

CoLeaF: A Contrastive-Collaborative Learning Framework for Weakly Supervised Audio-Visual Video Parsing

Faegheh Sardari¹, Armin Mustafa¹, Philip J. B. Jackson¹, and
Adrian Hilton¹

Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, UK
{f.sardari, armin.mustafa, p.jackson, a.hilton}@surrey.ac.uk

Abstract. Weakly supervised audio-visual video parsing (AVVP) methods aim to detect audible-only, visible-only, and audible-visible events using only video-level labels. Existing approaches tackle this by leveraging unimodal and cross-modal contexts. However, we argue that while cross-modal learning is beneficial for detecting audible-visible events, in the weakly supervised scenario, it negatively impacts unaligned audible or visible events by introducing irrelevant modality information. In this paper, we propose CoLeaF, a novel learning framework that optimizes the integration of cross-modal context in the embedding space such that the network explicitly learns to combine cross-modal information for audible-visible events while filtering them out for unaligned events. Additionally, as videos often involve complex class relationships, modelling them improves performance. However, this introduces extra computational costs into the network. Our framework is designed to leverage cross-class relationships during training without incurring additional computations at inference. Furthermore, we propose new metrics to better evaluate a method’s capabilities in performing AVVP. Our extensive experiments demonstrate that CoLeaF significantly improves the state-of-the-art results by an average of 1.9% and 2.4% F-score on the LLP and UnAV-100 datasets, respectively. Code is available at: <https://github.com/faeghehsardari/coleaf>.

Keywords: Unaligned audio-visual learning · Audio-visual video parsing · Weakly supervised learning · Event detection

1 Introduction

Audio-visual learning has emerged as a crucial field within computer vision, enabling AI systems to develop a more nuanced understanding of the world. Most audio-visual tasks (*e.g.*, audio-visual event localization [4, 18, 23] and audio-visual question answering [10, 14, 25]) assume that audio and visual data are temporally aligned and share a common concept. However, in practice, this assumption is not always valid. For instance, a baby might cry out of the camera’s view. The Audio-Visual Video Parsing (AVVP) task [17] instead introduces an unaligned




Fig. 1: An example of weakly supervised AVVP task. The method learns to temporally detect audible-only, visible-only, and audible-visible events in a video, using only video-level labels during training.

setting of audio-visual learning to recognize and localize audible-only, visible-only, and audible-visible events in a video. Since training a fully supervised AVVP model requires modality-specific densely annotated data, making it a costly endeavour, the AVVP task is tackled in a weakly supervised scenario where only video-level labels are available, see Fig.1. To address the weakly supervised AVVP task, all existing approaches, such as [3,17,24], incorporate unimodal and cross-modal contexts. However, we argue that while extracting cross-modal information is beneficial for multi-modal learning tasks, in the weakly supervised scenario, it can negatively impact the detection of audible-only and visible-only events (i.e., unaligned audible-visible events) by introducing irrelevant modality information. We investigated this in Table 1 by comparing the performance of CMPAE [3] with and without using cross-modal contexts. The results show that leveraging cross-modal contexts improves the performance on average in detecting audible-visible events, while in detecting unaligned events, a network focusing solely on unimodal information yields better results. These findings motivated us to raise this question: **Can we optimize the integration of cross-modal context in the embedding space such that the network explicitly learns to leverage cross-modal information for audible-visible events, while filtering them out for unaligned events?** To achieve this, we introduce CoLeaF, a novel learning framework for the weakly supervised AVVP task. CoLeaF relies on two network branches: Reference and Anchor. Reference focuses on unimodal information, while Anchor incorporates unimodal and cross-modal contexts. Both branches process the same audio and visual input signals and are trained simultaneously using the video-level labels. Additionally, during training, Reference is leveraged by a novel contrastive objective to optimize cross-modal learning in Anchor based on the event types present in the input video. The proposed objective encourages Anchor to learn embedding representations similar to those learned in Reference, but the strength of the encouragement adapts to the degree of unaligned audible-visible events present in the input video. More unaligned events lead to a stronger encouragement, while fewer unaligned events lead to a weaker encouragement. The unalignment degree is computed from the Reference’s output predictions. Since the optimization

Table 1: Performance of CMPAE [3] when using only unimodal (U) information versus exploiting both unimodal and cross-modal contexts (U+C) in detecting audible-only (A), visible-only (V) and audible-visible (AV) event types, in terms of true positive (TP), true negative (TN), false positive (FP), and false negative (FN) rates.

Metric	A		V		AV	
	U/(U+C)		U/(U+C)		U/(U+C)	
TP	43.5	/ 42.7	67.8	/ 66.4	80.8	/ 84.6
TN	97.6	/ 97.4	97.3	/ 97.0	95.2	/ 95.6
FP	0.4	/ 0.6	0.9	/ 0.7	3.3	/ 3.8
FN	94.7	/ 95.5	97.7	/ 99.2	33.2	/ 29.9

process relies on the performance of Reference, its effectiveness may diminish if Reference is solely trained on ‘modality-agnostic’ video-level labels. To address this without requiring additional annotations, Anchor collaboratively distills free modality-aware pseudo-labels for Reference. After training, Anchor is deployed for inference.

To effectively perform the AVVP task, we should exploit unimodal and cross-modal contexts while aware of class dependencies. For example, this is often seen when musical instruments are played alongside people singing. Authors in [12,13] design networks to explicitly learn cross-class relationships. However, their approaches are computationally expensive. For example, in [12], to capture cross-class dependencies, they extend their model to map each temporal segment of the video into the maximum number of event classes, and apply attention layers within the event classes for each segment individually. For a video with a length of T segments and C class categories, this led to a significant computational overhead with a ratio of $T \times C \times C$. This key challenge motivated us to raise this novel question: **Can we take advantage of explicitly modelling co-occurrence class relationships without imposing the overhead of heavy computational costs?** To achieve this, in CoLeaF, Reference that is employed only during training is designed to explicitly learn co-occurrence class relationships and distil this knowledge to Anchor. This novel design allows us to benefit from cross-class relationships during training without introducing computational overhead on Anchor that is employed at inference.

