Large Scale Audiovisual Learning of Sounds with Weakly Labeled Data

Haytham M. Fayek* and Anurag Kumar*

Facebook Reality Labs, Redmond, WA, USA {haythamfayek, anuragkr}@fb.com

Abstract

Recognizing sounds is a key aspect of computational audio scene analysis and machine perception. In this paper, we advocate that sound recognition is inherently a multi-modal audiovisual task in that it is easier to differentiate sounds using both the audio and visual modalities as opposed to one or the other. We present an audiovisual fusion model that learns to recognize sounds from weakly labeled video recordings. The proposed fusion model utilizes an attention mechanism to dynamically combine the outputs of the individual audio and visual models. Experiments on the large scale sound events dataset, AudioSet, demonstrate the efficacy of the proposed model, which outperforms the single-modal models, and state-of-theart fusion and multi-modal models. We achieve a mean Average Precision (mAP) of 46.16 on Audioset, outperforming prior state of the art by approximately +4.35 mAP (relative: 10.4%).

1 Introduction

Sound recognition is key for intelligent agents that perceive and interact with the world, similar to how the perception of different sounds is critical for humans to understand and interact with the environment. Sounds are primarily produced by actions on or interactions between objects, and hence we often learn to immediately associate sounds with visual objects and entities. This visual association is often necessary for recognizing and understanding the acoustic activities occurring around us. For example, it is easier to distinguish between sounds produced by a *Vacuum Cleaner* and sounds produced by a *Hair Dryer* by listening to and seeing the appliances, as opposed to one or the other.

Sound recognition is inherently a multi-modal audiovisual task. Thus, we ought to build machine learning models for sound recognition that are multi-modal, inspired by how humans perceive sound. In addition to reducing uncertainties in certain situations as instantiated above, audiovisual learning of sounds can lead to a more holistic understanding of the environment. For example, an alarm clock can be designed to produce *Church Bell* sounds; while an audio only sound recognition model can perhaps label the sound as *Church Bells*, which would not be incorrect, it nevertheless does not represent the actual event. Hence, systems designed to do audiovisual perception of sounds would lead to better and more complete understanding of these phenomena. However, most prior work on sound recognition focuses on learning to recognize sounds only from audio [Virtanen *et al.*, 2017; Kumar, 2018].

Multi-modal learning approaches, which aim to train a single model by using the audio and visual modalities simultaneously are one class of well established audiovisual learning approaches. However, these methods present their own set of challenges. In these methods, it is often necessary to appropriately synchronize the audio and visual components and ensure temporal consistencies for a single model to consume both modalities. Moreover, training a single multi-modal model which uses both modalities is a difficult feat due to different learning dynamics for the audio and visual modalities [Wang *et al.*, 2019a]. Further, training and inference in multi-modal models is computationally cumbersome, particularly for large datasets [Gao *et al.*, 2019].

Another class of audiovisual learning methods can be described as *Fusion* approaches. Fusion approaches aim to combine two independent single-modal models. This offers several advantages. We can build strong individual models for each modality separately, with the freedom to design modal-specific models without requiring to factor in the unnecessary complexities arising out of multi-modal models and training. Furthermore, fusion approaches are often more interpretable than multi-modal models as it is easier to qualitatively analyze and understand the contributions of each modality through the individual models. Finally, fusion approaches allow combining multiple models per modality, which can be advantageous in improving performance.

In this paper, we propose fusion based approaches for audiovisual learning of sounds. We propose methods that learn to combine individual audio and visual models that were respectively trained on each modality separately. We first build state-of-the-art audio and visual sound recognition systems. We then propose attention fusion models to dynamically combine these audio and visual models. Specifically, our fusion models learn to pay attention to the appropriate modality in a *sample-specific class-specific* manner. Our models are de-

^{*}Equal contribution.

signed for weakly supervised learning, where we train models using weakly labeled data [Kumar and Raj, 2016a]. We analyze our proposed models on the largest dataset for sound events, Audioset [Gemmeke *et al.*, 2017], and show that the proposed models outperform state-of-the-art single-modal models, baseline fusion models, and multi-modal models. The results and ensuing analysis attest to the importance of using both the audio and visual modalities for sound recognition and efficacy of our attention fusion models.

2 Related Work

Learning Sound Events. Numerous prior works have proposed a variety of supervised methods for detecting and classifying sound events [Virtanen *et al.*, 2017; Kumar, 2018]. Large-scale sound event detection has been possible primarily through weakly supervised learning [Kumar and Raj, 2016a] and the release of large-scale weakly labeled sound events datasets, such as *Audioset* [Gemmeke *et al.*, 2017]. Most of the recent methods rely on deep neural networks, in particular, deep Convolutional Neural Networks (ConvNets) [Kumar and Ithapu, 2020; Adavanne *et al.*, 2019; Kong *et al.*, 2019; Ford *et al.*, 2019]. While *mean* and *max* pooling have been shown to be effective in handling weakly labeled data [Kumar, 2018], more involved methods such as adaptive pooling [McFee *et al.*, 2018] or attention based pooling [Wang *et al.*, 2019b; Kong *et al.*, 2019] have also been proposed.

