

Deep Learning Based Multimodal with Two-phase Training Strategy for Daily Life Video Classification

Lam Pham* Lam.Pham@ait.ac.at Austrian Institute of Technology Austria

> Dat Ngo dn22678@essex.ac.uk University of Essex UK

Trang Le* trang.le2019@qcf.jvn.edu.vn JVN Institute-VNU Viet Nam

Weißenfeld Axel Axel.Weissenfeld@ait.ac.at Austrian Institute of Technology Austria Cam Le cam.levt123@hcmut.edu.vn HCM University of Technology Viet Nam

Alexander Schindler Alexander.Schindler@ait.ac.at Austrian Institute of Technology Austria

ABSTRACT

In this paper, we present a deep learning based multimodal system for classifying daily life videos. To train the system, we propose a two-phase training strategy. In the first training phase (Phase I), we extract the audio and visual (image) data from the original video. We then train the audio data and the visual data with independent deep learning based models. After the training processes, we obtain audio embeddings and visual embeddings by extracting feature maps from the pre-trained deep learning models. In the second training phase (Phase II), we train a fusion layer to combine the audio/visual embeddings and a dense layer to classify the combined embedding into target daily scenes. Our extensive experiments, which were conducted on the benchmark dataset of DCASE (IEEE AASP Challenge on Detection and Classification of Acoustic Scenes and Events) 2021 Task 1B Development, achieved the best classification accuracy of 80.5%, 91.8%, and 95.3% with only audio data, with only visual data, both audio and visual data, respectively. The highest classification accuracy of 95.3% presents an improvement of 17.9% compared with DCASE baseline and shows very competitive to the state-of-the-art systems.

CCS CONCEPTS

- Computing methodologies \rightarrow Scene understanding.

KEYWORDS

Deep learning model, scene classification (SC), data augmentation.

ACM Reference Format:

Lam Pham, Trang Le, Cam Le, Dat Ngo, Weißenfeld Axel, and Alexander Schindler. 2023. Deep Learning Based Multimodal with Two-phase Training Strategy for Daily Life Video Classification. In 20th International Conference on Content-based Multimedia Indexing (CBMI 2023), September 20–22, 2023, Orléans, France. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/ 3617233.3617248

(*) Lam Pham and Trang Le made equal contribution to this paper.



This work is licensed under a Creative Commons Attribution International 4.0 License.

CBMI 2023, September 20–22, 2023, Orléans, France © 2023 Copyright held by the owner/author(s). ACM ISBN 979-8-4007-0912-8/23/09. https://doi.org/10.1145/3617233.3617248 1 INTRODUCTION

Recently, applying deep learning techniques to analyze videos has achieved many successes and opened a variety of real-life applications. Indeed, a wide range of deep learning based systems have been proposed for various human-relevant tasks of emotion recognition [22], lip-reading [3], or detecting human activities [5, 16, 17], etc. Recently, a dataset of daily-scene videos [20], which was proposed by DCASE challenge [4] for a new task of audio-visual scene classification (AVSC), was published and attracted attention from the video research community. Similar to the systems proposed for analyzing videos of human activities [17, 22], the state-of-theart systems proposed for AVSC task also leveraged deep learning based models and presented joined audio-visual analysis. For instances, the proposed systems in [1, 19] used convolutional based models to extract audio embeddings from audio data and leveraged pre-trained deep learning models for extracting visual embeddings from visual data. Then, the audio embeddings and the visual embeddings are concatenated and fed into dense layers for classification. To further enhance audio/visual embeddings, the authors in [23] proposed a graphed based model which was used to learn the audio/visual feature maps extracted from middle layers of deep learning backbone models. The graph based model then generates a graph based audio/visual embedding. The graph based audio/visual embeddings are finally fused with audio/visual embeddings before going to dense layers for classification. Meanwhile, authors in [6] improved the audio/visual embeddings by proposing a contrastive event-object alignment layer. The contrastive event-object alignment layer, which is based on the contrastive learning technique, helps to explore the relationship between audio and visual information by learning relative distances of event-object pairs occurring in both audio and visual scenes.

