

Learnable Irrelevant Modality Dropout for Multimodal Action Recognition on Modality-Specific Annotated Videos

Saghir Alfasly^{1,2} Jian Lu^{1,3,*} Chen Xu^{1,2} Yuru Zou¹

¹Shenzhen Key Laboratory of Advanced Machine Learning and Applications, Shenzhen University, China

²Guangdong Key Laboratory of Intelligent Information Processing, Shenzhen, China ³Pazhou Lab, Guangzhou, China

{saghiralfasly, jianlu, yuruzou}@szu.edu.cn, chenxuszu@sina.com

Abstract

With the assumption that a video dataset is multimodality annotated in which auditory and visual modalities both are labeled or class-relevant, current multimodal methods apply modality fusion or cross-modality attention. However, effectively leveraging the audio modality in vision-specific annotated videos for action recognition is of particular challenge.

To tackle this challenge, we propose a novel audio-visual framework that effectively leverages the audio modality in any solely vision-specific annotated dataset. We adopt the language models (e.g., BERT) to build a semantic audio-video label dictionary (SAVLD) that maps each video label to its most K -relevant audio labels in which SAVLD serves as a bridge between audio and video datasets. Then, SAVLD along with a pretrained audio multi-label model are used to estimate the audio-visual modality relevance during the training phase. Accordingly, a novel learnable irrelevant modality dropout (IMD) is proposed to completely drop out the irrelevant audio modality and fuse only the relevant modalities. Moreover, we present a new two-stream video Transformer for efficiently modeling the visual modalities. Results on several vision-specific annotated datasets including Kinetics400 and UCF-101 validated our framework as it outperforms most relevant action recognition methods.

1. Introduction

One of the deep neural network (DNN) learning schemes to improve video understanding is to leverage as many input modalities as available, such as audio, RGB frames, motion, textual data, the visible text on videos, and human skeleton joints. Therefore, multimodal learning has shown a remarkable improvement in video-based action recognition. These methods process either each modality with an independent

*Corresponding author

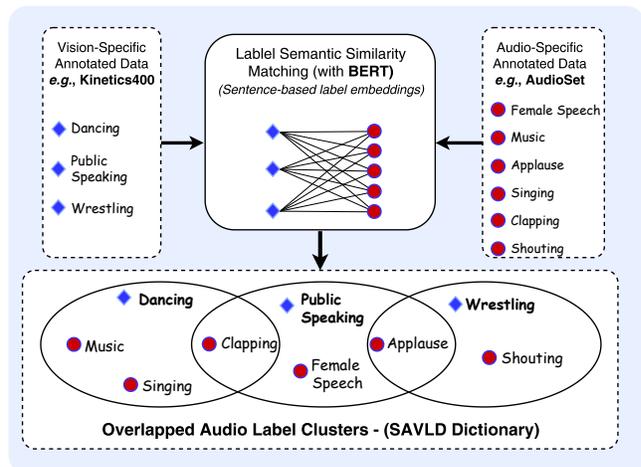


Figure 1. **Conceptual Overview:** Vision-Audio Label Mapping. In an overlapped manner, our method performs a cross-dataset textual labels mapping in which audio labels are mapped to their most closely relevant video-based human activity labels by using language models, e.g., BERT. The resulting clusters compose a SAVLD dictionary that serves as a bridge between video and audio datasets. Since our framework trains human activity multimodal models on vision-specific annotated datasets, we use SAVLD to leverage the auditory modality in videos.

DNN or all with a single shared DNN. The modalities or their feature representations are fused either in an early fusion, middle fusion, or late fusion manner [1–3]. Another fusion scheme is proposed by adopting cross-modality attention or gated units [4–7].

Motivation. Most of the attention and fusion methods [2] boost the audio-visual models’ performance on audio-vision correspondence-based videos in which auditory and visual modalities are corresponding or at least relevant as in Kinetics-Sounds [8] and EPIC-Kitchen dataset [9]. However, these fusion methods do not provide similar performance boost on the vision-specific annotated datasets, such as Kinetics400 [10] and UCF-101 [11]. This condition

is due to the high irrelevance and non-correspondence between visual and auditory modalities as concluded in [1]. In this case, multimodal methods do not effectively gain maximum benefit from unlabeled audio modality with noisy relevance on vision-specific datasets. For example, `music` can largely indicate that the human activity is `dancing`. However, music audio can be associated with several other human activities, such as car driving, or simply can be found as background audio in any edited video. This interprets the findings in an interesting study [7], which showed that unimodal models consistently outperform multimodal DNN in modality-specific datasets because multimodal networks are more prone to overfitting with their increase in capacity. This finding motivated us to look for a method that leverages the relevant audio modality while completely dropping out the irrelevant modality in training and inference phases.

Contributions. To tackle the aforementioned challenge, we present a novel multimodal training framework that trains action recognition networks with the best audio-visual modality combination on visual modality-specific datasets. In our method, we automatically estimate the audio-visual modalities relevance by leveraging the pre-trained audio classification models on large audio-specific datasets such as AudioSet [12] and VGGSound [13]. This task is effectively achieved in two phases. In the first phase, we map audio labels of an audio-specific dataset to their most semantically relevant labels in the vision-specific dataset. This process is achieved using the language pre-trained models, *e.g.*, BERT [14] or GloVe [15]. After obtaining the semantic sentence-based embedding of each label, we perform cross-dataset label matching and generate overlapped label clusters each of which contains a single label from the video dataset along with a set of its most semantically similar labels in the audio dataset as shown in Fig. 1. For example, `Dancing` label in Kinetics400 and `Music`, `Singing`, and `Clapping` labels in AudioSet are in the same cluster. Overall, the resulting overlapped clusters compose a semantic audio-video label dictionary (SAVLD) that serves as a bridge between video and audio datasets for the training supervision in our next phase. In the second phase, as a transfer learning task, our framework adopts a pre-trained audio Transformer to generate highly semantic audio features and multi-label audio predictions. The audio dataset used for building the SAVLD should be the same dataset on which the audio Transformer has been trained. Overall, the outputs of the audio Transformer and the SAVLD are then used by our proposed framework to annotate audio, as depicted in Fig. 2.

