

# MULTI-SOURCE TRANSFORMER ARCHITECTURES FOR AUDIOVISUAL SCENE CLASSIFICATION

## Technical Report

*Wim Boes, Hugo Van hamme*

ESAT, KU Leuven

wim.boes@esat.kuleuven.be, hugo.vanhamme@esat.kuleuven.be

### ABSTRACT

In this technical report, the systems we submitted for subtask 1B of the DCASE 2021 challenge, regarding audiovisual scene classification, are described in detail. They are essentially multi-source transformers employing a combination of auditory and visual features to make predictions. These models are evaluated utilizing the macro-averaged multi-class cross-entropy and accuracy metrics.

In terms of the macro-averaged multi-class cross-entropy, our best model achieved a score of 0.620 on the validation data. This is slightly better than the performance of the baseline system (0.658).

With regard to the accuracy measure, our best model achieved a score of 77.1% on the validation data, which is about the same as the performance obtained by the baseline system (77.0%).

**Index Terms**— DCASE 2021, audiovisual scene classification, transformer

## 1. INTRODUCTION

Subtask 1B of DCASE 2021 [1] is dedicated to audiovisual scene classification. Models are ranked using the macro-averaged multi-class cross-entropy, which is further explained in Section 3.

For this work, we were influenced by the prior use of transformers in the context of environmental event classification with audiovisual data [2]. We also were inspired by the employment of multi-source transformer architectures in the context of machine translation with multiple languages [3].

In Section 2, the submitted systems are described. In Section 3, we go into the experimental setup. Next, in Section 4, we report the obtained results, and finally, we draw a conclusion in Section 5.

## 2. MODELS

In this section, the submitted models are elaborated upon.

### 2.1. Architecture

The architecture of the models is visualized in Figure 1.

The input to the system is a combination of three inputs: spectral auditory features, pretrained auditory features and pretrained visual features. How these are obtained from audiovisual scene recordings is further explained in Section 3.

The spectral auditory features are first processed by a convolutional neural network (CNN), consisting of four blocks. Each block comprises five layers: a convolutional layer, a batch normalization layer [4], a ReLU activation layer, a dropout layer [5] (with a drop rate of 33%) and an average pooling layer.

Each convolutional layer uses a square kernel of size 3 and a stride of 1. The amount of output channels are equal to 12, 24, 48 and 96 for the first, second, third and fourth blocks respectively.

Table 1: Kernel sizes and strides of pooling layers in CNN

Block	Kernel size = stride
0	(3, 4)
1-2	(2, 4)
3	(1, 2)

The kernel sizes and strides of the pooling operations are listed in Table 1. The first and second numbers of each tuple relate to the time and frequency axes respectively.

At the end of the last block, the frequency-related dimension of the spectral input has been reduced to one and the corresponding axis can therefore be discarded.

The sequences of convolutional auditory, pretrained auditory and pretrained visual features are linearly mapped to embeddings of size 96. The outputs of these operations are additionally passed through dropout layers with a drop rate of 33%.

The embeddings extracted from the convolutional auditory and pretrained visual features are then (separately) passed through a transformer encoder consisting of three layers.

Next, a multi-source serial transformer decoder with three layers is used. This module is a simple extension of the regular transformer decoder allowing for more than two inputs by using multiple multi-head cross-attention blocks (in series) instead of just one.

In the considered model, the embeddings extracted from the pretrained auditory features constitute the queries of the multi-source serial transformer decoder. The keys and values of the first and second multi-head cross-attention blocks in this decoder come from the transformer encodings based on the convolutional auditory and pretrained visual features respectively.

All transformer components described above use three attention heads, fully connected layers with 96 units and dropout rates of 0.1.

As explained in Section 3, the pretrained auditory features fed as input to the architecture are “sequences” of one vector. Thus, the output of the transformer decoder also consists of just one element.

This vector is passed through a dropout layer, a linear map is computed, and finally, the softmax activation function is applied to obtain the output audiovisual scene probabilities.