In this paper, we also leverage deep learning techniques, propose a deep learning based multimodal system for the task of AVSC. We present our main contributions: (a) We propose a mechanism, which combines a multi-model fusion and a two-phase training strategy, to generate an audio-visual embedding representing for one video input. (b) We evaluate our proposed deep learning based multimodal system on the DCASE 2021 Task 1B Development set which is the benchmark and largest dataset for the task of audiovisual daily scene classification. Results reveal that our proposed system is very competitive to the state-of-the-art systems. CBMI 2023, September 20-22, 2023, Orléans, France



Figure 1: The high-level architecture of the proposed deep learning based multimodal for AVSC task

2 PROPOSED DEEP LEARNING BASED MULTIMODAL FOR AVSC TASK

As Figure 1 shows, the high-level architecture of our proposed deep learning based multimodal for audio-visual scene classification (AVSC) comprises two individual branches: the audio branch and the visual branch, which focus on either audio or visual data extracted from the input video. Regarding the audio branch, the audio is first transformed into spectrograms which are then fed into three Audio Backbones to extract audio embeddings. Meanwhile, images are fed into two Visual Backbones to extract image embeddings. The audio and image embeddings are then combined by a Fusion Layer to generate an audio-visual embedding (i.e. The Fusion Layer is denoted by the function f). The audio-visual embedding is finally classified into target categories by a Dense Layer. From results shown in recent papers [1, 6, 19], we can see that the visual data contributes to the AVSC performance more significantly than the audio data. If we train our proposed AVSC system with an end-toend training process, it possibly causes an overfitting on the visual branch and reduces the role of the audio branch. We therefore propose a two-phase training strategy to train our proposed AVSC system. While the first training phase (Phase I) is used to train and achieve the individual Audio and Visual Backbones, the Fusion Layer and the Dense Layer are trained in the second phase (Phase II).

2.1 Phase I: Train deep learning models on individual audio or visual data to achieve audio and visual backbones

In Phase I, we aim to achieve individual high-performance Audio and Visual Backbones as shown in Figure 1. To this end, we consider the AVSC task as a combination of two independent tasks of Acoustic Scene Classification (ASC) and Visual Scene Classification (VSC). To deal with the ASC task, we leverage multiple input spectrograms, which proves powerful to improve the ASC performance [9, 11, 13, 14]. In particular, we propose deep learning based systems as shown in Figure 2 to train audio data. The audio is firstly re-sampled to 32,000 Hz, then transformed into three types of spectrogram: Mel spectrogram (MEL), Gammatone (GAM), and Constant-Q-Transform (CQT) where both temporal and spectral Lam Pham et al



Figure 2: Deep learning based models with audio data



Figure 3: Deep learning based models with visual data

features are presented. By using two channels and setting parameters of the filter number, the window size, the hop size to 128, 80 ms, 14 ms, respectively, we generate MEL, GAM, and CQT spectrograms of 128×309×2 from one 10-second audio segment. Delta and delta-delta are then applied to the three-dimensional spectrograms to obtain six-dimensional spectrograms of 128×305×6. Next, the Mixup [18] data augmentation method is applied on the sixdimensional spectrograms before feeding into a residual-inception based network for classification. Regarding the residual-inception based network for training audio spectrograms, it is separated into two main parts: A Residual-Inception block and a Dense block. The Residual-Inception block in this paper is the CNN-based backbone of the novel residual-inception deep neural network architecture which is reused from our previous works in [12]. Meanwhile, the Dense block comprises two dense layers with detailed configuration shown in the lower part of Figure 2. As we apply three types of spectrogram transformation (e.g. MEL, GAM, and CQT), we obtain three individual deep learning based models for audio input, referred to as the Aud-MEL, Aud-GAM, Aud-CQT models, respectively.

As regards the visual data, we propose deep learning models as shown in Figure 3. As Fig 3 shows, the original images (i.e. two images from each second) extracted from scene videos are first scaled into the tensor of 224×224×3 with RGB format. Then, the Mixup [18] data augmentation method is applied on the scaled images before feeding into classification models. To construct the classification models, we are inspired by [21] which shows that a combination of Inception based and ConvNet based models proves effective to improve the performance of VSC tasks. We, therefore, select InceptionV3 and ConvNeXtTiny networks from Keras library [2] for evaluating the VSC task in this paper. As both InceptionV3 and ConvNeXtTiny networks were trained with ImageNet [15] in advance, we reuse the trainable parameters from the first layer to the global pooling layer of these networks. We then connect these pre-trained layers with a two dense layers as shown in the lower part in Figure 3 to perform the InceptionV3 and ConvNeXtTiny based classification models for the VSC task in this paper. The InceptionV3 and ConvNeXtTiny based classifiers,

which are finetuned on the downstream VSC task, are referred to as Vis-CONV and Vis-INC models, respectively.

