



Full length article

# Multimodal Attentive Fusion Network for audio-visual event recognition

Mathilde Brousmiche<sup>a,b,\*</sup>, Jean Rouat<sup>a</sup>, Stéphane Dupont<sup>b</sup>

<sup>a</sup> Université de Sherbrooke, NECOTIS Lab, 2500 boul. de l'Université, Sherbrooke, Canada

<sup>b</sup> University of Mons, Numediart Institute, 20 Place du Parc, Mons, Belgium

## ARTICLE INFO

### Keywords:

Audio-visual fusion  
Modality conditioning  
Attention  
Multimodal deep learning  
Event recognition

## ABSTRACT

Event classification is inherently sequential and multimodal. Therefore, deep neural models need to dynamically focus on the most relevant time window and/or modality of a video. In this study, we propose the Multimodal Attentive Fusion Network (MAFnet), an architecture that can dynamically fuse visual and audio information for event recognition. Inspired by prior studies in neuroscience, we couple both modalities at different levels of visual and audio paths. Furthermore, the network dynamically highlights a modality at a given time window relevant to classify events. Experimental results in AVE (Audio-Visual Event), UCF51, and Kinetics-Sounds datasets show that the approach can effectively improve the accuracy in audio-visual event classification. Code is available at: <https://github.com/numediart/MAFnet>

## 1. Introduction

Event recognition is an active research area in machine learning. It has numerous potential applications such as video surveillance [1], autonomous driving [2], sports analysis [3], and content-based retrieval [4]. Thanks to the collection of large video datasets such as YouTube-8M [5], Kinetics [6] and Sports-1M [7], event recognition performance has improved during recent years. However, compared to the success of image classification, event recognition is still a challenging task due to the high computational complexity. The current state-of-the-art performance remains not as accurate as human performance.

Inspired by the success of image classification [8,9], current models for visual event recognition use convolutional neural networks (CNN). Several methods have been explored to process the temporal information such as two-streams network [10], recurrent neural network (RNN) [11], and three-dimensional convolution (3D CNN) [12], to name a few. However, all these works ignore an important information present in video: acoustic features.

Video understanding is a natural human ability. From an early age, humans are able to understand events or actions based on image, video as well as sound. However, most contemporary approaches ignore the acoustic information. It is obvious that the acoustic signal can be useful for event recognition. For example, two musical instruments may be difficult to distinguish based on the visual information but can produce distinct sounds or the object of interest may be occluded. Furthermore, some actions such as *Whistling* are visually subtle but can be recognized based on acoustic features.

Given the potential of sound to facilitate event recognition, researchers have attempted to combine the audio and visual signals [13–21]. Still, audio-visual event recognition is not easy as it is difficult to effectively use the audio information. In fact, audio can be corrupted with irrelevant background noise and sounds. Moreover, some event such as *Shaking hands* does not produce a particular sound signature.

Visual and audio information are different types of signals. Therefore, how to exploit the maximum relevant information coming from both modalities? There is not a simple answer. In real life situations, scenes are dynamics and the respective contribution of audio and video to the scene (or object) understanding evolves through time. Effective solutions should be adaptive and take into consideration the fact that the dynamics of visual and acoustics inputs are very different and change with the actions in the scene.

In the context of deep learning, visual and audio paths do not have the same complexity and therefore the same learning dynamics during training. Indeed, the visual and audio paths do not have the same learning speed, the number of iterations necessary to train a network based on visual information is not the same as for a network based on audio. This can lead to a generalization problem when training modalities together.

Therefore, we propose the Multimodal Attentive Fusion network (MAFnet) to dynamically fuse visual and audio information for event recognition. Respective contributions of the modalities are dynamically weighted through time via an attention mechanism that weights visual and audio features. Modalities interact with each other at different levels (different layers) of the network. A score highlights a modality

\* Corresponding author at: University of Mons, Numediart Institute, 20 Place du Parc, Mons, Belgium.

E-mail addresses: [mathildebrousmiche@gmail.com](mailto:mathildebrousmiche@gmail.com), [mathilde.brousmiche@usherbrooke.ca](mailto:mathilde.brousmiche@usherbrooke.ca) (M. Brousmiche).

at a given time window that may be effective to recognize the event. We also propose to go further than the simple fusion by coupling modalities with a lateral connection between visual and audio paths of MAFnet. Moreover, to overcome the incompatibility of learning dynamics between visual and audio paths, we propose to randomly drop the update of the visual path during training. We evaluate our architecture on multiple datasets: Kinetics-Sound [6], UCF51 [22] and AVE [19]. In addition, we evaluate the contribution of each module of MAFnet with an ablation study.

## 2. Related work

### 2.1. Visual event recognition

Inspired by the success of image classification, convolutional neural networks (CNN) have also been applied in visual event recognition. Several methods have been proposed to take advantage of the temporal information. The usage of recurrent neural networks (RNN) on top of 2D convolutional layers is investigated in [11,23,24] to take into account long-term dependencies. Li et al. went further by proposing a convolutional long-short term memory (LSTM) [25]. Another approach was to extend 2D convolution kernels to 3D convolution kernels to learn spatio-temporal features [6,7,12]. In addition, to reduce complexity, the 3D convolution is decomposed into two convolutions: a spatial 2D convolution and a temporal 1D convolution [26,27]. Another strategy was to capture fine low-level motion by calculating optical flow [10,28,29]. However, all these techniques do not exploit an important part of the video classification: the acoustic information.

### 2.2. Audio-visual event recognition

In recent years, only few works exploited the information present in the audio signal. The concatenation is used in [13,30] to fuse the visual and audio paths to exploit the information from the two modalities. Long et al. went further by testing different levels of fusion in the network [21]. These networks integrated the visual and audio information with hard fusion without exploiting a possible interaction between visual and audio paths. These works did not study more complex fusion techniques such as multimodal compact bilinear pooling (MCB) [31] or dual multimodal residual fusion (DMR) [19]. Furthermore, they did not take into account the different learning dynamics of the different modalities.