After obtaining the audio modality predictions using any audio pre-trained model (*e.g.*, AST), our framework guided by the SAVLD drops the irrelevant audio modality. This is performed by using a novel trainable irrelevant modality dropout (IMD) that consists of two main modules: The first

module is a neural network termed relevance network (RN) that receives auditory and visual modalities and decides whether they are relevant. The output of this network is sigmoidal predictions representing the relevance level that is used by the second module to decide whether to fuse the two modalities for the final video classification or to drop the audio modality. We further improve our framework learning by proposing a new intra-class cross-modality augmentation in which it randomly pairs auditory and visual modalities of the same class from different videos. This augmentation method may show a negative effect on some applications of audio-visual networks when the audio and video are required to be corresponding as the case of speech recognition. However, in our case, the audio modality is required to be relevant but is not necessarily to be aligned or accurately correspond to the visual modality because we tackle the problem of human activity recognition in this work. This motivates us to propose this augmentation method that empirically provides a reasonable performance boost.

Generally, our framework is an entirely convolution-free Transformer-based network. It leverages three video modalities: RGB frames, optical flow, and audio. The RGB and optical flow modalities are processed by using our efficient two-stream video Transformer, whereas the audio modality is processed by any audio pre-trained Transformer.

The key contributions of this work are as follows:

- A novel multimodal human activity recognition framework that leverages the power of the NLP BERT model and the pre-trained audio classification models is proposed to automatically annotate audio modality. It effectively trains audio-visual action recognition models on any vision-specific annotated dataset.
- A novel learnable IMD network is proposed to completely drop out the irrelevant audio modality, whereas the relevant modalities are fused on the basis of their relevance level.
- An efficient two-stream video Transformer is designed to learn the visual modality with few parameters as compared to the relevant video Transformers.
- An intra-class cross-modality augmentation method is proposed for generating more training samples by allowing each audio modality sample to be paired with any visual modality sample among the same class.

2. Related Work

Multimodal Action Recognition. Multimodal action recognition aims to exploit multimodality input for better human activity recognition. DNNs have empowered this learning scheme, which enriches the global feature learning [1, 3, 5–7, 16]. On the basis of the modality combination, some methods use RGB frames along with motion modality as in the two-stream networks [4, 17, 18], with the skeleton modality [19], with detected objects [2], with textual

modality [20], or with the audio modality [1, 8, 19, 21–23] which is the most common multimodality approach to learn global representation for video understanding. More than two modalities are exploited by some methods [2, 5, 7, 16, 24] to improve action recognition. Another approach is followed for obtaining video modalities as in [5, 25], which uses the compressed videos as a source of four modalities, I-frames, motion vectors, residuals, and the audio modalities. Using these modalities, the 2D/3D convolution networks have shown substantial progress in global feature modeling [2, 4, 8, 19, 21–23, 26]. However, multimodal video Transformers [1, 3, 5, 16] have shown a competitive performance as a normal reflection of the remarkable Vision Transformer (ViT) success witnessed on image recognition [27]. Although, supervised multimodal learning provides good performance [1–3, 5–8, 19, 21], some methods tend to leverage the huge unlabeled video data to learn multimodal networks in self-/weakly supervised scheme [16, 26, 28, 29]

Modality Fusion. One of the main components of any multimodal algorithm is its modality fusion module that is used to fuse and derive the cross-modality representations for the final prediction. Several fusion methods have been proposed and can be categorized into early, mid, and late fusion. These fusion approaches are adopted and empirically studied in more detail in [1–3, 6, 7, 18]. The fusion methods vary between a simply modality feature aggregation, concatenation, and cross-modality attention [1, 26]. Recently, several multimodal methods [1, 4, 5] leverage the attention mechanism to perform cross-modality feature modeling resulting in a remarkable improvement in action recognition.

Modality Dropout and Gating. Several methods, including the well-known dropout [30], have been proposed to improve the training process of the deep networks, which prevents DNNs from overfitting. Model overfitting becomes more challenging when it comes to multimodal training particularly for video understanding, as concluded in [7]. This is due to the small video dataset size as compared to image-based datasets and due to the bias towards majority classes resulting by highly imbalanced datasets. Therefore, several methods use random modality dropout, which is first used in [31], for improving the multimodal gesture recognition training. By contrast, a drop input token is proposed in [16] for reducing the training time at the cost of the prediction quality. Other studies are presented in [6, 32], which propose a gated multimodal unit for text-vision fusion in case two modalities are relevant.

Our work differs from previous methods in the following aspects. First, our vision model is a convolution-free two-stream Transformer that derives spatiotemporal features with few parameters as compared to [1, 5, 16]. Second, unlike [6, 16, 32], our proposed IMD is learnable that learns to completely drop the irrelevant modality whereas the relevant modalities are passed for the fusion step in both

Table 1. Samples of video labels in Kinetics400 dataset and their most relevant labels in AudioSet dataset selected by KNN after a semantic label embedding mapping with BERT, when $k = 5$.

Label	Relevant AudioSet Labels
applauding	speech;applause;whistling;chime;clapping
clean and jerk	fill with liquid;pump liquid;filing rasp;rumble;rustle
feeding birds	wild animals;insect;mosquito;bird;patter
sniffing	whimper;growling;cheering;whispering;rattle
sneezing	gurgling;snoring;babbling;gargling;raping
tickling	whispering;rustle;cheering;growling;screaming
yawning	babbling;raping;frying food;gurgling;snoring
writing	writing;speech;typing;chatter;mechanisms

training and inference phase. IMD aims to learn multi-modality representations from modality-specific annotated datasets. Thus, the proposed IMD is different in its structure, purpose, and training approach. Third, to the best of our knowledge, our work is the first attempt to exploit the power of both NLP and audio Transformers to automatically annotate audio modality that significantly boosted the action recognition performance. Fourth, the audio-visual correspondence (AVC) [8], which considers the dataset as audio-visual corresponding, learns to align auditory and visual modalities and decides whether they are extracted from the same time track of the same video, on the basis of only positive and negative labels. However, our relevance network (RN) does not learn the modality correspondence or alignment. Instead, it learns to simply decide whether the audio modality is class-relevant to the visual modality or not based on the generated label dictionary SAVLD regardless of which videos are used to sample the input modalities. Finally, our proposed intra-class cross-modality augmentation is different from [33], which is an inter-class sample mixing. It is also different from [26], which exchanges audio and video tracks for one time as a preprocessing step to ensure each input sample involves auditory and visual modalities.

