

A COMPARATIVE ASSESSMENT OF MULTI-VIEW FUSION LEARNING FOR CROP CLASSIFICATION

Francisco Mena^{1,2}, Diego Arenas², Marlon Nuske², and Andreas Dengel^{1,2}

¹University of Kaiserslautern-Landau (RPTU)
Department of Computer Science
Kaiserslautern, Germany

²German Research Center for Artificial Intelligence (DFKI)
Smart Data and Knowledge Services
Kaiserslautern, Germany

ABSTRACT

With a rapidly increasing amount and diversity of remote sensing (RS) data sources, there is a strong need for multi-view learning modeling. This is a complex task when considering the differences in resolution, magnitude, and noise of RS data. The typical approach for merging multiple RS sources has been input-level fusion, but other - more advanced - fusion strategies may outperform this traditional approach. This work assesses different fusion strategies for crop classification in the CropHarvest dataset. The fusion methods proposed in this work outperform models based on individual views and previous fusion methods. We do not find one single fusion method that consistently outperforms all other approaches. Instead, we present a comparison of multi-view fusion methods for three different datasets and show that, depending on the test region, different methods obtain the best performance. Despite this, we suggest a preliminary criterion for the selection of fusion methods.

Index Terms— Crop Classification, Remote Sensing, Data Fusion, Multi-view Learning, Deep Learning.

1. INTRODUCTION

Many phenomena in our environment are studied through multiple sources, e.g. a farm that could be observed by satellites with different sensors. The idea is to corroborate and complement the information between observations. Deep learning models have been widely used to model complex relationships between input data and target tasks. However, the situation described above poses a machine learning scenario of combining information coming from multiple sources, named multi-view or multi-modal fusion learning. In the context of Remote Sensing (RS), relevant challenges arise regarding the different types of sensor resolution. For instance, RS views could have different spatial or temporal resolutions, even different spectral bands (bandwidth or numbers of channels), calibration, and amount of noise.

F. Mena acknowledges the financial support from the chair of Prof. Dr. Prof. h.c. Andreas Dengel with RPTU in Kaiserslautern.

Data	View	AA	AUC	Entropy
Kenya	DEM	48.0 ± 0.9	41.1 ± 0.9	83.1 ± 3.8
	Weather	50.0 ± 0.0	37.1 ± 0.0	72.2 ± 8.7
	Radar	63.0 ± 1.0	66.8 ± 2.9	69.2 ± 3.8
Togo	DEM	61.2 ± 2.6	65.0 ± 1.8	94.6 ± 1.3
	Weather	55.6 ± 3.5	59.3 ± 5.4	98.6 ± 0.7
	Optical	80.0 ± 1.3	88.6 ± 0.7	62.7 ± 2.1
Global	DEM	64.9 ± 0.1	69.7 ± 0.1	91.1 ± 1.1
	Weather	72.7 ± 0.8	81.2 ± 0.6	74.1 ± 1.4
	Optical	78.7 ± 0.8	86.7 ± 0.8	59.6 ± 2.0

Table 1: Predictive performance of a model trained with individual views. The view with the best individual performance for the task is highlighted in **bold**.

In RS, there could be cases when a model trained with individual views (single-view model) gives predictions similar to a random classifier. See AA scores in Tab. 1 of DEM and Weather views for an example. We remark this as the main difference between multi-view learning in RS [1] to conventional multi-modal learning [2] (e.g. with image, text, and audio). This suggests that individual RS views may not contain enough information to achieve an optimal classification for some tasks, and that multi-view fusion is required.

The most common approach used in the literature to merge the information on multi-view RS is the input-level fusion [3, 4], i.e. align the input views to the same resolution (e.g. spatial and temporal interpolation) and concatenate them. However, some works have shown that the heterogeneity of RS views could be better exploited by using other fusion approaches [3, 5]. In this work, we compare different multi-view fusion methods for the crop classification task. The key contributions are as follows:

- A new assessment of different fusion strategies on the recent CropHarvest dataset [4].
- Achieve state-of-the-art results without pre-training in the African testing regions (Kenya and Togo).
- Open-source code to ensure reproducibility of our experiments and findings.¹