Given the individual pre-trained models of Aud-MEL, Aud-GAM, Aud-CQT for audio input and Vis-CONV and Vis-INC for visual input, we remove header layers of these pre-trained models (i.e. The header layers of the pre-trained models are either the softmax layer or the final dense layer) to perform the Audio and Visual Backbones as shown in Figure 1. In the other words, when we feed an audio or visual data into the pre-trained models of Aud-MEL, Aud-GAM, Aud-CQT, Vis-CONV, Vis-INC, the feature maps extracted at the first fully connected FC(1024) or at second fully connected FC(10) are considered as the audio and visual embeddings as shown in Figure 1.

2.2 Phase II: Train the Fusion Layer and the Dense Layer

In this Phase II, we aim to train the Fusion Layer and the Dense Layer as shown in the lower part of the Figure 1. Regarding the Fusion Layer, it is used to combine audio and visual embeddings, which are extracted from the Audio and Visual Backbones, to generate an audio-visual embedding representing for one video input. In this paper, we proposed three combination methods for the Fusion Layer. Additionally, we have two types of audio/visual embeddings: The first type of audio/visual embeddings are extracted from the first fully connected layer FC(1024) of the pre-trained deep learning based models: Aud-MEL, Aud-GAM, Aud-CQT, Vis-CONV, Vis-INC; and the second type of audio/visual embeddings are extracted from the second fully connected layer FC(10) of these pre-trained deep learning based models. As a result, we totally evaluate six types of Fusion Layer, referred to as f_1 , f_2 , f_3 , f_4 , f_5 , and f_6 . While f_1 , f_2 , f_3 are three types of combinations for the first type of audio/visual embeddings, f_4 , f_5 , f_5 are for the second type of audio/visual embeddings.

Let consider { $\mathbf{ae_g}, \mathbf{ae_m}, \mathbf{ae_c}, \mathbf{ve_i}, \mathbf{ve_c}$ } $\in \mathbb{R}^{1024}$ as the the first type of audio and visual embeddings extracted from the the first fully connected layer FC(1024) of the Audio and Visual Backbones, the fusion functions of f_1, f_2, f_3 representing for the Fusion Layer are described by

 $f_1 = \mathbf{a}\mathbf{e}_{\mathbf{g}}.\mathbf{w}_1 + \mathbf{a}\mathbf{e}_{\mathbf{m}}.\mathbf{w}_2 + \mathbf{a}\mathbf{e}_{\mathbf{c}}.\mathbf{w}_3 + \mathbf{v}\mathbf{e}_{\mathbf{i}}.\mathbf{w}_4 + \mathbf{v}\mathbf{e}_{\mathbf{c}}.\mathbf{w}_5 + \mathbf{b}, \quad (1)$

 $f_2 = (\mathbf{ae_g.w_1} + \mathbf{ae_m.w_2} + \mathbf{ae_c.w_3}).\mathbf{w_a} + (\mathbf{ve_i.w_4} + \mathbf{ve_c.w_5}).\mathbf{w_v} + \mathbf{b}, (2)$

 $f_3 = concat[(\mathbf{ae_g.w_1} + \mathbf{ae_m.w_2} + \mathbf{ae_c.w_3}), (\mathbf{ve_i.w_4} + \mathbf{ve_c.w_5})], (3)$

where $\{w_1, w_2, w_3, w_4, w_5, w_a, w_v, b\} \in \mathbb{R}^{1024}$ are trainable parameters.

Regarding the fusion function f_1 , we assume that individual audio/visual embeddings have a linear relation across each dimension. Therefore, we apply the element-wise product between each trainable vector of $\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3, \mathbf{w}_4, \mathbf{w}_5$ and each individual embedding before adding a bias **b**. By this way, a linear function, which helps to learn the relation of audio/visual embeddings across 1024 dimension, is established. Meanwhile, in the fusion function f_2 , we first apply the linear combination for only audio embeddings and for visual embeddings independently. Then, we again apply the linear combination for both audio and visual embeddings using trainable vectors of $\mathbf{w}_a, \mathbf{w}_v$ and **b**. For the fusion function f^3 , we also first apply the linear combination for only audio embeddings and only visual embeddings independently. We then concatenate two audio and visual embeddings to perform one audio-visual embedding. The fusion functions f_4 , f_5 , f_6 share the same equation as f_1 , f_2 , f_3 respectively with the second type of audio/visual input embeddings of { $ae_g, ae_m, ae_c, ve_i, ve_c$ } $\in \mathbb{R}^{10}$ and the trainable parameters of { $w_1, w_2, w_3, w_4, w_5, w_a, w_v, b$ } $\in \mathbb{R}^{10}$.