Audiovisual Learning. There is a profusion of work on multi-modal audiovisual learning for video recognition and understanding in recent years [Baltrušaitis et al., 2018]. Progress in multi-modal learning was facilitated by the availability of large-scale video datasets [Kay et al., 2017; Gemmeke et al., 2017; Gu et al., 2018], advances in selfsupervised learning of multi-modal representations that exploit cross-modal information [Aytar et al., 2016; Arandjelovic and Zisserman, 2017; Owens and Efros, 2018; Korbar et al., 2018], custom loss functions that synchronize learning multiple modalities [Wang et al., 2019a], and others. Despite the recent progress, multi-modal audiovisual learning remains a challenging problem. Primary reasons include difficulties in designing multi-modal neural network architectures that work well across various datasets and tasks, jointly learning multiple modalities, training and inference efficiency, etc.

Audiovisual Fusion. Late fusion approaches are a popular to combine audio and visual models [Ghanem *et al.*, 2018]. Unweighted late fusion such as averaging the outputs of the audio and visual models can demonstrate competitive performance in addition to not requiring any additional parameters and training beyond the audio and visual single-modal models [Lan *et al.*, 2012]. Nevertheless, weighted late or mid-level fusion has been extensively studied as it provides avenues to adaptively combine information between the different modalities [Lan *et al.*, 2012; Wu *et al.*, 2016; Kazakos *et al.*, 2019]. Further, attention mechanisms [Bahdanau *et al.*, 2014] were used to dynamically combine the outputs of the single-modal models [Long *et al.*, 2018; Sterpu *et al.*, 2018; Hori *et al.*, 2018; Zhou *et al.*, 2019].

Audiovisual Learning of Sounds Prior work on audiovisual learning of sounds, whether through single-model multimodal learning or through fusion of individual models independently trained on the respective modalities, is scarce. [Wang et al., 2019a] describes a multi-modal learning approach in which a single model is trained jointly for both the audio and visual modalities and applies this to sound recognition. Such methods are susceptible to several difficulties as outlined in Section 1. We, on the other hand, propose fusion models for audiovisual learning of sounds and show in Section 2 (Table 3) the effectiveness of our fusion models compared with the multi-modal models in [Wang et al., 2019a]. Another work related to ours is [Parekh et al., 2019], which attempts to learn representations from weakly labeled audiovisual data and use these representations in different tasks including sound event recognition; therein, the audio and visual models are only combined through average fusion. Perhaps worth mentioning are other works such as as [Aytar et al., 2016; Arandjelovic and Zisserman, 2017; Alwassel et al., 2019]; they attempted to use the visual modality to learn representations for the audio modality that can be used for downstream sound event classification tasks. However, in all prior work, the downstream sound classification tasks are done on small-scale datasets, which does not provide in-depth understanding of audiovisual learning for sounds at scale.

3 Audiovisual Models for Sounds

We learn sound events in the weakly supervised setting, where only recording-level labels are provided for each training recording, with no temporal information. The fundamentals of weakly supervised learning of sound events are based on Multiple Instance Learning (MIL) [Kumar and Raj, 2016a]. In MIL, the learning problem is formulated in terms of *bags* and respective labels, (\mathcal{B}, Y) ; each *bag* is a collection of instances [Andrews et al., 2003]. A bag is labeled as positive for a class, Y = 1, if at least one instance within the bag is positive. On the other hand, if all instances within the bag are negative for a given class, then the *label* is negative, Y = 0. For weakly labeled sound recognition, recordings which are tagged with the presence of a sound class become a positive bag for that class and negative otherwise. The audio and visual models in Sections 3.1 and 3.2 follow this formulation.

3.1 Audio Model

In the MIL setting for weakly labeled sound event recognition, each recording becomes a *bag* as required in the MIL framework. The instances in the *bag* are segments of the recording. Assuming that each instance is represented by a feature vector \mathbf{x} ; the training data takes the form of (\mathcal{B}_i, Y_i) , i = 1 to *n*, where the *i*th bag consists of m_i instances $\mathcal{B}_i = {\mathbf{x}_{i1} \cdots \mathbf{x}_{im_i}}$.

In the audio model, the key idea is that the learner can be trained by first mapping the instance-level predictions to the corresponding bag-level predictions, and then these bag-level predictions can be used to compute a loss function to be minimized, akin to the classic MI-SVM approach for MIL [Andrews *et al.*, 2003]. Herein, the *max* function was used to aggregate the instance-level predictions.

Let f be the function we wish to learn; then the training involves minimizing the following loss function.

$$\mathcal{L}(\Theta, \Phi) = \sum_{i=1}^{n} l(g_{\Phi}(f_{\Theta}(\mathbf{x}_{i1}), \cdots, f_{\Theta}(\mathbf{x}_{im_i})), Y_i) \quad (1)$$

where, l is the loss function measuring divergence between predictions and targets, and Θ represents learnable parameters of f. Note that the function g, used to aggregate the instance-level predictions from f to the bag-level predictions, can itself contain learnable parameters Φ .

