The combination of crop invariance and noise in Aug3 also increased the prediction results for E2 to 0.71 OA and E4 to 0.87 OA. For E4, the highest improvement of 16% was achieved here. Aug1 has only one added value with 10% samples. The noise-invariant representation also stabilises the results with 10% no longer collapsing as in Table 3 or Table 2. Nevertheless, the robustness of RF proved to be evident as well.

14 ts	E1 (OA)	E2 (OA)	E3 (OA)	E4 (OA)
Aug1	0.92	0.54	0.68	0.85
Aug3	0.93	0.71	0.87	0.87

Table 5. SSL with 14 time steps: Overall accuracy (OA) with all data (E1); OA with 0% labels for 2018 (E2); with 5% labels (E3) and 10% for 2018 (E4).

In contrast, Table 6 addressed the results with 11 time steps for a growing season from the end of March to the end of August. Thus, noise caused by clouds was removed in the first 3 time steps in 2018 and changed the presentation of the results. Both, Aug1 and Aug2 performed well in this scenario. Especially the test with 5% labels reached an accuracy of 0.9 OA compared to 0.85 OA supervised.

11 ts	E1 (OA)	E2 (OA)	E3 (OA)	E4 (OA)
Aug1	0.90	0.85	0.89	0.89
Aug2	0.93	0.81	0.90	0.84

Table 6. SSL with 11 time steps: Overall accuracy (OA) with all data; OA with 0% labels for 2018; with 5% labels and 10% for 2018.

One observation with Aug1 was the lower accuracy using all data compared with supervised TF (0.9 vs 0.93 OA). Looking at the embeddings visualised with PCA, it is evident that winter wheat and winter barley were difficult to separate. While fine-tuning thereafter separated the two classes well, supervised learning was able to distinguish these two classes more clearly. Figure 3 visualises examples with raw data, pre-training and fine-tuning. Figure 4 shows the collapse level of embeddings during the training. Our observations confirmed that convergence near $1/\sqrt{dim}$ is important to assess whether the embeddings collapse in the pre-training. Evaluation based on training loss alone was not always sufficient. We sometimes observed a collapse of the embeddings during pre-training, which prevented fine-tuning from improving accuracy. Therefore, it was important to observe the collapse level from 300 epochs.

Intermediate evaluation of the backbone

The direct evaluation of the SimSiam network without using a dedicated classification layer, can be understood as an intermediate evaluation of the quality of the latent representation. Table 7 shows the crop-wise prediction accuracy of the SimSiam network when evaluated according to equation (1).

In both E1 and E2, we observed a decrease in accuracy for




Figure 3. Examples of embeddings generated with principal component analysis (PCA). The six different colors represent the crop types. A) gives an overview of the raw data. B) shows the result after pre-training with SimSiam and Aug1. Here it is already possible to differ between clusters. C) shows the last step and the embeddings after fine-tuning with labels.




Figure 4. Standard deviation of the l_2 -normalized output and training loss.

11 steps	c_1	c_2	c_3	c_4	c_5	c_6
E1	0.94	0.84	0.97	0.97	0.96	1.00
E2	0.94	0.53	0.85	0.64	0.85	0.72

Table 7. Classification accuracy resolved by crop type based on 11 time steps for E1 and E2.

crop type 2 (winter barley), which was more pronounced in E2. Since the evaluation for E1 is based on data that included samples in training, a decline in a single crop type may indicate a systematic problem either in the framework or in the data themselves. It is hence of interest to resolve the actual predictions when the network is confronted with crop type 2 samples from 2018 (Table 8).

$\mathcal{P}(c_2)$	c_1	c_2	c_3	c_3	c_5	c_6
E1	0	84	1	5	10	0
E2	1	53	2	3	41	0

Table 8. Predictions for all 100 crop type 2 samples from 2018 (correct predictions in bold numbers).

We have noticed a systematic tendency for crop type 2 to be incorrectly identified as crop type 5 (winter wheat). This suggested that crop type 2 and 5 cannot be sufficiently separated by the framework.

4. DISCUSSION & CONCLUSION

Methods in the field of supervised learning enabled several applications for precision farming (Yang et al., 2019; Waldner and Diakogiannis, 2019; Mazzia et al., 2019). However, supervised learning is in general associated with a large effort since labeled data are required. So far, few experiments were performed with SSL for precision agriculture. In our work, we explored the potential to use SSL for domain adaptation in crop classification.