### 2.2. Loss function

The loss used to train these models is the categorical cross entropy.

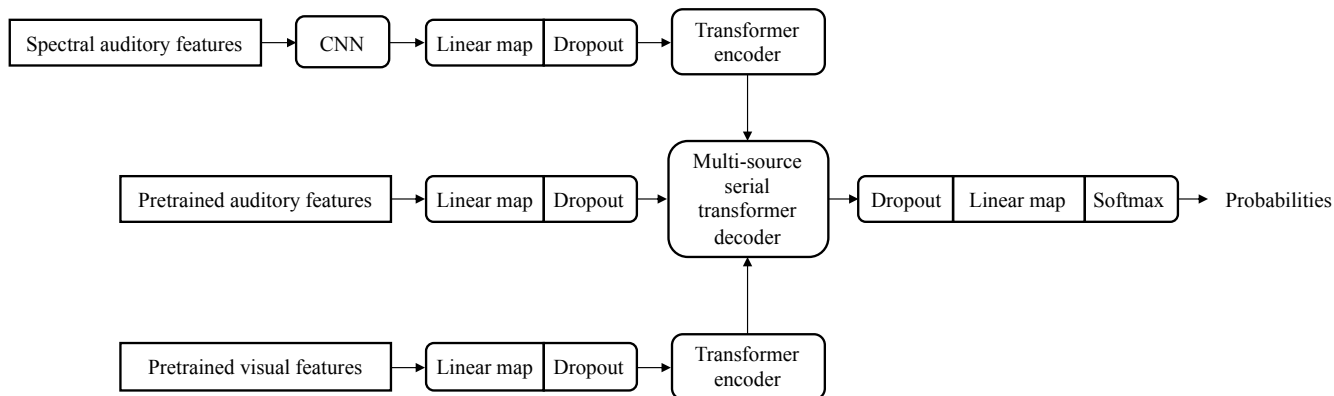


Figure 1: Architecture of submitted models.

### 2.3. Mean teacher ensemble

In this work, we employ an ensembling method based on the mean teacher principle [6]: The architecture described above is duplicated but is not trained in the regular way. Instead, its parameters are computed as the exponential moving average of the weights of the original structure with a multiplicative decay factor of 0.999 per training iteration. This “teacher” model can effectively be regarded as an averaged version of the “student” system.

## 3. EXPERIMENTAL SETUP

### 3.1. Data

The multi-class data set for problem 1B of DCASE 2021 [7] consists of synchronized auditory and visual scene recordings. There are 10 possible, mutually exclusive audiovisual scenes.

The training partition consists of 8696 files of 10 seconds, which are cut into 86960 samples of one second. The validation set contains 72000 examples with a length of one second.

### 3.2. Preprocessing

To get spectral maps of the audio segments, we resampled them to 22.050 kHz and performed peak amplitude normalization. Next, log mel spectrograms with 128 bins were computed using a window of 2048 samples and a hop length of 368 samples. Per-frequency bin standardization was done as a final step. For the sound snippets of one second, this resulted in sequences of 600 features.

To get pretrained auditory features, we resampled the audio to 16 kHz and extracted log mel spectrograms with 128 frequency bins utilizing a window of 400 samples and a hop length of 160 samples. The resulting spectral maps were then fed to vggish [8], a model for video tag classification based on audio, pretrained using a preliminary version of the YouTube-8M data set [9]. The 128-dimensional outputs of the last fully connected layer of this network were used as pretrained auditory features. For an audio recording of one second, this resulted in a “sequence” of one vector.

To get pretrained visual features, the frames of the visual files (recorded at 29.97 fps) were passed through VGG16 [10], a network for image classification, pretrained on the ImageNet data set [11]. The 4096-dimensional outputs of the last fully connected layer of this model were used as pretrained visual features. For a video of one second, this resulted in a sequence of 30 features.

### 3.3. Data augmentation

In order to diminish the risk of overfitting, we employed data augmentation during the training of the considered models. More specifically, we used mixup [12] with an activation probability of 50%. The mixing ratios were randomly sampled from a beta distribution with shape parameters set to 0.2.

