







# Label-anticipated Event Disentanglement for Audio-Visual Video Parsing

Jinxing Zhou<sup>1</sup>,  Dan Guo<sup>1,2,3</sup>, Yuxin Mao<sup>4</sup>, Yiran Zhong<sup>5</sup>,  
Xiaojun Chang<sup>6,7</sup>, and  Meng Wang<sup>1,3</sup>

<sup>1</sup>Hefei University of Technology, <sup>2</sup>Anhui Zhonghuitong Technology Co., Ltd.,  
<sup>3</sup>Institute of Artificial Intelligence, Hefei Comprehensive National Science Center,  
<sup>4</sup>Northwestern Polytechnical University, <sup>5</sup>Shanghai AI Laboratory,  
<sup>6</sup>University of Science and Technology of China, <sup>7</sup>MBZUAI


**Abstract.** Audio-Visual Video Parsing (AVVP) task aims to detect and temporally locate events within audio and visual modalities. Multiple events can overlap in the timeline, making identification challenging. While traditional methods usually focus on improving the early audio-visual encoders to embed more effective features, the decoding phase – crucial for final event classification, often receives less attention. We aim to advance the decoding phase and improve its interpretability. Specifically, we introduce a new decoding paradigm, label semantic-based projection (LEAP), that employs labels texts of event categories, each bearing distinct and explicit semantics, for parsing potentially overlapping events. LEAP works by iteratively projecting encoded latent features of audio/visual segments onto semantically independent label embeddings. This process, enriched by modeling cross-modal (audio/visual-label) interactions, gradually disentangles event semantics within video segments to refine relevant label embeddings, guaranteeing a more discriminative and interpretable decoding process. To facilitate the LEAP paradigm, we propose a semantic-aware optimization strategy, which includes a novel audio-visual semantic similarity loss function. This function leverages the Intersection over Union of audio and visual events (EIoU) as a novel metric to calibrate audio-visual similarities at the feature level, accommodating the varied event densities across modalities. Extensive experiments demonstrate the superiority of our method, achieving new state-of-the-art performance for AVVP and also enhancing the relevant audio-visual event localization task.

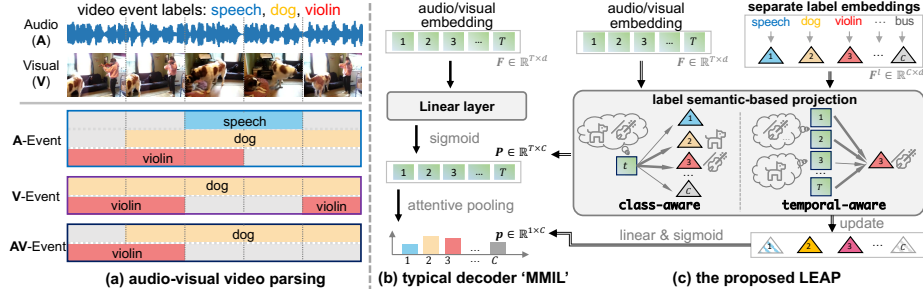
**Keywords:** Audio-visual video parsing · Event disentanglement · Audio-visual event localization

## 1 Introduction

Human perception involves the remarkable ability to discern various types of events in real life through their intelligent auditory and visual sensors [7, 26]. We

---

 Corresponding authors ({guodan, wangmeng}@hfut.edu.cn)



**Fig.1: Illustration of the AVVP task and different event decoding paradigms.** (a) The AVVP task is required to parse audio events, visual events, and audio-visual events within the video. Each segment may contain multiple overlapping events. Given the latent audio/visual embeddings, (b) the typical decoding paradigm ‘MMIL’ directly predicts multiple event classes by using simple linear layers. (c) We propose to elucidate the disentanglement of the potentially overlapping semantics through the projection of latent features into multiple, semantically separate label embeddings.

can even recognize multiple events simultaneously when they occur at the same time. For instance, we can witness one musician playing the guitar and another playing the piano at a concert (visual events), or we can hear the sounds of a TV show and a baby crying (audio events). Audio-Visual Video Parsing (AVVP) [23] task aims to identify all the events in the respective audio and visual modalities and localize the temporal boundaries of each event. To avoid extensive annotation cost, the pioneer work [23] performs this task under a weakly supervised setting where only the event label of the entire video is provided for model training. As shown in Fig. 1(a), we only know that this video contains events of *speech*, *dog*, and *violin*, and the AVVP task requires temporally parsing the audio events, visual events, and audio-visual events (both audible and visible). Moreover, multiple events may occur in the same segment, *i.e.*, *overlapping events* in the timeline, adding challenges for accurate event parsing.

