Audiovisual Transformer Architectures for Large-Scale Classification and Synchronization of Weakly Labeled Audio Events

Wim Boes
Hugo Van hamme
wim.boes@esat.kuleuven.be
hugo.vanhamme@esat.kuleuven.be
ESAT, KU Leuven
Leuven, Belgium

ABSTRACT

We tackle the task of environmental event classification by drawing inspiration from the transformer neural network architecture used in machine translation. We modify this attention-based feedforward structure in such a way that allows the resulting model to use audio as well as video to compute sound event predictions.

We perform extensive experiments with these adapted transformers on an audiovisual data set, obtained by appending relevant visual information to an existing large-scale weakly labeled audio collection. The employed multi-label data contains clip-level annotation indicating the presence or absence of 17 classes of environmental sounds, and does not include temporal information.

We show that the proposed modified transformers strongly improve upon previously introduced models and in fact achieve state-of-the-art results. We also make a compelling case for devoting more attention to research in multimodal audiovisual classification by proving the usefulness of visual information for the task at hand, namely audio event recognition.

In addition, we visualize internal attention patterns of the audiovisual transformers and in doing so demonstrate their potential for performing multimodal synchronization.

CCS CONCEPTS

• Computing methodologies → Neural networks.

KEYWORDS

audio event classification, audiovisual classification, audiovisual synchronization, transformer neural network architecture

ACM Reference Format:

Wim Boes and Hugo Van hamme. 2019. Audiovisual Transformer Architectures for Large-Scale Classification and Synchronization of Weakly Labeled Audio Events. In *Proceedings of the 27th ACM International Conference on Multimedia (MM '19, October 21–25, 2019, Nice, France)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3343031.3350873>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '19, October 21–25, 2019, Nice, France

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6889-6/19/10...\$15.00

<https://doi.org/10.1145/3343031.3350873>

1 INTRODUCTION

Human beings heavily depend on both auditory and visual cues to detect environmental events and to infer how to act in response. In order to create situation-aware machines, it is therefore vital to study and design models which are capable of utilizing both modalities to accurately classify such events.

Image classification has been a red hot topic ever since the advent of the ImageNet challenge and data set. Much more recently, research into audio event classification has also taken off because of among other the introduction of the Audio Set data set and challenges such as DCASE 2016, DCASE 2017 and DCASE 2018. However, in contrast to the unimodal classification tasks, multimodal classification tasks are much less prevalent in the literature, even though studying those endeavours could naturally lead to important practical consequences [4, 6, 11, 15, 19].

In this work, we draw inspiration from the machine translation community to tackle classification of environmental events. As is the case with many machine learning tasks, the domain of machine translation is currently dominated by the deep learning paradigm. In particular, recently the transformer neural network has gained a lot of popularity. This type of model attains state-of-the-art performance while also achieving comparatively quick training times and allowing for interpretation relatively easily [18].

The data used in the field of machine translation is similar to audiovisual data in the sense that two modalities are involved, namely a source language and a target language. The comparison can be taken even further as audio and vision could be regarded as two abstract languages capable of describing environmental events. For example, a train passing by could be expressed in the abstract audio language by the sound of a train horn, while it could simultaneously be represented in the abstract vision language by a picture of a train. This loose analogy indicates that it could be useful to apply machine translation techniques to audiovisual data.

The approach taken in this work is therefore the following: we modify the transformer neural network architecture used in machine translation to be viable for the task of audiovisual classification of environmental events. We then evaluate and analyze the results obtained by these models.

In Section 2, we discuss the original transformer architecture and how it can be adapted for the task of audiovisual classification of environmental events. In Section 3, we discuss the setup for the conducted experiments. Next, in Section 4 we discuss and interpret the obtained results. Finally, in Section 5 we draw a conclusion.

2 AUDIOVISUAL TRANSFORMER NEURAL NETWORKS

In this section, we discuss some key concepts of the original transformer used for machine translation and explain how this model can be adapted to become suitable for audiovisual classification.

