
Towards Global Crop Maps with Transfer Learning

Hyun-Woo Jo ^{*1} **Alkiviadis Koukos** ^{*2} **Vasileios Sitokonstantinou**²

Woo-Kyun Lee¹ **Charalampos Kontoes**²

¹ Department of Environmental Science and Ecological Engineering, Korea University

²BEYOND Centre, IAASARS, National Observatory of Athens

{akoukos, vsito, kontoes}@noa.gr

endeavor4a1@gmail.com

leewk@korea.ac.kr

Abstract

The continuous increase in global population and the impact of climate change on crop production are expected to affect the food sector significantly. In this context, there is need for timely, large-scale and precise mapping of crops for evidence-based decision making. A key enabler towards this direction are new satellite missions that freely offer big remote sensing data of high spatio-temporal resolution and global coverage. During the previous decade and because of this surge of big Earth observations, deep learning methods have dominated the remote sensing and crop mapping literature. Nevertheless, deep learning models require large amounts of annotated data that are scarce and hard-to-acquire. To address this problem, transfer learning methods can be used to exploit available annotations and enable crop mapping for other regions, crop types and years of inspection. In this work, we have developed and trained a deep learning model for paddy rice detection in South Korea using Sentinel-1 VH time-series. We then fine-tune the model for i) paddy rice detection in France and Spain and ii) barley detection in the Netherlands. Additionally, we propose a modification in the pre-trained weights in order to incorporate extra input features (Sentinel-1 VV). Our approach shows excellent performance when transferring in different areas for the same crop type and rather promising results when transferring in a different area and crop type.

1 Introduction

Food security, but also social and economic development are at high risk due to the population growth and climate change and the pressure they put on agriculture. Several recent studies indicate that the changes in climate cause substantial yield losses at the global level [20, 19]. At the same time, the production should be increased by 24% until 2030 - compared to 2022 - to achieve zero hunger, while reducing emissions by 6% [14]. Apart from climate-friendly policies and practices to ensure and promote the increase of agricultural productivity [11], there is also a need of global, timely and precise crop type mapping systems to assist in the monitoring and management of agricultural fields [17], the prediction of crop production [21, 9] and the effective spatial allocation of the agricultural practices [5].

Earth Observation (EO) data have been extensively used to train Machine Learning (ML) and Deep Learning (DL) models to produce crop maps [4, 18, 22, 2, 3, 24, 7]. Nevertheless, most approaches require large labeled datasets for training [16, 10, 23]. In reality, ground samples that capture spatio-temporal differences of crops worldwide are scarce and this remains one of the main barriers for global applications [8], limiting most studies to small or homogeneous areas. The most trustworthy

^{*}Equal contribution.

way to acquire such ground observations is field surveys that are time-consuming, expensive and cannot cover every part of the world (e.g. inaccessible/remote areas). Transfer Learning (TL) [15] has been successfully applied to overcome this issue, by improving the learning performance on reduced datasets while decreasing the computational complexity [26]. To this direction, a handful of works have been published the past couple of years that apply TL for crop classification using EO data [6, 1, 13, 25].

In this paper, we apply TL on paddy rice mapping using only Sentinel-1 time series, by transferring knowledge from South Korea to European areas, for which there is a small amount of ground samples. We implement and apply a recurrent U-net model to detect paddy in South Korea using Sentinel-1 VH backscatter time-series as input. Then, we explore the capability of transferring the knowledge captured from the paddy rice model to effectively predict i) paddy rice in Spain and France and ii) summer barley in the Netherlands. Moreover, we explore fine-tuning the model by augmenting the input space with Sentinel-1 VV backscatter coefficients.

2 Data & Problem Formulation

Data. We used Synthetic Aperture Radar (SAR) Sentinel-1 data and computed the 20-day-mean backscattering coefficient (VH|VV) for each pixel for a time-series throughout the growing period. Then we extracted patches of 256x256 pixels for the areas of interest, using Google Earth Engine. All input data were scaled using max-min normalization. The datasets of rice in South Korea (2017-2019), Spain (2021), and France (2020) consist of 12,942, 88, and 134 patches, respectively. The dataset of summer barley in Netherlands (2018) includes 2,280 patches, however with few barley pixels in the patches. In each dataset, 60% of randomly selected patches were used for training and the rest 40% were used for testing.

Problem Formulation. The time series of EO data are denoted by $x_{r,t}^s$ and the annotated crop data by y_r^c where r, t, s, c represent region, time, feature, and crop type, respectively. In this study, the crop mapping was performed with a recurrent U-net model (h) exploiting time-series during the growing period:

$$\hat{y}_{r,t} = h(x_{r,1}, x_{r,2}, \dots, x_{r,t}) \quad (1)$$

where $t = 1 \dots 8$ indicates the relative time instance, i.e., 20-day feature vector, within the cultivation period. The model was pre-trained (h^p) to classify paddy rice in South Korea by using Sentinel-1 VH backscatter, in patches extended all over the country and for the years of 2017-2019. However, unlike South Korea where rice is a staple food, the number of x in the other cases of r, c hardly suffice to efficiently train deep neural networks.

3 Methodology

3D Recurrent U-net. A custom recurrent U-Net (Fig. 2) was designed to exploit both spatial and temporal context of the EO time-series in order to produce timely paddy rice segmentation maps. The model follows a standard U-Net architecture; the encoder consists of a series of recurrent modules including convolutional layers, drop-out, and spatial max-pooling. In the recurrent module, each time step shares the convolution layers, and the weighted output of the previous time step is added to the next time step; thus, the phenological context can be passed to the later calculation. Considering both the preservation of temporal features and computational efficiency, the max-pooling layers at the skip connections were applied to the time axis so that the adjacent time steps at the same developing phase were pooled, and half the size of the features were passed to the decoder.