The output of the Fusion Layer, known as the audio-visual embedding, is finally classified by a Dense Layer performed by a fully connected layer FC(10) and a Softmax layer as shown in the Figure 1. Notably, as we freeze the Audio and Visual Backbones in the Phase II training process, the model is forced to learn the Fusion Layer and the Dense Layer.

3 EXPERIMENTAL RESULTS AND DISCUSSIONS

3.1 Implementation of deep learning models

We apply Tensorflow framework for implementing all deep learning based models in this paper. As mixup [18] data augmentation is used for audio spectrograms, image frames, and audio/visual embeddings to enforce classifiers, the labels of the augmented data are no longer one-hot. We therefore use Kullback-Leibler (KL) divergence loss to train back-end classification models:

$$LOSS_{KL}(\Theta) = \sum_{n=1}^{N} \mathbf{y}_n \log\left(\frac{\mathbf{y}_n}{\hat{\mathbf{y}}_n}\right) + \frac{\lambda}{2} ||\Theta||_2^2, \tag{4}$$

where *N* is the training samples, Θ present the trainable network parameters, and λ denotes the ℓ_2 -norm regularization coefficient. y_n and \hat{y}_n denote the ground-truth and the network output, respectively. All the training processes in this paper are run on two GeForce RTX 2080 Titan GPUs using Adam method [7] for optimization.

3.2 Datasets and evaluation metric

This dataset is referred to as DCASE Task 1B Development, which was proposed for DCASE 2021 challenge [4]. The dataset is slightly imbalanced and contains both acoustic and visual information, recorded from 12 large European cities: Amsterdam, Barcelona, Helsinki, Lisbon, London, Lyon, Madrid, Milan, Prague, Paris, Stockholm, and Vienna. It consists of 10 scene classes: airport, shopping mall (indoor), metro station (underground), pedestrian street, public square, street (traffic), traveling by tram, bus and metro (underground), and urban park, which can be categorized into three meta-class of indoor, outdoor, and transportation. The dataset was recorded by four recording devices simultaneously with the same setting of 48000 Hz sampling rate and 24 bit resolution. We obey the DCASE 2021 Task 1B challenge [4], separate this dataset into training (Train.) subset for the training process and evaluation (Eval.) subset for the inference. As regards the evaluation metric, the Accuracy (Acc.%), which is commonly applied in classification tasks [4] and is also proposed for DCASE Task 1B challenge, is used to evaluate the AVSC task in this paper.

3.3 Experimental results and discussion

We first evaluate the performance of our proposed systems with different types of fusion methods mentioned in Section 2.2. As the

CBMI 2023, September 20-22, 2023, Orléans, France



Figure 4: The confusion matrix results of the propose systems using f_4 fusion method with only audio data (a), with only visual data (b), and with both audio and visual data (c)

Table 1: The performance of the proposed system (Acc.%) with different types of fusion function: $f_1, f_2, f_3, f_4, f_5, f_6$.

Category	f_1	f_2	f_3	f_4	f_5	f_6
Airport	87.9	82.2	85.4	87.9	86.8	83.3
Bus	96.9	96.3	99.1	97.6	99.1	97.6
Metro	97.5	94.2	96.9	97.2	98.1	96.9
Metro Stattion	98.2	97.7	98.2	98.7	99.0	97.9
Park	95.6	92.7	93.8	97.9	99.0	98.2
Public square	92.0	91.7	89.7	91.5	87.9	88.1
Shopping mall	91.0	90.7	88.9	93.5	92.5	92.0
Street pedestrian	99.3	99.5	99.5	99.5	99.5	99.0
Street traffic	96.3	96.5	96.8	96.0	95.8	95.5
Tram	83.1	91.2	84.1	90.6	90.3	85.7
Overall	94.1	93.7	93.6	95.3	95.0	93.9

results show in Table 1, fusion methods of f_4 , f_5 , f_6 outperform f_1 , f_2 , f_3 respectively. In the other words, the fusions of audio/visual embeddings extracted from the second fully connected layer FC(10) are more effective rather than the fusions of audio/visual embeddings from the first fully connected layer FC(1024). We also see that the best accuracy score of 95.3% is achieved from f_4 method which presents a linear combination of all five audio/visual embeddings.