### 2.3. Audio-visual event detection

The release of the AVE dataset [19] has stimulated research in audio-visual event detection. For example, an audio-guided visual attention mechanism is introduced in [19] to learn which visual region to look at based on the visual and audio information. Lin et al. proposed to learn global and local event information in a sequence to sequence manner [32]. Finally, Wu et al. extracted the global representation of one modality and found the local segments that are relevant to the event in the other modality and vice versa [33]. In our work, we propose to better integrate audio-visual information by computing a global feature with an attention module to include only the relevant information present in the modalities. In addition, instead of visual spatial attention, we propose to use an attention mechanism on the audio feature maps based on the visual information, called modality conditioning. Indeed, we observed that the visual modality has more information to contribute to the audio path of MAFnet than the reverse.

### 2.4. Modality conditioning

Modality conditioning is the influence of a modality on another modality. It is the interaction between the paths of each modality inside the neural network. Interactions can be created by simple operations between paths such as an element-wise multiplication [20] or a sum [34] at different levels of the network.

More complex approaches to condition modalities have been explored, for example, the attention mechanism. Attention models were first proposed for object detection [35] and then used for other applications such as natural language processing with the self-attention mechanism [36]. Attention has been applied to video classification under the form of temporal and/or spatial attention [21,37,38]. However, these models did not include an interaction between modality paths. Only Tian et al. proposes to realize the interaction between paths by implementing visual attention guided by audio to condition vision [19].

Another approach to condition one modality with the other is the conditional normalization (CN). Instead of focusing attention on a particular region of space or a particular time window, the CN highlights some feature maps based on a given input. Various forms of CN have proven to be highly effective across a number of domains and modalities: image stylization [39], speech recognition [40], visual question answering [41] and audio question answering [42]. As presented in our previous work [43], we propose to condition the audio path with the visual information by using the Feature-wise Linear Modulation (FiLM) method [44]. The FiLM layer highlights audio feature maps based on visual information.

## 3. Multimodal Attentive Fusion Network

Inspired by the ability of humans to pay attention to different regions, instants and modalities [47], we propose to compute a score for each modality and for each time window with the modality & temporal attention module. The attention module combines modality and temporal information to create a global feature containing the relevant multimodal and temporal information.

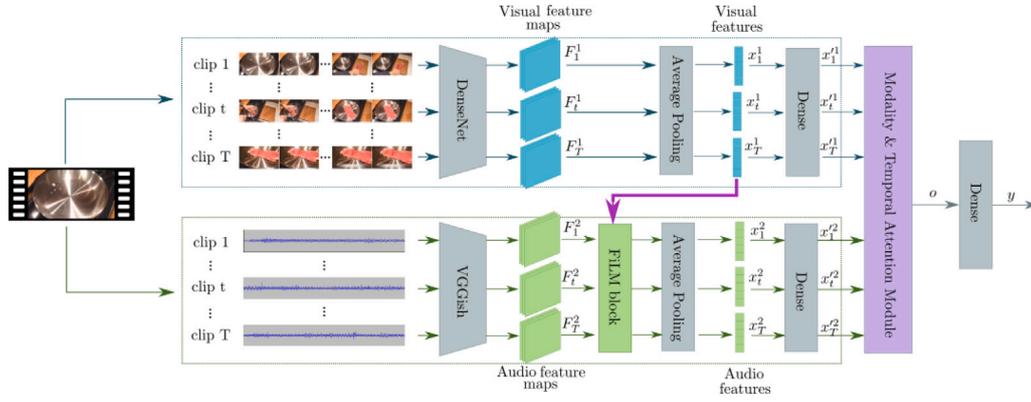
In addition to the fusion with the attention module, we propose to go further than modality fusion at high level and include interaction between visual and audio paths with a FiLM layer. In this section, we overview the Multimodal Attentive Fusion Network (MAFnet) and then detail the different components of the network.

### 3.1. Overview of the Multimodal Attentive Fusion Network (MAFnet)

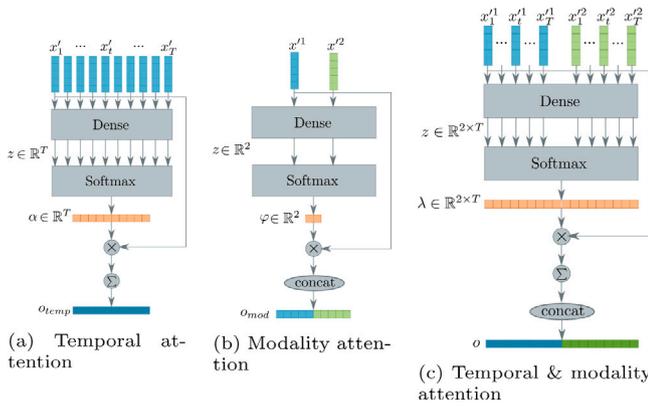
Fig. 1 presents the architecture of MAFnet. As in [19], we split each video into  $T$  non-overlapping clips, where each clip is 1s long. We extract information for  $K = 2$  modalities (visual and audio information). For each clip, we extract visual and audio feature maps with pretrained convolutional neural networks. So, we have 2 input sequences:  $\{F_1^1, \dots, F_T^1\}, F_t^1 \in \mathbb{R}^{H_v \times W_v \times D_v}$  for the visual information and  $\{F_1^2, \dots, F_T^2\}, F_t^2 \in \mathbb{R}^{H_a \times W_a \times D_a}$  for the audio information.  $H$ ,  $W$  and  $D$  are respectively the height, the width and the number of feature maps.