The audio model is a ConvNet that maps Log-scaled Melfilter-bank feature representations of the entire audio recording to (multiple) labels, trained by minimizing the loss function defined in Equation 1. The sampling rate of all audio recordings is 16 kHz; 64 Mel-filter-bank representations are obtained for 16 ms windows in the audio recording, shifted by 10 ms, which amounts to 100 frames per second of audio.

The overall architecture of the network is detailed in Table 1. The network produces *instance* and *bag* representations for any given input at Block B5. The bag consists of 2048-dimensional instance representations. The network is designed for a receptive field of ~1 second (96 frames). Hence, it outputs 2048-dimensional features for ~1 second of audio, every ~0.33 seconds (32 frames). Three 1×1 convolutional layers (B6 to B8) are then used to predict outputs for each segment, which are then aggregated via the layer *G* that uses global average pooling. The number of segments depends on the size of the input. As shown in Table 1, for an input with 1024 frames, 30 segments are produced.

The loss function for the i^{th} training audio recording is defined as:

$$l(\mathcal{B}_{i}, Y_{i}) = \frac{1}{C} \sum_{c=1}^{C} h(Y_{c}^{i}, p_{c}^{i})$$
(2)

where, p_c^i is the output for the c^{th} class and $h(Y_c^i, p_c^i) = -Y_c^i * \log(p_c^i) - (1 - Y_c^i) * \log(1 - p_c^i)$ is the binary cross entropy loss function. Furthermore, our audio model training also incorporates the sequential co-supervision method described in [Kumar and Ithapu, 2020].

3.2 Visual Model

Similar to the audio model, the visual model for sounds is based on the MIL framework. Herein, the entire video is the bag and instances of the bag are the frames of the video. Specifically, we sample 64 frames from the videos to form the bags. Each frame (instance) is then represented by 2048feature representations obtained from a ResNet-152 [He *et al.*, 2016], pre-trained on Imagenet [Deng *et al.*, 2009]. This yields the video bag representations for each recording. We obtain a single 2048-dimensional vector representation for each bag by averaging the feature representations for each instance. This idea of mapping bags to a single vector representations was previously employed in [Kumar and Raj, 2016b; Wei *et al.*, 2014] to make MIL algorithms scalable.

We then train a 4 hidden layer Deep Fully Connected Neural Network (DNN) to recognize sound events. The number

| Stage | Layers | Output Size |
|-----------------------------|------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------|
| Input | Unless specified – (S)tride = 1, (P)adding = 1 | $1\times 1024\times 64$ |
| Block B1 | Conv: $64, 3 \times 3$ Conv: $64, 3 \times 3$ Pool: 4×4 (S:4) | $\begin{array}{c} 64\times1024\times64\\ 64\times1024\times64\\ 64\times256\times16 \end{array}$ |
| Block B2 | Conv: 128, 3 × 3 Conv: 128, 3 × 3 Pool: 2 × 2 (S:2) | $\begin{array}{c} 128 \times 256 \times 16 \\ 128 \times 256 \times 16 \\ 128 \times 128 \times 8 \end{array}$ |
| Block B3 | Conv: 256, 3 × 3 Conv: 256, 3 × 3 Pool: 2 × 2 (S:2) | $\begin{array}{c} 256 \times 128 \times 8 \\ 256 \times 128 \times 8 \\ 256 \times 64 \times 4 \end{array}$ |
| Block B4 | Conv: 512, 3 × 3 Conv: 512, 3 × 3 Pool: 2 × 2 (S:2) | $\begin{array}{c} 512\times 64\times 4\\ 512\times 64\times 4\\ 512\times 32\times 2\end{array}$ |
| Block B5 | Conv: 2048, 3 × 2 (P:0) | $2048\times 30\times 1$ |
| Block B6 | Conv: 1024, 1×1 | $1024\times 30\times 1$ |
| Block B7 | Conv: 1024, 1×1 | $1024\times 30\times 1$ |
| Block B8 | Conv: C, 1×1 | $C\times 30\times 1$ |
| $\mathrm{G}\left(g() ight)$ | Global average pooling | $C \times 1$ |

Table 1: Audio Convolutional Neural Network (ConvNet) architecture. The number of filters and size of each filter varies between convolutional layers; e.g., Conv: 64, 3×3 denotes a convolutional layer of 64 filters, each of size 3×3 . C denotes the number of classes. All convolutional layers except in B8 are followed by batch normalization and Rectified Linear Units (ReLU); the convolutional layer in B8 is followed by a sigmoid function.

of units in each layer are 2048, 2048, 1024, 1024, C respectively, where C is the number of classes. The first and second hidden layers are followed by a dropout layer with a dropout rate of 0.3. Rectified Linear Units (ReLU) activation is used in all layers except for the last layer, where sigmoid activation is used. Similar to the audio model, the network is trained with the binary cross entropy loss function.

3.3 Baseline Audiovisual Fusion

Our audiovisual learning of sounds focuses on developing fusion models that combine outputs from models trained individually on the audio and visual modalities. The following is a description of a number of baseline fusion methods.