For a deeper understanding of how our proposed framework models unaligned events, we introduce novel metrics for AVVP. Current metrics evaluate the method’s performance in detecting audible-only and visible-only events by comparing the estimated audio and visual event proposals individually with audio and visual ground-truth labels. In our proposed metrics, we consider both audio and visual modalities, as comparing them individually can lead to the mistaken interpretation of an audible-visible event as both audible-only and visible-only. Our key contributions are summarized as: (i) we pioneer the first exploration in explicitly optimizing the integration of cross-modal context for the weakly supervised AVVP task, (ii) we propose an approach that allows the network to explicitly learn cross-class relationships without adding computational overhead during inference. This can benefit various computer vision tasks that demand a deeper understanding of interactions between classes, (iii) we introduce novel metrics for effective evaluation of AVVP which are beneficial for future development of this task, (iv) we evaluate CoLeaF on two challenging publicly available datasets and improve the state-of-the-art results on average by 1.9% and 2.4% F-score on LLP and UnAV-100 datasets, respectively, and (v) we perform extensive ablation studies to evaluate our proposed framework’s design.

2 Related Works

Tian et al. [17] defined the AVVP task and also introduced the LLP dataset for this purpose. Due to the laborious nature of the labelling process for this task,

for the training set, they obtained only video-level labels, so the learning process is performed in a weakly supervised setting. To perform the task, they design the Hybrid Attention Network (HAN), where they extract and combine unimodal and cross-modal contexts from audio and visual data by applying self-attention and cross-attention mechanisms and introduce Multimodal Multiple Instance Learning (MMIL) to facilitate network training under weakly supervised setting. Following HAN, the proposed approaches can be divided into the two categories:

Architectural Improvements – To better explore unimodal and cross-modal contexts, the authors in [8,24] expand HAN to a multi-scale pyramid network. Alternatively, Gao et al. [3] attempt to decrease the uncertainty of HAN’s predictions under weakly supervised learning by introducing a modality presence-absence confidence term in the cross-entropy loss function based on Subjective Logic Theory [16]. Similar to the network architecture, their loss confidence term is also explicitly derived from the integration of unimodal and cross-modal information. However, as stated in Section 1, the cross-modal learning in these architectures negatively affects the detection of unaligned events. In contrast, in our proposed framework, the network explicitly learns to filter out cross-modal contexts for unaligned events. On the other hand, [12,13] additionally learn cross-class dependencies. Although they improve HAN’s performance, they suffer from computational complexity. In contrast, our framework leverages the class relationships during training without incurring additional costs in inference.

Label Refinement – The core goal of these approaches [2,3,8,21,26] is to enhance the performance of existing network architectures by obtaining modality-specific knowledge for training. For example, the authors in [3] devised a two-stage approach in which they initially extract offline modality-specific annotations by applying large-scale pre-trained models such as CLAP [22] and CLIP [15] on the audio and visual data. Subsequently, they train a network backbone from scratch using the additional labels and a cross-entropy loss function. However, the training process of such approaches is complicated and inefficient. Alternatively, [8] propose an online method where the estimated probabilities of a network, focusing on individual modalities, guide the final estimation of the AVVP task through the mean square error loss. These approaches may optimize the integration of cross-modal information as they obtain hard/soft modality-aware labels for training. However, their optimization process is implicit and suboptimal. In contrast, our proposed approach explicitly optimizes cross-modal learning in the embedding space. Furthermore, in these approaches, the pseudo labels are obtained from ‘modality-agnostic’ video-level labels, making them unreliable. However, our online collaborative solution mitigates the impact of weak labels on our optimization process.

3 Methodology

In this section, we first define the weakly supervised audio-visual video parsing task, followed by a brief review of the HAN network [17]. Then, we elaborate on our proposed framework CoLeaF.

3.1 Preliminaries

Problem Formulation – In weakly supervised AVVP, our goal is to detect audible-only, visible-only, and audible-visible events in a given video by training from video-level labels, as defined in [17]. For the video sequence $S = \{A_t, V_t\}_{t=1}^T$ with a length of T segments, A_t and V_t represent the audio and visual signals at the t^{th} segment, respectively. At the inference, the network task is to provide the multi-class event labels $Y_t = \{y_{t,i}^a, y_{t,i}^v, y_{t,i}^{av}\}_{i=1}^C$ for each time segment of S , where $y_{t,i}^a, y_{t,i}^v, y_{t,i}^{av} \in \{0, 1\}$ indicate audible-only, visible-only, and audible-visible events, respectively, and C represents the maximum number of event classes in the dataset. However, during training, we lack access to the ‘modality-specific segment-level’ labels Y_t , having only the ‘video-level’ labels that are ‘modality-agnostic’ $Y = \{y_i\}_{i=1}^C$, where $y_i \in \{0, 1\}$.

Hybrid Attention Network (HAN) – HAN [17] is a popular baseline for the weakly supervised AVVP task. In HAN, first, two pre-trained models are applied on the audio and visual signals of the input video S to extract segment-level audio and visual input tokens, $F^a = \{f_t^a\}_{t=1}^T$ and $F^v = \{f_t^v\}_{t=1}^T$, respectively. Then, it aggregates unimodal and cross-modal information through applying and combining self-attention and cross-attention mechanisms [20] on the input tokens

$$\hat{f}_t^a = f_t^a + \text{Self-attn}(f_t^a, F^a) + \text{Cross-attn}(f_t^a, F^v), \quad (1)$$

$$\hat{f}_t^v = f_t^v + \text{Self-attn}(f_t^v, F^v) + \text{Cross-attn}(f_t^v, F^a), \quad (2)$$

Finally, HAN uses a share fully connected layer FC to predict multi-class segment-level audio and visual event probabilities from the audio and visual output tokens as $P^\phi = \text{Sigmoid}(FC(\hat{F}^\phi))$, where $P^\phi = \{p_t^\phi \in \mathbb{R}^{1 \times C}\}_{t=1}^T$ and $\phi \in \{a, v\}$. During the inference P^ϕ is utilized to predict segment-level audible-only, visible-only, and audible-visible events. However, during the training, as we only have video-level labels, HAN applies a Multimodal Multiple Instance Learning (MMIL) module to provide video-level probabilities as

$$W^\Delta = \text{Soft-max}_t(FC_t([\hat{F}^a; \hat{F}^v])), \quad W^\blacktriangle = \text{Soft-max}_{av}(FC_{av}([\hat{F}^a; \hat{F}^v])), \quad (3)$$

$$\mathbb{P} = \sum_{t=1}^T \sum_{m=1}^M (W_{t,m}^\Delta \odot W_{t,m}^\blacktriangle \odot P_{t,m}), \quad \text{where } P = [P^a; P^v], \quad (4)$$

and $M=2$ refers to audio and visual modalities, $P, W^\Delta, W^\blacktriangle \in \mathbb{R}^{T \times 2 \times C}$, $\mathbb{P} \in \mathbb{R}^{1 \times C}$, and $[\cdot]$ and \odot indicates the concatenation and Hadamard product operations, respectively. Soft-max_t and Soft-max_{av} represent the Softmax layers along the temporal and audio-visual dimensions, respectively. With the video-level probabilities \mathbb{P} , and the video-level labels Y , HAN is optimized through Binary Cross Entropy (BCE) loss $\mathcal{L}_{video} = BCE(Y, \mathbb{P})$.