We reduce the feature maps with average pooling and feed the visual features ( $\{x_1^1, \dots, x_T^1\}, x_t^1 \in \mathbb{R}^{D_v}$ ) and audio features ( $\{x_1^2, \dots, x_T^2\}, x_t^2 \in \mathbb{R}^{D_a}$ ) in the modality & temporal attention module. This module is the combination of temporal and modality attentions. It attempts to learn the attention scores  $\lambda_t^k$  with  $t = 1, \dots, T$  and  $k = 1, \dots, K$  to weight temporal and modality dimensions. We therefore obtain a temporal-multimodal representation of the entire video. The output of the network is  $y \in \mathbb{R}^N$  with  $N$  the number of classes.

To go further than a simple fusion, we implement a lateral connection between visual and audio paths with the FiLM layer [44]. With a FiLM layer, the visual modality influences the audio modality. Greater importance is given to some of the audio feature maps based on the visual information. The FiLM layer is placed directly at the output of the audio feature extractor before reducing feature maps into vectors.



**Fig. 1.** Multimodal Attentive Fusion Network (MAFNet): one video is split into  $T$  non-overlapping clips. Then, audio and visual information are extracted with two pretrained CNNs: DenseNet [45] for visual features and VGGish [46] for audio features. The clip features are further fed into the modality & temporal attention module to build a global feature comprising multimodal and temporal information. This global feature is then used to predict the label of the video. A lateral connection between visual and audio paths is created through a FiLM layer [44].



**Fig. 2.** Attention mechanisms. (a) Temporal attention: a score  $\alpha$  is computed for each time window and the video-level feature representation  $o_{temp}$  with the sum. (b) Modality attention: a score  $\varphi$  is computed for each modality and the multimodal feature representation  $o_{mod}$  with the concatenation. (c) Temporal & modality attention: a score  $\lambda$  is computed for each time window AND modality and the global feature representation  $o$  with the combination of the sum over time windows and the concatenation over modalities.

### 3.2. Temporal attention

The aim of temporal attention [21,33] is to assign a positive weight score to each clip descriptors extracted from the video (Fig. 2(a)). The score can be interpreted as the relative contribution of each clip to the recognition of the target action, or the relative importance of each clip to generate an accurate global video representation.

Technically, for a given modality, given the input feature  $X' = \{x'_1, \dots, x'_T\}$ ,  $x'_t \in \mathbb{R}^D$ , the corresponding score  $\alpha = \{\alpha_1, \dots, \alpha_T\}$  over the  $T$  feature vectors is computed by

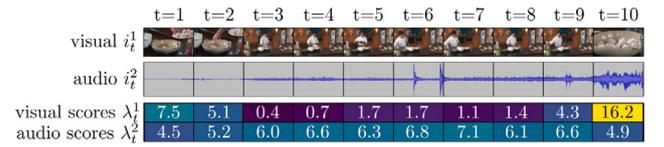
$$z_t = g_{att}(x'_t; \theta_{att}) = \text{ReLU}(W_{temp}^T x'_t + b) \quad (1)$$

$$\alpha_t = \frac{\exp(z_t)}{\sum_{j=1}^T \exp(z_j)} \quad (2)$$

where  $g_{att}$  is the temporal attention network parameterized by  $\theta_{att}$ .  $g_{att}$  can take different forms such as a perceptron.  $z_t$  is an intermediate attention score, normalized with the softmax function.

We compute the video-level feature representation  $o_{temp}$  with the weighted sum of the clip features. The weights are the scores computed by the temporal attention module:

$$o_{temp} = \sum_{t=1}^T \alpha_t x'_t \quad (3)$$



**Fig. 3.** Visualization of the scores  $\lambda_t^k$  determined by the modality & temporal attention module for a video labeled *Frying (food)* of the AVE dataset.

### 3.3. Modality attention

In the context of speech recognition, Zhou et al. proposed a modality attention mechanism [48]. The attention mechanism fuses input from multiple modalities into a single representation by weighted summing the information from individual modalities. We propose to use a similar mechanism but use the concatenation of the weighted modalities instead of the sum (Fig. 2(b)). In Section 4.5, we discuss the choice of the modality fusion.

The attention module computes a score for each modality, the score is proportional to the importance of the modality for the video classification.

Technically, at a given time, given the input feature  $X' = \{x'^1, \dots, x'^K\}$ ,  $x'^k \in \mathbb{R}^{D^k}$  with  $K$  the number of modalities. The score for each modality is computed by:

$$z^k = h_{att}(x'^k; \theta_{att}) = \text{ReLU}(W_{mod}^T x'^k + b) \quad (4)$$

$$\varphi^k = \frac{\exp(z^k)}{\sum_{j=1}^K \exp(z^j)} \quad (5)$$

where  $h_{att}$  is the attention network parameterized by  $\theta_{att}$  and  $z^k$  is an intermediate attention score, normalized with the softmax function.

The multimodal feature  $o_{mod}$  is obtained by fusing the weighted unimodal features with a concatenation:

$$o_{mod} = \text{concat}([\varphi^1 x'^1, \dots, \varphi^K x'^K]) \quad (6)$$

The modality attention module can dynamically choose the most relevant modality for a better classification of the events. Indeed, we can imagine that *Frying (food)* or *Truck* have strong visual information while *Violin* or *Flute* have strong audio information.

### 3.4. Modality & temporal attention module

We can combine the temporal and the modality attention modules to constitute the modality & temporal attention module (Fig. 2(c)). The aim of the modality & temporal attention module is to assign a positive

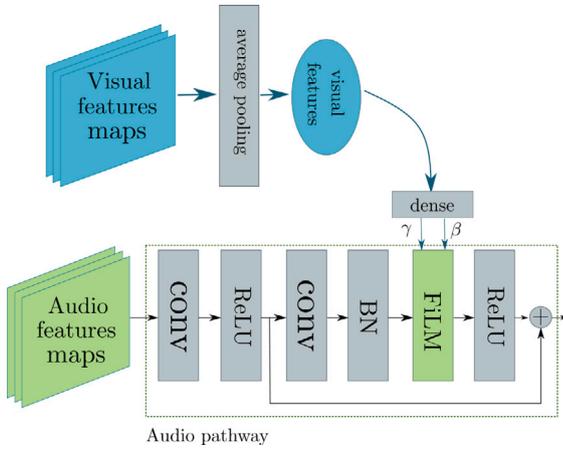


Fig. 4. Lateral connection between visual and audio paths through FiLM layer: The FiLM layer inside the residual block uses the visual features to modulate the audio feature maps.  $\gamma$  and  $\beta$  parameters are computed from a dense layer having its input from the visual features.

score for each modality and clip. Indeed, for example in Fig. 3, we see that most of the time the audio information is more relevant than the visual information except for the last clip where you can clearly see the food frying. This visual clip has the largest score and can be identified as the most relevant to classify the video as *Frying (food)*.