Average Fusion. We simply compute the arithmetic mean of the outputs of the audio and visual models. Despite its simplicity, average fusion serves as a strong baseline. [Parekh *et al.*, 2019] used average fusion to combine audio and visual predictions.

Regression. We train a linear layer with a sigmoid activation function that ingests the (mean-standard-deviation) normalized outputs of the audio and visual models, and predicts the probability of the sound classes. The model is L2 regularized with the regularization parameter set to $\lambda = 10^{-5}$.

Multi-Layer Perceptron (MLP). We train a fully connected neural network with a single hidden layer that ingests concatenated outputs of the audio and visual models, and predicts the probability of the sound classes. The hidden layer comprises 512 units with batch normalization and ReLU activation. Furthermore, a dropout layer with a dropout rate of 0.5 is used. The output layer is a fully connected layer with *C* units and a sigmoid activation function.



Figure 1: Audiovisual fusion for sound recognition. The audio and visual models map the respective input to segment-level representations which are then used to obtain single-modal predictions, \mathbf{h}_a and \mathbf{h}_v , respectively. The attention fusion function, n_{attn} , ingests the single-modal predictions, \mathbf{h}_a and \mathbf{h}_v , to produce weights for each modality, $\boldsymbol{\alpha}_a$ and $\boldsymbol{\alpha}_v$. The single-modal audio and visual predictions, \mathbf{h}_a and \mathbf{h}_v , are mapped to $\tilde{\mathbf{h}}_a$ and $\tilde{\mathbf{h}}_v$ via functions, n_a and n_v , respectively, and fused using the attention weights, $\boldsymbol{\alpha}_a$ and $\boldsymbol{\alpha}_v$. The fused output, $\tilde{\mathbf{h}}_{av}$, is then mapped to multi-modal multi-label predictions via function n_{av} .

3.4 Attention Audiovisual Fusion

The audio and visual models have been separately trained to map the audio and visual modalities, respectively, to the various sound classes. The performance of each class varies between modalities. This variation can be sample-specific as well. For a given sound class, in some videos, the audio modality might be sufficient to correctly recognize the sound, whereas in other videos, the visual modality might be more suited. Thus, we ought to fuse the outputs from the audio and visual models in a sample-specific and class-specific manner. To this end, we design an attention mechanism to combine the audio and visual models.

Figure 1 shows the general schema of the attention fusion mechanism. Let \mathbf{h}_a and \mathbf{h}_v represent the outputs of the audio and visual models for a given video recording respectively. We learn an attention function, $a(\cdot)$, that generates weights α_a and α_v for the audio and visual modalities respectively. We constrain the attention weights to sum to 1 for each class, i.e., $\alpha_a + \alpha_v = 1$, where 1 is a vector of ones.

The attention function is implemented as a neural network, n_{attn} :

$$\boldsymbol{\alpha}_a = n_{attn}([\mathbf{h}_a, \mathbf{h}_v], W_{attn}) \tag{3}$$

where, W_{attn} denotes the learnable parameters of network n_{attn} , and $[\mathbf{h}_a, \mathbf{h}_v]$ denotes the concatenation of \mathbf{h}_a and \mathbf{h}_v . The output layer in n_{attn} has a sigmoid activation function.

The final output that combines the audio and visual models is obtained as follows:

$$\tilde{\mathbf{h}}_a = n_a(\mathbf{h}_a, W_a^n)$$
, $\tilde{\mathbf{h}}_v = n_v(\mathbf{h}_v, W_v^n)$ (4)

$$\mathbf{h}_{av} = \boldsymbol{\alpha}_a \odot \mathbf{h}_a + (1 - \boldsymbol{\alpha}_a) \odot \mathbf{h}_v \tag{5}$$

$$\mathbf{o}_{av} = n_{av}(\mathbf{h}_{av}, W_{av}^n) \tag{6}$$

where, n_a and n_v are neural networks that process inputs \mathbf{h}_a and \mathbf{h}_v into $\tilde{\mathbf{h}}_a$ and $\tilde{\mathbf{h}}_v$, respectively, which are then combined through attention weights (α_a and $\alpha_v = 1 - \alpha_a$), as in Equation 5. W_a^n and W_v^n are the learnable parameters of n_a and n_v respectively. The combined outputs, $\tilde{\mathbf{h}}_{av}$, are passed through another neural network n_{av} with a sigmoid activation function to produce the final outputs \mathbf{o}_{av} .