¹<https://github.com/fmenat/MultiviewCropClassification>

2. CASE STUDY

For this case study, we use a recent benchmark dataset on crop classification called CropHarvest [4]. Each sample corresponds to a yearly time-series at the pixel-level. The samples are labeled as positive or negative for specific crop categories, i.e. the task corresponds to identifying whether a pixel contains a specific (or any) growing crop during a season. There are 5 views as input data coming from 4 RS sources: Sentinel-2, Sentinel-1, ERA5, and SRTM. The multi-band optical view with 11 spectral channels was obtained from Sentinel-2 at 10-60m/px spatial resolution and 5 days temporal resolution. The 2-band radar view (VV and VH) was obtained from Sentinel-1 at 10m/px and a variable revisit time. The 2-band weather view (precipitation and temperature) was obtained from ERA5 at 31km/px and hourly temporal resolution. The 1-band Normalized Difference Vegetation Index (NDVI) was calculated from the optical view. The 2-band Digital Elevation Map (DEM) view (elevation and slope) was obtained from SRTM at 30m/px and contains topography information (static across time). The views are aligned by monthly averaging (temporal) and 10-m interpolating (spatial).

Out of the total 95186 samples in the benchmark, we only use the 65243 samples that have all views available. The dataset contains two geographical testing regions in Africa: Kenya and Togo. We also include a global testing region, which consists of all the training samples across the globe. Tab. 2 displays the distribution of the samples per region.

Data	Total	Training	Testing
Kenya	2217 (37.8%)	1319 (20.0%)	898 (64.0%)
Togo	1596 (51.1%)	1290 (55.0%)	306 (34.6%)
Global	65243 (66.3%)	45723 (66.4%)	19520 (66.0)

Table 2: Number of samples in each data region. The percentage of positive data is shown in parentheses.

3. METHODS

We compare three main fusion strategies highlighted in [1, 3, 5]: input, feature, and decision level fusion. In addition, multiple losses model [6] and an ensemble-based aggregation [7] are compared. An illustration of these methods is in Fig. 1.

Input: a concatenation of *input views* with a resolution alignment is directly fed to a single model. **Feature:** it uses view-encoder models to map each view to a new *feature space* (named view-representation), followed by a merge function and fully connected layers. We test different merge functions. Simple summary functions (**Feature-S**) and with gated modules (**Feature-G**). The gated module adaptively fuses the view-representations through a weighted sum, where the (attention) weights are computed for each sample [8]. The

weights could be modeled for each view or for each feature and view combination (feature-specific weight). **Decision:** it uses parallel models that process each view, outputs the crop probability (*decision*), and then averages to yield the aggregated prediction. **Multi-Loss:** based on the feature-level fusion, it includes one auxiliary predictive model that is fed with a single view representation and yields the crop classification. For training the auxiliary model for each view, a loss function is added to the optimization (with a weight of 0.3, [6]), so the model has to learn the task based on the individual views in addition to the fused information. **Ensemble:** similar to decision-level fusion, but on a two-step basis. The first step corresponds to training a model for each view (without fusion), while the second step aggregates (averages) the predicted probabilities from individual models at test time.

We use a recurrent neural network with 2 GRU layers and 64 units as view-encoders. While for prediction, we stack 1 fully connected layer with 64 hidden units. We include 20% of dropout during training and batch-normalization on all layers, as suggested [9]. The optimization is carried out with the cross entropy loss and ADAM optimizer with a 256 batch size. We train the models for 1000 epochs maximum with an early stopping criterion, patience of 5, and delta loss of 0.01 on the validation set (a 10% randomly selected from training).

4. EXPERIMENTS

We use two classification metrics [4], Average Accuracy (AA), and Area Under the Curve (AUC). We also include the entropy of the predicted crop probability to analyze the uncertainty of the model. Based on the variability of the weights initialization and stochastic optimization, we repeat the experiment 10 times and report the average and std. deviation.

By selecting the best combination of the merge function and the gate type for Feature methods (see Section 4.2 for the corresponding experiment), we compare all fusion methods in Tab. 3. We observe that methods using multiple models often obtain better classification performance than a single model (Input fusion). Besides, fusions at the feature-level are always in the top 3 best performances and have the lowest std. deviation. The evidence shows the effectiveness of multi-view fusion, since the top 3 fusion methods have better classification performance than individual views (from Tab. 1) in all cases. In terms of *relative improvement*, $100(a - b)/b$, when we compare the best fusion method (a) with the best model on an individual view (b), a *relative improvement* of 6%, 5%, and 4% is obtained in AA for Kenya, Togo, and Global, respectively. While comparing the best fusion (a) with the worst individual model (b), we get a relative improvement of 39%, 37%, and 27% in AA for Kenya, Togo and Global. These results suggest that significant improvements can be obtained by exploring different fusion strategies.