3. Proposed Multimodal Framework

3.1. Label Dictionary SAVLD

SAVLD is a cross-dataset label dictionary in which each label in an audio-specific dataset is semantically mapped into one or more labels on the vision-specific dataset. We achieved this preprocessing step by using the well-known NLP model BERT [14]. We obtain the sentence-based semantic embeddings of each textual label in audio and video datasets because most human activity video and audio labels involve more than one word. Each video label embedding is matched with all audio label embeddings. The most k relevant audio labels are picked by using the k nearest neighbor algorithm. However, the matching process does not require much time because the number of labels in two datasets is small, that is, Kinetics400 has 400 class labels

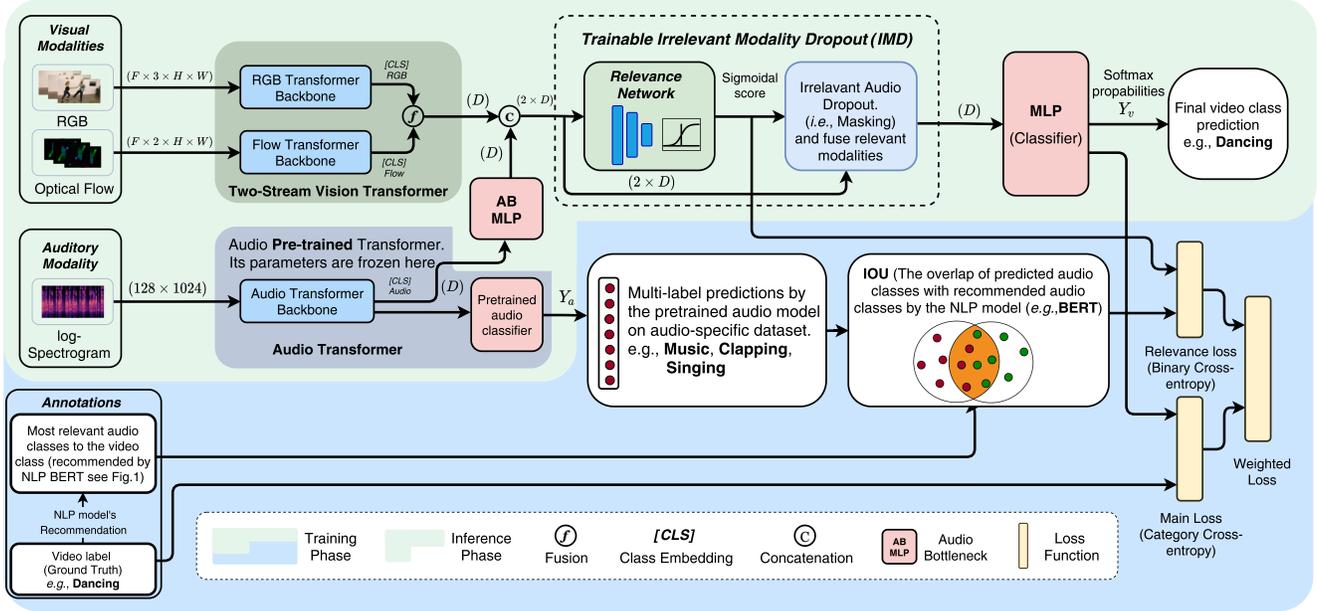


Figure 2. **Proposed Multimodal Learning Framework.** It uses three modalities: RGB, flow, and audio. Audio modality features are obtained by an off-the-shelf Transformer trained on the audio dataset AudioSet. Here, audio Transformer parameters are not trainable except the added AB which works as a bridge between our trainable network and the pre-trained frozen audio model. The trainable IMD is the main part of this framework, it estimates the auditory-visual modality relevance of the input video using the RN, and then it decides whether to fuse the audio modality or to completely drop it out using thresholding and masking layers (Fig. 4). IMD is optimized with binary cross-entropy. The IOU is computed between the audio Transformer multi-label predictions and the corresponding k relevant audio labels generated by SAVLD. The IOU is normalized into the [0–1] range to match the relevant network (RN) output range.

and AudioSet has 527 class labels. Thus, finding the closest k similar labels can be done simply by using any distance metric. We obtain a dictionary in which each video label has k audio labels of which they all are considered semantically similar. This dictionary can be considered U_v overlapped clusters, where U_v denotes the number of video dataset classes. The cluster overlap is represented here in terms of audio labels only. A conceptual overview of building SAVLD is illustrated in Fig.1, whereas a set of real Kinetics400-AudioSet mapped labels is shown in Table 1.

3.2. Visual-Modality Two-Stream Transformer

To leverage the pre-trained knowledge on large image datasets [34], we adopt several parts from the image-based ViT [27] and the video Transformer [35] to build our two-stream video Transformer.