Note that n_a and n_v are generic functions representing all possible transforms: linear, non-linear, or even the identity of \mathbf{h}_a and \mathbf{h}_v respectively. For the identity case, $n(\mathbf{h}) = \mathbf{h}$, and n does not contain any learnable parameters. For the linear transform case, n does not contain any non-linear activation functions. Similarly, n_{av} might just be a sigmoid function, i.e., $n_{av}(\tilde{\mathbf{h}}_{av}) = 1/(1 + \exp^{-\tilde{\mathbf{h}}_{av}})$, and hence, containing no learnable parameters. Otherwise, it can be a neural network with one or more layers with a sigmoid activation at the output layer. We also attempted to learn the attention weights with a single modality, either \mathbf{h}_a or \mathbf{h}_v instead of the combined $[\mathbf{h}_a, \mathbf{h}_v]$. In this case, α_a is either obtained through self-attention $\alpha_a = n_{attn}(\mathbf{h}_a, W_{attn})$ or through cross-modal attention $\alpha_a = n_{attn}(\mathbf{h}_v, W_{attn})$. Similarly, α_v is obtained, and both α_a and α_v are normalized to sum to 1 for each class. However, attention weights learned from single modality was found to be inferior to learning them through both modalities as in Equation 3.

4 **Experiments**

Experiments are conducted on AudioSet which is described in Section 4.1. The experiments and results are presented in Section 4.2.

4.1 Audioset Dataset

Audioset [Gemmeke *et al.*, 2017] is the largest dataset for sound events. The dataset provides YouTube videos for 527 sound events. Each video clip is approximately 10s in length annotated by a human with multiple labels denoting the sound events present in the video clip. The average number of labels per video clip is 2.7. The dataset is weakly labeled as the labels of each video clip denote the absence or presence of sound events but do not contain any temporal information. The distribution of sound event classes in the training set is severely unbalanced, ranging from around 1 million videos for the most represented class, *Music*, to approximately 120 videos for the least represented class, *Screech*.

We use the predefined *Unbalanced* set for training and *Eval* set for performance evaluation. The training set comprises

| Model | mAP | mAUC |
|-------------------|-------|--------------|
| Audio | 38.35 | 97.12 |
| Visual | 25.73 | 91.30 |
| Average Fusion | 42.84 | 97.63 |
| Regression Fusion | 43.10 | 95.35 |
| MLP Fusion | 45.60 | 96.83 |
| Attention Fusion | 46.16 | 97.51 |

Table 2: Comparison of mAP and mAUC for different fusion methods for combining audio and visual models.

| Model | mAP | mAUC |
|----------------------------------------------------|--------------|--------------|
| Audio (ResNet-50) [Ford <i>et al.</i> , 2019] | 38.0 | 97.10 |
| Audio (Ours) | 38.35 | 97.12 |
| Visual (R(2+1)D-101) [Wang et al., 2019a] | 18.80 | 91.80 |
| Visual (Ours) | 25.73 | 91.30 |
| AudioVisual (G-Blend) [Wang <i>et al.</i> , 2019a] | 41.80 | 97.50 |
| AudioVisual (Ours) | 46.16 | 97.51 |

Table 3: mAP and mAUC for state-of-the-art audio, visual, and audiovisual sound recognition models on AudioSet.

approximately 2 million videos, whereas the evaluation set comprises approximately 20,000 videos. The evaluation set contains at least 59 videos for each class. In practice, the total video count is slightly less than the numbers listed above due to the unavailability of links when this work was carried out. We sample approximately 25,000 videos from the training set to use as the validation set.

4.2 Performance Evaluation of Fusion Methods

Experimental Setup. The outputs (predictions) of the audio and visual models are used as input to all fusion models. These outputs are normalized to zero mean and unit standard deviation prior to feeding them into the fusion models (except for average fusion). The mean and standard deviation are computed using the training set only.

The details of the audio model training can be found in [Kumar and Ithapu, 2020]. The network in the video model is trained for 20 epochs using the Adam optimizer [Kingma and Ba, 2014] with a mini-batch size of 144. Hyperparameters such as the learning rate are selected using the validation set, which is also used to select the best model during training.

In the fusion experiments, all neural networks are trained using Adam for 100 epochs. The mini-batch size is set to 256. The validation set is used to select the best hyperparameters, such as the learning rate for each model, and the best model during training. Similar to prior work on Audioset, we use Average Precision (AP) and Area Under Receiver Operating Characteristic (ROC) Curve (AUC) to measure the performance per class. Mean Average Precision (mAP) and mean Area Under ROC Curve (mAUC) computed over all 527 classes are used for comparison.

Comparison of Fusion Methods. Table 2 summarizes the results. The audio model achieves 38.35 mAP and 97.12

| Model | mAP | mAUC |
|----------------------------------------------|-------|-------|
| Attention (self unimodal) | 46.04 | 97.48 |
| Attention (cross unimodal) | 46.06 | 97.47 |
| Attention (multi-modal w/o n_a and n_v) | 38.84 | 95.34 |
| Attention (multi-modal w/o n_{av}) | 40.84 | 95.15 |
| Attention (multi-modal) | 46.16 | 97.51 |

Table 4: Ablation of attention fusion models.

mAUC, whereas the visual model achieves 25.73 mAP and 91.30 mAUC. The superiority of the audio model over the visual model is expected due to the nature of the task.

Average fusion of the audio and visual outputs achieves 42.84 mAP, an absolute improvement of 4.49 mAP (relative: 11.7%) over the audio model, and an absolute improvement of 17.11 mAP (relative: 66.5%) over the visual model. The regression fusion model leads to a minor improvement over average fusion: +0.26 mAP. The MLP fusion model leads to considerable improvement over average fusion: +2.76 mAP.