Transfer Learning. We implement different scenarios of TL to identify an optimal application according to data availability and similarity. As a baseline, we transferred only the architecture with randomly initialized weights (RI). The others include initializing with the pre-trained weights and then fine-tuning ($f^{r,c,s}(\cdot)$) to adapt to the target r, c, s . Considering that the model (h) consists of an encoder (h_E), which extracts the crop's phenological characteristics, and it is followed by a decoder (h_D), the applications were to fine-tune the entire networks (FT), fine-tune only h_E while freezing h_D (FT_E) or fine-tune h_D while freezing h_E (FT_D).

$$RI = f(h_E \cdot h_D) \quad FT = f(h_E^p \cdot h_D^p) \quad FT_E = f(h_E^p) \cdot h_D^p \quad FT_D = h_E^p \cdot f(h_D^p) \quad (2)$$

Incorporation of additional feature types. In crop classification, diverse characteristics of each crop (e.g., texture, reflection) raise the need of an extended application of TL, such as using different sources of data as input. In this direction, we adapt h^p , pre-trained on Sentinel-1 VH backscatter, to take as input both Sentinel-1 VH and VV features. To do this, the pre-trained weights at the first layer of the encoder ($W_{E_0}^P$) are divided by the total number of input layers (Eq.3). Therefore, a similar scale of signal intensity is transferred to the activation functions (σ) that is invariant to the number of inputs, and ensures that h^p maintains the trained feature extraction process.

$$h_{E_0}^P = \sigma((W_{E_0}^P \cdot x^s + W_{E_0}^P \cdot x^{s'})/2 + b) \quad (3)$$

4 Experiments and Results

Experiments. We implemented 10 scenarios by combining different r, c, s , where fine-tuning was conducted in r_1 and performance was tested in r_2 . The main goal of this study is to investigate the effect of TL for the same target labels (e.g., paddy rice) in different areas. Therefore, we run the following experiments: r_1-r_2-c-s : 1) Spain-Spain-rice-VH, 2) Spain-Spain-rice-VH|VV, 3) Spain-France-rice-VH, 4) Spain-France-rice-VH|VV, 5) France-France-rice-VH, 6) France-France-rice-VH|VV, 7) France-Spain-rice-VH, 8) France-Spain-rice-VH|VV. Additionally, we explore the efficiency of TL in different regions and different crop types (summer barley in the Netherlands) to examine if the knowledge of paddy rice could contribute in mapping other crops; r_1-r_2-c-s : 9) Netherlands-Netherlands-summer barley-VH, 10) Netherlands-Netherlands-summer barley-VH|VV. It is worth mentioning that the data of each region have been acquired from different years, which makes the application of TL even more challenging.

Based on our experiments, we found FT_E achieved better performance than FT and FT_D . Fine-tuning the decoder did not converge, whereas by fine-tuning only its last (or 2-3 last) layers the model was successfully trained, but provided suboptimal performance. Table 1 presents the Intersection over Union (IoU) of the positive class, for the RI , FT and FT_E and the 10 different scenarios mentioned earlier. We also compare their performance against locally trained Random Forest (RF) models. Visual maps of predictions for each scenario and method can be found in the Appendix (Figures 6-13)

Table 1: Mean IoU for the different scenarios and methods

Fine-tuning		Spain				France				The Netherlands	
Test	Spain	France		France		Spain		The Netherlands			
Feature	VH	VH VV	VH	VH VV							
RF	0.87	0.90	0.63	0.66	0.76	0.84	0.77	0.78	0.26	0.40	
RI	0.86	0.69	0.52	0.36	0.76	0.74	0.70	0.73	0.31	0	
FT	0.89	0.90	0.57	0.63	0.82	0.83	0.82	0.83	0.40	0.45	
FT_E	0.90	0.90	0.63	0.66	0.86	0.86	0.79	0.84	0.42	0.54	

Using only the model’s architecture without transferring the pre-trained weights (RI), we observe a poor performance in most cases, which is even poorer in the case of VH|VV input. This can be explained by the fact that augmenting the input with more layers results in more parameters which, in combination with the few labels and no transferred knowledge, prevents the model from learning. On the other hand, fine-tuning the pre-trained U-net works very well for paddy rice mapping, both in the case of freezing the decoder and updating only the parameters of the encoder (FT_E) and in the case of updating the whole network’s weights (FT). As expected, when fine-tuning and testing in the same area the performance is better. But when we fine-tune in Spain and test in France, we notice a significant drop in the IoU.

Figure 1 shows the mean VH time-series of the True Positive (TP) and the False Negative (TN) of the predictions of the model fine-tuned in Spain and tested in France, together with the corresponding mean VH time-series of rice pixels in Spain. TP (i.e., correctly predicted rice pixels) have very similar VH signature with that of the rice in Spain, whereas for FN predictions (i.e., rice instances that the model failed to identify) backscatter coefficients differ significantly, as compared to both the TP and the Spain rice pixels.

It seems that in France there are two different types of rice, out of which only one shares the same growing phases with those in Spain, and therefore the model is not able to successfully predict both of them. Moreover, the incorporation of the VV feature, especially in France, enhances slightly the performance. This is a strong indicator that TL not only works with the augmentation of the input with more features, but it can also provide better results.

Transferring the paddy rice model to predict summer barley does not perform as well. Paddy rice fields are intentionally flooded at the start of the cultivation period; SAR data have a great ability of identifying water content, which makes them ideal in classifying paddy rice. However, this is not the case for summer barley, and thus the discrimination of it using SAR data is much more challenging. Nevertheless, using both VH and VV backscatter coefficients we are able to achieve an IoU of 0.54 - also recall, precision and f1-score of 0.7 (Tables 2, 3 and 4) - which is interestingly high given the nature of the problem. In this case, we notice not only an improvement by using the extra input of VV, but rather a significantly better performance by fine-tuning only the encoder.

It is also worth mentioning that RF also achieves great performance in almost every paddy rice experiment. However, it fails in predicting summer barley, even in the case of using both VV and VH, with an IoU of 0.4. Identifying paddy rice using EO data is not a particularly difficult problem, thanks to the flooding in the early vegetation period that was mentioned above. On the contrary, prediction of summer barley is a much more complex task, since it could share common phenological characteristics with other summer crop types (e.g., maize, summer wheat).

5 Conclusion

Precise, dynamic and detailed global crop type maps are essential for monitoring crop production that is under pressure. Such maps are powerful datasets that enable the timely identification of food security challenges and the large-scale, yet local-specific, rural planning to mitigate climate change. In this context, we propose a transfer learning method that leverages a pre-trained recurrent U-net model for paddy rice mapping in South Korea and fine-tunes it in other areas (France, Spain and the Netherlands) and/or crop types (paddy rice and summer barley) with a few available annotated data. TL for paddy rice mapping yielded excellent results both in Spain and in France. Based on our experiments, fine-tuning the encoder or the entire network provided the best performance, whereas fine-tuning the decoder did not converge. Additionally, the incorporation of an additional feature (i.e., VV backscatter coefficient) boosts the performance in almost every scenario. Finally, TL for barley in the Netherlands exhibits promising results, especially in the case of fine-tuning the encoder and incorporating the VV input, which outperforms significantly the corresponding RF model.