2.1 Machine translation transformer

The original transformer is a neural network based entirely on feed-forward layers, residual connections and attention. It is depicted in Figure 1. This encoder-decoder model converts a sequence of input tokens in a source language into a sentence in a target language by iteratively calculating conditional probabilities of each output word and combining them through beam search. In what follows, we will further describe some of the key components of this model. However, we will not discuss elements which are not used in this work such as the masking and shifting of output tokens and implementation details like the used regularization methods. For full details on the machine translation transformer, we refer the reader to the original work on this model [7, 18].

Figure 1: Transformer model for machine translation [18]

2.1.1 Multi-head attention. The transformer in Figure 1 uses multi-head scaled dot product attention blocks. Scaled dot product attention involves computing linear combinations of a set of values. The so-called attention weights of these sums are determined by comparing a number of queries to a collection of keys corresponding to the values: they are obtained by calculating scaled dot products and applying the softmax function to normalize w.r.t. the keys and values. Mathematically, this can be expressed as follows [18]:

$$A(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d_k}} \right) V \quad (1)$$

In Equation (1), $Q \in \mathbb{R}^{n_q \times d_k}$, $K \in \mathbb{R}^{n_k \times d_k}$ and $V \in \mathbb{R}^{n_k \times d_v}$ denote the matrices containing respectively the queries, the keys and the values. The number of queries and the number of keys, which is equal to the number of corresponding values, are symbolized by n_q and n_k respectively. The dimension of the keys, which is the same as the dimension of the queries, and the dimension of the values are indicated by d_k and d_v respectively. The matrix $A \in \mathbb{R}^{n_q \times d_v}$ contains the resulting weighted sums [18].

The multi-head attention blocks add flexibility on top of this bare attention mechanism in two ways. Firstly, linear transformations are introduced before and after the computation of the scaled dot products. Secondly, instead of one linear combination, multiple attention heads, i.e., weighted sums, are calculated and concatenated per query, explaining the name of the component [18].

The multi-head attention block is zoomed in on in Figure 2.

Figure 2: Multi-head scaled dot product attention block [18]

In the transformer architecture for machine translation, the queries, keys and values consist of abstract word token representations. It has been demonstrated that applying the multi-head attention mechanism to these word embeddings in both the encoder and decoder allows the model to capture a variety of interesting intra- as well as interlingual word relationships [18].

2.1.2 Positional encodings. Because of the absence of recurrent and convolutional connections, transformers have no real perception of order, which is obviously an undesirable property for neural machine translation models. This issue is alleviated by injecting sequential information via the introduction of positional encodings.

This comes down to adding abstract vector representations of word positions to both the input and output word embeddings. In the machine translation transformer, sinusoidal positional encodings are employed: for this type of encodings, fixed position differences can be represented by simple linear transformations [5, 18].

2.2 Adaptations to original architecture

Due to the similarities between machine translation and audiovisual classification, fitting the original transformer to the task at hand is relatively easy. Below, we discuss the relevant changes.

The resulting audiovisual transformer architecture incorporating all of these adaptations is depicted in Figure 3.

2.2.1 Data adaptation. Linguistic data is sequential in nature, defined at discrete word positions. In this work, we ensure that the used audiovisual data complies to this format by converting the audio into sequences of frames and sampling images from the videos. More detailed information is provided in Section 3.

In machine translation, the data from the source language is served as inputs, while the data from the target language constitutes the desired outputs. For the task at hand, all data from the considered modalities, i.e., audio and vision, can be used as inputs.

2.2.2 Introduction of extra linear maps. For the machine translation transformer, the size of the word embeddings constrains the amount of neurons used by the linear transformations in the multi-head attention blocks. This is a consequence of the residual connections in the model. To relax this restriction, the audiovisual transformer contains additional linear maps projecting the employed auditory and visual embeddings to vectors of some predefined dimensionality.

2.2.3 Change of output layer. In contrast to machine translation, for classification of environmental events the prediction categories are not mutually exclusive. Indeed, while word positions can only be occupied by a single token, multiple environmental events can occur at the same time. The softmax layer at the end of the transformer is therefore changed into a sigmoid layer as it is more appropriate.