Regarding entropy, the Feature-G method obtains the lowest uncertainty in its classification predictions. It is possible

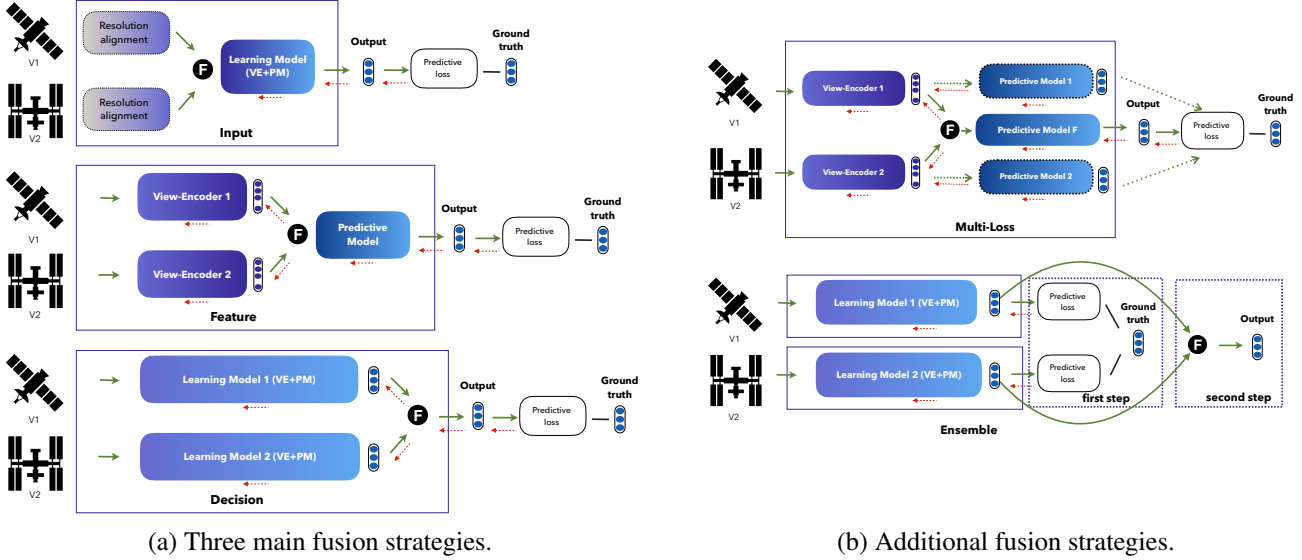


Fig. 1: Fusion methods compared in this manuscript. Green arrows represent the forward pass of the models (left to right), while red arrows represent the backward pass. VE stands for view-encoder and PM for predictive model. Images from [1].

that, by letting the model choose how it will combine the views for each sample (i.e. the adaptive fusion in the gated module proposal [8]), it becomes more confident in its decision. In most methods, fusing multiple views decreases the uncertainty regarding individual views. However, with some methods the entropy increases, e.g. with Ensemble. Interestingly, when merging towards the output-level, the model increases the entropy, e.g. the order of increasing entropy goes like Input, Feature-S, and then Decision. This suggests that the method becomes more uncertain when it has few layers to exchange information between the views.

4.1. State-of-the-Art comparison

The Tab. 4 compares our results with the state-of-the-art including the binary F1 score. We show that, by selecting an appropriate model for fusion, we outperform previous work. Indeed, the best method is different for each region and metric. In Kenya, we obtain the best AUC with Feature-G, and the best binary F1 with Feature-S, while in Togo with the Ensemble method.