Tokenization. Similar tokenization part of [35] is used in which the input visual-modality sample $\mathbf{V} \in \mathbb{R}^{T \times C \times W \times H}$, of T frames and $C = 3$ for RGB modality and $C = 2$ for optical flow modality, is tokenized into $T \times N$ patches in which each frame is projected to $N = HW/P^2$ patches each of which has a spatial dimension of P^2 . After organizing the obtained $T \times N$ patches, each patch is mapped to $\mathbf{x}_{j,i} \in \mathbb{R}^D$, where $j = 1, \dots, T$, $i = 1, \dots, N$, and d denotes the embedding size to be used throughout the visual

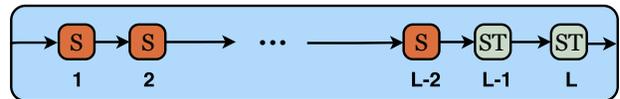


Figure 3. **Visual-Modality Transformer Encoder.** It consists of $L - 2$ Spatial **S** blocks and 2 Spatio-Temporal **ST** blocks. **ST**s are spatial-temporal factorized-based attention blocks to preserve the Transformer efficiency

two-stream transformer.

$$\mathbf{z}_{j,i}^0 = \mathbf{x}_{j,i} \mathbf{E} + \mathbf{E}_{j,i}^{pos}, \quad (1)$$

where for each patch $\mathbf{x}_{j,i}$, a spatiotemporal positional embedding $\mathbf{E}_{j,i}^{pos} \in \mathbb{R}^D$ is added. Additionally, a learnable classification embedding $\mathbf{z}_{0,0}^0 \in \mathbb{R}^D$ is added to learn the visual-modality semantic video representation.

Visual-Modality Transformer Encoder. Our visual two-stream Transformer encoder involves L stacked blocks, each of which is a spatial-based module except the last two blocks $\{L, L - 1\}$ which involve a spatiotemporal factorized self-attention to better model the spatiotemporal knowledge. Using this approach, our Transformer provides competitive performance with few parameters as compared to [35] which involves spatiotemporal attention in each block. Each Transformer block consists of a multi-

headed self-attention (MSA) [36], LayerNorm (LN) [37], and multilayer perceptron (MLP) modules. The first $L - 2$ blocks of the proposed Transformer encoder can be formulated as:

$$\mathbf{z}_{j,i}^\ell = \text{MSA}(\text{LN}(\mathbf{z}_{j,i}^{\ell-1})) + \mathbf{z}_{j,i}^{\ell-1}, \quad (2)$$

where $\ell \in \{0, \dots, L - 2\}$. The last two blocks can be formulated as:

$$\begin{aligned} \hat{\mathbf{z}}_{j,i}^\ell &= \text{LR}(\text{MSA}^{time}(\text{LN}(\mathbf{z}_{j,i}^{\ell-1}))) + \mathbf{z}_{j,i}^{\ell-1}, \\ \mathbf{z}_{j,i}^\ell &= \text{MSA}^{space}(\text{LN}(\hat{\mathbf{z}}_{j,i}^\ell)) + \hat{\mathbf{z}}_{j,i}^\ell, \end{aligned} \quad (3)$$

where LR refers to a linear layer that added after each temporal module. The classification embeddings of RGB and optical flow streams are fused with two fully-connected layers each of which is followed by ReLU activation.

3.3. Audio Transformer

We choose the pretrained AST [38] because it is the most closely similar Transformer to our video Transformer. We adopt AST to extract high semantically audio features and to perform a multi-label prediction. AST has been built on the basis of ViT [27] in which it treats audio modality processing as a vision task because it converts input audio modality into $128 \times 100t$ spectrogram. AST is mostly identical to ViT except for a few changed Transformer parameters, where AST uses patch size $P = 16$ with a stride of 10. Overall, we follow the audio input normalization and inference settings in [38] because we adopt AST as a downstream task with frozen parameters. We obtain the output predictions of AST on its audio dataset (*i.e.*, AudioSet) to obtain the video label using the label dictionary SAVLD. Additionally, we obtain the learned class embedding which is then passed to a learnable audio bottleneck.

Audio Bottleneck (AB). We add a learnable audio bottleneck as a semantic representation bridge between audio and video datasets, such as AudioSet-Kinetics400. This condition is because all audio model parameters are frozen while training our video-based action recognition model. This bottleneck plays a significant role by transforming the audio features from audio-dataset-oriented predictions to video-oriented predictions in the training phase. Simply, it consists of one LN and two LR each of which is followed by ReLU activation.

3.4. Trainable Irrelevant Modality Dropout (IMD)

IMD is the main part of our framework that automatically drops out the irrelevant audio modality. IMD consists of two main modules: The first module is a simple relevance network that receives the concatenated modalities of auditory and visual Transformers and decides whether they are relevant. The output of this network is sigmoidal predictions representing the relevance level. The sigmoidal predictions are forwarded to a binary thresholding layer to decide whether audio-modality embedding can be fused with

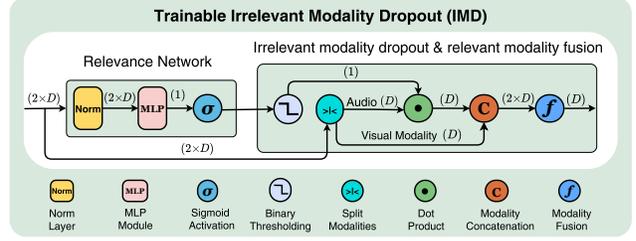


Figure 4. **Proposed Irrelevant Modality Dropout (IMD).** It is composed of two main networks. Relevance network (RN) which estimates the class-based audio-vision relevance of the input clip. The IMD network performs two steps: the output of the RN is thresholded and masked as in Eqs. 5-7. Consequently, the audio modality is completely dropped or fused as in Eq. 8.

visual-modality embedding for the final classification or to drop it. The thresholding and dot product layers compose a masking layer in which the irrelevant modality is multiplied with 0. The proposed IMD can be formulated as follows:

$$rev = \sigma(\text{MLP}(\text{LN}(\mathbf{z}_{0,0}^L)) + \mathbf{z}_{0,0}^L), \quad (4)$$

$$\delta = \begin{cases} 0, & \text{if } z < \alpha \\ rev, & \text{otherwise} \end{cases}, \quad (5)$$

$$z_{0,0}^{(a)L}, z_{0,0}^{(v)L} = \text{split}(z_{0,0}^L), \quad (6)$$

$$z_{0,0}^L = \text{Concat}((z_{0,0}^{(a)L} \cdot \delta), z_{0,0}^{(v)L}), \quad (7)$$

$$z_{av} = \text{MLP}(\text{LN}(z_{0,0}^L)), \quad (8)$$

where α denotes the defined threshold for the audio-visual modality relevance and z_{av} is the fused output. Figure 4 illustrates the IMD network and its main layers. Although it is designed for audio-modality dropout, it may be adopted for other modalities in the case of textual-vision applications whenever possible to easily build a cross-dataset label dictionary as the SAVLD we built for audio-visual training.