Our attention fusion model achieves 46.16 mAP, which is an absolute improvement of 7.81 mAP (relative: 20.4%) over the audio model. It also outperforms all baseline fusion methods: an absolute improvement of 3.32 mAP (relative: 7.7%) over average fusion. In this case, n_{attn} takes in the concatenated normalized outputs from the audio (h_a) and visual (h_v) models. n_{attn} is a neural network with a single hidden layer of 512 units; batch normalization and ReLU activation is applied on this layer which is followed by a dropout layer with a dropout rate of 0.5. The output layer has 527 units and a sigmoid activation function. n_a , n_v , and n_{av} are all single layer networks with 512 units and sigmoid activations.

We also study different variations of the proposed fusion model. The *unimodal* approaches uses only a single modality in the attention network, either \mathbf{h}_a or \mathbf{h}_v as opposed to using \mathbf{h}_a and \mathbf{h}_v combined through concatenation. When α_a is obtained by using \mathbf{h}_a and α_v using \mathbf{h}_v , we denote it as *self unimodal* attention. However, if α_a uses \mathbf{h}_v and vice versa, we call it *cross unimodal* attention. In these cases, there are two attention networks, n_{attn}^a and n_{attn}^v but other aspects of the architecture remains identical to the one described above.

We also consider another variation where n_a and n_v are simply identity functions that do not contain learnable parameters. Similarly, in another variation, we make n_{av} an identity mapping. The performance of these fusion models are listed in Table 4. The multi-modal attention model is slightly better than the unimodal attention models. However, removing neural networks n_a , n_v , and n_{av} leads to significant deterioration in performance.

Comparison with State-of-the-Art. Table 3 shows comparison with state-of-the-art models on Audioset. Our audio model is slightly better than the state-of-the-art performance on Audioset. Note that [Ford *et al.*, 2019] also shows a performance of 39.2 mAP. However, this is obtained by a ensembeling outputs from multiple models and the best single model performance (which is fairer to compare with) is 38.0. To the best of our knowledge, [Wang *et al.*, 2019a] is the only prior work that reported visual and audiovisual models for sound



Figure 2: Audio, visual, and audiovisual models Average Precision (AP). The horizontal and vertical axes denote the audio and visual models APs respectively for all classes. The diameter of each bubble denotes the absolute change in AP over the single-modal models (audio or visual, whichever is better). Blue denotes improvement over the single-modal model, whereas red indicates deterioration over the single-modal model, as depicted in the color bar. Several sound classes of interest are annotated accordingly.

events on Audioset. Our visual model is +6.93 mAP (relative: 36.8%) better than [Wang *et al.*, 2019a]. More importantly, our audiovisual model is **+4.35 mAP** (relative: 10.4%) compared with [Wang *et al.*, 2019a] and sets a new state-of-the-art on Audioset.

5 Discussions

5.1 Class Specific Analysis

Figure 2 provides a per class depiction of the audio, visual, and audiovisual models. As expected, the audio model outperforms the video model for the majority of sound events. For sounds such as *Heartbeat*, *Jingle Bell*, and *Whispering*, audio recognition of sounds is considerably better than visual recognition of sounds. However, for classes such as *Patter*, *Shuffling Cards*, and *Bus*, the visual model outperforms the audio model by a large margin. Overall, the visual model outperforms the audio model for 117 sound classes by a small or large margin. For some classes, both audio and visual models perform well (e.g., *Turkey, Train*, and *Music Accordion*), or both fail to achieve good performance (e.g., *Squish* and *Hum*).

For the audiovisual fusion model, we see that for most classes, audiovisual learning leads to an improvement over the better of the audio or visual models, as observable by mostly blue bubbles in Figure 2. Figure 3 shows a histogram of the absolute deterioration and improvement in performance through audiovisual learning compared to the better of the audio or visual models. For over 83% of classes (439 of 527), we see an improvement, whereas for the remaining



Figure 3: Histogram of change in AP for our audiovisual attention fusion model with respect to single modality models (audio or visual whichever is better).



Figure 4: Mean of attention weights for all classes in AudioSet.

classes, audiovisual learning leads to deterioration in performance, compared with the better of the audio or visual models. Furthermore, for over 17% of classes (91 out of 527), audiovisual learning achieves an absolute improvement of more than +10 mAP. On the other hand, for only 15 classes audiovisual learning leads to an absolute deterioration of more than 5 mAP, with only 3 classes with more than 10 deterioration in mAP. For most cases, the deterioration is minor.

From Figure 2, we can see some examples of classes with large improvements in the audiovisual model, e.g., *Explosion*, *Ringtone Clicking*, *Vacuum Cleaner*, and *Bathtub*. Several of these classes with large improvement have similar AP using audio and visual models.