To tackle this task, the majority of previous works [1, 5, 11, 20, 30, 32] try to develop more robust audio-visual *encoders* for embedding more effective audio-visual features, thus facilitating late event decoding. Meanwhile, to ease this weakly supervised task, some works attempt to provide additional supervision by generating audio and visual pseudo labels at either the video-level [2, 27] or segment-level [31, 34, 35]. While these efforts have achieved significant improvements, they typically employ a conventional event *decoding* paradigm – Multi-modal Multi-Instance Learning (MMIL) [23] strategy. As illustrated in Fig. 1(b), the encoded audio/visual embeddings are simply processed through linear layers, which directly transform the features from the latent space into the event category space. Then, the transformed logits are activated using the sigmoid function to obtain the segment-level event probabilities, which are then attentively averaged over timeline to predict video-level events. The MMIL successfully achieves event prediction through simple linear functions, yet it is not very intuitive in

demonstrating *how the semantics of potentially overlapped events are decoded from the latent features*. To approach this goal, we seek to improve the event decoding phase by exploring a more explicit category semantic-guided paradigm.

Inspired by that the natural language can convey specific and independent semantics, we try to utilize explicit label texts of all the event classes in the event decoding stage. Specifically, we propose a **label semantic-based projection (LEAP) strategy**, which iteratively projects the encoded audio and visual features into semantically separate label embeddings. The projection is realized by modeling the cross-modal relations between audio/visual segments and event texts using a straightforward Transformer architecture. This enables each audio/visual segment to clearly perceive and interact with distinct label embeddings. As shown in Fig. 1(c), if one segment contains overlapping events, then multiple separate label embeddings corresponding to the events are enhanced through the higher cross-modal attention weights (*class-aware*), indicated by thicker arrows in the figure. In other words, the semantics mixed within hidden features are clearly separated or disentangled into multiple independent label embeddings, which makes our event decoding process more interpretable and traceable. The intermediate cross-modal attention matrix reflecting the similarity between audio/visual with label texts can be used to generate segment-level event predictions. Afterwards, each label embedding is refined by aggregating matched event semantics from all the relevant temporal segments (*temporal-aware*). Those label embeddings of events that actually occur in the video are enhanced to be more discriminative. The updated label embeddings can be utilized for video-level event predictions.

To facilitate the above LEAP process, we explore a **semantic-aware optimization strategy**. The video-level weak label and segment-level pseudo labels [31] are used as the basic supervision to regularize predictions. Moreover, we propose a **novel audio-visual semantic similarity loss function**  $\mathcal{L}_{avss}$  to further enhance the audio-visual representation learning. Given that each audio/visual segment may contain multiple events, we propose the use of *Intersection over Union of audio Events and visual events* as a metric (abbreviated as EIoU) to assess cross-modal semantic similarity. The more identical events the audio and visual modalities contain, the higher the EIoU will be. Then  $\mathcal{L}_{avss}$  computes the EIoU matrix for all audio-visual segment pairs and employs it to regulate the similarity between the early encoded audio and visual features.

In summary, the main contributions of this paper are:

- We propose a label semantic-based projection (LEAP) method as a new event decoding paradigm for the AVVP task. Our LEAP utilizes semantically independent label embeddings to disentangle potentially overlapping events.
- We develop a semantic-aware optimization strategy that considers both uni-modal and cross-modal regularizations. Particularly, the EIoU metric is introduced to design a novel audio-visual semantic similarity loss function.
- Extensive experiments confirm the superiority of our LEAP method compared to the typical paradigm MMIL in parsing events across different modalities and in handling overlapping cases.

- Our method is compatible with existing AVVP backbones and achieves new state-of-the-art performances. Besides, the proposed LEAP is beneficial for the related AVEL [24] task, demonstrating its generalization capability.

## 2 Related Work

**Audio-Visual Learning** focuses on exploring the relationships between audio and visual modalities to achieve effective audio-visual representation learning and understanding of audio-visual scenarios. Over the years, various research tasks have been proposed and investigated [26], such as the sound source localization [10, 17, 36, 37], audio-visual event localization [24, 29, 33, 38], audio-visual question answering and captioning [14, 15, 22]. While a range of sophisticated networks have been proposed for solving these tasks, most of them emphasize establishing correspondences between audio and visual signals. However, audio-visual signals are not always spatially or temporally aligned. As exemplified by the studied audio-visual video parsing task, the events contained in a video may be modality-independent and temporally independent. Consequently, it is essential to explore the semantics of events within each modality.