3.5. Framework Optimization

Our framework has two learning targets; optimizing the modality Transformers to learn the final human activity and optimizing the relevance network to accurately estimate the audio-video relevance. Therefore, we adopt two cross-entropy losses. The first loss is a binary cross-entropy for optimizing the RN by using the label dictionary SAVLD. The second loss is the default category cross-entropy for the classification learning, which is weighted with the RN loss to form the final loss. Before the category cross-entropy is applied to the output fused class embedding, it is first fed into the final classification MLP as follows:

$$Y_v = \text{Softmax}(\text{MLP}(\text{LN}(z_{av}))) \quad (9)$$

3.6. Intra-Class Cross-Modality Augmentation

We propose an intra-class cross-modality augmentation for human activity recognition, which has shown a positive impact in the Transformer training because human activity recognition on video does not require audio-video correspondence as in the case of speech-visual recognition. Thus, in our work, we focus on relevance rather than correspondence in which the audio modality should be relevant to the visual class label regardless of its video source. Therefore, we run a cross-modality intra-class augmentation by pairing audio and visual modalities from different videos but belong to the same class.

4. Experimental Results

4.1. Experiments Setup

Datasets and Evaluation Metrics. Since our aim is to tackle the problem of multimodal video understanding in vision-specific datasets, we have conducted an extensive experiments on two well-known video datasets. **Kinetics400** [10] is downloadable from YouTube, where we used in this work 241, 722 videos of 400 classes for training, and 19, 877 videos for validation. The videos of Kinetics400 are trimmed in 10 seconds. **UCF-101** [11] contains 13, 320 videos of 101 classes. Its video length about 7.2 seconds. It is split into three training/validation splits. We adopted AST audio model pretrained on **AudioSet** [12], which is considered one of the largest available audio dataset. It contains 1.8M sound clips of 10-second length. We followed the most used evaluation metric Top-1 and Top-5. We sampled 3 spatial crops and 4 temporal clips from each video and averaged their obtained accuracy. We compared our multimodal video Transformer with state-of-the-art methods in terms of computational cost floating-point operation and the number of parameters

Implementation Details. The entire proposed framework is built in Pytorch [39]. The training and evaluation phases are conducted on two machines (*i.e.*, $8 \times$ TITAN RTX and $2 \times$ RTX 2080 Ti). Unless otherwise specified, we ensure to run each set of the framework training instances by using the same settings on the same machine. The GPU’s memory is saved for the visual RGB and flow stream because AST is frozen during training and inference, except its bottleneck AB. The input RGB modality is sampled as a clip with spatial resolution of $3 \times 224 \times 224$ and temporal length of 8-frames. Similarly, flow modality is sampled with the same RGB spatial and temporal settings except for the channel of 2 to represent the x and y flow components. We trained each instance with a global input batch of 64 for 15 epochs. SGD optimizer is utilized with an initial learning rate of 0.005 and weight decay of $1e - 4$. Our visual video Transformer uses a patch of 16×16 size, block depth 12, embedding size $D = 768$, and self-attention of 12 heads.

Table 2. Impact of embedding extractor and similarity distance methods on the AudioSet-Kinetics400 SAVLD quality.

Embed. Extractor	Similarity	Top-1 (%)	Top-5 (%)
word2vec	Euclidean	80.9	94.3
GloVe	Euclidean	81.2	94.5
BERT	Euclidean	82.8	95.7
BERT	Manhattan	82.6	95.5
BERT	Cosine	82.8	95.9

The label dictionary SAVLD is saved as a text file, where each row represents the video dataset label and its k relevant audio labels.

4.2. Ablation Study

In this section, we conduct a set of extensive ablation studies on the proposed framework variants to explore their optimal settings. We start with SAVLD building techniques and then, the modality combination on the visual-modality specific dataset performance. Finally, the effect of IMD elements is studied.

4.2.1 Semantic Audio-Video Dictionary Quality

Embedding Extractor. We chose the well-known NLP model termed BERT to extract the best sentence-based semantic features for each textual label due to two main reasons: first, BERT has been trained in an unsupervised manner and has undergone millions of unorganized text from real-world sources. Thus, it can provide better semantic features that can detect the most relevant text. Secondly, unlike word2vec [40] and GloVe [15], BERT as a text feature extractor is a context-aware text analyzer that provides different embedding for the same word depending on its position in the sentence. The performance of the proposed framework on Kinetics400 using SAVLD constructed with the three embedding extractors is reported in Table 2.

Embedding Matching. Three vector similarity matching methods, namely, Euclidean, Manhattan, and cosine similarity distance are evaluated. The lower part of Table 2 reports their performance on BERT embedding matching. Notably, the similarity distance function does not show large effect on the framework performance.

4.2.2 Modality Combination

Modality combination on human activity recognition with the proposed framework is evaluated. Table 3 reports three modality combination schemes, namely, unimodal, bimodal, and multimodal on all the possible three modalities, audio, RGB, and optical flow. The use of more than one modality can remarkably improve video-based action recognition. However, RGB remains the modality with the largest effect on recognition accuracy.

Table 3. Performance of the proposed framework variants in terms of modality combination on Kinetics400.

Scheme	Modality	Top-1 (%)	Top-5 (%)
Uni-modal	Audio	26.8	45.1
	RGB	79.1	93.9
Bi-modal	Audio-Flow	56.7	72.8
	Audio-RGB	81.6	94.8
	RGB-Flow	81.1	94.3
Multi-modal	Audio-RGB-Flow	82.8	95.7

Table 4. Impact of different IMD settings on the proposed framework performance on Kinetics400.