For several sound classes, one of the modalities is much more superior. One may argue for using only the appropriate single-modal models. However, we see that in several of these cases the other modality, despite being inferior on its own, can provide complimentary information, which can lead to improved performance through audiovisual learning. Consider for example, classes such as *Angry Music* and *Plop*, as shown in Figure 2. While the audio modality is much more useful for these classes compared to the visual modality, audiovisual learning still leads to considerable improvement as the information from the visual modality compliments the audio modality. The same inference applies for classes such as *Bus* and *Motorboat*, where the visual modality is much more superior compared to the audio modality. However, there are exceptions to the above observations. On the audio side, *Whispering* and *Police Siren*, and on the video side, *Shuffling Cards* and *Writing*, are examples of such exceptions.

One final observation we would like to point out from Figure 2 is that, in general for classes with low performance for the audio and visual models (less than 10 AP), audiovisual fusion also does not lead to any improvement. We believe that a multi-modal framework in which the audio and visual modalities are jointly used to train a single model might be more suitable in these situations.

5.2 Analysis of Attention

We analyze the attention weights obtained for each sample for each class. For each class, we compute the mean attention weight $(\bar{\alpha}_a^c)$ by averaging over all evaluation examples which contains the c^{th} class. This gives us an estimate of the importance of the audio modality for the c^{th} class (and consequently for the visual modality as well). Figure 4 is a histogram plot of the mean attention weights. The right side of the plot with higher α_a represents dominance of the audio modality, whereas on the left side the visual modality is the dominant modality. The distribution is clearly skewed towards the audio modality, which is expected since the audio model is better than the visual model for more than 75% of the sound classes. There are more than 110 sounds for which on average, the audio modality gets all the weight ($\bar{\pmb{\alpha}}_a^c \approx 1$). Interestingly, there are around 38 sound classes where the visual modality has complete dominance ($\bar{\alpha}_a^c \approx 0$). Furthermore, for more than 70 sound classes the visual modality on average gets more than 90% weight ($\bar{\alpha}_a^c \leq 0.1$). This signifies that for several sounds the visual modality can play a vital role in their recognition.

6 Conclusions

We propose audiovisual models for learning sounds. Instead of jointly using audio and visual modalities in a single model, we propose to combine the individual models that were trained separately on each modality. To this end, we use a state-of-the-art audio based sound event recognition model and propose a novel vision based model for recognizing sound events. On the Audioset dataset of sound events, it outperforms a prior vision model for sounds by +6.93 mAP (relative: 36.9%). We then propose an attention mechanism for audiovisual fusion that dynamically fuses the outputs from the audio and visual models in a sample-specific and class-specific manner. Our proposed audiovisual learning model improves state-of-the-art performance on Audioset by approximately +4.35 mAP (relative: 10.4%). We also provide thorough analysis of the role of the audio and visual modalities for various sound classes. The outcomes of this analysis can play a crucial role in the development of future models for audiovisual learning of sounds.

References

[Adavanne *et al.*, 2019] Sharath Adavanne, Haytham M. Fayek, and Vladimir Tourbabin. Sound event classification and detection with weakly labeled data. In *Detection* and Classification of Acoustic Scenes and Events Workshop (DCASE), pages 15–19, New York, USA, October 2019.

- [Alwassel *et al.*, 2019] Humam Alwassel, Dhruv Mahajan, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Selfsupervised learning by cross-modal audio-video clustering. *arXiv preprint arXiv:1911.12667*, 2019.
- [Andrews *et al.*, 2003] Stuart Andrews, Ioannis Tsochantaridis, and Thomas Hofmann. Support vector machines for multiple-instance learning. In *NIPS*, pages 577–584, 2003.
- [Arandjelovic and Zisserman, 2017] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *CVPR*, pages 609–617, 2017.
- [Aytar *et al.*, 2016] Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. In *NIPS*, pages 892–900, 2016.
- [Bahdanau *et al.*, 2014] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [Baltrušaitis *et al.*, 2018] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. Multimodal machine learning: A survey and taxonomy. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 41(2):423–443, 2018.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009.
- [Ford *et al.*, 2019] Logan Ford, Hao Tang, François Grondin, and James Glass. A deep residual network for large-scale acoustic scene analysis. In *Interspeech*, 2019.
- [Gao *et al.*, 2019] Ruohan Gao, Tae-Hyun Oh, Kristen Grauman, and Lorenzo Torresani. Listen to look: Action recognition by previewing audio. *arXiv preprint arXiv:1912.04487*, 2019.
- [Gemmeke *et al.*, 2017] J Gemmeke, D Ellis, D Freedman, A Jansen, W Lawrence, C Moore, M Plakal, and M Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *ICASSP*, pages 776–780. IEEE, 2017.
- [Ghanem *et al.*, 2018] B Ghanem, J Niebles, C Snoek, F Heilbron, H Alwassel, V Escorcia, R Khrisna, S Buch, and C Dao. The activitynet large-scale activity recognition challenge 2018 summary. *arXiv preprint arXiv:1808.03766*, 2018.
- [Gu et al., 2018] C Gu, C Sun, D Ross, C Vondrick, C Pantofaru, Y Li, S Vijayanarasimhan, G Toderici, S Ricco, R Sukthankar, et al. AVA: A video dataset of spatiotemporally localized atomic visual actions. In *CVPR*, pages 6047–6056, 2018.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.