If we have the input:

$$X' = \begin{bmatrix} x_1^1 & \dots & x_t^1 & \dots & x_T^1 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ x_1^K & \dots & x_t^K & \dots & x_T^K \end{bmatrix} \quad (7)$$

The equations of the attention module become:

$$z_t^k = f_{att}(x_t^k; \theta_{att}) = \text{ReLU}(W_{mod+temp}^T x_t^k + b) \quad (8)$$

$$\lambda_t^k = \frac{\exp(z_t^k)}{\sum_{j=1}^T \sum_{l=1}^K \exp(z_j^l)} \quad (9)$$

$$o = \text{concat}\left(\left[\sum_{t=1}^T \lambda_t^1 x_t^1, \dots, \sum_{t=1}^T \lambda_t^K x_t^K\right]\right) \quad (10)$$

We add a dense layer in the path of each modality before the attention module because the attention module needs each modality to have the same dimension.

### 3.5. Lateral connection

We propose to go further than the “simple” fusion at high level by including a lateral connection to condition audio with vision. Indeed most approaches do not exploit a possible interaction between the different paths. As presented in our previous work [43], the Feature-wise Linear Modulation (FiLM) layer can create a lateral connection between visual and audio paths. We use visual features as input to the FiLM layer to highlight feature maps of the audio modality (Fig. 4).

More formally, FiLM learns functions  $f$  and  $h$  to compute  $\gamma_{t,c}$  and  $\beta_{t,c}$  as a function of input  $x_t^1$ :

$$\gamma_{t,c} = f_c(x_t^1) \quad \beta_{t,c} = h_c(x_t^1) \quad (11)$$

$\gamma_{t,c}$  and  $\beta_{t,c}$  modulate the activations  $\mathbf{F}_{t,c}^2$ , whose subscripts refer to the  $t$ th input and  $c$ th audio feature map, via a feature-wise affine transformation:

$$\text{FiLM}(\mathbf{F}_{t,c}^2 | \gamma_{t,c}, \beta_{t,c}) = \gamma_{t,c} \mathbf{F}_{t,c}^2 + \beta_{t,c} \quad (12)$$

$f$  and  $h$  can be arbitrary functions which are typically implemented with neural networks. FiLM layers allow to manipulate feature maps of

a target according to an input by scaling them up or down, negating them, shutting them off, selectively thresholding them (when followed by a ReLU).

### 3.6. Audio-visual training

Wang et al. noticed in [49] that multi-modal networks are prone to overfitting due to their increased capacity and different modalities overfit and generalize at different rates. So, they proposed a complex Gradient-Blending training. Xiao et al. noticed also different dynamic of training depending on the modality [34] and propose to randomly drop the audio path during the training. Unlike [34], when training unimodal networks, we notice that the audio network need more epoch to reach overfitting compared to the visual network. Therefore, we follow the idea of [34] and randomly drop the upgrade of the visual path, to train more the audio path.

## 4. Experimental results

### 4.1. Datasets

We evaluate our network on three public datasets: AVE [19], UCF51 [22] and Kinetics-Sounds [14].

AVE is a subset of AudioSet [52]. The dataset consists of 4143 videos from 28 event classes. Each video lasts 10 s. It covers a wide range of audio-visual events from different domains, e.g., human activities, animal activities, music performances, and vehicle sounds.

UCF51 is the second part of the UCF101 dataset [22]. Only the videos of the new 51 classes have sound information. UCF51 dataset consists of 6836 videos from 51 event classes. It concentrates on human actions. The mean video length is 7.0 s. The dataset is partitioned into three splits for training and testing.

Kinetics-Sounds is a subset of the Kinetics dataset [6] and consists of only action classes that are potentially recognizable both visually and aurally. It consists of 21945 videos from 32 events categories. The mean video length is 9.7 s.

### 4.2. Feature extraction

Audio and visual features can easily be extracted from a new video using trained models [19,21,32,33]. The extracted features are significantly smaller in size than the raw RGB frame and audio data and allow working with smaller networks.

**Visual feature extraction.** We use an ImageNet pre-trained deep learning model named DenseNet [45] to extract visual features from video. The video is split into  $T$  clips. As in [19], we choose  $T = 10$ , so each clip is one second long without overlapping. For each clip, we extract the output of the DenseNet last convolutional layer for 16 RGB video frames with a global average pooling over the 16 frames to generate one  $7 \times 7 \times 1920$  dimensional feature map.

**Audio feature extraction.** We use a VGG-like network [46] pre-trained on AudioSet to extract audio features. Again, the video is split into  $T = 10$  clips of one second each without overlapping. For each clip, we extract the output of the last convolutional layer of the network to generate one  $12 \times 8 \times 512$  dimensional feature map.

### 4.3. Implementation details

The number of filters in the Residual Block and the number of neurons of the hidden dense layer is 512. Batch normalization is used after hidden dense layers. The network is trained with cross-entropy loss and Adam optimizer with an initial learning rate of 0.001. Early stopping based on the validation accuracy is done, the training is stopped when the validation accuracy does not improve since 50 epochs. During training, we randomly do not update the weights of the visual path. The model is implemented in Tensorflow [53].