3.2 CoLeaF: Contrastive-Collaborative Learning Framework

In this section, we introduce CoLeaF, a novel learning framework designed to address two key challenges in the weakly supervised AVVP task: (i) effectively

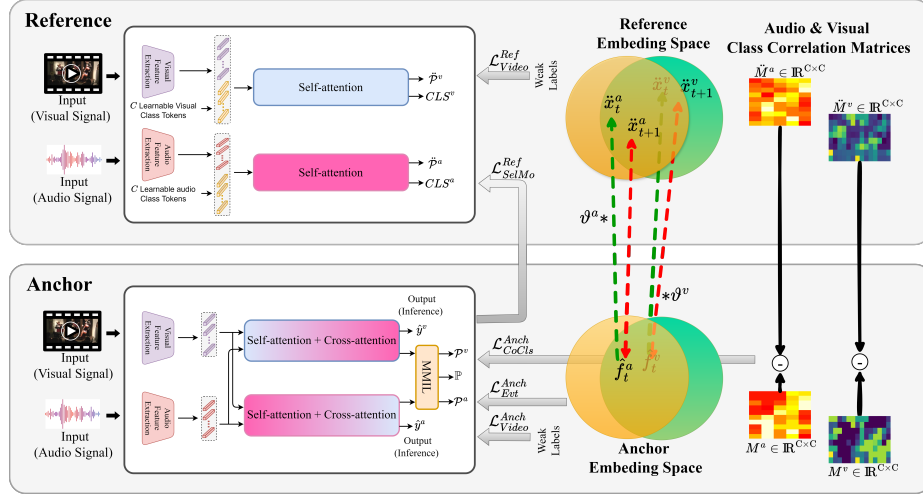


Fig. 2: The overall scheme of CoLeaF for the weakly supervised AVVP task. CoLeaF has two network branches, Reference and Anchor. The Reference branch focuses on unimodal information and cross-class relationships, while the Anchor branch exploits both unimodal and cross-modal contexts. The branches are trained simultaneously through the conventional video-level losses $\{\mathcal{L}_{video}^{Ref}, \mathcal{L}_{video}^{Anchr}\}$ and our novel contrastive and knowledge distillation losses $\{\mathcal{L}_{Evt}^{Anch}, \mathcal{L}_{SelfMo}^{Ref}, \mathcal{L}_{CoCls}^{Anchr}\}$. In inference, Anchor is deployed for the AVVP task.

integrating the cross-modal contexts with the lack of modality-specific annotations, and (ii) efficiently learning cross-class relationships. CoLeaF relies on two network branches, Reference and Anchor, as shown in Fig. 2. The Reference branch, used only during training, is dedicated to exploiting unimodal and co-occurrence class relationships. The Anchor branch exploits unimodal and cross-modal contexts and serves as the main network branch deployed during inference. The branches receive the same audio and visual input data and are trained simultaneously using two types of losses: (a) the conventional video-level weakly supervised loss, as in HAN, and (b) our proposed collaborative losses, including event-aware contrastive loss, co-occurrence class knowledge distillation loss, and self-modality-aware knowledge distillation loss. In the following, we explain the details of each branch and the losses.

Reference Branch – This branch addresses the weakly supervised AVVP task by primarily focusing on unimodal information. This can be achieved by applying two self-attention layers on the audio and visual input tokens individually. However, as videos with multi-label events contain complex relationships between the event classes, we simultaneously leverage this information to improve the method’s performance. To model explicitly the cross-class relationships, inspired by [13], we first concatenate the audio and visual input tokens, $F^a = \{f_t^a\}_{t=1}^T$ and $F^v = \{f_t^v\}_{t=1}^T$, with learnable audio and visual class tokens, $C^a = \{c_i^a\}_{i=1}^C$ and $C^v = \{c_i^v\}_{i=1}^C$,

$$X^\phi = \{x_n^\phi\}_{n=1}^{T+C} = [\{f_t^\phi\}_{t=1}^T; \{c_i^\phi\}_{i=1}^C], \quad (5)$$

where $\phi \in \{a, v\}$ and $[\cdot]$ indicates the concatenation operations. We then forward them into our self-attention backbones to exploit both intra-modality and class relationships,

$$\{\ddot{x}^\phi\}_{n=1}^{T+C} = \{\text{Self-attn}_\phi(x_n^\phi, X^\phi)\}_{n=1}^{T+C}. \quad (6)$$

We take the outputs of self-attention layers which correspond to the audio and visual input tokens, $\ddot{F}^\phi = \{\ddot{x}_n^\phi\}_{n=1}^T$, to predict segment-level audio and visual event probabilities through a fully connected layer as $\{\ddot{p}_t^\phi\}_{t=1}^T = \text{Sigmoid}(FC_\phi(\ddot{F}^\phi))$, where $\ddot{p}_t^\phi \in \mathbb{R}^{1 \times C}$. Subsequently, the audio and visual event probabilities are computed as $\ddot{P}^\phi = \sum_{t=1}^T (\ddot{w}_t^\phi \odot \ddot{p}_t^\phi)$

$$\ddot{P}^\phi = \sum_{t=1}^T (\ddot{w}_t^\phi \odot \ddot{p}_t^\phi), \quad (7)$$

where $\ddot{P}^\phi, \ddot{w}_t \in \mathbb{R}^{1 \times C}$, and \ddot{w}_t are temporal learnable weight. With the audio and visual event probabilities and video-level annotations $Y \in \mathbb{R}^{1 \times C}$, the Reference branch is optimized through the video-level loss $\mathcal{L}_{video}^{Ref} = \sum_{\phi \in \{a, v\}} BCE(Y, \ddot{P}^\phi)$.

Optimizing the Reference branch by relying only on the \ddot{F}^ϕ output tokens does not guarantee that the input class tokens learn to model the class events. To ensure that they represent the classes effectively, we also utilize the output tokens of self-attention layers which correspond to the input class tokens, $\{\ddot{c}_i^\phi\}_{i=1}^C = \{\ddot{x}_n^\phi\}_{n=T+1}^{T+C}$, for the network optimization. To achieve this, we apply a global average pooling layer on these tokens along their embedding dimension to convert them into a 1-D dimensional vector, followed by a sigmoid activation function to convert them into range 0 and 1 as $CLS^\phi = \{\text{Sigmoid}(\text{AvgPool}(\ddot{c}_i^\phi))\}_{i=1}^C$. Then, they are supervised by video-level annotations and cross entropy loss. Therefore, the video-level loss is extended as $\mathcal{L}_{video}^{Ref} = \sum_{\phi \in \{a, v\}} BCE(Y, \ddot{P}^\phi) + BCE(Y, CLS^\phi)$.