2.2.4 Addition of aggregation block. The goal of the machine translation transformer is to compute the probability of a sentence. This requires iteration as conditional probabilities have to be calculated for each word in order to eventually multiply them. However, for the task at hand and the used data (see Section 3), we only need a single probability vector per audiovisual clip. Thus, we add an aggregation block at the end of the model which combines frame-level scores by performing a mean or max pooling operation. Consequently, unlike in the original transformer, there is no more need for iteration in the adapted variant: inputs need to be passed through this model just once to obtain sound event predictions.

2.2.5 Modification of attention function. Originally, the softmax function was used to calculate attention values as is evident from Equation (1). We also test some other straightforward attention functions, namely the sigmoid function σ and the normalized sigmoid function $\bar{\sigma}$. The latter function is characterized as follows:

$$\bar{\sigma}(\mathbf{x}) = \frac{\sigma(\mathbf{x})}{\sum_{i=1}^n \sigma(x_i) + \epsilon} \quad (2)$$

In Equation (2), $\mathbf{x} \in \mathbb{R}^n$ is a vector and ϵ is a very small positive real number used to avoid overflow.

2.2.6 Optional deletion of positional encodings. Lastly, we turn the addition of positional encodings at the inputs into an optional operation and perform experiments with and without these abstract representations as it is unclear how important ordering information is for the task and data at hand.

Figure 3: Audiovisual transformer model

3 EXPERIMENTAL SETUP

In this section, we expand upon the experimental setup. We provide details on the audiovisual data set used to train the multimodal transformers and discuss implementation details such as the used data preprocessing and regularization methods.

3.1 Data set

The large-scale multi-label data set for task 4 of DCASE 2017 consists of 51172 training, 488 validation and 1103 test audio samples originated from YouTube. Most of these audio files are 10 seconds long, some are shorter. They are weakly labeled: only clip-level annotations are provided, indicating the presence or absence of 17 audio event classes, which all represent environmental sounds. Temporal knowledge is not included, meaning that there is no information about the time boundaries of the sounds [12].

Taking an approach similar to a procedure used in prior work, YouTube downloads allow us to readily append visual data to this audio data set. Unfortunately, due to availability issues, videos could not be obtained for all of the audio files in the original data set. Ultimately, we ended up creating a slightly reduced collection of audio and video files consisting of 48883 training, 465 validation and 1050 test samples: about 5% of samples of the original data was discarded in the process [14].

3.2 Implementation details

The audiovisual transformers are built using the TensorFlow toolkit. In this subsection, some important implementation details are elaborated upon [1].

3.2.1 Data preprocessing. To obtain auditory frame embeddings, first logarithmic mel spectrograms are extracted from the audio samples. Next, a sliding window is applied to these spectrograms to divide the signal into frames. The length of this window and its hop length are set to 960 ms. The spectral frames are passed through vggish, a model for video tag classification using audio, pretrained using a preliminary version of the YouTube-8M data set. The 128-dimensional outputs of the last feedforward layer of this deep convolutional network are used as frame embeddings [2, 8].

Visual frame embeddings are obtained in the following way: firstly, still images are sampled from the videos. The sampling times are chosen to coincide with the centers of the possible positions of the sliding window used in the audio preprocessing described above. The resulting video frames are then fed into VGG16, a deep convolutional model for image classification, pretrained on the ImageNet data set. The 4096-dimensional outputs of the last feedforward layer of this neural network are used as visual frame embeddings [4, 16].

Most of the preprocessing steps described above are included in the TensorFlow toolkit and are therefore easy to implement [1].

3.2.2 Data sampling. The audiovisual data used for training is strongly unbalanced: the ratio of number of samples containing a sound of the most common event class to the number of examples comprising the least likely event is approximately equal to 130.

To alleviate this issue, we use a sampling scheme proposed in prior work: for each training batch, class labels are non-uniformly sampled and appropriate data examples are fetched. This happens in a way that limits the maximum class imbalance ratio to 5 [20].