Supervised learning achieved very good results through noise removal when trained on bands B4, B5, B6, B7, B8, B8A, B9, B11 and B12. Excluding samples from the target year, an overall accuracy of 0.86 was achieved. This resulted in a decrease of about 8% compared to the test with all the data. We hypothesize that the difference is mainly due to climatological variations (Marszalek et al., 2022). The transformer had difficulties with noise caused by clouds and improved significantly with the reduction of bands and noisy time steps. RF showed robust performance in all tests regardless of noise.

Pre-training was used to learn robust representations for each crop type. Augmentation is an important part of SSL and crucial for performance. The introduction of 5% samples from the target year improved accuracies and outperformed results using supervised learning. While previous attempts were mainly based on typical augmentations such as rotation or jitter (Güldenring and Nalpantidis, 2021; Chen and He, 2021), we took a different approach with Aug1. Learning a crop-invariant representation separated most crop types in the embedding space. Only winter wheat and winter barley were not separable. By directly evaluating the backbone, we showed that these two crop types could not be well distinguished by the embeddings. Therefore, we hypothesize that this is the reason why pre-training combined with fine-tuning led to slightly worse results than supervised training in some cases. With 10% samples we recommend to proceed supervised because here already OA values of around 0.9 were achieved. Aug3 has proven itself in the test with the noise caused by clouds. Here, a significant increase in accuracy was achieved. Successful denoising has also been confirmed in previous studies (Dalsasso et al., 2021). The choice of augmentation provides several opportunities to learn a better representation that could also be invari-

ant to regional differences. For example, Nyborg et al. (2021) introduced the temporal shift of crop time series in different regions which inspired us to use a drift (Aug2) as augmentation. However, this was only partially successful in learning a shift-invariant representation. Further experiments will be necessary to find a domain-invariant representation for crop types. We also plan to further develop SimSiam based on the lessons learned to more efficiently incorporate unlabeled data into training, which could be important for domain adaptation without labels. Dwibedi et al. (2021) and Zell et al. (2022) confirmed that incorporation of unlabeled data improved SSL.

The **suitability of SSL** was considered for crop classification in this study. We showed that it is essential to introduce few labels for a new year. Although the prediction of crop types in a deviating year was possible, we would like to emphasise that the collection of few samples improved the accuracy and reduced uncertainties. The study targeted crop classification, but we also see potential for other applications in precision agriculture (e.g. yield prediction). We showed that the pre-training improved the results and that a costly collection of labels for a region with already existing data from previous years is only partly necessary. Originally, we assumed that the influence of climatological conditions would be much more significant. Nevertheless, supervised learning and SSL provided good results. Although e.g. a pronounced drought could strongly influence the spectral response of plants, one can question the influence of climate for the classification of crops in this study. The advantage of this test was that we used data from one region and thus minimised the influence of soil and fertilisation. We assumed that farming practices were similar over the years. However, for a change of the target region, these two factors need to be taken into account as they are likely to have an impact on the spectral response. While this study looked at domain adaptation, other studies have also identified the benefits for agriculture (Güldenring and Nalpantidis, 2021; Agastya et al., 2021). In general, this study can be an argument for exploring other applications of SSL where data availability is a limitation.

ACKNOWLEDGEMENT

We would like to thank the LIGHTLY team for the support during this research project and the Institute of Advanced Research in Artificial Intelligence (IARAI) for providing the computational resources. We would also like to thank Maximilian Lössch for the crop type data set.