Anchor Branch – This branch tackles the weakly supervised AVVP task by exploiting both unimodal and cross-modal contexts. To achieve this, a backbone that performs the AVVP task using self-attention and cross-attention mechanisms (*e.g.*, HAN) can serve as Anchor to predict segment-level and video-level probabilities $P^\phi = \{p_t^\phi \in \mathbb{R}^{1 \times C}\}_{t=1}^T$ and $\mathbb{P} \in \mathbb{R}^{1 \times C}$ (Eqs 3-4), respectively, from the audio and visual input tokens F^ϕ . The segment-level labels are employed for inference phase, as well as later in this section for our proposed self-modality-aware knowledge distillation loss, while the video-level probabilities are utilized to compute the video-level loss $\mathcal{L}_{video}^{Anch} = BCE(Y, \mathbb{P})$, as in HAN network.

Event-Aware Contrastive Learning – We argue that while leveraging cross-modal information is beneficial in the detection of audible-visible events, in the weakly supervised scenario, it can negatively impact the detection of audible-only and visible-only events by introducing irrelevant modality information. Our findings in Table 1 support this argument. To tackle this challenge, we introduce contrastive learning to encourage the learned representations in the Anchor branch to become similar to those of the Reference branch for unaligned audio-visible event types. To achieve this, for each audio and visual output tokens of the Anchor branch \hat{f}_t^ϕ , the output tokens of the corresponding segment from the

Reference branch \ddot{x}_t^ϕ can be chosen as positive samples, and the output tokens of the rest of its segments can be considered as negative set. Subsequently, the Anchor branch can be optimized through Noise Contrastive Estimation (NCE) loss [5]. In this formulation, regardless of the event types in a video, Anchor is encouraged to align with Reference. However, a video can contain multiple events spanning audible-only, visible-only, and audible-visible, and our goal is to get close to the tokens' representations of the Reference branch for only unaligned audio-visible event types. We address this issue by introducing a novel event-aware NCE loss that adjusts the strength of contrastive learning encouragement based on the degree of audible-only and visible-only events present in the input video

$$\mathcal{L}_{Evt}^{Anch} = -\frac{1}{T} * \sum_{\phi \in \{a,v\}} \vartheta^\phi * \sum_{t=1}^T \log \frac{\exp(\hat{f}_t^{\phi^\top} \cdot \ddot{x}_t^\phi / \tau)}{\sum_{n=1, n \neq t}^T \exp(\hat{f}_t^{\phi^\top} \cdot \ddot{x}_n^\phi / \tau)}, \quad (8)$$

where weight ϑ^ϕ is computed from a set of pseudo-labels $\ddot{G}^\phi = \{\ddot{g}_i^\phi\}_{i=1}^C$ distilled from the Reference branch

$$\ddot{g}_i^\phi = \begin{cases} 1 & \text{if } \ddot{\mathcal{P}}_i^\phi > \theta \\ 0 & \text{else} \end{cases}, \quad (9)$$

$$N^a = \sum_{c=1}^C \ddot{g}_i^a \odot (1 - \ddot{g}_i^v), N^v = \sum_{c=1}^C \ddot{g}_i^v \odot (1 - \ddot{g}_i^a), N^{av} = \sum_{c=1}^C \ddot{g}_i^a \odot \ddot{g}_i^v, \quad (10)$$

$$\vartheta^a = \frac{N^a}{N^a + N^{av}}, \text{ and } \vartheta^v = \frac{N^v}{N^v + N^{av}}. \quad (11)$$

N^a , N^v , and N^{av} indicate the number of audible-only, visible-only, and audible-visible events in the video, respectively. ϑ^ϕ reflects the degree of encouragement in the contrastive learning and takes a value between 0 and 1 where 0 indicates there is no event or all the events are audible-visible, resulting in no encouragement, and 1 indicates all the events are unaligned, resulting in maximum encouragement.

Self-Modality-Aware Knowledge Distillation – If the Reference branch relies solely on ‘modality-agnostic’ video-level labels for training, we face two key limitations: (i) Reference cannot acquire rich representations necessary to teach Anchor, and (ii) the proposed event-aware contrastive loss may not effectively optimize cross-modal learning since it relies on the predicted event types by Reference. To overcome these *without additional annotations*, we propose to leverage *free modality-aware knowledge* from the Anchor branch to improve the performance of the Reference branch. To perform this, we first distil pseudo labels $G^\phi = \{g_i^\phi\}_{i=1}^C$ from the estimated audio and visual event probabilities in the Anchor branch,

$$g_i^\phi = \begin{cases} 1 & \text{if } \mathcal{P}_i^\phi > \theta \\ 0 & \text{else} \end{cases} \quad (12)$$

Then, G^ϕ are used to compute our proposed self-modality-aware knowledge distillation loss for the Reference branch as $\mathcal{L}_{SelfMo}^{Ref} = \sum_{\phi \in \{a,v\}} BCE(G^\phi, \ddot{\mathcal{P}}^\phi)$.

Co-occurrence Class Knowledge Distillation – The primary goal of modeling the co-occurrence class relationships explicitly in the Reference branch is to transfer them to the Anchor branch so that we can benefit from this knowledge during inference without suffering from the associated computation overhead. Optimizing the Anchor branch through our proposed contrastive learning may implicitly transfer class relationships from the Reference branch to the Anchor branch, but cannot fully transfer them because it is not designed for this purpose. To explicitly transfer the cross-class relationships, we introduce the co-occurrence class knowledge distillation loss, where we first obtain the audio and visual class correlation matrices, $\dot{M}^\phi \in \mathbb{R}^{C \times C}$ and $M^\phi \in \mathbb{R}^{C \times C}$, for the Reference and Anchor branches as

$$\dot{M}_{i,j}^\phi = \ddot{\mathcal{P}}_i^\phi * \ddot{\mathcal{P}}_j^\phi, \text{ and } M_{i,j}^\phi = \mathcal{P}_i^\phi * \mathcal{P}_j^\phi. \quad (13)$$

By computing the correlations, we quantify the relationships between each pair of classes. Then, we encourage the Anchor branch to mimic the audio and visual cross-class correlations of the Reference branch through the mean square error (MSE) as $\mathcal{L}_{COCls}^{Anch} = \sum_{\phi \in \{a,v\}} MSE(\dot{M}^\phi, M^\phi)$.

