EGU24-3127, updated on 23 Jan 2025
https://doi.org/10.5194/egusphere-egu24-3127
EGU General Assembly 2024

Cracking Ground Truth Barriers: Harnessing the Power of Transfer Learning for Crop Mapping

Miloš Pandžić[1], Dejan Pavlović[1], Sanja Brdar[1], Milan Kilibarda[2], and Oskar Marko[1]
[1]BioSense Institute, University of Novi Sad, Novi Sad, Serbia
[2]University of Belgrade, Faculty of Civil Engineering, Belgrade, Serbia

Transfer learning (TL) is rapidly gaining popularity in recent years in various research disciplines due to its practicality, temperate need for resources, and often quite promising results. The same principle applies for Earth observation, especially for tasks such as crop mapping where TL already showed a certain potential. Our focus in this research was on temporal transfer learning for a single agricultural region. In our study we built an initial CNN-1D crop mapping model for Vojvodina province, Serbia, using SAR satellite imagery and ground truth (GT) data collected for 2017-2020. We did it using a leave-one-year-out approach where each year served only once as a validation dataset. The top-performing model was further employed for transfer learning analysis, utilising a limited set of target season ground truth data. The aim was to diminish reliance on labour-intensive and time-consuming large-scale ground truth data collection, typically carried out through hands-on field inspections. Instead of collecting it all over Vojvodina for the 2021 season, we tried to focus on a limited area around the departure point. Three options were analysed, i.e., approximately 20, 25, and 30 km radius around the departure point for which the province capital Novi Sad was taken. From the practical standpoint, labels of these parcels are easier to record than those more distant (distributed), so it seems reasonable to visit only these locations to reduce the costs of ground truth collection. Visited parcels that fell within these radiuses served for retraining the model, and the remaining parcels (those outside 30 km radius) served for testing and accuracy assessment. For each parcel 50 randomly selected pixels were used for the analysis. After 5 retraining cycles, the average F1 score for transfer learning approach of the CNN-1D model for 20, 25 and 30 km buffer zones was 74%, 79% and 83%, respectively. Training the same CNN-1D model from scratch reached 69%, 73% and 78% respectively, i.e., approximately 5% lower score on average. Inferencing using the pre-trained model as such (without adaptations) achieved F1 score of 78%, which set 20 km radius case to the irrational use of TL, while the use of other two buffer areas were justified as they achieved comparably better results. Also, three buffer cases achieved between 3% and 9% lower F1 score than their respective pairs when the same number of parcels used for retraining were randomly distributed all over the test area. This was likely in a relationship with the restricted sampling region characteristics (uniform soil type, management practice, weather conditions) and distribution of classes in that region which may not have properly represented the entire test area. In addition, the comparison of these two approaches showed that adding more samples for retraining scaled down the difference. The viability of the presented approach was confirmed within the experiment and, therefore, practitioners are

encouraged to weigh the trade-off between practicality and accuracy in their future work.