- [Hori *et al.*, 2018] Chiori Hori, Takaaki Hori, Gordon Wichern, Jue Wang, Teng-Yok Lee, Anoop Cherian, and Tim K Marks. Multimodal attention for fusion of audio and spatiotemporal features for video description. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 2528–2531, 2018.
- [Kay et al., 2017] W Kay, J Carreira, K Simonyan, B Zhang, C Hillier, S Vijayanarasimhan, F Viola, T Green, T Back, P Natsev, et al. The kinetics human action video dataset. arXiv preprint arXiv:1705.06950, 2017.
- [Kazakos *et al.*, 2019] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *ICCV*, October 2019.
- [Kingma and Ba, 2014] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Kong et al., 2019] Qiuqiang Kong, Changsong Yu, Yong Xu, Turab Iqbal, Wenwu Wang, and Mark D Plumbley. Weakly labelled audioset tagging with attention neural networks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing (TASLP)*, 27(11):1791–1802, 2019.
- [Korbar et al., 2018] Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. In *NeurIPS*, pages 7763–7774, 2018.
- [Kumar and Ithapu, 2020] Anurag Kumar and Vamsi Krishna Ithapu. Secost: Sequential co-supervision for weakly labeled audio event detection. In *ICASSP*. IEEE, 2020.
- [Kumar and Raj, 2016a] Anurag Kumar and Bhiksha Raj. Audio event detection using weakly labeled data. In *ACM Multimedia*, pages 1038–1047. ACM, 2016.
- [Kumar and Raj, 2016b] Anurag Kumar and Bhiksha Raj. Weakly supervised scalable audio content analysis. In 2016 IEEE International Conference on Multimedia and Expo (ICME), pages 1–6. IEEE, 2016.
- [Kumar, 2018] Anurag Kumar. *Acoustic Intelligence in Machines*. PhD thesis, Carnegie Mellon University, 2018.
- [Lan et al., 2012] Zhen-zhong Lan, Lei Bao, Shoou-I Yu, Wei Liu, and Alexander G Hauptmann. Double fusion for multimedia event detection. In *International Conference* on *Multimedia Modeling*, pages 173–185. Springer, 2012.
- [Lin *et al.*, 2019] Yan-Bo Lin, Yu-Jhe Li, and Yu-Chiang Frank Wang. Dual-modality seq2seq network for audio-visual event localization. In *ICASSP*, pages 2002–2006. IEEE, 2019.
- [Long *et al.*, 2018] Xiang Long, Chuang Gan, Gerard De Melo, Jiajun Wu, Xiao Liu, and Shilei Wen. Attention clusters: Purely attention based local feature integration for video classification. In *CVPR*, pages 7834–7843, 2018.
- [McFee *et al.*, 2018] Brian McFee, Justin Salamon, and Juan Pablo Bello. Adaptive pooling operators for weakly

labeled sound event detection. *IEEE/ACM Transactions* on Audio, Speech and Language Processing (TASLP), 26(11):2180–2193, 2018.

- [Owens and Efros, 2018] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *ECCV*, pages 631–648, 2018.
- [Parekh et al., 2019] S Parekh, S Essid, A Ozerov, N Duong, P Pérez, and G Richard. Weakly supervised representation learning for audio-visual scene analysis. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* (TASLP), 28:416–428, 2019.
- [Sterpu *et al.*, 2018] George Sterpu, Christian Saam, and Naomi Harte. Attention-based audio-visual fusion for robust automatic speech recognition. In *Proceedings of the* 20th ACM International Conference on Multimodal Interaction, pages 111–115, 2018.
- [Virtanen et al., 2017] Tuomas Virtanen, Mark D Plumbley, and Dan Ellis. *Computational analysis of sound scenes and events*. Springer, 2017.
- [Wang *et al.*, 2019a] Weiyao Wang, Du Tran, and Matt Feiszli. What makes training multi-modal networks hard? *arXiv preprint arXiv:1905.12681*, 2019.
- [Wang *et al.*, 2019b] Yun Wang, Juncheng Li, and Florian Metze. A comparison of five multiple instance learning pooling functions for sound event detection with weak labeling. In *ICASSP*, pages 31–35, 2019.
- [Wei *et al.*, 2014] Xiu-Shen Wei, Jianxin Wu, and Zhi-Hua Zhou. Scalable multi-instance learning. In *IEEE International Conference on Data Mining*, pages 1037–1042. IEEE, 2014.
- [Wu et al., 2016] Zuxuan Wu, Yu-Gang Jiang, Xi Wang, Hao Ye, and Xiangyang Xue. Multi-stream multi-class fusion of deep networks for video classification. In ACM Multimedia, pages 791–800. ACM, 2016.
- [Zhou et al., 2019] Pan Zhou, Wenwen Yang, Wei Chen, Yanfeng Wang, and Jia Jia. Modality attention for endto-end audio-visual speech recognition. In *ICASSP*, pages 6565–6569. IEEE, 2019.