Acknowledgments and Disclosure of Funding

This work was supported by the International Research and Development Program of the National Research Foundation of Korea (NRF) funded by the Ministry of Science and ICT [2021K1A3A1A78097879] and by the CALLISTO project, funded by EU's Horizon 2020 research and innovation programme under grant agreement No. 101004152.

References

- [1] P. Bosilj, E. Aptoula, T. Duckett, and G. Cielniak. Transfer learning between crop types for semantic segmentation of crops versus weeds in precision agriculture. *Journal of Field Robotics*, 37(1):7–19, 2020.
- [2] P. Defourny, S. Bontemps, N. Bellemans, C. Cara, G. Dedieu, E. Guzzonato, O. Hagolle, J. Inglada, L. Nicola, T. Rabaute, et al. Near real-time agriculture monitoring at national scale at parcel resolution: Performance assessment of the sen2-agri automated system in various cropping systems around the world. *Remote sensing of environment*, 221:551–568, 2019.
- [3] Z. Du, J. Yang, C. Ou, and T. Zhang. Smallholder crop area mapped with a semantic segmentation deep learning method. *Remote Sensing*, 11(7):888, 2019.
- [4] R. d’Andrimont, A. Verhegghen, G. Lemoine, P. Kempeneers, M. Meroni, and M. Van der Velde. From parcel to continental scale—a first european crop type map based on sentinel-1 and lucas copernicus in-situ observations. *Remote sensing of environment*, 266:112708, 2021.
- [5] C. Folberth, N. Khabarov, J. Balkovič, R. Skalský, P. Visconti, P. Ciais, I. A. Janssens, J. Peñuelas, and M. Obersteiner. The global cropland-sparing potential of high-yield farming. *Nature Sustainability*, 3(4):281–289, 2020.
- [6] P. Hao, L. Di, C. Zhang, and L. Guo. Transfer learning for crop classification with cropland data layer data (cdl) as training samples. *Science of The Total Environment*, 733:138869, 2020.
- [7] H.-W. Jo, S. Lee, E. Park, C.-H. Lim, C. Song, H. Lee, Y. Ko, S. Cha, H. Yoon, and W.-K. Lee. Deep learning applications on multitemporal sar (sentinel-1) image classification using confined labeled data: The case of detecting rice paddy in south korea. *IEEE Transactions on Geoscience and Remote Sensing*, 58(11):7589–7601, 2020.
- [8] A. Kamaras and F. X. Prenafeta-Boldú. Deep learning in agriculture: A survey. *Computers and electronics in agriculture*, 147:70–90, 2018.
- [9] S. Khaki and L. Wang. Crop yield prediction using deep neural networks. *Frontiers in plant science*, 10:621, 2019.
- [10] R. M Rustowicz, R. Cheong, L. Wang, S. Ermon, M. Burke, and D. Lobell. Semantic segmentation of crop type in africa: A novel dataset and analysis of deep learning methods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 75–82, 2019.
- [11] G. S. Malhi, M. Kaur, and P. Kaushik. Impact of climate change on agriculture and its mitigation strategies: A review. *Sustainability*, 13(3):1318, 2021.
- [12] R. M. French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999.
- [13] A. Nowakowski, J. Mrziglod, D. Spiller, R. Bonifacio, I. Ferrari, P. P. Mathieu, M. Garcia-Herranz, and D.-H. Kim. Crop type mapping by using transfer learning. *International Journal of Applied Earth Observation and Geoinformation*, 98:102313, 2021.
- [14] OECD, Food, and A. O. of the United Nations. *OECD-FAO Agricultural Outlook 2022-2031*. 2022. doi: <https://doi.org/https://doi.org/10.1787/f1b0b29c-en>. URL <https://www.oecd-ilibrary.org/content/publication/f1b0b29c-en>.
- [15] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- [16] M. Rußwurm, C. Pelletier, M. Zollner, S. Lefèvre, and M. Körner. Breizhcrobs: A time series dataset for crop type mapping. *arXiv preprint arXiv:1905.11893*, 2019.
- [17] J. Segarra, M. L. Buchaillot, J. L. Araus, and S. C. Kefauver. Remote sensing for precision agriculture: Sentinel-2 improved features and applications. *Agronomy*, 10(5):641, 2020.
- [18] V. Sitokonstantinou, A. Koukos, T. Drivas, C. Kontoes, I. Papoutsis, and V. Karathanassi. A scalable machine learning pipeline for paddy rice classification using multi-temporal sentinel data. *Remote Sensing*, 13(9):1769, 2021.
- [19] B. Sultan, D. Defrance, and T. Iizumi. Evidence of crop production losses in west africa due to historical global warming in two crop models. *Scientific reports*, 9(1):1–15, 2019.

- [20] R. Tito, H. L. Vasconcelos, and K. J. Feeley. Global climate change increases risk of crop yield losses and food insecurity in the tropical andes. *Global Change Biology*, 24(2):e592–e602, 2018.
- [21] T. Van Klompenburg, A. Kassahun, and C. Catal. Crop yield prediction using machine learning: A systematic literature review. *Computers and Electronics in Agriculture*, 177:105709, 2020.
- [22] K. Van Tricht, A. Gobin, S. Gilliams, and I. Picard. Synergistic use of radar sentinel-1 and optical sentinel-2 imagery for crop mapping: A case study for belgium. *Remote Sensing*, 10(10):1642, 2018.
- [23] G. Weikmann, C. Paris, and L. Bruzzone. Timesen2crop: A million labeled samples dataset of sentinel 2 image time series for crop-type classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:4699–4708, 2021.
- [24] N. You, J. Dong, J. Huang, G. Du, G. Zhang, Y. He, T. Yang, Y. Di, and X. Xiao. The 10-m crop type maps in northeast china during 2017–2019. *Scientific data*, 8(1):1–11, 2021.
- [25] W. Zhang, H. Liu, W. Wu, L. Zhan, and J. Wei. Mapping rice paddy based on machine learning with sentinel-2 multi-temporal data: Model comparison and transferability. *Remote Sensing*, 12(10):1620, 2020.
- [26] F. Zhuang, Z. Qi, K. Duan, D. Xi, Y. Zhu, H. Zhu, H. Xiong, and Q. He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020.