**Audio-Visual Video Parsing** aims to recognize the event categories and their temporal locations for both audio and visual modalities. The pioneering work [23] performs this task in a weakly supervised setting and frames it as a Multi-modal Multi-Instance Learning (MMIL) problem, demanding the model to be modality-aware and temporal-aware. To tackle this challenging task, subsequent works primarily focus on designing more effective audio-visual encoders [1, 18, 30, 32]. For instance, MM-Pyr [30] utilizes a pyramid unit to constrain the unimodal and cross-modal interactions to occur in adjacent segments, improving the temporal localization. Additionally, some approaches try to generate pseudo labels for audio and visual modalities from the video level [2, 27] and segment level [31, 35]. However, prior works [5, 20, 28, 30, 31] mainly adopt the typical strategy MMIL proposed in [23] as the decoder for final event prediction. The MMIL approach directly regresses multiple classes based on the semantic-mixed hidden feature. In contrast, we further introduce the textual modality as an intermediary and disentangle the semantics of potentially overlapping events contained in audio/visual features by projecting them into semantically separate label embeddings.

## 3 Audio-Visual Video Parsing Approach

### 3.1 Task Definition

The AVVP task aims to recognize and temporally localize all types of events that occur within an audible video. Those events encompass audio events, visual events, and audio-visual events. Specifically, an audible video is divided into  $T$  temporal segments, each spanning one second. The audio and visual streams at  $t$ -th segment are denoted as  $X_t^a$  and  $X_t^v$ , respectively. A video parsing model needs to classify each audio/visual segment  $X_t^m$  ( $m \in \{a, v\}, t = 1, \dots, T$ ) into

predefined  $C$  event categories, **being aware of the events from the perspectives of class, modality, and temporal timeline.**

The AVVP task, initially introduced in [23], is conducted under a weakly-supervised setting, where only the event label for the entire video is provided for model training, denoted as  $\mathbf{y}^{a||v} \in \mathbb{R}^{1 \times C}$ . Here,  $\mathbf{y}_c^{a||v} \in \{0, 1\}$ , with ‘1’ indicating the presence of an event in the  $c$ -th category in the video. However, it does not specify which modality (audio or visual) or which temporal segments contain events of this category. The most recent advance of the field [31] has introduced more explicit supervision by generating high-quality segment-level audio and visual pseudo labels, denoted as  $\{\mathbf{Y}^a, \mathbf{Y}^v\} \in \mathbb{R}^{T \times C}$ . It is important to note that  $\sum \mathbf{Y}_{t,:}^m \geq 0$  ( $m \in \{a, v\}$ ), indicating that each audio/visual segment may carry overlapping events of multiple classes, potentially occurring simultaneously.

### 3.2 Typical Event Decoding Paradigm – MMIL

As introduced in Sec. 1, prior works [5, 11, 20, 30] usually rely on the Multi-modal Multi-Instance Learning (MMIL) [23] strategy as the late decoder used for final event prediction. We briefly outline the main steps of MMIL.

First, an audio-visual encoder  $\Phi$  is employed to obtain audio and visual features:  $\mathbf{F}^a, \mathbf{F}^v = \Phi(X^a, X^v)$ , where  $\mathbf{F} \in \mathbb{R}^{T \times d}$  and  $d$  is the feature dimension. Then, a linear layer is used to transform the obtained features, and the sigmoid activation is directly used to generate the segment-wise event probabilities:

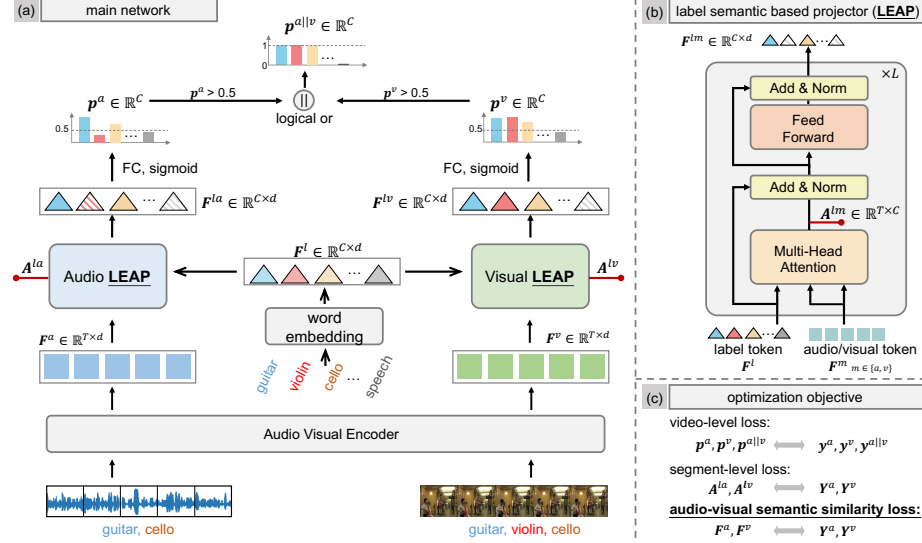
$$\begin{cases} \mathbf{P}^a = \text{sigmoid}(\mathbf{F}^a \mathbf{W}^a), \\ \mathbf{P}^v = \text{sigmoid}(\mathbf{F}^v \mathbf{W}^v), \end{cases} \quad (1)$$