4 New Metrics for AVVP

Current approaches utilize 5 metrics, under segment-level and event-level, to evaluate the method’s performance for the AVVP task. These metrics include F-score for audible-only (A), visible-only (V), and audible-visible (AV) events separately, as well as the average F-score of the A, V, and AV (Type@AV), and the overall F-score for all A and V events within each sample (Event@AV). To compute these metrics, they first employ a threshold on the estimated audio and visual event probabilities to identify audible and visible events, \hat{y}_t^a and \hat{y}_t^v , respectively. Then, the F-score for A and V are computed by comparing \hat{y}_t^a and \hat{y}_t^v with their corresponding audio and visual ground-truth labels, y_t^a and y_t^v , respectively. The F-score for AV instead is computed by simultaneously considering both audible and visible events and their ground-truth labels.

We argue that the A and V metrics fail to accurately evaluate the method’s performance in detecting audible-only and visible-only events because they consider only their corresponding modalities. In contrast, we believe that both modalities should be considered when computing A and V, similar to AV. For instance, consider a scenario where the network predicts $\hat{y}_t^a = 1$ and $\hat{y}_t^v = 1$ for an audible-only event with ground-truth labels $y_t^a = 1$ and $y_t^v = 0$. According to the A metric, this prediction is considered a true positive. However, when both modalities are considered, it becomes evident that the network’s prediction is incorrect, as it predicted an audible-visible event instead of an audible-only event. Thus, this should be considered a false positive for the audible-only metric. Consequently, a network that excels at detecting audible-visible events may falsely exhibit high true positive rates for the A and V metrics. This issue becomes particularly prominent when using cross-modal context in a weakly supervised scenario, as methods under these settings inherently tend to detect events as

audible-visible. To overcome this, we introduce ‘audible-only’ and ‘visible-only’ event proposals, \hat{y}_t^{ao} and \hat{y}_t^{vo} , as well as ‘audible-only’ and ‘visible-only’ ground-truth labels, y_t^{ao} and y_t^{vo} , that are derived by considering both modalities as

$$\hat{y}_t^{ao} = \hat{y}_t^a \odot (1 - \hat{y}_t^v), \text{ and } \hat{y}_t^{vo} = \hat{y}_t^v \odot (1 - \hat{y}_t^a), \quad (14)$$

$$y_t^{ao} = y_t^a \odot (1 - y_t^v), \text{ and } y_t^{vo} = y_t^v \odot (1 - y_t^a). \quad (15)$$

These new event proposals and their corresponding ground-truth labels are used to compute the F-score for audible-only (Ao) and visible-only (Vo) events. Subsequently, we have Type@AVo, which computes the average F-score of Ao, Vo, and AV, and Event@AVo, which computes the overall F-score for all Ao and Vo events within each sample.

5 Experimental Results

Dataset – LLP [17] is the only benchmark dataset for the AVVP task. It comprises 11,849 10-second YouTube video clips spanning 25 event classes, offering a diverse range of video events. Notably, 7,202 of these videos contain multiple event categories. In the LLP dataset, videos are partitioned into training, validation, and testing sets. While the training set includes only video-level annotations, the validation and test sets feature annotations for individual audio and visual events with second-wise temporal boundaries.

Implementation Details – We implement the Reference branch by employing two self-attention layers, *i.e.*, one for each modality, followed by two fully connected layers that obtain the temporal weights and estimate class probabilities, respectively. For the Anchor branch, we employ CMPAE [3] as the backbone. Following previous works [3, 17], audio input tokens are extracted through pre-trained VGGish [7], and visual tokens are obtained through the pre-trained models ResNet152 [6] and R(2+1)D [19]. The feature dimension for all input tokens, including the class tokens, is set to 512, as in previous approaches. To train CoLeaF, we utilized the Adam optimizer with an initial learning rate of 5×10^{-4} and a batch size of 128 for 15 epochs. The learning rate was decayed by a factor of 0.25 every 6 epochs. Our experiments were performed using PyTorch on an NVIDIA GeForce RTX 3090 GPU.

5.1 Ablation Studies

In this section, we examine the impact of key components of CoLeaF, using the Ao, Vo, and AV metrics on LLP.

Impact of Event-Aware Contrastive Learning – To evaluate the effect of our proposed contrastive loss, \mathcal{L}_{Evt}^{Anch} , we compare the performance of our framework against three variants: one without \mathcal{L}_{Evt}^{Anch} , and two where it is replaced with the standard NCE loss \mathcal{L}_{NCE} and BCE loss \mathcal{L}_{BCE} , the latter using the pseudo-labels employed by \mathcal{L}_{Evt}^{Anch} . Table 2 demonstrates that standard contrastive learning is ineffective for the AVVP task. While \mathcal{L}_{BCE} yields average

Table 2: Ablation studies on \mathcal{L}_{Evt}^{Anch} .

Loss	Segment-level			Event-level		
	Ao	Vo	AV	Ao	Vo	AV
\mathcal{L}_{NCE}	46.1	60.8	56.3	40.5	60.6	50.1
\mathcal{L}_{BCE}	48.5	61.2	57.4	43.5	61.1	51.2
-	47.3	61.8	56.7	41.6	61.7	50.3
\mathcal{L}_{Evt}^{Anch}	49.3	62.4	58.6	44.1	62.2	52.1
	(+2.1)	(+0.6)	(+0.9)	(+2.5)	(+0.5)	(+0.8)

Table 4: Ablation studies on class tokens (Tkn) and $\mathcal{L}_{CoCls}^{Anch}$.

Tkn	$\mathcal{L}_{CoCls}^{Anch}$	Segment-level			Event-level		
		Ao	Vo	AV	Ao	Vo	AV
✓	✗	47.5	60.3	57.0	41.9	60.2	50.4
✗	✓	48.2	62.7	57.3	42.5	62.5	52.0
✗	✗	47.8	60.8	56.6	42.2	60.6	49.8
✓	✓	49.3	62.4	58.6	44.1	62.2	52.1
		(+1.1)	(+1.6)	(+2.0)	(+1.9)	(+1.6)	(+2.3)

improvement, its effectiveness is not consistent across all event types, i.e., it deteriorates performance in detecting the visible-only events in both segment and event levels. Our contrastive loss tailored for this task increases the method’s performance across all event types (see blue numbers in Table 2). Specifically, it leads to a significant improvement (**over 2%** F-score) in detecting audible-only events in both segment and event levels.

Impact of Self-Modality-Aware Knowledge – In Table 3, we ablate our proposed self-modality-aware knowledge distillation loss $\mathcal{L}_{SelfMo}^{Anch}$. The results demonstrate that the additional free modality-aware knowledge obtained by Anchor for Reference during training eventually enhances the performance of Anchor itself across all event types, specifically resulting in **over a 1%** improvement in F-score for detecting unaligned audible-visible event types.