We train the models for a maximum of 30 epochs. In each training epoch, 300 batches consisting of 40 data samples are presented to the networks, resulting in a total of 12000 examples per epoch.

3.2.3 Number of encoder and decoder blocks. The number of encoder and decoder blocks is referred to with the letter N in Figure 3. We set this hyperparameter equal to three.

3.2.4 Number of attention heads. We use three attention heads in all multi-head attention blocks in both the encoder and decoder of the audiovisual transformers.

3.2.5 Feedforward ReLU subnetworks. The feedforward ReLU subnetworks referenced in Figure 3 match the ones used in the original transformer shown in Figure 1: each subnet consists of one ReLU layer followed by a linear layer without activation function [18].

3.2.6 Number of neurons. All feedforward layers with and without activation functions, including all linear transformations used, consist of 128 neurons.

3.2.7 Regularization methods. Identically to the original transformer architecture, computation of the output of each multi-head attention block of the audiovisual transformer is followed by addition with its corresponding residual connection and layer normalization, a popular regularization technique which performs sample-wise scaling. We also apply dropout with a rate of 0.1 to the outputs of the following parts of the model [3, 7, 17, 18]:

- Linear layers at the input of the encoder and decoder
- Attention function in the multi-head attention blocks
- Final linear map in the multi-head attention blocks
- Both layers of the feedforward ReLU subnetworks

3.2.8 Loss function and optimization algorithm. We use the well-known binary cross entropy loss function in conjunction with the Adam optimization algorithm (using standard settings as defined in the TensorFlow toolkit) to train the transformers [1, 10].

Early stopping is applied if at any point during training the cross entropy calculated on the validation data is higher than all validation losses obtained after the last 7 epochs.

3.2.9 Epoch-level ensembling. Instead of only using the model obtained after the final training epoch to obtain audio event class probabilities, we employ the epoch-level ensembling method proposed in the original work on transformers: after each of the last 7 training epochs the model is saved. Eventually, these saved states are used to obtain 7 sets of probabilities, which are then averaged class-wise to obtain a more robust output [18].

3.2.10 Prediction thresholding. The final goal of the audiovisual transformer models is to make binary decisions: they should predict whether or not certain audio event classes are present in a given clip. The continuous probabilities at the output of the transformer systems thus should be converted into binary values via thresholding. Preliminary experiments have shown that using a global threshold which maximizes the performance of the models on the validation data is favorable to utilizing class-dependent optimal thresholds, hence the former approach is taken.

4 EXPERIMENTAL RESULTS

In this section, we discuss the results we obtained by training the audiovisual transformer neural networks. We split this discussion into two parts: the first is the quantitative analysis, where the performance of the models is numerically evaluated and compared to other approaches, the second part is the qualitative analysis which provides interesting, interpretable insights into the workings of the audiovisual transformers.

Table 1: F1 scores of audiovisual transformer models

modality of first input	modality of second input	modalities utilizing positional encodings	F1 score					
			mean pooling as aggregation method			max pooling as aggregation method		
			softmax attention	sigmoid attention	normalized sigmoid attention	softmax attention	sigmoid attention	normalized sigmoid attention
audio	video	audio/video	66.1%	66.0%	66.3%	66.2%	66.3%	66.4%
		video	69.5%	69.8%	69.4%	69.3%	69.5%	69.4%
		none	69.2%	69.4%	69.3%	69.1%	69.3%	68.9%
video	audio	audio/video	67.7%	67.5%	67.8%	67.3%	67.1%	67.6%
		video	69.7%	70.1%	69.7%	69.7%	69.8%	69.7%
		none	69.6%	69.8%	69.5%	69.7%	69.7%	69.7%
video	video	video	59.1%	59.3%	59.3%	59.0%	59.1%	59.0%
		none	58.8%	59.1%	59.0%	58.8%	58.8%	58.8%
audio	audio	audio	57.0% (57.0%)	56.9% (57.0%)	56.8% (57.1%)	56.2% (56.4%)	56.4% (56.7%)	56.4% (56.7%)
		none	58.6% (59.3%)	59.1% (59.5%)	58.8% (59.3%)	58.5% (59.0%)	58.9% (59.4%)	58.7% (59.0%)

4.1 Quantitative analysis

We quantitatively evaluate the audiovisual transformer models by measuring micro-averaged F1 scores on the test samples of the employed data. This metric is chosen as it is very often utilized in the field of audio event recognition, e.g., in task 4 of the DCASE 2017 challenge, where the used data originates from [12, 13].