As UCF51 and Kinetics-Sounds datasets have different video lengths, feature vectors are zero padded to obtain the same length.

**Table 1**

Comparison with state-of-the-art on AVE, UCF51 and Kinetics-Sound datasets. Each model was trained based on code available online. Models are split into two types: end-to-end training and feature extraction. End-to-end training models are trained on larger datasets and then fine-tuned on a smaller dataset. By contrast, feature extraction models are trained on feature previously extracted from the video. Depending on the model, input can be visual frame (V) and/or audio (A).

	Model	Inputs	Pretrained dataset	Accuracy [%]		
				AVE	UCF51	Kinetics-Sound
End-to-end training	I3D [6]	V	Kinetics	73.28	86.92	80.22
	R(2+1)D [26]	V	Sports1M + Kinetics	79.19	95.54	79.10
	SlowFast [50]	V	Kinetics	80.41	91.78	81.91
	MARS [51]	V	Kinetics	79.44	<b>97.83</b>	–
	Model	Inputs	Extraction network	AVE	UCF51	Kinetics-Sound
Feature extraction	Attention Cluster [30]	V + A	DenseNet (V) + VGGish (A)	80.71	84.79	73.91
	DMRN [19] (our feat.)	V + A	DenseNet (V) + VGGish (A)	80.96	82.93	77.5
	DMRN [19] (their feat.)	V + A	VGG19(V) + VGGish-PCA(A)	85.02	81.04	–
	MAFnet (our)	V + A	DenseNet (V) + VGGish (A)	<b>90.86</b>	86.72	<b>83.94</b>

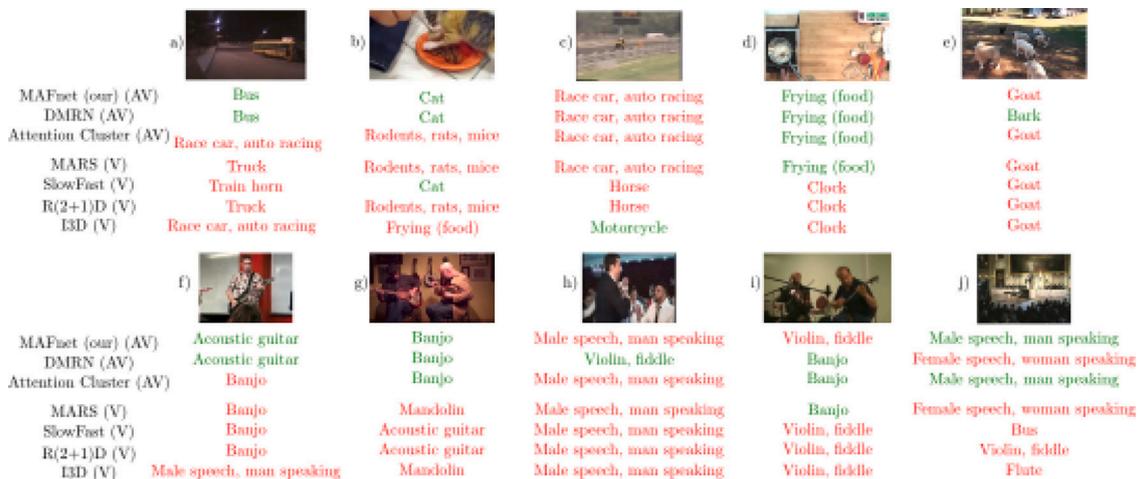


Fig. 5. Output prediction of different visual only (V) models and audio-visual (AV) models for some example of the AVE dataset. Each model predicts one class per video. (Green: correct prediction, Red: False prediction). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

#### 4.4. Event recognition performance

Table 1 presents event recognition results of MAFnet on AVE, UCF51 and Kinetics-Sounds datasets. We also compare our result with several state-of-the-art methods using different modalities, *i.e.* audio (A) and visual frames (V). For the UCF51 dataset, we report the average accuracy over three testing splits.

Unlike our method, the Attention Cluster network [30] does not exploit the modality interaction at several levels. It comprises an attention mechanism that works independently inside each modality, followed by a concatenation of the vector of each modality. DMRN [19] comprises a lateral connection that projects from audio to visual features before the concatenation. It highlights image pixels based on sound information, while MAFnet highlights audio features maps based on visual information.

MAFnet obtains the best accuracy performance on the AVE dataset among methods based on end-to-end training or feature extraction. End-to-end training methods have the advantage of being trained on larger datasets such as Sports1M or Kinetics to avoid overfitting and then the entire network is fine-tuned on smaller datasets. As the AVE dataset was built as an audio-visual set, the audio information is as important as the visual information. Therefore, as the end-to-end models take into account only the visual information, performances decreases. Furthermore, the AVE dataset includes classes from different events unlike Sports1M and Kinetics datasets which include classes from human activities only. Models based on feature extraction obtain slightly better results than end-to-end training methods due to the use of audio information.

The UCF51 dataset comprises fewer classes with relevant audio information. Indeed, it contains classes that do not produce a particular

sound signature or event video with irrelevant background noise. Our network is then not as good as end-to-end training models which take advantage of pretraining on larger datasets and fine-tuning the entire network. On the other hand, models using feature extraction do not fine-tune extractor networks. However, MAFnet is the best model among the architectures that use audio-visual features.

The Kinetics-Sound dataset as well as the UCF51 dataset is centered on human action but comprises classes potentially recognizable both visually and aurally. Therefore, when the dataset comprises relevant audio and visual information, our network provides the best result. It is capable to take advantage of both modalities. Moreover, it has better integration of the audio and visual information than the other audio-visual models.

Fig. 5 shows examples of output prediction from the different models. For examples (a) to (e), we observe that the background might impact the choice of the class. The raceway is classified as *Race car, auto racing* even if there is a bus (example a) or a motorcycle (example c). The field with a herd is classified as *Goat* (example e). Moreover, some specific elements in the video can fool models. The spoon and the plate (example b) may influence the I3D model in the choice of the *Frying (food)* class. In example (d), visual models may be fooled by the round shape of the pan. In the case of examples (c) and (e), the audio modality is not distinctive enough to help the network.