References

- Agastya, C., Ghebremusse, S., Anderson, I., Reed, C., Vahabi, H., Todeschini, A., 2021. Self-supervised Contrastive Learning for Irrigation Detection in Satellite Imagery. *ArXiv*, abs/2108.05484.
- Baevski, A., Hsu, W.-N., Xu, Q., Babu, A., Gu, J., Auli, M., 2022. data2vec: A general framework for self-supervised learning in speech, vision and language.
- Belgiu, M., Csillik, O., 2018. Sentinel-2 cropland mapping using pixel-based and object-based time-weighted dynamic time warping analysis. *Remote Sensing of Environment*, 204, 509–523.
- Biau, G., Scornet, E., 2015. A random forest guided tour. *TEST*, 25, 197–227.
- Castillo-Navarro, J., Le Saux, B., Boulch, A., Lefèvre, S., 2022. Energy-Based Models in Earth Observation: From Generation to Semisupervised Learning. *IEEE Transactions on Geoscience and Remote Sensing*, 60, 1–11.
- Chen, X., He, K., 2021. Exploring simple siamese representation learning. *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 15750–15758.
- Dalssasso, E., Denis, L., Tupin, F., 2021. SAR2SAR: A Semi-Supervised Despeckling Algorithm for SAR Images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14, 4321–4329.
- Dwibedi, D., Aytar, Y., Tompson, J., Sermanet, P., Zisserman, A., 2021. With a little help from my friends: Nearest-neighbor contrastive learning of visual representations.
- Fanzo, J., Haddad, L., Schneider, K. R. et al., 2021. Viewpoint: Rigorous monitoring is necessary to guide food system transformation in the countdown to the 2030 global goals. *Food Policy*, 104, 102163.
- GEE, 2012. Spectral bands. https://developers.google.com/earth-engine/datasets/catalog/COPERNICUS_S2. (Accessed on: 21.07.2020).
- Google, 2020. GEE data catalog. <https://earthengine.google.com/>. (Accessed on: 21.07.2020).
- Güldenring, R., Nalpantidis, L., 2021. Self-supervised contrastive learning on agricultural images. *Computers and Electronics in Agriculture*, 191, 106510.
- Leenstra, M., Marcos, D., Bovolo, F., Tuia, D., 2020. Self-supervised pre-training enhances change detection in sentinel-2 imagery. *ICPR Workshops*.
- Lucas, B., Pelletier, C., Schmidt, D., Webb, G., Petitjean, F., 2021. A Bayesian-inspired, deep learning-based, semi-supervised domain adaptation technique for land cover mapping. *Machine Learning*, 1–33.
- Marszalek, M., 2021. Crop type and yield data. <https://github.com/ESA-PhiLab/WorldCrops>. (Accessed on: 30.03.2022).
- Marszalek, M., Körner, M., Schmidhalter, U., 2022. Prediction of multi-year winter wheat yields at the field level with satellite and climatological data. *Computers and Electronics in Agriculture*, 194, 106777.
- Marszalek, M., Lössch, M., Körner, M., Schmidhalter, U., 2020. Early Crop-Type Mapping Under Climate Anomalies. *Preprints*.
- Mazzia, V., Khaliq, A., Chiaberge, M., 2019. Improvement in Land Cover and Crop Classification based on Temporal Features Learning from Sentinel-2 Data Using Recurrent-Convolutional Neural Network (R-CNN). *Applied Sciences*, 10, 238.
- Nowakowski, A., Spiller, D., Cremer, N., Bonifacio, R., Marszalek, M., Garcia-Herranz, M., Mathieu, P. P., Kim, D.-H., 2021. Ai opportunities and challenges for crop type mapping using Sentinel-2 and drone data. *2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*, 258–261.

- Nyborg, J., Pelletier, C., Assent, I., 2022. Generalized classification of satellite image time series with thermal positional encoding.
- Nyborg, J., Pelletier, C., Lefèvre, S., Assent, I., 2021. Timematch: Unsupervised cross-region adaptation by temporal shift estimation.
- Orynbaikyzy, A., Gessner, U., Conrad, C., 2022. Spatial Transferability of Random Forest Models for Crop Type Classification Using Sentinel-1 and Sentinel-2. *Remote Sensing*, 14.
- Rockström, J., Edenhofer, O., Gaertner, J., DeClerck, F., 2020. Planet-proofing the global food system. *Nature Food*.
- Rußwurm, M., Pelletier, C., Zollner, M., Lefèvre, S., Körner, M., 2020. BreizhCrops: A Time Series Dataset for Crop Type Mapping. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences ISPRS (2020)*.
- StMELF, 2020. Bavarian state ministry of food, agriculture and forestry. <https://www.stmelf.bayern.de>. [Online; accessed 21-July-2020].
- tsaug, 2020. Time series augmentation. <https://github.com/arundo/tsaug>. (Accessed on: 28.03.2022).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L. u., Polosukhin, I., 2017. Attention is all you need. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (eds), *Advances in Neural Information Processing Systems*, 30, Curran Associates, Inc.
- Waldner, F., Diakogiannis, F., 2019. Deep learning on edge: Extracting field boundaries from satellite images with a convolutional neural network.
- Yang, Q., Shi, L., Han, J., Zha, Y., Zhu, P., 2019. Deep convolutional neural networks for rice grain yield estimation at the ripening stage using UAV-based remotely sensed images. *Field Crops Research*, 235, 142-153.
- Zell, A., Sumbul, G., Demir, B., 2022. Deep metric learning-based semi-supervised regression with alternate learning.