We report F1 scores which are averaged over 25 training runs to ensure the results are reliable. This is practically feasible as a consequence of the use of fixed pretrained embeddings and due to the fact that transformers are very quick to train, in contrast to neural models containing recurrent connections.

Table 1 contains the measured F1 scores on the audiovisual data set created in this work which is described in Section 3. The referred settings are described in detail in Section 2. This table contains results for multimodal transformers employing both audio and video as well as transformers utilizing only a single modality. For the models using only audio, performance on the slightly larger (complete) data set of task 4 of the DCASE 2017 challenge also mentioned in Section 3 is reported between parentheses [12].

In the rest of this section, we analyze these numerical results and compare to approaches presented in prior works.

4.1.1 Evaluation of positional encodings. The unimodal video transformers benefit from the utilization of positional encodings, albeit only slightly. However, employing these abstract position representations is strongly detrimental to the models utilizing only audio.

The micro-averaged F1 scores obtained by multimodal transformers employing both audio and video show a similar trend: adding positional encodings to the used visual frame embeddings provides minor improvements, adding these abstract position representations to both the auditory and visual embeddings causes a severe deterioration in performance.

4.1.2 Evaluation of aggregation methods and attention functions. Generally, audiovisual transformers employing mean pooling as aggregation method seem to perform slightly better than those using max pooling. The disparity in F1 scores is the most pronounced for the unimodal models utilizing only audio.

Remarkably, transformers utilizing different types of attention functions perform very similarly. In particular, we note that there is no meaningful difference in F1 scores between models employing normalized (softmax, normalized sigmoid) and unnormalized attention functions (sigmoid).

4.1.3 Comparison to prior work in audio classification. We list relevant prior audio event classification results measured on the test samples of the data set of task 4 of the DCASE 2017 challenge in Table 2. These micro-averaged F1 scores can most directly be compared to the outcomes reported between parentheses in Table 1, as their computation is based on the same data [12].

Table 2: F1 scores of prior audio classification models

model	F1 score
fusion of gated convolutional recurrent networks [21]	55.6%
capsule-based gated convolutional network [9]	58.6%

The transformers outdo the models in Table 2. We speculate that this can be attributed to the power of the multi-head attention mechanism and the fact that they employ rich, externally pretrained embeddings instead of training audio features from scratch.

4.1.4 Comparison to prior work in audiovisual classification. The multimodal transformers employing both audio and video can be fairly compared to the two-stream audiovisual neural network introduced in prior work: this model performs environmental audio event classification based on audiovisual data acquired in the same manner as the procedure used in this work. The micro-averaged F1 score achieved by this model is reported in Table 3.

Table 3: F1 score of prior audiovisual classification model

model	F1 score
two-stream audiovisual neural network [14]	64.2%

Figure 4: Visualization for transformer using video/audio as first/second input modalities (sound produced by train)

The multimodal audiovisual transformers strongly improve upon the two-stream audiovisual neural network. As both of these types of models use nearly the same data and employ similar externally pretrained embeddings, this truly illustrates the power of the transformer architecture for the task at hand.

As can be inferred from Table 1, the best performing multimodal transformer type achieves a micro-averaged F1 score of 70.1%. To the best of our knowledge, this is the best classification result achieved for the considered data.

4.1.5 Comparison of unimodal and multimodal transformers. A comparison between the transformers using both audio and video and the transformers utilizing only audio unveils the usefulness of visual information for environmental audio event classification: on average, there is a difference in F1 scores of about 10%.