We also note that some instruments can be difficult to distinguish (example f and g) or are occluded (example h). Some videos include several classes but are annotated with only one class (example i). Others are visually indistinctive (example j) but can be classified thanks to the audio modality.

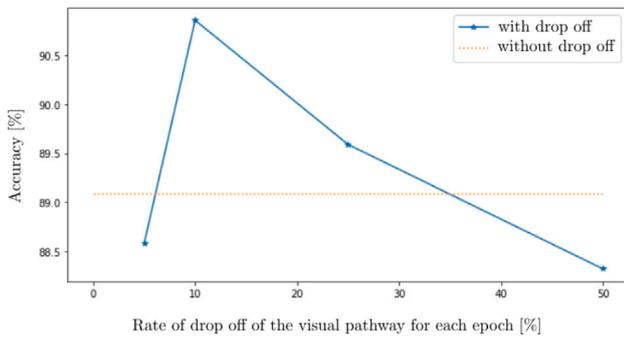


Fig. 6. Accuracy of the event recognition of the AVE dataset when using different rates of dropping the weight update of the visual path during training.

4.5. Model analysis and discussion

In this section, we report studies to identify the impact of each module of the MAFnet. We work with the AVE dataset as the dataset assures the presence of the two modalities. We analyze the training method, the impact of the temporal attention, the modality attention and the combination of the two attentions, compare different fusion techniques and the impact of the modality conditioning.

**Training method: Drop off.** Visual and audio paths do not have the same learning speed. Even without the additional convolutional layers comprised in FiLM, the training of the audio path needs more iterations than for the visual training. Inspired from [34], we investigate a new multimodal training technique by randomly dropping the update of the visual weights to allow the audio path to train longer. In Fig. 6, we report the accuracy in function of the dropping rate of the weight update of the visual path. Dropping too often the visual path decreases results compared to training without dropping. We suppose that the visual path is not trained enough. It is also observed that not dropping enough the visual path gives also slightly poorer results. We suppose that the network cannot exploit enough sound information. Furthermore, it loses visual information as the visual path is not trained as much as needed.

**Fusion techniques comparison.** MAFnet creates a multimodal feature by concatenating the information coming from the visual and audio paths. We analyze the importance of using unimodal information versus multimodal information in the case of event recognition. We also test different fusion techniques present in the literature to determine the best fusion method: addition, concatenation, multimodal compact bilinear pooling (MCB) [31] and the multimodal residual fusion (DMR) [19].

As we want to test the fusion techniques, the experiments are made without the FiLM layer. In the case of the unimodal network, the network comprises only the temporal attention module without the modality attention module as only one modality is present.

From the results in Table 2, and as expected, we conclude that the dataset is easier to classify using visual information only than sound information only. Multimodal information increases the performance compared to unimodal. Concatenation has the best result and is even slightly better than more complex fusion techniques like MCB or DMR.

**Attention analysis.** In this section, we analyze the impact of each attention module. Table 3 presents the event recognition results without attention, with temporal attention only and with modality attention only. Again, the network does not comprise the FiLM layer for this ablation study. The temporal attention allows to take into account the temporal context and dynamically highlights particular time windows. Not each time window comprises relevant information for the classification. The modality attention highlights a modality. Indeed, depending

Table 2

Comparison unimodal versus multimodal event recognition and the use of different fusion techniques on AVE dataset.

Fusion technique	Accuracy [%]
Visual	75.63
Audio	69.29
Addition	84.77
Concatenation	<b>89.34</b>
MCB	88.83
DMR	87.56

Table 3

Ablation study of the modality & temporal attention module on the AVE dataset.

Attention type	Accuracy [%]
Without attention	87.82
Temporal attention	88.92
Modality attention	88.66
Modality & temporal attention	<b>89.34</b>

Table 4

Evaluation of the lateral connection between visual and audio paths with FiLM layer on the AVE dataset.

Modality conditioning	Accuracy [%]
Add residual block without modality conditioning	86.55
Conditioning of both modalities	87.62
Conditioning of audio by vision	<b>90.86</b>
Conditioning of vision by audio	90.61

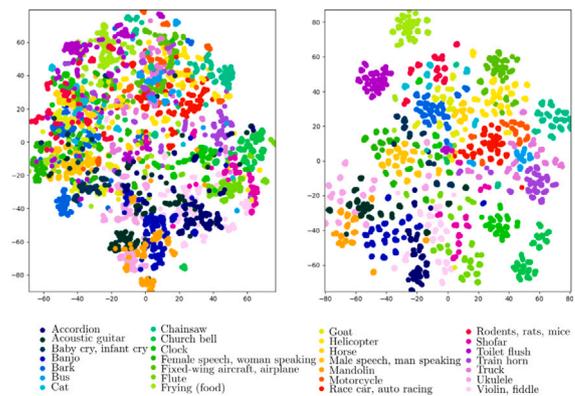


Fig. 7. T-SNE visualization of the embedding of the residual block just before (left) and after (right) the FiLM layer in the audio path.

on the video a modality can have more importance than the other. Each attention module has a positive impact on the accuracy and the combination of both attentions gets the best result.

**Modality conditioning analysis.** In this section, we analyze the impact of the lateral connection, the modality conditioning ( Table 4). It is observed that adding FiLM in visual and audio path provides better results than without any conditioning. However, conditioning only one modality is better than conditioning both modalities whatever the conditioned modality. The best result is obtained when visual conditions audio. We suppose that vision contributes more to audio than the other way. This is in accordance with previous observations in [43]. When one modality can statistically drive the other, it might be more efficient to use only one FiLM layer that projects from the driving modality to the other. In that situation, a second FiLM layer to drive attention from the dependent modality to the driving one might be counterproductive.