### 3.4. Training and evaluation

All of the models were trained and evaluated using PyTorch [13].

#### 3.4.1. Training hyperparameters

All systems were trained for 30 epochs with a batch size of 128. Every sample in each batch originated from a different training file.

Adam [14] was employed to train the weights of the models. Learning rates were ramped up exponentially from 0 to 0.00025 for the first 675 optimization iterations. Subsequently, they decayed multiplicatively at a rate of 0.999 per training step.

#### 3.4.2. Evaluation metrics

For this challenge, two evaluation metrics are used: macro-averaged multi-class (categorical) cross-entropy and accuracy.

After each training epoch, these metrics were calculated using the probabilities produced by the teacher models. The best scores were retained and are reported in Section 4.

## 4. RESULTS

The baseline model for this particular challenge [7] achieved a macro-averaged multi-class cross-entropy measure of 0.658 and an accuracy of 77.0% on the validation data.

Our best model achieved a cross-entropy of 0.620 on the validation data, which is a slight improvement over the baseline.

With regard to the accuracy measure, our best model achieved a score of 77.1% on the validation data, which is about the same as the performance obtained by the baseline system.

## 5. CONCLUSION

In this technical report, we described the systems we submitted for subtask 1B of the DCASE 2021 challenge and reported their results.

The best macro-averaged cross-entropy and accuracy scores achieved by our models were equal to 0.620 and 77.1% respectively.

## 6. ACKNOWLEDGMENT

This work is supported by a PhD Fellowship of Research Foundation Flanders (FWO-Vlaanderen).

## 7. REFERENCES

- [1] Wang, Shanshan and Heittola, Toni and Mesaros, Annamaria and Virtanen, Tuomas, “Audio-visual scene classification: analysis of DCASE 2021 Challenge submissions,” *arXiv preprint arXiv:2105.13675*, 2021.
- [2] Boes, Wim and Van hamme, Hugo, “Audiovisual transformer architectures for large-scale classification and synchronization of weakly labeled audio events,” in *Proceedings of the 27th ACM International Conference on Multimedia*, 2019, pp. 1961–1969.
- [3] J. Libovický, J. Helcl, and D. Mareček, “Input Combination Strategies for Multi-Source Transformer Decoder,” in *Proceedings of the Third Conference on Machine Translation: Research Papers*, 2018, pp. 253–260.
- [4] S. Ioffe and C. Szegedy, “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift,” in *Proceedings of ICML*, 2015, pp. 448–456.
- [5] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: A Simple Way to Prevent Neural Networks from Overfitting,” *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [6] A. Tarvainen and H. Valpola, “Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results,” in *Advances in Neural Information Processing Systems*, 2017, pp. 1195–1204.
- [7] Wang, Shanshan and Mesaros, Annamaria and Heittola, Toni and Virtanen, Tuomas, “A Curated Dataset of Urban Scenes for Audio-Visual Scene Analysis,” in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [8] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, *et al.*, “CNN Architectures for Large-Scale Audio Classification,” in *2017 IEEE International Conference on Acoustics, Speech and Signal processing (ICASSP)*, 2017, pp. 131–135.
- [9] S. Abu-El-Haija, N. Kothari, J. Lee, P. Natsev, G. Toderici, B. Varadarajan, and S. Vijayanarasimhan, “Youtube-8M: A Large-Scale Video Classification Benchmark,” *arXiv preprint arXiv:1609.08675*, 2016.
- [10] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” in *International Conference on Learning Representations*, 2015.
- [11] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, “ImageNet: A Large-Scale Hierarchical Image Database,” in *CVPR09*, 2009.
- [12] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” *arXiv preprint arXiv:1710.09412*, 2017.
- [13] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, “PyTorch: An Imperative Style, High-Performance Deep Learning Library,” in *Advances in Neural Information Processing Systems*, 2019, pp. 8024–8035.
- [14] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” *arXiv preprint arXiv:1412.6980*, 2014.