Initially, it is reasonable to expect unimodal transformers employing visual information to perform worse than those utilizing audio, since the annotation of the data at hand only indicates the presence or absence of audio event classes as outlined in Section 3. Additionally, the used audiovisual data set was created by simply extending labeled audio files originated from YouTube with their corresponding videos, as also discussed in Section 3, and therefore there is no guarantee that the appended visual material contains valuable information. By conducting a manual inspection of the data set, we were even able to confirm that not all of the downloaded videos are informative for the considered classification task. Despite these remarks, we still find that the unimodal visual models slightly outperform the unimodal models based on audio.

This peculiar observation can be explained by the fact that in many cases, images or videos are more useful than auditory signals.

For example, multiple classes in the used data set pertain to similar vehicle sounds such as police car, fire truck and ambulance sirens: even for humans, differentiating these events solely based on audio can be quite difficult and visual information may be of great help. Another plausible reason for this strange outcome is a disparity in richness of the used pretrained auditory and visual embeddings, which are described in Section 3.

4.2 Qualitative analysis

We qualitatively analyze the transformers utilizing both audio and video by visualizing cross-modal attention weights in a manner similar to a procedure followed in prior work, to investigate the ability of these models to perform multimodal synchronization [18].

4.2.1 Models using video/audio as first/second inputs respectively. Figures 4, 5 and 6 contain visualizations for transformers without positional encodings which utilize video as first and audio as second input, mean pooling for aggregation and employ the softmax attention function. A full description of these configurations can be found in Section 2. Each figure contains a number of auditory (spectral) and visual frames for a test sample of the audiovisual data set described in Section 3, as well as a heat map of the cross-modal attention weights obtained for the circled audio query frame from the cross-modal multi-head attention component in the last decoder block of the corresponding model. Each row in each heat map pertains to a different attention head.

In all of the cases visualized in Figures 4, 5 and 6, the models are able to detect visual frames relevant to the environmental sounds present in the circled audio query frames.

Figure 5: Visualization for transformer using video/audio as first/second input modalities (sound produced by motorcycle)

Figure 6: Visualization for transformer using video/audio as first/second input modalities (police car and fire truck sirens)

Figure 7: Visualization for transformer using audio/video as first/second input modalities (sound produced by skateboard)

For models using video as first input and audio as second input, we were able to discover similarly interesting attention patterns for many other validation and test samples. Strangely, this was only possible for audiovisual transformers employing the softmax attention function: the architectures utilizing the sigmoid and normalized sigmoid attention functions did not allow for meaningful visualizations, even though they achieve comparable performance as reported in Section 4.1.2.

4.2.2 Models using audio/video as first/second inputs respectively. In Figure 7, we present an example for an audiovisual transformer with visual positional encodings which utilizes audio and video as first and second input modalities respectively, mean pooling as aggregation method and uses softmax attention. The figure contains 10 (spectral) auditory and visual frames for a validation sample of the data set described in Section 3, as well as a heat map of the attention weights obtained for the circled video query frame from the cross-modal multi-head attention component in the last decoder block of the associated model. Each row in the heat map pertains to a different attention head.

The circled video query frame of the sample visualized in Figure 7 is a picture of someone skateboarding. Through its attention heads, the corresponding model is able to detect the relevant auditory frames containing skateboarding sounds, which are circled.

For this type of transformer, discovering synchronization patterns in the attention weights proved to be quite difficult. This can likely be explained by the fact that sounds are naturally much less localized compared to visual events. Indeed, also for humans, pointing to specific audio frames as containing valuable information is much harder than pointing to specific video frames, i.e., images.

5 CONCLUSION

We tackled the task of environmental event classification by drawing inspiration from machine translation: we changed the popular transformer architecture into a model which can use auditory as well as visual data to compute multi-label audio event predictions.

We experimented with these transformers on an audiovisual data set, obtained by extending an existing weakly labeled auditory data set with visual information. The used multi-label data contains clip-level annotation indicating the presence of 17 classes of environmental sounds, without disclosure of temporal information.