Fig. 7 compares the embedding of the residual block just before and after the FiLM layer in the audio path (Fig. 4). We use average

pooling and t-SNE [54] to reduce the embedding dimension to 2D. We observe a better clustering of the different classes after including the visual information in the audio path.

## 5. Conclusion

We proposed the Multi-level Attention Fusion (MAFnet) in the context of event recognition task. Our network includes a modality & temporal attention module. It dynamically associates a score to each modality at each time window to highlight the relevant modality and time window. To go further than a simple late fusion, we condition one modality with the other with a FiLM layer. It highlights selected audio feature maps based on visual information. Finally, to take into account the different learning dynamics of each modality, we randomly drop the weight update of the visual path. We evaluate our network on three datasets and achieve better accuracy than the current audio-visual models.

MAFnet shows promising results for audiovisual data fusion in the context of event classification. When exploiting audio-visual data, the fusion of the two modalities is not the only important element, the conditioning between the modalities paths is necessary to make the best use of the audio-visual information. Conditioning with the FiLM layer modifies the hidden representation of audio modalities based on visual information. In view of the conditioning results, future research should investigate the interaction between the audio and visual path at different levels of the architecture as well as study different conditioning methods.

## CRedit authorship contribution statement

**Mathilde Brousmiche:** Conceptualization, Methodology, Investigation, Software, Writing – original draft. **Jean Rouat:** Resources, Writing – review & editing, Supervision, Funding acquisition. **Stéphane Dupont:** Resources, Writing – review & editing, Supervision, Funding acquisition.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work was funded by CHISTERA IGLU, Canada project and the European Regional Development Fund (ERDF). The authors would also like to thank the NVIDIA Corporation for donating GPUs.

## References

- [1] J.-Y. Dufour, *Intelligent Video Surveillance Systems*, John Wiley & Sons, 2012.
- [2] R. Kala, *On-Road Intelligent Vehicles: Motion Planning for Intelligent Transportation Systems*, Butterworth-Heinemann, 2016.
- [3] T. D'Orazio, M. Leo, A review of vision-based systems for soccer video analysis, *Pattern Recognit.* 43 (8) (2010) 2911–2926.
- [4] D.C. Gibbon, Z. Liu, *Introduction to Video Search Engines*, Springer Science & Business Media, 2008.
- [5] E. Real, J. Shlens, S. Mazzocchi, X. Pan, V. Vanhoucke, Youtube-boundingboxes: A large high-precision human-annotated data set for object detection in video, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2017*, pp. 5296–5305.
- [6] J. Carreira, A. Zisserman, Quo vadis, action recognition? a new model and the kinetics dataset, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2017*, pp. 6299–6308.
- [7] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, L. Fei-Fei, Large-scale video classification with convolutional neural networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2014*, pp. 1725–1732.
- [8] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, 2014, arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).

- [9] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, Z. Wojna, Rethinking the inception architecture for computer vision, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2016*, pp. 2818–2826.
- [10] K. Simonyan, A. Zisserman, Two-stream convolutional networks for action recognition in videos, in: *Advances in Neural Information Processing Systems, NIPS, 2014*, pp. 568–576.
- [11] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, T. Darrell, Long-term recurrent convolutional networks for visual recognition and description, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2015*, pp. 2625–2634.
- [12] D. Tran, L. Bourdev, R. Fergus, L. Torresani, M. Paluri, Learning spatiotemporal features with 3D convolutional networks, in: *Proceedings of the IEEE International Conference on Computer Vision, ICCV, 2015*, pp. 4489–4497.
- [13] E. Kazakos, A. Nagrani, A. Zisserman, D. Damen, Epic-fusion: Audio-visual temporal binding for egocentric action recognition, in: *Proceedings of the IEEE International Conference on Computer Vision, ICCV, 2019*, pp. 5492–5501.
- [14] R. Arandjelovic, A. Zisserman, Look, listen and learn, in: *Proceedings of the IEEE International Conference on Computer Vision, ICCV, 2017*, pp. 609–617.
- [15] R. Arandjelovic, A. Zisserman, Objects that sound, in: *Proceedings of the European Conference on Computer Vision, ECCV, 2018*, pp. 435–451.
- [16] Y. Aytar, C. Vondrick, A. Torralba, Soundnet: Learning sound representations from unlabeled video, in: *Advances in Neural Information Processing Systems, NIPS, 2016*, pp. 892–900.
- [17] A. Owens, A.A. Efros, Audio-visual scene analysis with self-supervised multisensory features, in: *Proceedings of the European Conference on Computer Vision, ECCV, 2018*, pp. 631–648.
- [18] A. Owens, J. Wu, J.H. McDermott, W.T. Freeman, A. Torralba, Ambient sound provides supervision for visual learning, in: *Proceedings of the European Conference on Computer Vision, ECCV, Springer, 2016*, pp. 801–816.
- [19] Y. Tian, J. Shi, B. Li, Z. Duan, C. Xu, Audio-visual event localization in unconstrained videos, in: *Proceedings of the European Conference on Computer Vision, ECCV, 2018*, pp. 247–263.
- [20] Y. Tian, Y. Cao, J. Wu, W. Hu, C. Song, T. Yang, Multi-cue combination network for action-based video classification, *IET Comput. Vis.* 13 (6) (2019) 542–548.
- [21] X. Long, C. Gan, G. De Melo, X. Liu, Y. Li, F. Li, S. Wen, Multimodal keyless attention fusion for video classification, in: *Thirty-Second AAAI Conference on Artificial Intelligence, 2018*.
- [22] K. Soomro, A.R. Zamir, M. Shah, UCF101: A Dataset of 101 human actions classes from videos in the wild, 2012, arXiv preprint [arXiv:1212.0402](https://arxiv.org/abs/1212.0402).
- [23] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, G. Toderici, Beyond short snippets: Deep networks for video classification, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2015*, pp. 4694–4702.
- [24] C.-Y. Ma, M.-H. Chen, Z. Kira, G. AlRegib, TS-LSTM And temporal-inception: Exploiting spatiotemporal dynamics for activity recognition, *Signal Process., Image Commun.* 71 (2019) 76–87.
- [25] Z. Li, K. Gavriljuk, E. Gavves, M. Jain, C.G. Snoek, VideoLSTM convolves, attends and flows for action recognition, *Comput. Vis. Image Underst.* 166 (2018) 41–50.
- [26] D. Tran, H. Wang, L. Torresani, J. Ray, Y. LeCun, M. Paluri, A closer look at spatiotemporal convolutions for action recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2018*, pp. 6450–6459.
- [27] S. Xie, C. Sun, J. Huang, Z. Tu, K. Murphy, Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification, in: *Proceedings of the European Conference on Computer Vision, ECCV, 2018*, pp. 305–321.
- [28] C. Feichtenhofer, A. Pinz, A. Zisserman, Convolutional two-stream network fusion for video action recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2016*, pp. 1933–1941.
- [29] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, L. Van Gool, Temporal segment networks: Towards good practices for deep action recognition, in: *European Conference on Computer Vision, Springer, 2016*, pp. 20–36.
- [30] X. Long, C. Gan, G. De Melo, J. Wu, X. Liu, S. Wen, Attention clusters: Purely attention based local feature integration for video classification, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2018*, pp. 7834–7843.
- [31] Y. Gao, O. Beijbom, N. Zhang, T. Darrell, Compact bilinear pooling, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016*, pp. 317–326.
- [32] Y.-B. Lin, Y.-J. Li, Y.-C.F. Wang, Dual-modality seq2seq network for audio-visual event localization, in: *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, IEEE, 2019*, pp. 2002–2006.
- [33] Y. Wu, L. Zhu, Y. Yan, Y. Yang, Dual attention matching for audio-visual event localization, in: *Proceedings of the IEEE International Conference on Computer Vision, ICCV, 2019*, pp. 6292–6300.
- [34] F. Xiao, Y.J. Lee, K. Grauman, J. Malik, C. Feichtenhofer, Audiovisual SlowFast networks for video recognition, 2020, arXiv preprint [arXiv:2001.08740](https://arxiv.org/abs/2001.08740).
- [35] V. Mnih, N. Heess, A. Graves, et al., Recurrent models of visual attention, in: *Advances in Neural Information Processing Systems, NIPS, 2014*, pp. 2204–2212.
- [36] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, in: *Advances in Neural Information Processing Systems, NIPS, 2017*, pp. 5998–6008.