where  $\mathbf{W}^a, \mathbf{W}^v \in \mathbb{R}^{d \times C}$  are learnable parameters and  $\mathbf{P}^a, \mathbf{P}^v \in \mathbb{R}^{T \times C}$ . To learn from the weak video label  $\mathbf{y}^{a||v}$ , the video-level event probability  $\mathbf{p}^{a||v} \in \mathbb{R}^{1 \times C}$  is obtained by an attentive pooling operation, which produces attention weights for both modality and temporal segments.

Therefore, the MMIL primarily relies on simple linear transformations of audio/visual features to directly classify the multiple event classes. However, this mechanism lacks clarity in demonstrating how potentially overlapped events are disentangled from the semantically mixed hidden features. To enhance the decoding stage, we introduce all  $C$ -class label embeddings, each representing separate event semantics, and iteratively project encoded audio/visual features onto them. Through the projection process, the overlapping semantics in the hidden features are gradually disentangled to improve the distinctiveness of the corresponding label embeddings, thereby enhancing the interpretability of our event decoding process. We elaborate on our method in the next subsections.

### 3.3 Our Label Semantic-based Projection

As shown in Fig. 2(a), we propose the label semantic-based projection (LEAP) to improve the *decoder* for final event parsing, serving as a new decoding paradigm.



**Fig. 2: Overview of our method.** (a) Our network for audio-visual video parsing. We focus on enhancing the later decoder with the proposed label semantic-based projection (LEAP) strategy. Specifically, we explicitly introduce the separate label embeddings of all event classes and then disentangle potentially overlapping events by projecting the audio or visual features into those label embeddings. (b) The illustration of LEAP. LEAP models the cross-modal relations between audio/visual with label embeddings. The label embeddings corresponding to the ground truth events are enhanced to be discriminative. The intermediate cross-attention matrix  $A^{lm}$  and the final enhanced label embedding  $F^{lm}$  is used for segment-level and video-level event predictions, respectively. (c) For effective projection and model optimization, we consider the supervision from uni-modal labels at both the video level and segment level ( $\mathcal{L}_{basic}$ ). We also design a new audio-visual semantic similarity loss function  $\mathcal{L}_{avss}$  to regularize the model by considering cross-modal relations at the feature level.

For the audio-visual *encoder*, prior typical backbones, such as HAN [23] and MM-Pyr [30], can be used to obtain the intermediate audio and visual features, denoted as  $\{F^a, F^v\} \in \mathbb{R}^{T \times d}$ . Then, we begin to establish the foundation for our LEAP method by acquiring the independent label embeddings. Given texts of all  $C$  event classes, *e.g.*, *dog* and *guitar*, we obtain their label embeddings using the pretrained Glove [19] model. The resulting label embeddings are then combined into one label-semantic matrix, denoted as  $F^l \in \mathbb{R}^{C \times d}$ .

The essence of our LEAP lies in discerning semantics within audio and visual latent features by projecting them into separate label embeddings. We achieve this goal by modeling the cross-modal (audio/visual-label) interactions using a Transformer block. As illustrated in Fig. 2(b), the label embeddings are used as *query*, and the audio/visual features serve as the *key* and *value*, formulated as,

$$\mathcal{Q}^{lm} = F^l W_Q^m, \mathcal{K}^m = F^m W_K^m, \mathcal{V}^m = F^m W_V^m, \quad (2)$$

where  $m \in \{a, v\}$  denotes the audio and visual modalities.  $\{\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V\} \in \mathbb{R}^{d \times d}$  are learnable parameters.  $\mathbf{Q}^{lm} \in \mathbb{R}^{C \times d}$  and  $\{\mathbf{K}^m, \mathbf{V}^m\} \in \mathbb{R}^{T \times d}$ . Then, the cross-modal (audio/visual-label) attention  $\mathbf{A}^{lm}$  can be obtained by computing the scaled-dot production. Based on  $\mathbf{A}^{lm}$ , the initial label embeddings are enriched by aggregating related semantics from the audio/visual temporal segments. A feed-forward network is finally used to update the label embeddings. This process can be formulated as,

$$\begin{cases} \mathbf{A}^{lm} = \text{softmax}(\frac{\mathbf{Q}^{lm} \mathbf{K}^m}{\sqrt{d}}), \\ \tilde{\mathbf{F}}^{lm} = \mathbf{F}^l + \text{LN}(\mathbf{A}^{lm} \mathbf{V}^m), \\ \mathbf{F}^{lm} = \tilde{\mathbf{F}}^{lm} + \text{LN}(\text{FF}(\tilde{\mathbf{F}}^{lm})), \end{cases} \quad (3)$$