Impact of Learning Co-occurrence Class Relationships – In CoLeaF, the class tokens and $\mathcal{L}_{CoCls}^{Anch}$ are introduced to explicitly model and transfer the cross-class relationships from Reference to Anchor. In Table 4, we conduct ablations on these two components. The results demonstrate that leveraging both components increases the method’s performance **over a 1%** F-score across all event types (see blue numbers in Table 4). Additionally, the results reveal that (i) the network can benefit from $\mathcal{L}_{CoCls}^{Anch}$ with or without the use of the class tokens and (i) adding the class tokens does not improve the performance on average unless the network leverages $\mathcal{L}_{CoCls}^{Anch}$. In these experiments, we also observed that using the class tokens increased the computational complexity of Reference by 41.7% Flops (i.e., from 12.7 to 18.0 GFlops) for 25 classes of LLP. However, our novel design in CoLeaF allows us to benefit from the knowledge learned from them without suffering from their computational overhead in inference. Notably, this novel design can benefit other video understanding tasks where exploiting class dependencies is crucial.

Table 3: Ablation studies on $\mathcal{L}_{SelfMo}^{Ref}$.

$\mathcal{L}_{SelfMo}^{Anch}$	Segment-level			Event-level		
	Ao	Vo	AV	Ao	Vo	AV
✗	48.1	60.9	57.9	42.8	60.6	51.4
✓	49.3	62.4	58.6	44.1	62.2	52.1
	(+1.2)	(+1.5)	(+0.7)	(+1.3)	(+1.6)	(+0.7)

Table 5: Framework generalization on LLP. ⊗ indicates that the network is embedded in CoLeaF as the Anchor branch.

	Segment-level			Event-level		
	Ao	Vo	AV	Ao	Vo	AV
HAN [18]	33.1	50.7	48.9	31.0	50.1	43.0
HAN⊗	36.3	54.1	49.4	34.2	53.5	43.7
	(+3.2)	(+3.4)	(+0.5)	(+3.2)	(+3.4)	(+0.7)
CMPAE [3]	48.2	57.9	57.5	43.6	57.5	49.6
CMPAE⊗	49.3	62.4	58.6	44.1	62.2	52.1
	(+1.1)	(+4.5)	(+1.1)	(+0.5)	(+4.7)	(+2.5)

Framework Generalization – CoLeaF is a general learning framework, which means that any existing AVVP method can be embedded in our framework to (i) explicitly optimize the integration of cross-modal information and (ii) benefit from modeling co-occurrence class relationships in an end-to-end manner. In Table 5, we present two examples: HAN [17], a common architecture for AVVP, and CMPAE [3]. The results demonstrate that our proposed framework significantly enhances the performance of HAN and CMPAE, *i.e.*, on average, the F-score is increased by **2.4%** in both networks. This improvement is achieved in an end-to-end manner without requiring additional annotations or altering their network architecture.

5.2 State-of-the-Art Comparison

In this section, we compare the performance of CoLeaF with current state-of-the-art approaches using both traditional and our proposed metrics. The results of the previous methods under the new metrics were reproduced using their available code whenever accessible. Note that to obtain the event proposals from the estimated output probabilities, CMPAE [3] utilizes a set of ‘selected thresholds’ for different event categories of LLP, while in the standard setting used by other state-of-the-art approaches, a fixed threshold is applied to all event classes. For a fair comparison, we also present the CMPAE’s results under the standard setting.

Table 6 presents comparative results on the LLP dataset when only video-level labels are used for training. The results demonstrate that our proposed approach CoLeaF achieves new state-of-the-art results on all new metrics, as well as AV under both standard and ‘selected thresholds’ settings. CoLeaF exhibits over **1.1%** F-score improvement on most new metrics. Specifically, it significantly increases Vo and Event@Avo by **3.6%** and **2.0%** F-score at the segment level, and Vo, AV, and Type@Avo by **2.2%**, **2.7%**, and **2.4%** F-score at the event level (under the standard setting), and Type@Avo by **2.1%** F-score at the event level (under the ‘selected thresholds’ setting). On average across new metrics and AV, our proposed method improves the state-of-the-art results by 1.9% and 1.2% F-score under standard and selected threshold settings, respectively.

Analysis of Metrics – In Section 4, we stated that the traditional A and V metrics do not adequately reflect the method’s ability to detect audible-only and visible-only events, respectively. As in the A and V metrics, apart from counting audible-only and visible-only events, respectively, they mistakenly include audible-visible events as true positive. This can have a detrimental impact on the analysis of the method’s performance. For example, using the traditional A metric, HAN and MA [21] appear to show nearly identical performance in detecting audible-only events at 60.1% and 60.3% F-score (at segment-level), respectively. However, when we evaluate them using our proposed Ao metric that counts correctly audible-only events, we observe that MA outperforms HAN with a margin of 10.4%. Similarly, based on the V metric (at event-level), CMPAE [3] outperforms JoMoLD [2] in detecting visible-only events, while our proposed

Table 6: Weakly supervised AVVP results on the LLP dataset when utilizing video-level labels for training. All methods use VGGish, ResNet152, and R(2+1)D to generate the input tokens. ‘T@’ and ‘E@’ refers ‘Type@AV’ and ‘Event@AV’, respectively. † indicates the results are under ‘selected thresholds’ setting. The best and the second-best results are in **Bold** and underlined, respectively.