Loss	$K-Y_\alpha$	α	M/W	Top-1 (%)	Top-5 (%)
IOU	10-10	0.50	M	81.7	95.4
Dice	10-10	0.50	M	81.6	95.5
IOU	20-10	0.50	M	81.7	94.8
IOU	10-20	0.50	M	82.3	95.5
IOU	20-20	0.50	M	81.9	95.1
IOU	10-20	0.25	M	82.8	95.7
IOU	10-20	0.75	M	80.9	94.4
IOU	10-20	0.25	W	80.4	94.3

4.2.3 Trainable Irrelevant Modality Dropout (IMD)

Overlapping Loss. The overlap between the output predictions of the pre-trained audio model AST and the corresponding K relevant textual audio labels in the SAVLD is computed as ground truth for the RN optimization. The overlap is then normalized between 0 and 1 to match the sigmoidal prediction of RN. We tested two overlap loss functions IOU and Dice overlap [41]. The upper part of Table 4 shows the performance of two functions. The overlap function does not show a large effect on the framework performance because this loss is performed only to build the ground truth for the binary cross-entropy rather than to optimize the framework itself.

Most k Audio Relevant Labels. The overlap between AST predictions and the relevant textual label vectors is computed with same size of two vectors. However, we tested the impact of using different sizes of k against Y_α . Table 4 shows that 10 – 20 is the best $k - Y_\alpha$ ratio for ground truth computation.

Relevance Thresholding (α). The sigmoidal predictions of the RN indicate the audio relevance to the visual modality on the input video sample. Thus, these sigmoidal predictions range between 0 and 1, which are going to be used by the thresholding layer in the dropout module. The effect of various threshold values on the dropout rate affecting the framework performance is reported in Table 4. Notably, when $\alpha = 0.25$, the framework shows better results.

Irrelevant Modality Masking Scheme. Two types of

masking schemes in IMD are evaluated. The first scheme is the default explained masking (**M**) method in this work in which the relevant sigmoidal score larger than α is simply multiplied with the audio embeddings received from the audio Transformer, whereas the relevance is less than α is multiplied by 0 to completely mask the audio embeddings from the fusion step. The second scheme is applying the same weighting (**W**) method in two cases when relevant sigmoidal scores are larger or less than α . The lower part of Table 4 shows that the weighting relevance method provides less performance compared with the masking method.

4.3. Comparison with State-of-the-art

Modality Fusion Methods. We compare our proposed IMD to other fusion methods in the literature. When we adopt the relevant fusion methods, we disable the RN learning to ensure the fusion process is performed on each audio-visual input sample. Several fusion methods, including SE Gate [7], NL Gate [7], late-concat [7], AVC [8], and GMU [6], are evaluated. Table 5 reports the performance of our framework when using each method of these fusion methods.

Table 5. Performance comparison of the proposed IMD against other relevant fusion methods on Kinetics400. IMD* refers to apply the proposed intra-class cross-modality augmentation.

Fusion Method	Top-1 (%)	Top-5 (%)
SE Gate [7]	79.2	94.1
NL Gate [7]	80.8	94.8
late-concat [7]	78.6	93.9
AVC [8]	79.5	94.2
GMU [6]	80.9	94.8
IMD (ours)	82.3	95.1
IMD* (ours)	82.8	95.7

Action Recognition. In this part, we evaluate our overall framework against the state-of-the-art multimodal methods for video-based action recognition on two datasets Kinetics400 [10] and UCF-101 [11]. Several uni/multimodal action recognition methods have shown good performance including the CNN-based models: TSN [50], I3D [10], AVS-lowFast [47], SlowFast [43], X3D [44], MoViNet-A5 [42], and G-Blend [7] along with the Transformer-based models: VATT [16], TimeSFormer [35], Swin Video Transformer [45], X-ViT [46], MBT [1], Mformer [48], ViViT-L [49], MM-ViT [5]. However, our multimodal Transformer-based framework outperforms all aforementioned methods on the two datasets with considerable margins. Table 6 reports the action recognition performance on Kinetics400 dataset, and Table 7 reports the results on UCF-101 dataset. Unlike our framework, most multimodal video-based deep models are usually evaluated on multimodality datasets. However, we compared our method with the best available models on

Table 6. Comparison to prior uni and multimodal action recognition methods on Kinetics400 dataset. Most compared methods are Transformer-based except the CNN-based methods: MoViNet-A5, SlowFast, and X3D. All methods are organized by their size in which the upper part involves the Transformers of base (**B**) size, whereas the lower part involves the ones with large (**L**) size.

Model	P-train	Top-1	Top-5	GFLOPs
MoViNet-A5 [42]	IN-1K	78.2	N/A	29
SlowFast 16x8 [43]	IN-1K	79.8	93.9	7,020
VATT-B [16]	IN-1K	79.6	94.9	9,090
X3D-XXL [44]	IN-1K	80.4	94.6	5,823
TimeSFormer-B [35]	IN-21K	78.0	93.7	590
Swin-B [45]	IN-1K	80.6	94.6	3,384
Swin-B [45]	IN-21K	82.7	95.5	3,384
X-ViT-16x [46]	IN-21K	80.2	94.7	850
IMD-B (ours)	IN-21K	82.8	95.7	4,464
AVSlowFast [47]	-	78.8	93.6	7,020
G-Blend [7]	Sport1M	80.4	94.8	3,303
MBT (AV) [1]	IN-21K	80.8	94.6	N/A
TimeSFormer-L [35]	IN-21K	80.7	94.7	7,140
Mformer-L [48]	IN-1K	80.2	94.8	35,553
ViViT-L [49]	IN-21K	80.6	94.7	17,352
VATT-L [16]	IN-1K	82.1	95.5	29,800
Swin-L [45]	IN-21K	83.1	95.9	7,248
IMD-L (ours)	IN-21K	84.2	96.3	8,232

Table 7. Comparison to prior multimodal action recognition methods on UCF-101. Methods are organized by pretrained data.