A Supplemental Material

A.1 Model Details

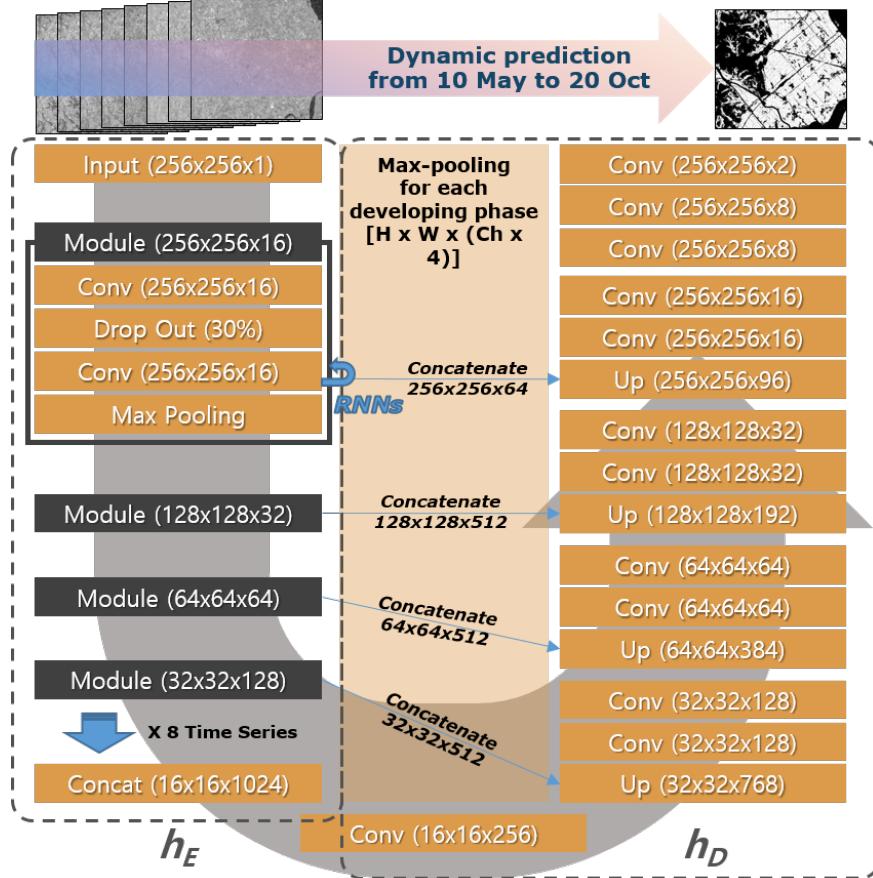


Figure 2: Recurrent U-net architecture

A.2 Training Details

Apart from making predictions with the entire time-series, we are able to acquire predictions early in the cultivation period using zero padding in the later time steps (Figure 3). The prediction using the full time-series ($t = 8$) and the early prediction using confined time steps ($1 \leq t < 8$) share a common feature extraction process but the signal intensity through the neural network can be greatly differed by σ . Therefore, if the input of the aforementioned time-series is provided in a random order during the training phase, the loss function is hardly optimized and overfitted to lastly seen training instances, which is associated with the problem of catastrophic forgetting [12].

Therefore, we manipulate the training order as described in Fig. 3, so that the model will be able to sequentially learn from each time-step's data. Specifically, in each epoch we firstly provide as input the cases of the very early prediction (e.g. only 1 time step) and lastly the ones of the full growing season. By training the model with chronologically ordered batches the parameters are updated gradually, with an additional time step after every new batch type, which is more likely to preserve the knowledge gained from the values of each new time step.

A.3 Additional results and plots

Additional metrics Table 2, 3 and 4 present the recall, the precision and the f1-score of the positive class for each of the different experiments, respectively.

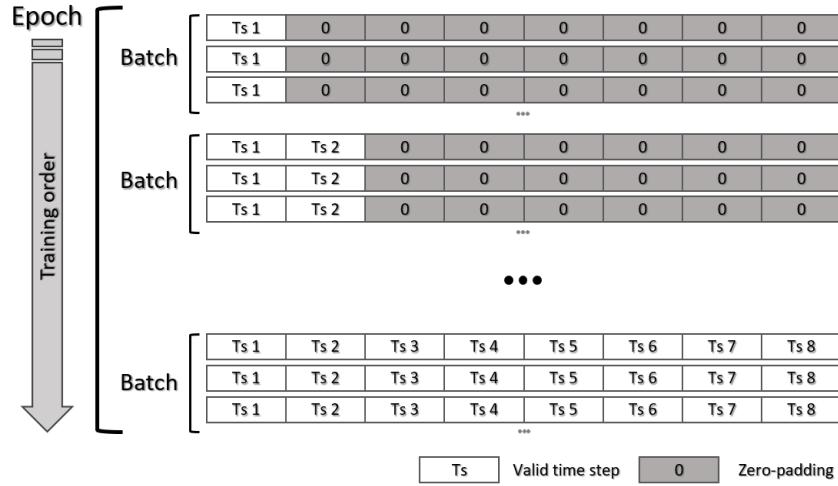


Figure 3: Chronologically ordered batch training

Table 2: Recall for the positive class in the different scenarios and methods

Fine-tuning		Spain				France				The Netherlands	
Test		Spain		France		France		Spain		The Netherlands	
Feature		VH	VH VV	VH	VH VV						
RF		0.957	0.964	0.665	0.698	0.831	0.881	0.878	0.846	0.293	0.446
RI		0.967	0.765	0.584	0.460	0.879	0.841	0.809	0.818	0.400	0.004
FT		0.957	0.962	0.601	0.674	0.896	0.894	0.954	0.950	0.490	0.614
FT _E		0.964	0.962	0.674	0.694	0.915	0.911	0.914	0.964	0.518	0.705

Table 3: Precision for the positive class in the different scenarios and methods

Fine-tuning		Spain				France				The Netherlands	
Test		Spain		France		France		Spain		The Netherlands	
Feature		VH	VH VV	VH	VH VV						
RF		0.910	0.926	0.926	0.931	0.901	0.943	0.866	0.906	0.718	0.781
RI		0.891	0.875	0.826	0.638	0.849	0.856	0.838	0.871	0.590	0.477
FT		0.932	0.928	0.919	0.915	0.905	0.920	0.858	0.872	0.686	0.631
FT _E		0.929	0.935	0.903	0.929	0.933	0.939	0.854	0.873	0.683	0.696

Table 4: F1-score of the positive class for the different scenarios and methods

Fine-tuning		Spain				France				The Netherlands	
Test		Spain		France		France		Spain		The Netherlands	
Feature		VH	VH VV	VH	VH VV						
RF		0.932	0.945	0.774	0.798	0.864	0.911	0.872	0.875	0.416	0.567
RI		0.927	0.816	0.684	0.534	0.864	0.848	0.824	0.844	0.477	0.007
FT		0.944	0.945	0.726	0.776	0.901	0.907	0.903	0.909	0.571	0.622
FT _E		0.947	0.948	0.772	0.794	0.924	0.925	0.883	0.916	0.589	0.700

Below we present visual predictions of the different scenarios and for the different methods of crop mapping. In every Figure below, the first column illustrates a composed of the first three Sentinel-1 VH images (Image), the second the ground truth labels (Label), the third the RF predictions, the fourth the RI predictions, the fifth the FT predictions and the sixth the FT_E predictions.