4.2. Ablation

Firstly, we compare different merge functions for the feature-level fusion with simple functions (Feature-S) in Tab. 5. The *average* operation has the best classification performance for Kenya and Togo, while the *concatenation* has it for the Global region. These best performing approaches also have the lowest entropy. In Tab. 6 we compare different gate types in the feature-level fusion with gated modules (Feature-G). The standard gated module [8] where the *concatenation* of

Data	Method	AA	AUC	Entropy
Kenya	Input	61.3 ± 7.5	70.0 ± 7.2	72.9 ± 6.1
	Feature-S	63.0 ± 5.1	71.6 ± 3.7	73.9 ± 4.4
	Feature-G	66.5 ± 3.6	71.8 ± 3.5	71.6 ± 5.8
	Decision	57.5 ± 7.6	63.2 ± 7.0	76.8 ± 7.6
	Ensemble	56.0 ± 5.3	69.5 ± 1.4	83.5 ± 2.2
	Multi-Loss	64.8 ± 5.2	71.6 ± 4.1	73.4 ± 4.3
Togo	Input	79.7 ± 1.5	89.0 ± 1.3	54.3 ± 4.2
	Feature-S	79.9 ± 1.0	88.9 ± 0.8	57.6 ± 4.1
	Feature-G	78.1 ± 1.8	87.6 ± 1.2	52.1 ± 7.8
	Decision	81.5 ± 1.6	89.5 ± 0.5	63.7 ± 2.5
	Ensemble	84.0 ± 1.0	90.9 ± 0.5	93.5 ± 0.6
	Multi-Loss	78.2 ± 6.1	88.6 ± 1.7	60.4 ± 12.2
Global	Input	81.2 ± 0.9	89.7 ± 0.8	52.7 ± 2.5
	Feature-S	81.3 ± 1.1	89.6 ± 1.1	53.8 ± 3.8
	Feature-G	82.1 ± 1.4	90.6 ± 1.1	51.1 ± 4.3
	Decision	79.1 ± 1.4	87.8 ± 1.3	59.4 ± 4.6
	Ensemble	79.2 ± 0.8	87.5 ± 1.5	85.4 ± 0.6
	Multi-Loss	81.6 ± 1.4	90.1 ± 1.1	52.3 ± 3.7

Table 3: Prediction performance of multi-view fusion methods. The top 3 best results for each test are underlined. The mean and std. deviation between multiple runs is shown.

the view-representation is used to compute the weights is Gated-C, while Gated-A uses the *average*. The GatedF uses a feature-specific weight in the gated module. Similarly to the Feature-S method, the average obtains better performance than concatenation. Nevertheless, including the feature-specific weight improves overall. Given that each feature might contain different information, the method allows more flexibility when combining the view’s information.

Data	Method	AUC	binary F1
Kenya	Random Forest [4]	57.8 ± 0.6	55.9 ± 0.3
	LSTM [4]	32.9 ± 1.1	78.2 ± 0.0
	Feature-S	71.6 ± 3.7	79.4 ± 3.3
	Feature-G	71.8 ± 3.5	77.2 ± 4.1
Togo	Random Forest [4]	89.2 ± 0.1	75.6 ± 0.2
	LSTM [4]	86.1 ± 0.2	72.0 ± 0.5
	Ensemble	90.9 ± 0.5	77.8 ± 1.3

Table 4: State-of-the-art prediction performance on the CropHarvest dataset. The best average result is in **bold**.

Data	Merge	AA	AUC	Entropy
Kenya	Average	63.0 ± 5.1	71.6 ± 3.7	73.9 ± 4.4
	Maximum	55.4 ± 8.0	58.9 ± 16.7	85.3 ± 7.7
	Product	50.0 ± 0.0	45.6 ± 12.8	99.7 ± 9.9
	Concatenate	64.0 ± 5.2	67.5 ± 9.8	79.8 ± 0.4
Togo	Average	<u>79.9 ± 1.0</u>	<u>88.9 ± 0.8</u>	57.6 ± 4.1
	Maximum	64.7 ± 9.0	76.0 ± 7.6	86.9 ± 9.9
	Product	50.0 ± 0.0	49.9 ± 8.0	99.5 ± 0.5
	Concatenate	79.5 ± 1.9	88.5 ± 1.8	59.2 ± 3.7
Global	Average	80.5 ± 1.8	89.1 ± 1.5	54.0 ± 3.3
	Maximum	81.1 ± 1.1	89.5 ± 1.0	53.6 ± 2.7
	Product	79.5 ± 1.2	88.8 ± 0.9	62.9 ± 4.6
	Concatenate	81.3 ± 1.1	89.6 ± 1.1	53.8 ± 3.8

Table 5: Merge function comparison in the **Feature-S** method. The best average result for each test is underlined.