We showed that the proposed models strongly improve upon prior approaches in the fields of audio as well as audiovisual classification and achieve state-of-the-art results.

We also proved the usefulness of visual information for the task at hand, i.e., audio event recognition. This shows that it certainly can be worth researching multimodal audiovisual models.

In addition, we visualized internal attention patterns of the audiovisual transformers. This way, we were able to demonstrate that they show potential for multimodal synchronization.

In the future, we will use these transformers to tackle some other audiovisual data collections and learn more about their properties. In particular, we are keen to analyze the impact of scale: we will experiment with the small-scale data set of task 4 of DCASE 2018, and Audio Set, which is humongous in size. In addition, we will look into how the models can be used for unsupervised tasks [6, 15].

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 265–283.
- [2] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. 2016. Youtube-8M: A Large-Scale Video Classification Benchmark. *arXiv preprint arXiv:1609.08675* (2016).
- [3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450* (2016).
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Ji Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- [5] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional Sequence to Sequence Learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. 1243–1252.
- [6] Jort F. Gemmeke, Daniel P.W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. 2017. Audio Set: An ontology and human-labeled dataset for audio events. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 776–780.
- [7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [8] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F. Gemmeke, Aren Jansen, R. Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, et al. 2017. CNN Architectures for Large-Scale Audio Classification. In *2017 IEEE International Conference on Acoustics, Speech and Signal processing (ICASSP)*. 131–135.
- [9] Turab Iqbal, Yong Xu, Qiuqiang Kong, and Wenwu Wang. 2018. Capsule Routing for Sound Event Detection. In *2018 26th European Signal Processing Conference (EUSIPCO)*. 2255–2259.
- [10] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [11] Annamaria Mesaros, Toni Heittola, Emmanouil Benetos, Peter Foster, Mathieu Lagrange, Tuomas Virtanen, and Mark D. Plumbley. 2018. Detection and Classification of Acoustic Scenes and Events: Outcome of the DCASE 2016 Challenge. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)* 26, 2 (2018), 379–393.
- [12] Annamaria Mesaros, Toni Heittola, Aleksandr Diment, Benjamin Elizalde, Ankit Shah, Emmanuel Vincent, Bhiksha Raj, and Tuomas Virtanen. 2017. DCASE 2017 challenge setup: Tasks, datasets and baseline system. In *DCASE 2017-Workshop on Detection and Classification of Acoustic Scenes and Events*.
- [13] Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen. 2016. Metrics for Polyphonic Sound Event Detection. *Applied Sciences* 6, 6 (2016), 162.
- [14] Sanjeel Parekh, Slim Essid, Alexey Ozerov, Ngoc QK Duong, Patrick Pérez, and Gaël Richard. 2018. Weakly Supervised Representation Learning for Unsynchronized Audio-Visual Events. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 2518–2519.
- [15] Mark D. Plumbley, Christian Kroos, Juan P. Bello, Gaël Richard, Daniel P.W. Ellis, and Annamaria Mesaros. 2018. *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2018 Workshop (DCASE2018)*.
- [16] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *International Conference on Learning Representations*.
- [17] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *The Journal of Machine Learning Research* 15, 1 (2014), 1929–1958.
- [18] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Advances in Neural Information Processing Systems*. 5998–6008.
- [19] Tuomas Virtanen, Annamaria Mesaros, Toni Heittola, Aleksandr Diment, Emmanuel Vincent, Emmanouil Benetos, and Benjamin Martinez Elizalde. 2017. *Proceedings of the Detection and Classification of Acoustic Scenes and Events 2017 Workshop (DCASE2017)*.
- [20] Yong Xu, Qiuqiang Kong, Wenwu Wang, and Mark D Plumbley. 2017. Surrey-CVSSP system for DCASE2017 challenge task4. *arXiv preprint arXiv:1709.00551* (2017).
- [21] Yong Xu, Qiuqiang Kong, Wenwu Wang, and Mark D Plumbley. 2018. Large-scale weakly supervised audio classification using gated convolutional neural network. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 121–125.