Method	Segment-level (%)										Event-level (%)									
	A	Ao	V	Vo	AV	T@	T@o	E@	E@o	A	Ao	V	Vo	AV	T@	T@o	E@	E@o		
AVE [18]	47.2	–	37.1	–	35.4	39.9	–	41.6	–	40.4	–	34.7	–	31.6	35.5	–	36.5	–		
AVSDN [11]	47.8	–	52.0	–	37.1	45.7	–	50.8	–	34.1	–	46.3	–	26.5	35.6	–	37.7	–		
CVCMS [12]	59.2	–	59.9	–	53.4	57.5	–	58.1	–	51.3	–	55.5	–	46.2	51.0	–	49.7	–		
MM-Pyr [24]	60.9	–	54.4	–	50.0	55.1	–	57.6	–	52.7	–	51.8	–	44.4	49.9	–	50.5	–		
DHNN [8]	61.3	–	58.3	–	52.9	57.5	–	58.1	–	54.0	–	55.1	–	47.3	51.5	–	51.5	–		
MGN [13]	60.8	–	55.4	–	50.0	55.5	–	57.2	–	51.1	–	52.4	–	44.4	49.3	–	49.1	–		
HAN [17]	60.1	33.1	52.9	50.7	48.9	54.0	44.3	55.4	24.1	51.3	31.0	48.9	50.1	43.0	47.7	41.4	48.0	22.0		
MA [21]	60.3	43.5	60.0	53.6	55.1	58.9	50.8	57.9	32.0	53.6	38.8	56.4	52.4	49.0	53.0	46.8	50.6	27.8		
JoMoLD [2]	61.3	46.2	63.8	58.8	57.2	60.8	54.4	59.9	35.4	53.9	40.9	59.9	59.0	49.6	54.5	50.3	52.5	31.1		
CMPAE [3]	63.9	48.2	63.3	57.9	57.5	61.2	54.2	62.1	35.7	57.1	43.6	60.6	57.5	49.6	55.8	50.3	55.6	32.0		
CoLeaF	63.5	49.3	64.5	62.4	58.6	62.1	56.7	61.8	37.7	56.2	44.1	62.0	62.2	52.1	56.7	52.7	54.5	33.4		
	(+1.1)	(+3.6)	(+1.1)	(+1.4)	(+1.4)	(+2.0)	(+2.0)	(+1.4)	(+2.0)	(+0.5)	(+2.2)	(+2.7)	(+2.4)	(+1.4)						
CMPAE† [3]	64.2	48.9	66.4	61.7	59.2	63.3	55.5	62.8	37.6	56.6	43.8	63.7	61.1	51.8	51.9	50.3	55.7	34.0		
CoLeaF†	64.2	49.6	67.1	63.1	59.8	63.8	57.0	61.9	38.4	57.1	45.0	64.8	62.7	52.8	58.2	52.4	55.5	34.6		
	(+0.7)	(+1.4)	(+0.6)	(+1.5)	(+0.8)	(+1.2)	(+1.6)	(+1.0)	(+2.1)	(+0.6)										

Table 7: Weakly supervised AVVP results on the LLP dataset when utilizing video-level labels and CLIP and CLAP annotations for training. All methods use VGGish, ResNet152, and R(2+1)D to generate the input tokens (see Table 6 for more definitions and terms).

Method	Segment-level (%)										Event-level (%)								
	A	Ao	V	Vo	AV	T@	T@o	E@	E@o	A	Ao	V	Vo	AV	T@	T@o	E@	E@o	
JoMoLD [2]	61.5	46.3	<u>64.7</u>	64.8	58.3	61.5	56.5	60.3	38.0	54.7	41.5	61.5	63.5	52.4	56.2	52.5	53.1	33.6	
CMPAE [3]	<u>64.1</u>	48.5	64.3	<u>66.9</u>	<u>58.9</u>	<u>62.4</u>	<u>58.1</u>	<u>62.0</u>	40.0	<u>57.0</u>	42.7	62.1	<u>66.0</u>	<u>53.5</u>	<u>57.6</u>	<u>54.1</u>	55.0	35.1	
VALOR [9]	61.8	<u>49.0</u>	65.9	66.7	58.4	62.0	57.9	61.5	<u>40.8</u>	55.4	<u>44.2</u>	<u>62.6</u>	65.0	52.2	56.7	53.8	<u>54.2</u>	<u>35.8</u>	
CoLeaF	64.2	51.2	64.4	67.7	59.3	62.6	59.4	62.5	41.3	57.6	45.5	63.2	67.0	54.2	57.9	55.6	55.6	36.4	
	(+2.2)			(+0.8)	(+0.4)			(+1.3)		(+0.5)		(+1.3)		(+1.0)	(+0.7)		(+1.5)	(+0.6)	

Vo metric reveals that JoMoLD achieves a better performance than CMPAE. Consequently, the traditional Type@AV and Event@AV have similar issues.

Qualitative Results – In Fig. 3, we qualitatively compare CoLeaF with the state-of-the-art JoMoLD [2] and CMPAE [3] on a test video sample of LLP. For the ‘Speech’ event, while JoMoLD and CMPAE cannot detect it, CoLeaF recognizes it correctly as audible-only. While all methods detect the audible-visible event ‘Singing’, CoLeaF’s event predictions have a better overlap with the ground-truth annotations.

CLIP and CLAP Annotations – Table 6 addresses the weakly supervised AVVP task without the need for additional annotations. The authors in [9] generate offline dense modality-specific annotations for the LLP dataset with large-scale pretrained models, CLIP [15], and CLAP [22]. In Table 7, we present comparative results with recent state-of-the-art approaches in which the networks are additionally trained on these annotations. The results of JoMoLD [2] and CMPAE [3] were obtained by adapting their available code. The results demonstrate that in this scenario with the modality-specific pseudo labels, CoLeaF still improves the state-of-the-art approaches consistently across all new metrics and AV F-score.

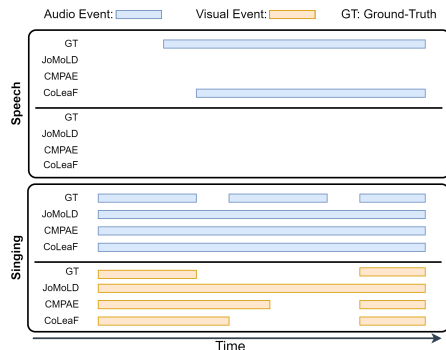


Fig. 3: Qualitative comparison to previous AVVP approaches (JoMoLD [2] and CMPAE [3]) training with weak labels.

Table 8: Weakly supervised DAVE results on the UnAV-100 dataset. All methods use pre-trained VGGish [7] and two-stream I3D [1] to generate the input tokens. The best and the second-best results are in **Bold** and underlined, respectively.

Method	AV (Seg)	AV (Evn)
HAN [17]	35.0	41.4
MA [21]	37.9	<u>44.8</u>
JoMoLD [2]	36.4	41.2
CMPAE [3]	<u>39.7</u>	43.8
CoLeaF	41.5 (+1.8)	47.8 (+3.0)

5.3 Generalization on the Large-Scale UnAV-100 Dataset

To demonstrate the superior generalization of our proposed framework on a different, yet challenging dataset, we present its results to the large-scale UnAV-100 dataset, a benchmark for the Dense Audio-Visual Event Localization (DAVE) task [4]. UnAV-100 includes videos with multiple events and encompasses a diverse range of 100 distinct audio-visible action types, four times more than the LLP dataset. Furthermore, UnAV-100 features untrimmed videos ranging from 0.2 to 60 seconds, significantly increasing the dataset’s complexity. To obtain weakly supervised results, we derived video-level labels from the videos’ temporal annotations (see supp. file for the implementation details for this task). The results of other state-of-the-art approaches were provided through their available code. Comparative results on UnAV-100 using AV F-score are available in Table 8. It is shown that CoLeaF outperforms the other approaches by **1.8%** and **3.0%** AV F-score at the segment and event levels, respectively, and achieves state-of-the-art results on this dataset.