Data	Type	AA	AUC	Entropy
Kenya	Gated-C	57.9 ± 6.1	64.9 ± 7.2	68.0 ± 7.0
	Gated-A	58.5 ± 7.5	66.4 ± 7.8	72.3 ± 6.6
	GatedF-A	<u>66.5 ± 3.6</u>	<u>71.8 ± 3.5</u>	71.6 ± 5.8
Togo	Gated-C	75.6 ± 5.8	86.3 ± 2.4	55.7 ± 13.6
	Gated-A	77.1 ± 5.3	87.3 ± 2.4	55.0 ± 13.6
	GatedF-A	78.1 ± 1.8	<u>87.6 ± 1.2</u>	52.1 ± 7.8
Global	Gated-C	80.3 ± 1.9	88.8 ± 1.8	53.3 ± 2.5
	Gated-A	81.6 ± 0.7	90.0 ± 0.6	53.1 ± 3.5
	GatedF-A	<u>82.1 ± 1.4</u>	<u>90.6 ± 1.1</u>	51.1 ± 4.3

Table 6: Gate types comparison in the **Feature-G** method. The best average result for each test is underlined.

5. FINAL REMARKS

We present an extensive comparison of fusion methods in a crop classification benchmark, achieving state-of-the-art performance in the African testing regions without pre-training. Even though our results are promising regarding different fusion strategies, there are some approaches that decrease the classification performance regarding individual views, while increasing prediction uncertainty. The best performing fusion methods usually depend on the region where they are applied. However, we suggest the following preliminary criterion. When testing in a region with a large positive area (a higher percentage of positively labeled pixels in the re-

gion, e.g. Kenya/Global), the methods with feature-level exchange between views (Feature-G, Feature-S, Multi-Loss) are the best. Whereas, when testing in a region with a small positive area (e.g. Togo), the methods with feature-level exchange get worse, and exchanging at the output level becomes the best (Ensemble, Decision). Anyhow, questions remain open whether the same results apply to other RS applications.

6. REFERENCES

- [1] Francisco Mena, Diego Arenas, Marlon Nuske, and Andreas Dengel, “Common practices and taxonomy in deep multi-view fusion for remote sensing applications,” *arXiv preprint arXiv:2301.01200*, 2023.
- [2] Xiaoqiang Yan, Shizhe Hu, Yiqiao Mao, Yangdong Ye, and Hui Yu, “Deep multi-view learning methods: A review,” *Neurocomputing*, vol. 448, pp. 106–129, 2021.
- [3] Stella Ofori-Ampofo, Charlotte Pelletier, and Stefan Lang, “Crop Type Mapping from Optical and Radar Time Series Using Attention-Based Deep Learning,” *Remote Sensing*, vol. 13, no. 22, pp. 4668, 2021.
- [4] Gabriel Tseng, Ivan Zvonkov, Catherine Lillian Nakalembe, and Hannah Kerner, “CropHarvest: A global dataset for crop-type classification,” *Proceedings of NIPS Datasets and Benchmarks Track*, 2021.
- [5] Danfeng Hong, Lianru Gao, N. Yokoya, J. Yao, J. Chanussot, Q. Du, and B. Zhang, “More Diverse Means Better: Multimodal Deep Learning Meets Remote-Sensing Imagery Classification,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 5, pp. 4340–4354, 2021.
- [6] Paola Benedetti, Dino Ienco, Raffaele Gaetano, Kenji Ose, R. Pensa, and S. Dupuy, “M3fusion: A Deep Learning Architecture for Multiscale Multimodal Multitemporal Satellite Data Fusion,” *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 11, no. 12, pp. 4939–4949, 2018.
- [7] Caleb Robinson, Kolya Malkin, Nebojsa Jojic, Huijun Chen, Rongjun Qin, C. Xiao, M. Schmitt, P. Ghamisi, R. Hänsch, and N. Yokoya, “Global Land-Cover Mapping With Weak Supervision: Outcome of the 2020 IEEE GRSS Data Fusion Contest,” *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, vol. 14, pp. 3185–3199, 2021.
- [8] John Arevalo, Tamar Solorio, Manuel Montes-y Gómez, and Fabio A. González, “Gated multimodal networks,” *Neural Computing and Applications*, vol. 32, no. 14, pp. 10209–10228, 2020.
- [9] Yushi Chen, Chunyang Li, Pedram Ghamisi, Xiuping Jia, and Y. Gu, “Deep Fusion of Remote Sensing Data for Accurate Classification,” *IEEE Geoscience and Remote Sensing Letters*, vol. 14, no. 8, pp. 1253–1257, 2017.