where ‘LN’ represents the layer normalization, and ‘FF’ denotes the feed-forward network mainly implemented using two linear layers. The outcome of the LEAP block is the cross-modal attention  $\mathbf{A}^{lm} \in \mathbb{R}^{C \times T}$  and updated label embedding  $\mathbf{F}^{lm} \in \mathbb{R}^{C \times d}$ . We summarize the above process as,

$$\mathbf{F}^{lm}, \mathbf{A}^{lm} = \text{LEAP}(\mathbf{F}^l, \mathbf{F}^m). \quad (4)$$

The LEAP block can be repeated iteratively. For the  $i$ -th iteration, the encoded audio/visual feature  $\mathbf{F}^m$  is repeatedly used to enhance the semantic-relevant label embeddings:

$$\mathbf{F}_i^{lm}, \mathbf{A}_i^{lm} = \text{LEAP}(\mathbf{F}_{i-1}^{lm}, \mathbf{F}^m), \quad (5)$$

where  $i = 1, \dots, N$  ( $N$  is the maximum iteration number) and  $\mathbf{F}_0^{lm} = \mathbf{F}^l$ .

It is worth noting that  $\mathbf{A}_i^{lm} \in \mathbb{R}^{C \times T}$  can act as an indicator of the similarity between each event class and every audio/visual segment. When an audio or visual segment (**modality-aware**) contains multiple overlapping events, the classes associated with those events occurring in the segment receive higher similarity scores compared to other classes (**class-aware**). Subsequently, the label embedding of each class traverses all the temporal segments and assimilates relevant semantic information from timestamps with high similarity scores for that class (**temporal-aware**). This mechanism effectively disentangles potential overlapping semantics, reinforcing label embeddings for classes present in the audio/visual segments.

### 3.4 Audio-Visual Semantic-aware Optimization

The cross-modal attention at the last LEAP cycle, *i.e.*,  $\mathbf{A}_N^{lm} \in \mathbb{R}^{C \times T}$ , indicates the similarity between all  $C$ -class label embeddings and all audio/visual segments. Therefore, we directly use  $\mathbf{A}_N^{lm}$  to generate *segment-level* event probabilities, written as,

$$\mathbf{P}^m = \text{sigmoid}((\mathbf{A}_N^{lm})^\top), \quad (6)$$

where  $\mathbf{P}^m = \{\mathbf{P}^a, \mathbf{P}^v\} \in \mathbb{R}^{T \times C}$ . Note that  $\mathbf{A}_N^{lm}$  in Eq. 6 is the raw attention logits without softmax operation. For *video-level* event prediction  $\mathbf{p}^m$ , it can

be produced by the obtained label embedding after LEAP, *i.e.*,  $\mathbf{F}_N^{lm}$ , since it indicates which event classes are finally enhanced:

$$\mathbf{p}^m = \text{sigmoid}(\mathbf{W}(\mathbf{F}_N^{lm})^\top), \quad (7)$$

where  $\mathbf{W} \in \mathbb{R}^{1 \times d}$  and  $\mathbf{p}^m = \{\mathbf{p}^a, \mathbf{p}^v\} \in \mathbb{R}^{1 \times C}$ . We use a threshold of 0.5 to identify events that happen in audio and visual modalities, then the event prediction of the *entire video*  $\mathbf{p}^{a||v}$  can be computed as follows,

$$\mathbf{p}^{a||v} = \mathbf{1}(\mathbf{p}^a \geq 0.5) \parallel \mathbf{1}(\mathbf{p}^v \geq 0.5), \quad (8)$$

where  $\mathbf{1}(\mathbf{z})$  is a boolean function that outputs ‘1’ when the  $\mathbf{z}_i \geq 0$ . ‘ $\parallel$ ’ is the *logical OR* operation, which computes the union of audio events and visual events.

For effective projection and better model optimization, we incorporate the segment-wise pseudo labels  $\mathbf{Y}^m \in \mathbb{R}^{T \times C}$  ( $m \in \{a, v\}$ ) generated in recent work [31] to provide fine-grained supervision. The video-level pseudo labels  $\mathbf{y}^m \in \mathbb{R}^{1 \times C}$  can also be easily obtained from  $\mathbf{Y}^m$ . The basic objective  $\mathcal{L}_{basic}$  constrains the audio and visual event predictions from both video-level and segment-level, computed by,

$$\mathcal{L}_{basic} = \sum_m \mathcal{L}_{bce}(\mathbf{p}^{a||v}, \mathbf{y}^{a||v}) + \mathcal{L}_{bce}(\mathbf{p}^m, \mathbf{y}^m) + \mathcal{L}_{bce}(\mathbf{P}^m, \mathbf{Y}^m), \quad (9)$$

where  $\mathcal{L}_{bce}$  is the binary cross entropy loss,  $m \in \{a, v\}$  denotes the modalities.