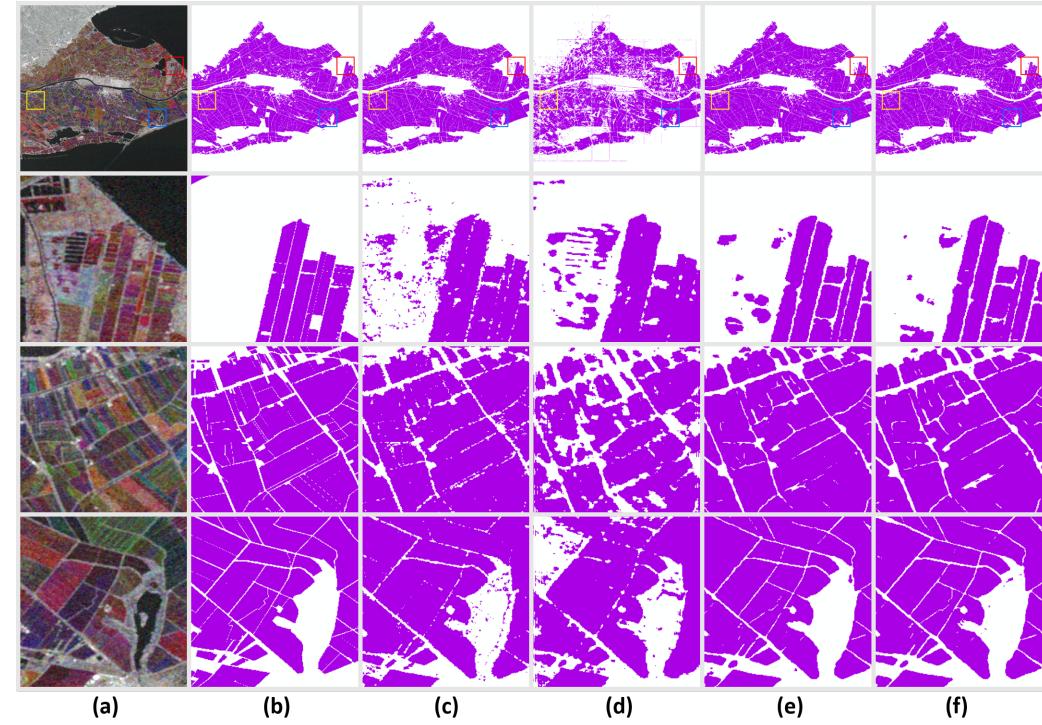
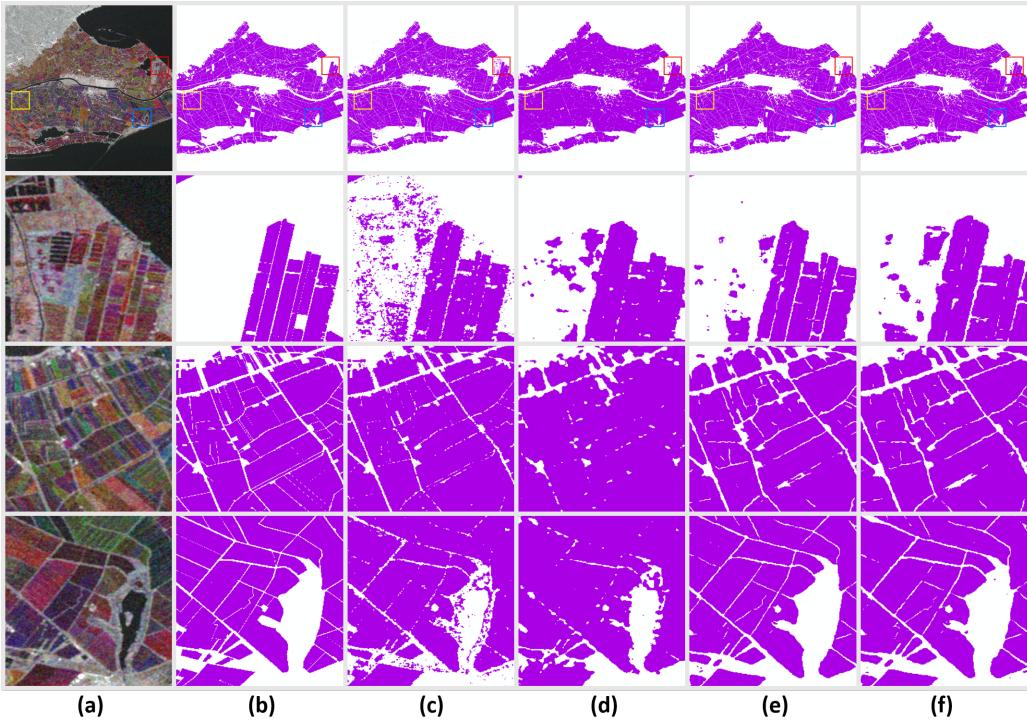


Figure 4: Visual comparisons on experiment 1 (Spain-Spain-rice-VH). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .

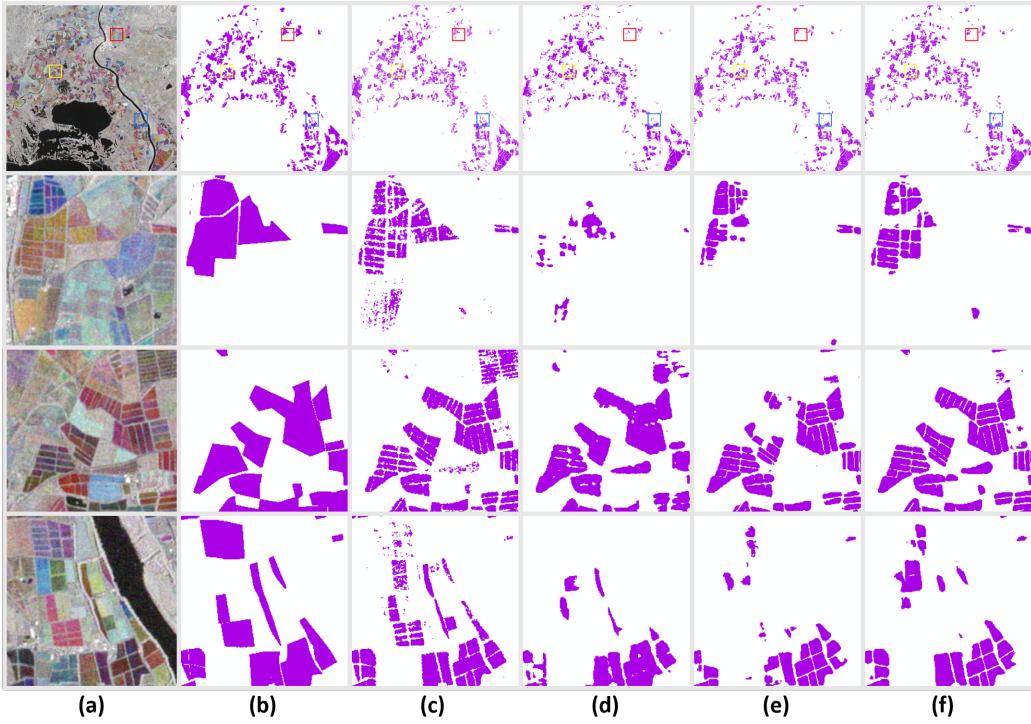


Figure 6: Visual comparisons on experiment 3 (Spain-France-rice-VH). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .

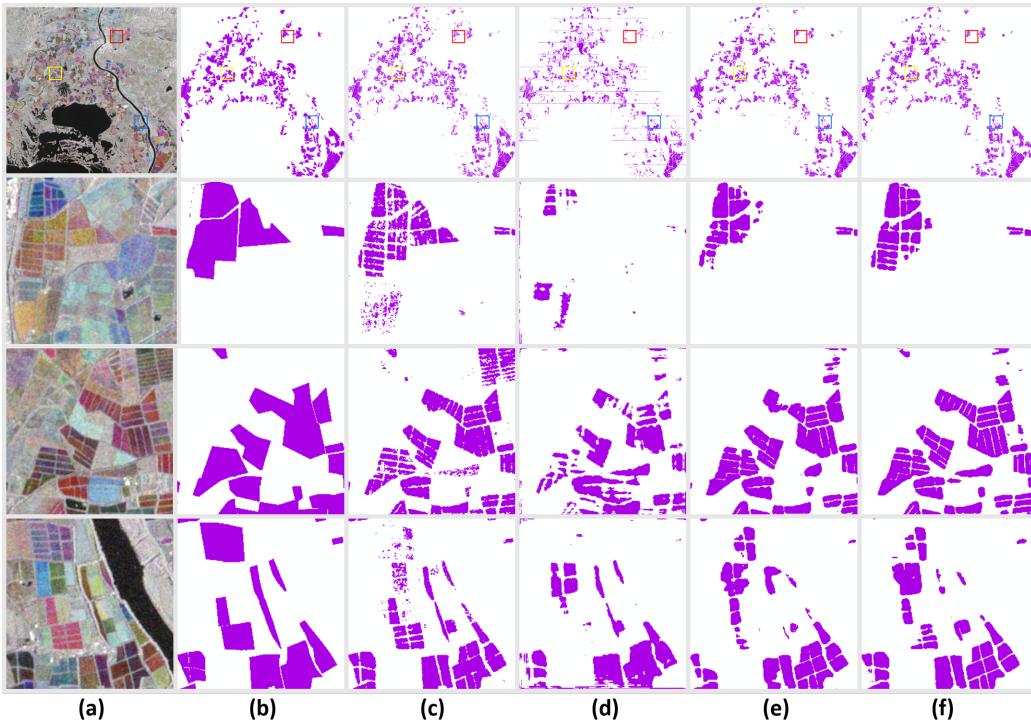


Figure 7: Visual comparisons on experiment 4 (Spain-France-rice-VH|VV). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .

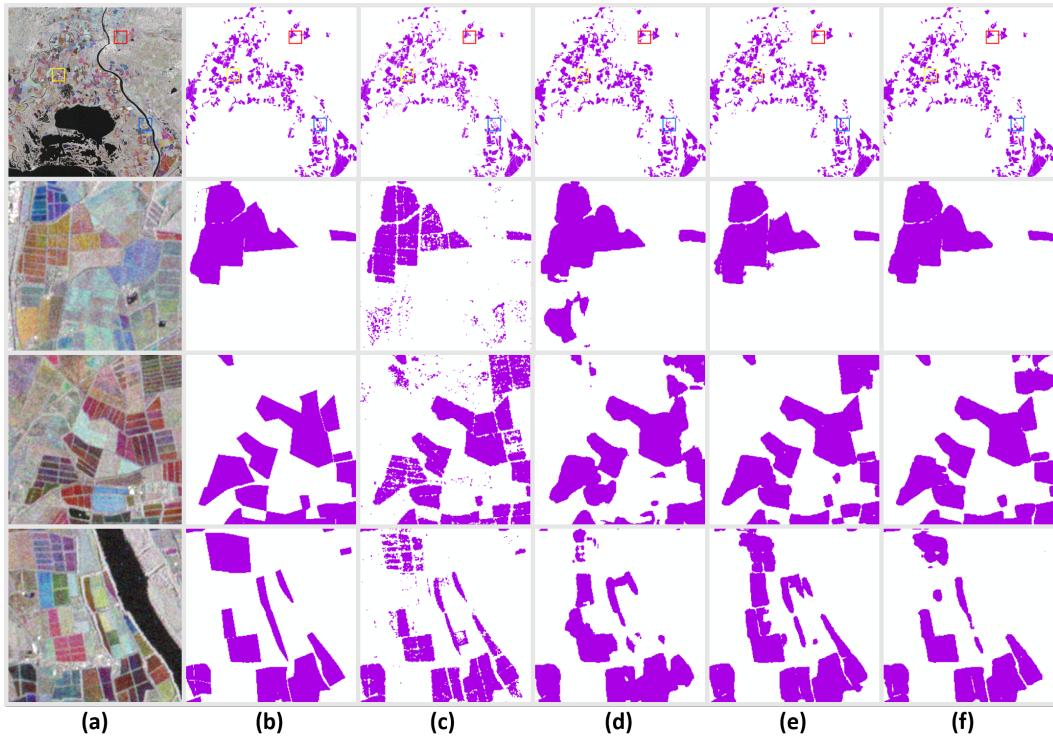


Figure 8: Visual comparisons on experiment 5 (France-France-rice-VH). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .

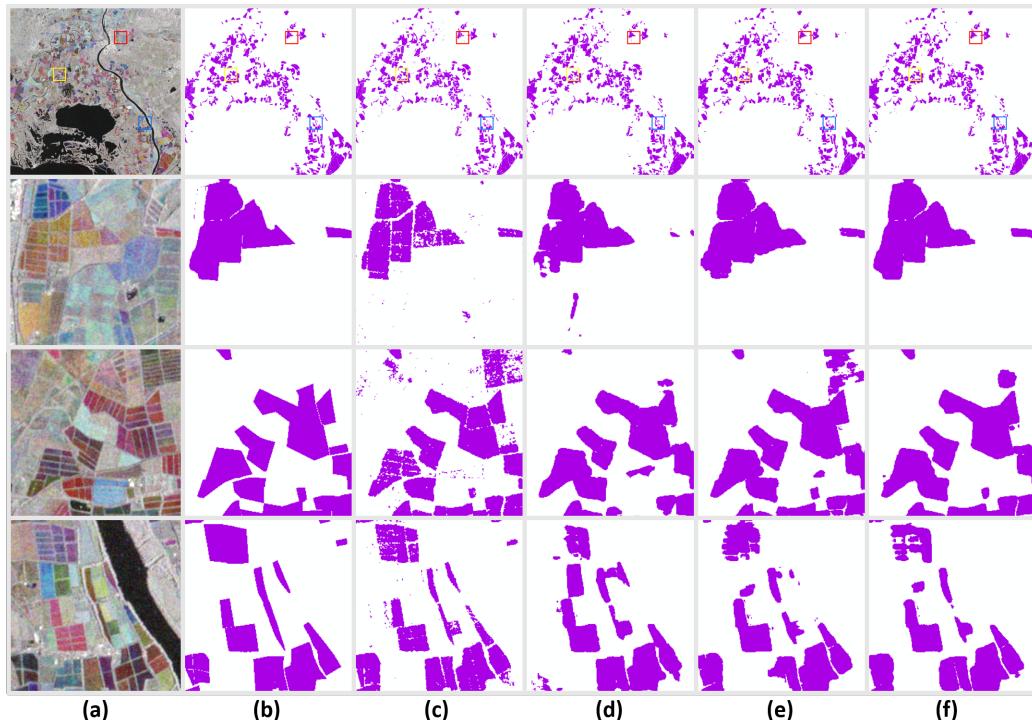


Figure 9: Visual comparisons on experiment 6 (France-France-rice-VH|VV). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .

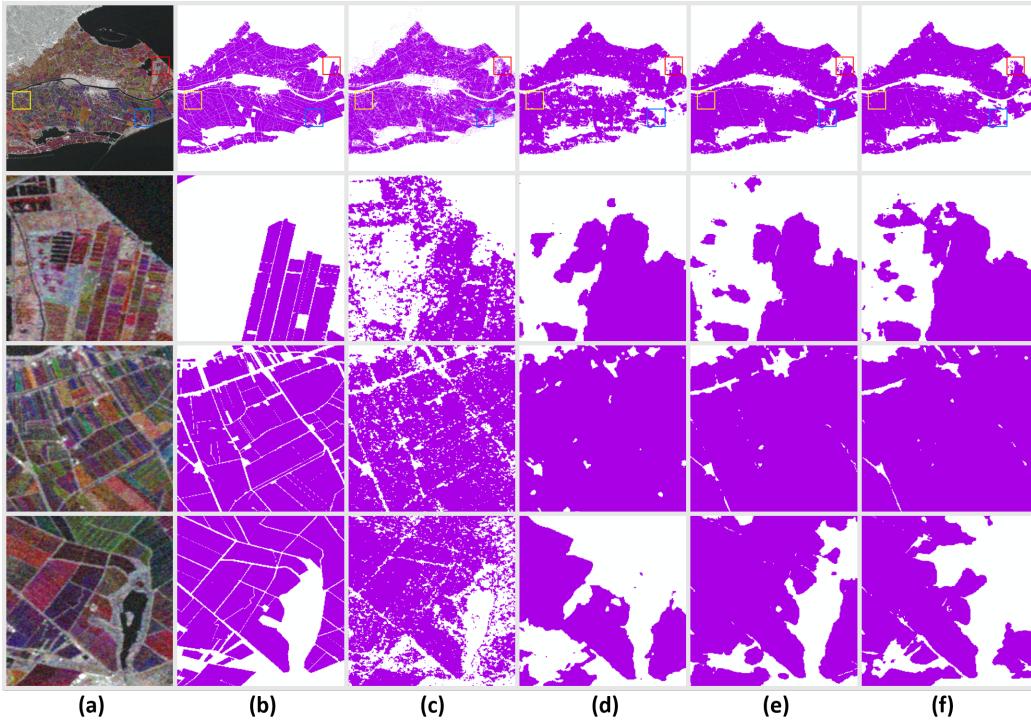


Figure 10: Visual comparisons on experiment 7 (France-Spain-rice-VH). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .

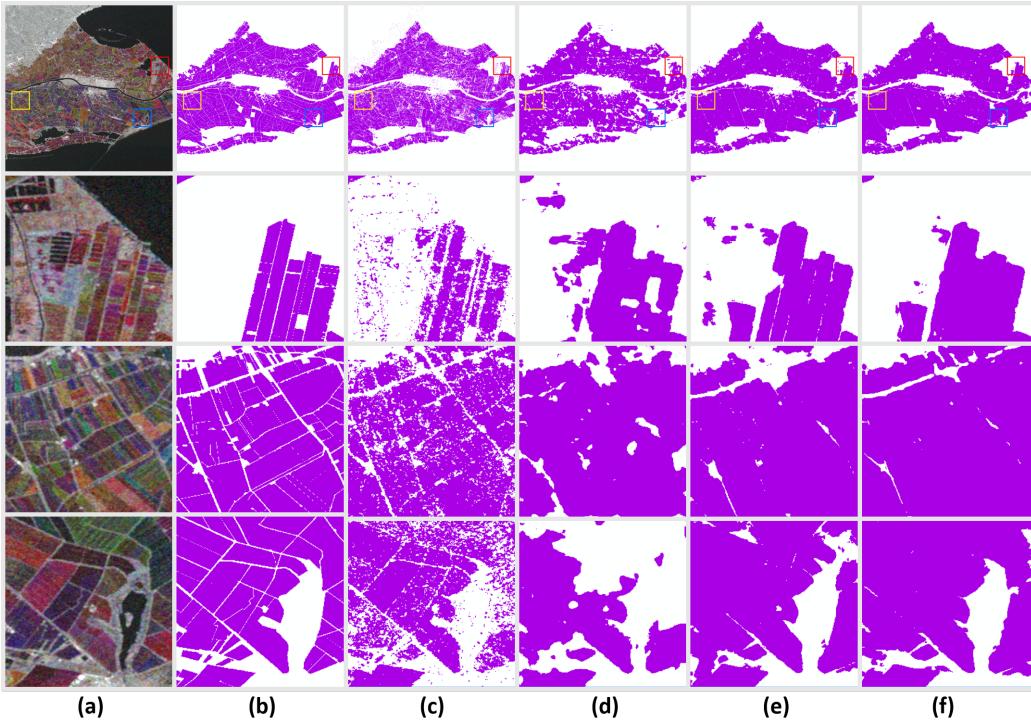


Figure 11: Visual comparisons on experiment 8 (France-Spain-rice-VH|VV). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .

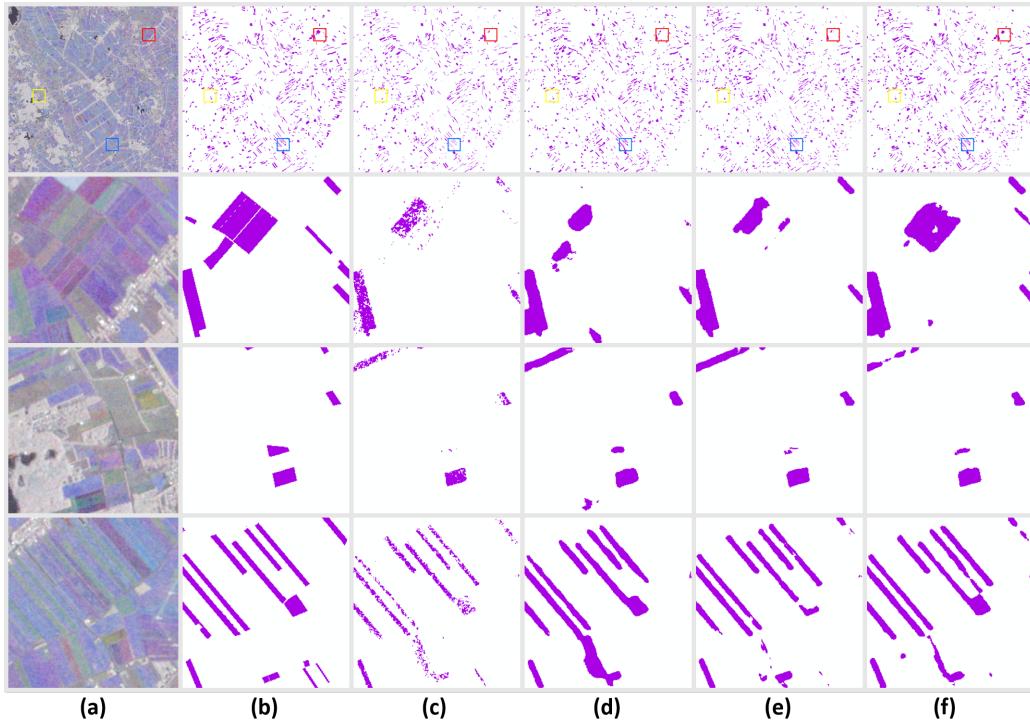


Figure 12: Visual comparisons on experiment 9 (Netherlands-Netherlands-barley-VH). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .

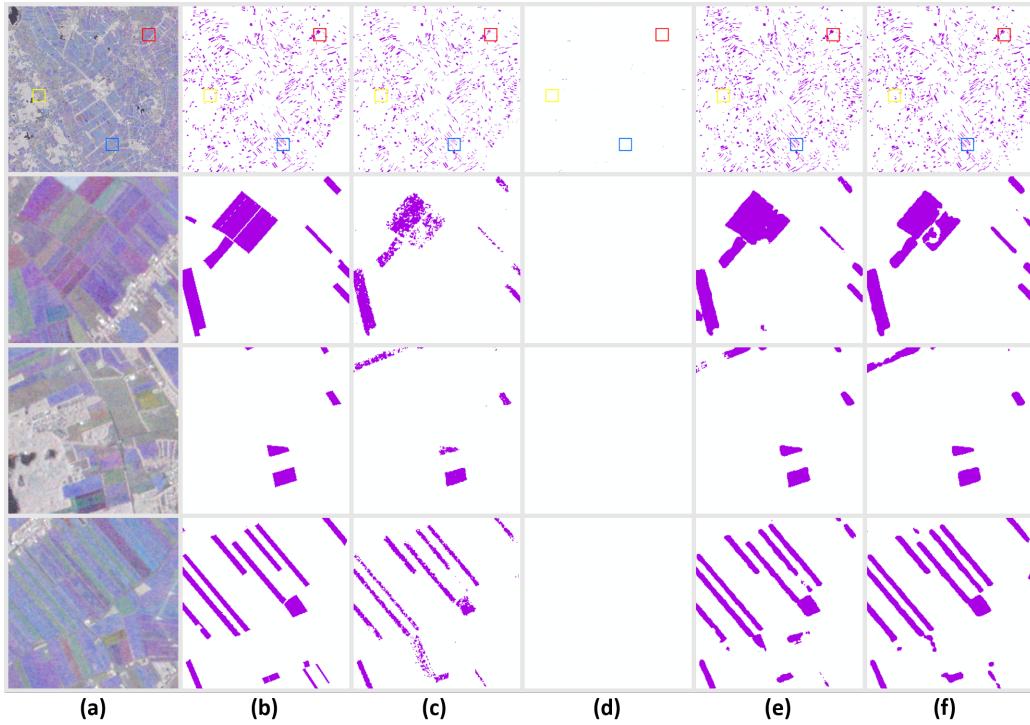


Figure 13: Visual comparisons on experiment 10 (Netherlands-Netherlands-barley-VH|VV). The first row shows the overall results of the test image and the following rows show three randomly selected test patches. (a) Image. (b) Label. (c) RF. (d) RI. (e) FT. (f) FT_E .