We then evaluate the performance comparison among the proposed systems using f_4 fusion of only audio data, of only visual data, of both audio and visual data. As the Figure 4 shows, the proposed AVSC system using only visual data (91.8%) outperforms the system with only audio data (80.5%) over almost categories, except of 'Tram' and 'Park'. When both audio and visual data are used, this helps to improve the performance in all categories (Most categories record have accuracy more than 90%, except 'Airport' with 88.0%).

We compare our best systems (i.e. using f_4 fusion) with the state-of-the-art systems. As the Table 2 shows, our proposed systems using only audio or using only visual data outperforms the state-of-the-art systems, records the accuracy of 80.5% and 91.8%, respectively. Our proposed system using both audio and visual data achieves the top-2 after the system from [10]. However, the top-1 system [10] presented an intensive ensemble of nine large deep learning models (EfficientNet, ResNeSt, and RegNet for directly training audio data; ResNet-6.4F, FesNetSt-50d, HRNet-W18 for directly training visual data; CLIP based networks of ResNet-101,

Table 2: Compare our AVSC system to the state-of-the-art
systems (Acc.%) with only using audio data, with only using
visual data, and with using both audio & visual data

Authors	Audio	Visual	Audio & Visual
DCASE baseline [20]	65.1	77.0	77.4
Javier [8]	69.0	87.0	90.0
Zhou et al. [23]	-	89.5	91.6
Hou et al. [6]	73.6	88.9	94.1
Wang et al. [19]	75.2	80.3	94.2
Chen et al. [1]	78.0	90.9	94.6
Soichiro [10]	78.1	90.9	96.1
Our best system	80.5	91.8	95.3

ResNet-50x4 ViT-B32 for extracting visual embeddings), which requires nine individual processes as well as a post processing method for an inference. Meanwhile, our proposed system combines five lighter models (3 residual-inception based models for audio data (36 M trainable parameters), InceptionV3 and ConvTiny based models for visual data (69.4 M trainable parameters)) and presents an end-to-end inference process.

4 CONCLUSION

We have proposed a deep learning based multimodal system with the two-phase training strategy for classifying daily life videos. Our proposed model, which makes use of a multi-spectrogram approach for audio data (i.e. MEL, GAM, and CQT) and multiple networks for visual data (InceptionV3 and ConvNeXtTiny), achieves the best performance of 95.3% on the benchmark dataset of DCASE 2021 Task 1B. The experimental results prove that our proposed AVSC system is very competitive to the state-of-the-art systems and potential for applying to real-life applications.

ACKNOWLEDGMENTS

The work described in this paper is performed in the H2020 project STARLIGHT ("Sustainable Autonomy and Resilience for LEAs using AI against High priority Threats"). This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 101021797.



Deep Learning Based Multimodal with Two-phase Training Strategy for Daily Life Video Classification