$\mathcal{L}_{basic}$  directly acts on final event predictions and constrains the audio/visual semantic learning through *uni-modal* event labels. In addition, we further propose a novel **audio-visual semantic similarity loss function** to explicitly explore the *cross-modal* relations, which provides extra regularization on audiovisual representation learning. We are motivated by the observation that the audio and the visual segments often contain different numbers of events. A video example has been shown in Fig. 1(a). An AVVP model should be aware of the semantic relevance and difference between audio events and visual events to achieve a better understanding of events contained in the video.

To quantify the cross-modal semantic similarity, we introduce the **Intersection over Union of audio Events and visual events** (EIoU, symbolized by  $r$ ). EIoU is computed for each audio-visual segment pair, illustrating the degree of overlap between their respective event classes. For instance, consider an audio segment  $a_1$  containing three events with classes  $\{c_1, c_2, c_3\}$ , a visual segment  $v_1$  with events of classes  $\{c_1\}$ , and another visual segment  $v_2$  with events  $\{c_1, c_2\}$ . In this scenario, the union event sets for these two audio-visual segment pairs are identical, consisting of  $\{c_1, c_2, c_3\}$ . However, the intersection event sets differ: for  $a_1$  and  $v_1$ , the intersection set is  $\{c_1\}$ , whereas for  $a_1$  and  $v_2$ , it is  $\{c_1, c_2\}$ . By calculating the ratio of the intersection set size to the union set size, we can obtain the EIoU values for these two audio-visual pairs, *i.e.*,  $r_{11} = 1/3$  and  $r_{12} = 2/3$ . This calculation extends to all combinations of  $T$  audio segments and  $T$  visual segments, resulting in the EIoU matrix  $\mathbf{r} \in \mathbb{R}^{T \times T}$ . Each entry  $r_{ij}$  in this matrix quantifies the semantic similarity between the  $i$ -th audio segment



and  $j$ -th visual segment. Notably, when two segments share precisely the same events,  $\mathbf{r}_{ij}$  equals 1. Conversely, if they contain entirely dissimilar events,  $\mathbf{r}_{ij}$  equals 0. Therefore,  $\mathbf{r}$  serves as an effective measure to assess the semantic similarity between audio and visual segments, particularly when segments contain multiple overlapping events.

Given the encoded audio and visual features  $\{\mathbf{F}^a, \mathbf{F}^v\} \in \mathbb{R}^{T \times d}$ , we compute the cosine similarity of all audio-visual segment pairs, denoted as  $\mathbf{s}$ , as below,

$$\mathbf{s} = \frac{\mathbf{F}^a}{\|\mathbf{F}^a\|_2} \otimes \left( \frac{\mathbf{F}^v}{\|\mathbf{F}^v\|_2} \right)^\top, \quad (10)$$

where  $\mathbf{s} \in \mathbb{R}^{T \times T}$ ,  $\otimes$  denotes the matrix multiplication operation. Then, the audio-visual semantic similarity loss  $\mathcal{L}_{avss}$  measures the discrepancy between the feature similarity matrix  $\mathbf{s}$  and the EIoU matrix  $\mathbf{r}$ , formulated as,

$$\mathcal{L}_{avss} = \mathcal{L}_{mse}(\mathbf{s}, \mathbf{r}), \quad (11)$$

where  $\mathcal{L}_{mse}$  denotes the mean squared error loss.

The overall semantic-aware objective  $\mathcal{L}$  is the combination of  $\mathcal{L}_{avss}$  and the basic loss  $\mathcal{L}_{basic}$ , computed by,

$$\mathcal{L} = \mathcal{L}_{basic} + \lambda \mathcal{L}_{avss}, \quad (12)$$

where  $\lambda$  is a hyperparameter to balance the two loss items. In this way, our LEAP model is optimized to be aware of the event semantics not only through *uni-modal* (audio or visual) label supervision but also by considering the *cross-modal* (audio-visual) semantic similarity.

## 4 Experiments

### 4.1 Experimental Setups

**Dataset.** We conduct experiments of AVVP on the widely used *Look, Listen, and Parse* (LLP) [23] dataset which comprises 11,849 YouTube videos across 25 categories, including audio/visual events related to everyday human and animal activities, vehicles, musical performances, *etc.* Following the standard dataset split [23], we use 10,000 videos for training, 648 for validation, and 1,200 for testing. The LLP dataset exclusively provides weak video event labels for the training set. We employ the strategy proposed in [31] to derive segment-wise audio and visual pseudo labels. For the validation and test sets, the segment-level labels are already available for model evaluation.