Model	Pretrained Weights	Top-1 (%)
CoViAR [25]	IN-1K	90.4
TSN [50]	IN-1K	94.2
CoViAR + OF [25]	IN-1K	94.9
Two-stream fusion [18]	IN-1K	92.5
MM-ViT [5]	IN-21K	95.4
IMD-B (ours)	IN-21K	95.9
I3D [10]	Kinetics	95.4
MM-ViT [5]	Kinetics	98.9
IMD-B (ours)	Kinetics	99.1

these modality-specific datasets in which our IMD models provides a Top-1 performance boost of the most similar multimodal methods (highlighted in LightBlue color) G-Blend and MBT with $\sim 3.8\%$ and $\sim 3.4\%$, respectively.

4.4. Discussion and Limitations

On the basis of the reported results and overall framework building, few worth-mentioning points need to be declared. *First*, our framework is a convolution-free Transformer-based neural network. It shows a larger performance enhancement on large datasets compared with its performance on small datasets, as shown in Tables 6 and 7. *Second*, our designed visual-modality two-stream Transformer attempts to cut the number of Transformer param-

eters while preserving the best spatiotemporal representation learning. *Third*, audio-modality masking based on the sigmoidal scores provided by the RN is relayed on the audio Transformer accuracy on its audio dataset (*e.g.*, AudioSet). For example, as reported in Table 4, the best $\alpha = 0.25$. This finding is because the AST and most audio recognition models in the literature still do not show as high recognition accuracy as visual-based models because audio modality is more challenging compared with visual data. AST [38] provides mAP of 0.485. In this manner, we think that threshold α can be higher when a better pretrained audio model is used. *Fourth*, the proposed intra-class cross-modality augmentation method shows a positive effect on the framework training. However, its main effectiveness is related to the framework design, which considers that the audio modality is noisy and unlabeled. Thus, intra-class cross-modality augmentation improves the performance by pairing visual and auditory modalities. *Finally*, even though we evaluate the quality of the SAVDL dictionary indirectly using its impact on the action recognition performance. However, it is of future interest to build a new evaluation method that can effectively evaluate the quality of text-to-text similarity matching.

5. Conclusion

In this work, we introduced a novel audio-visual action recognition framework. It leverages the power of NLP and audio recognition models to automatically annotate audio modality in the input unlabeled audio for video-based action recognition in an online training manner. The proposed framework features a novel IMD that completely drops the irrelevant audio modality based on a recommendation provided by an RN, which in turn learns to estimate the audio modality relevance to the input video category. Our framework aims to create cross-modality high semantic features that can help to bridge the relevance gap between the visual and auditory modalities in videos. As long as a strong pre-trained model on a large dataset is available to be used as a transfer learning task, this training scheme can be adopted by other bimodality and multimodality models. Thus, the same idea can be further investigated for other modalities, such as textual-auditory and textual-visual model training, which is part of our future interest.

Acknowledgments This work was supported in part by the National Natural Science Foundation of China under grants U21A20455, 61972265, 61872429 and 11871348, the Natural Science Foundation of Guangdong Province of China under grant 2020B1515310008, Educational Commission of Guangdong Province of China under grant 2019KZDZX1007, and the Shenzhen Basis Research Project under grant JCYJ20180305125521534.

References

- [1] Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. Attention bottlenecks for multimodal fusion. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. 1, 2, 3, 7, 8
- [2] Olga Zatsarynna, Yazan Abu Farha, and Juergen Gall. Multi-modal temporal convolutional network for anticipating actions in egocentric videos. In *CVPR Workshops*, pages 2249–2258, June 2021. 1, 2, 3
- [3] Sangho Lee, Youngjae Yu, Gunhee Kim, Thomas Breuel, Jan Kautz, and Yale Song. Parameter efficient multimodal transformers for video representation learning. In *International Conference on Learning Representations*, 2021. 1, 2, 3
- [4] Lu Chi, Guiyu Tian, Yadong Mu, and Qi Tian. Two-Stream Video Classification with Cross-Modality Attention. In *ICCVW*, pages 4511–4520. IEEE, oct 2019. 1, 2, 3
- [5] Jiawei Chen and Chiu Man Ho. MM-ViT: Multi-Modal Video Transformer for Compressed Video Action Recognition. *arXiv*, aug 2021. 1, 2, 3, 7, 8
- [6] John Arevalo, Thamar Solorio, Manuel Montes-Y-Gómez, and Fabio A. González. Gated multimodal units for information fusion. *5th International Conference on Learning Representations, ICLR 2017 - Workshop Track Proceedings*, 2019. 1, 2, 3, 7
- [7] Weiyao Wang, Du Tran, and Matt Feiszli. What Makes Training Multi-Modal Classification Networks Hard? *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 12692–12702, 2020. 1, 2, 3, 7, 8
- [8] Relja Arandjelovic and Andrew Zisserman. Look, Listen and Learn. In *ICCV*, volume Octob, pages 609–617, oct 2017. 1, 3, 7
- [9] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018. 1
- [10] Joao Carreira and Andrew Zisserman. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*, pages 4724–4733. IEEE, jul 2017. 1, 6, 7, 8
- [11] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *arXiv*, dec 2012. 1, 6, 7
- [12] Jort F. Gemmeke, Daniel P.W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter. Audio Set: An ontology and human-labeled dataset for audio events. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, pages 776–780, 2017. 2, 6
- [13] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. VGG SOUND : A large-scale audio-visual dataset. In *ICASSP*, pages 721–725, 2020. 2
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova Google, and A I Language. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Naacl-Hlt 2019*, 2018. 2, 3
- [15] Paul M. Brennan, James J.M. Loan, Neil Watson, Pragnesh M. Bhatt, and Peter Alwyn Bodkin. Pre-operative obesity does not predict poorer symptom control and quality of life after lumbar disc surgery. *British Journal of Neurosurgery*, 31(6):682–687, 2017. 2, 6
- [16] Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. VATT: Transformers for Multimodal Self-Supervised Learning from Raw Video, Audio and Text. In *NeurIPS*, 2021. 2, 3, 7, 8
- [17] Karen Simonyan and Andrew Zisserman. Two-Stream Convolutional Networks for Action Recognition in Videos. *Advances in Neural Information Processing Systems*, jun 2014. 2
- [18] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional Two-Stream Network Fusion for Video Action Recognition. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2016-Decem(i):1933–1941, 2016. 2, 3, 8
- [19] Hamid Reza Vaezi Joze, Amirreza Shaban, Michael L. Iuzolino, and Kazuhito Koishida. MMTM: Multimodal transfer module for CNN fusion. *CVPR*, pages 13286–13296, 2020. 2, 3
- [20] Hu Xu, Gargi Ghosh, Po-yao Huang, Prahal Arora, Masoumeh Aminzadeh, Christoph Feichtenhofer, Florian Metze, and Luke Zettlemoyer. VLM: Task-agnostic Video-Language Model Pre-training for Video Understanding. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4227–4239, Stroudsburg, PA, USA, 2021. Association for Computational Linguistics. 3
- [21] Relja Arandjelović and Andrew Zisserman. Objects that Sound. In *European Conference on Computer Vision (ECCV)*, 2018. 3
- [22] Ahmed Hussen Abdelaziz, Barry John Theobald, Paul Dixon, Reinhard Knothe, Nicholas Apostoloff, and Sachin Kajareker. Modality Dropout for Improved Performance-driven Talking Faces. *ICMI 2020 - Proceedings of the 2020 International Conference on Multimodal Interaction*, pages 378–386, 2020. 3
- [23] Alejandro Cartas, Jordi Luque, Petia Radeva, Carlos Segura, and Mariella DiMiccoli. Seeing and hearing egocentric actions: How much can we learn? *ICCVW*, pages 4470–4480, 2019. 3
- [24] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. EPIC-Fusion: Audio-Visual Temporal Binding for Egocentric Action Recognition. In *ICCV*, volume 2019-Octob, pages 5491–5500. IEEE, oct 2019. 3
- [25] Chao-Yuan Wu, Manzil Zaheer, Hexiang Hu, R. Manmatha, Alexander J. Smola, and Philipp Krahenbuhl. Compressed Video Action Recognition. In *CVPR*, pages 6026–6035. IEEE, jun 2018. 3, 8