- [37] J. Wang, X. Peng, Y. Qiao, Cascade multi-head attention networks for action recognition, *Comput. Vis. Image Underst.* (2020) 102898.
- [38] J. Li, X. Liu, W. Zhang, M. Zhang, J. Song, N. Sebe, Spatio-temporal attention networks for action recognition and detection, *IEEE Trans. Multimed.* (2020).
- [39] X. Huang, S. Belongie, Arbitrary style transfer in real-time with adaptive instance normalization, in: *Proceedings of the IEEE International Conference on Computer Vision, ICCV, 2017*, pp. 1501–1510.
- [40] T. Kim, I. Song, Y. Bengio, Dynamic layer normalization for adaptive neural acoustic modeling in speech recognition, in: *Proceedings of Interspeech, 2017*, pp. 2655–2659.
- [41] H. De Vries, F. Strub, J. Mary, H. Larochelle, O. Pietquin, A.C. Courville, Modulating early visual processing by language, in: *Advances in Neural Information Processing Systems, NIPS, 2017*, pp. 6594–6604.
- [42] J. Abdelnour, G. Salvi, J. Rouat, From visual to acoustic question answering, 2019, *arXiv preprint arXiv:1902.11280*.
- [43] M. Brousmiche, J. Rouat, S. Dupont, Audio-visual fusion and conditioning with neural networks for event recognition, in: *IEEE International Workshop on Machine Learning for Signal Processing, MLSP, IEEE, 2019*, pp. 1–6.
- [44] E. Perez, F. Strub, H. De Vries, V. Dumoulin, A. Courville, Film: Visual reasoning with a general conditioning layer, in: *Thirty-Second AAAI Conference on Artificial Intelligence, 2018*.
- [45] G. Huang, Z. Liu, L. Van Der Maaten, K.Q. Weinberger, Densely connected convolutional networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017*, pp. 4700–4708.
- [46] 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: *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, IEEE, 2017*, pp. 131–135.
- [47] E.B. Goldstein, J. Brockmole, *Sensation and Perception*, Cengage Learning, 2016.
- [48] P. Zhou, W. Yang, W. Chen, Y. Wang, J. Jia, Modality attention for end-to-end audio-visual speech recognition, in: *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, IEEE, 2019*, pp. 6565–6569.
- [49] W. Wang, D. Tran, M. Feiszli, What makes training multi-modal networks hard? 2019, *arXiv preprint arXiv:1905.12681*.
- [50] C. Feichtenhofer, H. Fan, J. Malik, K. He, SlowFast Networks for video recognition, in: *Proceedings of the IEEE International Conference on Computer Vision, ICCV, 2019*, pp. 6202–6211.
- [51] N. Crasto, P. Weinzaepfel, K. Alahari, C. Schmid, MARS: MOTion-augmented RGB stream for action recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2019*, pp. 7882–7891.
- [52] J.F. Gemmeke, D.P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R.C. Moore, M. Plakal, M. Ritter, Audio set: An ontology and human-labeled dataset for audio events, in: *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, IEEE, 2017*, pp. 776–780.
- [53] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al., Tensorflow: A system for large-scale machine learning, in: *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*, 2016, pp. 265–283.
- [54] L.v.d. Maaten, G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* 9 (Nov) (2008) 2579–2605.