**Evaluation metrics.** Following prior works [2, 5, 23, 27], we evaluate the model performance by using the F1-scores on all types of event parsing results, including the audio event (**A**), visual event (**V**), and audio-visual event (**AV**). For each event type, the F1-score is computed at both the segment level and event level. For the former, the event prediction and the ground truth are segment-wisely compared. As for the event-level metric, the consecutive segments in the same

event are regarded as one entire event. Then, the F1-score is computed with  $mIoU = 0.5$  as the threshold. In addition, two metrics are used to evaluate the overall audio-visual video parsing performance: “**Type@AV**” computes the average F1-scores of audio, visual, and audio-visual event parsing; “**Event@AV**” calculates the F1-score considering all audio and visual events in each video together.

**Implementation details.** We adopt the same backbones for feature extraction as previous works [2, 23, 27]. Specifically, we downsample video frames at 8 FPS and use the ResNet-152 [8] pretrained on ImageNet [3] and the R(2+1)D network [25] pretrained on Kinetics-400 [12] to extract 2D and 3D visual features, respectively. The concatenation of these two features is used as the initial visual feature. The audio waveform is subsampled at 16 KHz and we use the VGGish [9] pretrained on AudioSet [6] to extract the 128-D audio features. The loss balancing hyperparameter  $\lambda$  in Eq. 12 is empirically set to 1. We train our model for 20 epochs using the Adam [13] optimizer with a learning rate of  $1e-4$  and a mini-batch of 32.

## 4.2 Ablation Study

**Ablation studies of LEAP.** We begin by investigating the impacts of **1)** the maximum number of LEAP blocks ( $N$  in Eq. 5) and **2)** different label embedding generation strategies. In this part, we use the MM-Pyr [30] as the early audio-visual encoder of our LEAP-based method. **1)** As shown in the upper part of Table 1, the average video parsing performance increases with the number of LEAP blocks. The highest performance is 61.4%, achieved when using four LEAP blocks. This is slightly better than using two LEAP blocks but also doubles the computation cost in projection. Considering the trade-off of performance and computation cost, we finally utilize *two* LEAP blocks for constructing AVVP models. **2)** We test three commonly used word embedding strategies, *i.e.*, the Glove [19], Bert [4], and CLIP [21]. As shown in the lower part of Table 1, our LEAP method exhibits robustness to these three types of label embedding generation strategies. The highest average parsing performance is achieved with the Glove embedding. Therefore, we employ the pretrained *Glove* model to generate the label embeddings for our approach.

**Table 1: Ablation results of the LEAP block.** We explore the impacts of the maximum number  $N$  of LEAP blocks and different Label Embedding Generation strategies (LEG). “*Avg.*” is the average result of all ten metrics.

Setups		Segment-level					Event-level					Avg.
$N$	LEG	A	V	AV	Type@AV	Event@AV	A	V	AV	Type@AV	Event@AV	
1		63.6	<b>67.2</b>	60.6	63.8	62.8	57.5	64.6	55.1	59.1	56.0	61.0
2	Glove [19]	63.7	67.0	<b>61.3</b>	<b>64.0</b>	<b>62.8</b>	58.2	63.9	<b>56.2</b>	59.5	56.6	<u>61.3</u>
4		<b>63.8</b>	67.1	60.8	63.9	62.8	<b>58.4</b>	<b>64.7</b>	55.8	<b>59.7</b>	<b>56.7</b>	<b>61.4</b>
2	Glove [19]	63.7	<b>67.0</b>	<b>61.3</b>	<b>64.0</b>	<b>62.8</b>	<b>58.2</b>	<b>63.9</b>	<b>56.2</b>	<b>59.5</b>	<b>56.6</b>	<b>61.3</b>
	Bert [4]	63.4	66.7	60.2	63.4	62.7	58.1	63.5	55.5	59.0	56.2	60.9
	CLIP [21]	<b>64.4</b>	66.6	60.3	63.8	63.5	58.1	63.7	54.9	58.9	56.4	61.1

**Table 2: Effectiveness of the proposed LEAP and the loss function  $\mathcal{L}_{avss}$ .** We compare our LEAP with the typical video decoder – MMIL [23] by equipping them with two representative audio-visual encoders, *i.e.*, HAN [23] and MM-Pyr [30].