- [26] Yu Wu and Yi Yang. Exploring heterogeneous clues for weakly-supervised audio-visual video parsing. In *CVPR*, 2021. 3
- [27] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 3, 4, 5
- [28] Jean Baptiste Alayrac, Adrià Recasens, Rosalia Schneider, Relja Arandjelović, Jason Ramapuram, Jeffrey de Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. *Advances in Neural Information Processing Systems*, 2020-Decem(ii):1–19, 2020. 3
- [29] Andrew Owens and Alexei A. Efros. Audio-Visual Scene Analysis with Self-Supervised Multisensory Features. In *ECCV*, volume 11210 LNCS, pages 639–658, 2018. 3
- [30] Srivastava Nitish, Hinton Geoffrey, Krizhevsky Alex, Sutskever Ilya, and Salakhutdinov Ruslan. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 2014. 3
- [31] Natalia Neverova, Christian Wolf, Graham Taylor, and Florian Nebout. ModDrop: Adaptive Multi-Modal Gesture Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(8):1692–1706, aug 2016. 3
- [32] Feiran Huang, Kaimin Wei, Jian Weng, and Zhoujun Li. Attention-Based Modality-Gated Networks for Image-Text Sentiment Analysis. *ACM Transactions on Multimedia Computing, Communications, and Applications*, 16(3):1–19, sep 2020. 3
- [33] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. MixUp: Beyond empirical risk minimization. In *International Conference on Learning Representations (ICLR)*, 2018. 3
- [34] 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. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115:211–252, 2015. 4
- [35] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, volume 139, pages 813–824, Jul 2021. 4, 7, 8
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6000–6010, California, USA, 2017. .. 5
- [37] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer Normalization. *arXiv*, 2016. 5
- [38] Yuan Gong, Yu-An Chung, and James Glass. AST: Audio Spectrogram Transformer. *arXiv*, 2021. 5, 8
- [39] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, 2019. 6
- [40] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed Representations of Words and Phrases and their Compositionality. In *NIPS*, oct 2013. 6
- [41] W.R. Crum, Oscar Camara, and D.L.G. Hill. Generalized Overlap Measures for Evaluation and Validation in Medical Image Analysis. *IEEE Transactions on Medical Imaging*, 25(11):1451–1461, nov 2006. 7
- [42] Dan Kondratyuk, Liangzhe Yuan, Yandong Li, Li Zhang, Mingxing Tan, Matthew Brown, and Boqing Gong. MoViNets: Mobile Video Networks for Efficient Video Recognition. pages 16020–16030, 2021. 7, 8
- [43] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 7, 8
- [44] Christoph Feichtenhofer. X3D: Expanding Architectures for Efficient Video Recognition. *CVPR*, pages 200–210, 2020. 7, 8
- [45] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video Swin Transformer. *arXiv*, pages 1–12, 2021. 7, 8
- [46] Adrian Bulat, Juan-Manuel Perez-Rua, Swathikiran Sudhakaran, Brais Martinez, and Georgios Tzimiropoulos. Space-time mixing attention for video transformer. In *NeurIPS*, 2021. 7, 8
- [47] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer. Audiovisual SlowFast Networks for Video Recognition. *arXiv*, 2020. 7, 8
- [48] Mandela Patrick, Dylan Campbell, Yuki M. Asano, Ishan Misra, Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and João F. Henriques. Keeping Your Eye on the Ball: Trajectory Attention in Video Transformers. In *Advances in Neural Information Processing Systems*, jun 2021. 7, 8
- [49] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViViT: A Video Vision Transformer. *arXiv*, 2021. 7, 8
- [50] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal Segment Networks: Towards Good Practices for Deep Action Recognition. In *ECCV 2016*, pages 20–36. 2016. 7, 8
- [51] Hilde Kuehne, Hueihan Jhuang, E. Garrote, T. Poggio, and T. Serre. HMDB: A large video database for human motion recognition. In *ICCV*, pages 2556–2563. IEEE, nov 2011.