6 Conclusion

We introduce CoLeaF, a novel learning framework designed to: (i) explicitly optimize integrating cross-modal contexts in the embedding space for the weakly supervised AVVP task; and (ii) benefit from explicitly modelling cross-class relationships during training without adding computational overhead at inference. Having introduced novel metrics to better understand an AVVP method’s capability in detecting audible-only and visible-only events, we evaluate CoLeaF on LLP, the only benchmark dataset for the weakly supervised AVVP task, and UnAV-100, a challenging benchmark dataset for the weakly supervised DAVE task. Our method achieves new state-of-the-art results on both datasets. Ablation studies demonstrate the effectiveness of CoLeaF’s key components. Future work will adapt our end-to-end framework to benefit from language models.

Acknowledgement

This research is supported by UKRI EPSRC Platform Grant EP/P022529/1, and EPSRC BBC Prosperity Partnership AI4ME: Future Personalised Object-Based Media Experiences Delivered at Scale Anywhere EP/V038087/1.

References

1. Carreira, J., Zisserman, A.: Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6299–6308 (2017)
2. Cheng, H., Liu, Z., Zhou, H., Qian, C., Wu, W., Wang, L.: Joint-Modal Label Denoising for Weakly-Supervised Audio-Visual Video Parsing. In: *European Conference on Computer Vision*. pp. 431–448. Springer (2022)
3. Gao, J., Chen, M., Xu, C.: Collecting Cross-Modal Presence-Absence Evidence for Weakly-Supervised Audio-Visual Event Perception. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 18827–18836 (2023)
4. Geng, T., Wang, T., Duan, J., Cong, R., Zheng, F.: Dense-Localizing Audio-Visual Events in Untrimmed Videos: A Large-Scale Benchmark and Baseline. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 22942–22951 (2023)
5. Gutmann, M., Hyvärinen, A.: Noise-Contrastive Estimation: A New Estimation Principle for Unnormalized Statistical Models. In: *Proceedings of the 13th International Conference on Artificial Intelligence and Statistics*. pp. 297–304. JMLR Workshop and Conference Proceedings (2010)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep Residual Learning for Image Recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 770–778 (2016)
7. Hershey, S., Chaudhuri, S., Ellis, D.P., Gemmeke, J.F., Jansen, A., Moore, R.C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., et al.: CNN Architectures for Large-Scale Audio Classification. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 131–135. IEEE (2017)
8. Jiang, X., Xu, X., Chen, Z., Zhang, J., Song, J., Shen, F., Lu, H., Shen, H.T.: DHHN: Dual Hierarchical Hybrid Network for Weakly-Supervised Audio-Visual Video Parsing. In: *Proceedings of the 30th ACM International Conference on Multimedia*. pp. 719–727 (2022)
9. Lai, Y.H., Chen, Y.C., Wang, Y.C.F.: Modality-Independent Teachers Meet Weakly-Supervised Audio-Visual Event Parser. *Advances in Neural Information Processing Systems* (2023)
10. Li, G., Wei, Y., Tian, Y., Xu, C., Wen, J.R., Hu, D.: Learning to Answer Questions in Dynamic Audio-Visual Scenarios. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19108–19118 (2022)
11. Lin, Y.B., Li, Y.J., Wang, Y.C.F.: Dual-Modality Seq2Seq Network for Audio-Visual Event Localization. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 2002–2006. IEEE (2019)
12. Lin, Y.B., Tseng, H.Y., Lee, H.Y., Lin, Y.Y., Yang, M.H.: Exploring Cross-Video and Cross-Modality Signals for Weakly-Supervised Audio-Visual Video Parsing. *Advances in Neural Information Processing Systems* **34**, 11449–11461 (2021)

13. Mo, S., Tian, Y.: Multi-Modal Grouping Network for Weakly-Supervised Audio-Visual Video Parsing. *Advances in Neural Information Processing Systems* **35**, 34722–34733 (2022)
14. Nadeem, A., Hilton, A., Dawes, R., Thomas, G., Mustafa, A.: CAD-Contextual Multi-modal Alignment for Dynamic AVQA. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (2023)
15. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Quantify Classification Models from Natural Language Supervision. In: *International Conference on Machine Learning*. pp. 8748–8763 (2021)
16. Sensoy, M., Kaplan, L., Kandemir, M.: Evidential Deep Learning to Quantify Classification Uncertainty. *Advances in Neural Information Processing Systems* **31** (2018)
17. Tian, Y., Li, D., Xu, C.: Unified Multisensory Perception: Weakly-Supervised Audio-Visual Video Parsing. In: *European Conference in Computer Vision*. pp. 436–454. Springer (2020)
18. Tian, Y., Shi, J., Li, B., Duan, Z., Xu, C.: Audio-Visual Event Localization In Unconstrained Videos. In: *European Conference on Computer Vision*. pp. 247–263 (2018)
19. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A Closer Look at Spatiotemporal Convolutions for Action Recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 6450–6459 (2018)
20. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention Is All You Need. *Advances in Neural Information Processing systems* **30** (2017)
21. Wu, Y., Yang, Y.: Exploring Heterogeneous Clues for Weakly-Supervised Audio-Visual Video Parsing. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1326–1335 (2021)
22. Wu, Y., Chen, K., Zhang, T., Hui, Y., Berg-Kirkpatrick, T., Dubnov, S.: Large-Scale Contrastive Language-Audio Pretraining with Feature Fusion and Keyword-to-Caption Augmentation. In: *IEEE International Conference on Acoustics, Speech and Signal Processing*. pp. 1–5 (2023)
23. Xia, Y., Zhao, Z.: Cross-Modal Background Suppression for Audio-Visual Event Localization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 19989–19998 (June 2022)
24. Yu, J., Cheng, Y., Zhao, R.W., Feng, R., Zhang, Y.: Mm-Pyramid: Multimodal Pyramid Attentional Network for Audio-Visual Event Localization and Video Parsing. In: *Proceedings of the 30th ACM International Conference on Multimedia*. pp. 6241–6249 (2022)
25. Yun, H., Yu, Y., Yang, W., Lee, K., Kim, G.: Pano-AVQA: Grounded Audio-Visual Question Answering on 360deg Videos. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2031–2041 (October 2021)
26. Zhou, J., Guo, D., Zhong, Y., Wang, M.: Improving Audio-Visual Video Parsing with Pseudo Visual Labels. *arXiv preprint arXiv:2303.02344* (2023)