CBMI 2023, September 20-22, 2023, Orléans, France

REFERENCES

- Chengxin Chen, Meng Wang, and Pengyuan Zhang. 2022. Audio-Visual Scene Classification Using A Transfer Learning Based Joint Optimization Strategy. arXiv preprint arXiv:2204.11420 (2022).
- [2] François Chollet et al. 2015. Keras. https://keras.io.
- [3] Joon Son Chung, A. Senior, Oriol Vinyals, and Andrew Zisserman. 2017. Lip Reading Sentences in the Wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3444–3453.
- [4] Detection, Classification of Acoustic Scenes, and Events Community. 2021. DCASE Challenges Task 1A. http://dcase.community/challenge2021.
- [5] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. 2015. ActivityNet: A large-scale video benchmark for human activity understanding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 961–970.
- [6] Yuanbo Hou, Bo Kang, and Dick Botteldooren. 2022. Audio-visual scene classification via contrastive event-object alignment and semantic-based fusion. In 2022 IEEE 24th International Workshop on Multimedia Signal Processing (MMSP). 1–6.
- [7] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. CoRR abs/1412.6980 (2015).
- [8] Javier Naranjo-Alcazar, Sergi Perez-Castanos, Aaron Lopez-Garcia, Pedro Zuccarello, Maximo Cobos, and Francesc J Ferri. 2021. Squeeze-excitation convolutional recurrent neural networks for audio-visual scene classification. In Proc. DCASE, 16–20.
- [9] Dat Ngo, Lam Pham, Anh Nguyen, Tien Ly, Khoa Pham, and Thanh Ngo. 2022. Sound Context Classification based on Joint Learning Model and Multi-Spectrogram Features. *International Journal of Computing* 21, 2 (2022), 258–270.
- [10] Soichiro Okazaki, Quan Kong, and Tomoaki Yoshinaga. 2021. A Multi-Modal Fusion Approach for Audio-Visual Scene Classification Enhanced by CLIP Variants.. In Proc. DCASE, 95–99.
- [11] Lam Pham, Huy Phan, Truc Nguyen, Ramaswamy Palaniappan, Alfred Mertins, and Ian Mcloughlin. 2021. Robust Acoustic Scene Classification using a Multi-Spectrogram Encoder-Decoder Framework. *Digital Signal Processing* 110 (2021), 102943.
- [12] Lam Pham, Dusan Salovic, Anahid Jalali, Alexander Schindler, Khoa Tran, Canh Vu, and Phu X Nguyen. 2022. Robust, General, and Low Complexity Acoustic Scene Classification Systems and An Effective Visualization for Presenting a Sound Scene Context. arXiv preprint arXiv:2210.08610 (2022).

- [13] Lam Pham, Khoa Tran, Dat Ngo, Hieu Tang, Son Phan, and Alexander Schindler. 2022. Wider or Deeper Neural Network Architecture for Acoustic Scene Classification with Mismatched Recording Devices. In Proceedings of the 4th ACM International Conference on Multimedia in Asia. Article 35, 5 pages.
- [14] Huy Phan, Huy Le Nguyen, Oliver Y. Chén, Lam Pham, Philipp Koch, Ian McLoughlin, and Alfred Mertins. 2021. Multi-View Audio And Music Classification. In Proc. ICASSP. 611–615.
- [15] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. ImageNet Large Scale Visual Recognition Challenge. International Journal of Computer Vision (IJCV) 3 (2015), 211–252.
- [16] Jien-De Sui, Wei-Han Chen, Tzyy-Yuang Shiang, and Tian-Sheuan Chang. 2020. Real-time wearable gait phase segmentation for running and walking. In 2020 IEEE International Symposium on Circuits and Systems (ISCAS). 1–5.
- [17] Naoya Takahashi, Michael Gygli, and L. Van Gool. 2018. AENet: Learning Deep Audio Features for Video Analysis. *IEEE Transactions on Multimedia* 20 (2018), 513–524.
- [18] Yuji Tokozume, Yoshitaka Ushiku, and Tatsuya Harada. 2018. Learning from between-class examples for deep sound recognition. In International Conference on Learning Representations (ICLR).
- [19] Qing Wang, Jun Du, Siyuan Zheng, Yunqing Li, Yajian Wang, Yuzhong Wu, Hu Hu, Chao-Han Huck Yang, Sabato Marco Siniscalchi, Yannan Wang, and Chin-Hui Lee. 2022. A Study on Joint Modeling and Data Augmentation of Multi-Modalities for Audio-Visual Scene Classification. In 2022 13th International Symposium on Chinese Spoken Language Processing (ISCSLP). 453–457.
- [20] Shanshan Wang, Annamaria Mesaros, Toni Heittola, and Tuomas Virtanen. 2021. A Curated Dataset of Urban Scenes for Audio-Visual Scene Analysis. In ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). 626–630.
- [21] Weihao Yu, Pan Zhou, Shuicheng Yan, and Xinchao Wang. 2023. InceptionNeXt: When Inception Meets ConvNeXt. arXiv preprint arXiv:2303.16900 (2023).
- [22] Shiqing Zhang, Shiliang Zhang, Tiejun Huang, W. Gao, and Qi Tian. 2018. Learning Affective Features With a Hybrid Deep Model for Audio-Visual Emotion Recognition. *IEEE Transactions on Circuits and Systems for Video Technology* 28 (2018), 3030-3043.
- [23] Liguang Zhou, Yuhongze Zhou, Xiaonan Qi, Junjie Hu, Tin Lun Lam, and Yangsheng Xu. 2023. Attentional Graph Convolutional Network for Structure-Aware Audiovisual Scene Classification. *IEEE Transactions on Instrumentation and Measurement* 72 (2023), 1–15.