Methods		Objective		Segment-level						Event-level					
Encoders	Decoders	$\mathcal{L}_{basic}$	$\mathcal{L}_{avss}$	A	V	AV	Type@AV	Event@AV	A	V	AV	Type@AV	Event@AV	A	V
HAN	MMIL [23]	✓	✗	61.5	65.5	58.8	61.9	60.6	55.2	61.7	52.3	56.4	53.5		
	LEAP (ours)	✓	✗	62.1	65.2	58.9	62.1	61.1	56.3	62.7	54.0	57.7	54.7		
		✓	✓	62.7	65.6	59.3	62.5	61.8	56.4	63.1	54.1	57.8	55.0		
MM-Pyr	MMIL [23]	✓	✗	61.0	66.3	59.3	62.2	60.6	54.5	63.0	53.9	57.1	53.0		
	LEAP (ours)	✓	✗	63.7	67.0	61.3	64.0	62.8	58.2	63.9	56.2	59.5	56.6		
		✓	✓	64.8	67.7	61.8	64.8	63.6	59.2	64.9	56.5	60.2	57.4		

**Ablation study of our semantic-aware optimization objective.** We ablate the total objective  $\mathcal{L}$  (Eq. 12) and evaluate its impacts on two models employing the HAN [23] and MM-Pyr [30] as audio-visual encoders. As shown in Table 2 (with rows highlighted in gray), models trained with  $\mathcal{L}_{basic}$  have considerable performance as  $\mathcal{L}_{basic}$  uses explicit segment-level labels as supervisions. Moreover,  $\mathcal{L}_{avss}$  further boosts the parsing performances. Its effectiveness is more pronounced when integrated with the more advanced encoder MM-Pyr, resulting in a 1.0% improvement in event-level metrics for both audio and visual event parsing. These results indicate the benefits of  $\mathcal{L}_{avss}$  as part of our comprehensive semantic-aware optimization strategy, further enhancing the regularization of audio-visual relations.

### 4.3 Comparison with the Typical MMIL

We comprehensively compare our event decoding paradigm, LEAP, against the typical MMIL. Two widely employed audio-visual backbones, specifically HAN [23] and MM-Pyr [30], are used as the early encoders unless specified otherwise.

**Comparison on parsing events across different modalities. 1)** As shown in Table 2 (with numbers highlighted in blue), AVVP models utilizing our LEAP exhibit overall improved performances across audio, visual, and audio-visual event parsing, in contrast to models using MMIL. The improvement is more obvious when integrating with the advanced encoder MM-Pyr [30]. For example, the “Event@AV” metrics, indicative of the comprehensive audio and visual event parsing performance, at the segment level and event level are significantly improved by 2.2% and 3.6%, respectively. **2)** Beyond this holistic dataset comparison, we detail the parsing performances across distinct audio and visual event categories. As shown in Fig. 3, the proposed LEAP surpasses MMIL in most of the event categories for both audio and visual modalities. In particular, the event-level F-score for the event *telephone* experiences a substantial 14.4% and 50.0% improvement for audio and visual modalities, respectively. The average performances of audio and visual event parsing are improved by 2.7%. These results demonstrate the superiority of the proposed LEAP over traditional MMIL in parsing event semantics across audio and visual modalities.

Category	Audio			Visual		
	MMIL	LEAP (Ours)	$\Delta$	MMIL	LEAP (Ours)	$\Delta$
Speech	32.8	41.7	8.9	0.0	0.0	0.0
Car	61.8	67.5	5.7	58.0	68.6	10.6
Cheering	26.5	26.9	0.4	26.4	23.5	-2.9
Dog	40.7	47.5	6.7	70.8	71.1	0.3
Cat	18.8	21.1	2.3	49.2	48.8	-0.4
Frying food	85.3	80.8	-4.6	83.0	84.4	1.4
Basketball bounce	24.5	25.3	0.9	23.7	22.4	-1.2
Fire alarm	57.5	62.1	4.6	50.0	47.6	-2.4
Chainsaw	73.6	74.7	1.2	73.8	72.2	-1.6
Cello	64.6	69.0	4.5	68.3	65.5	-2.7
Banjo	57.1	65.8	8.6	65.5	73.3	7.8
Singing	46.5	51.0	4.5	40.6	41.0	0.4
Chicken	47.5	49.1	1.5	83.1	84.4	1.3
Violin	69.0	66.8	-2.2	69.5	68.2	-1.3
Vacuum cleaner	81.7	81.4	-0.3	80.7	87.0	6.3
Baby laughter	33.7	28.6	-5.1	36.3	26.3	-10.0
Accordion	72.7	78.6	5.9	75.4	77.6	2.2
Lawn mower	59.1	58.3	-0.8	72.2	74.7	2.5
Motorcycle	72.9	74.3	1.4	69.3	68.8	-0.5
Helicopter	80.8	83.8	3.0	70.7	71.4	0.7
Acoustic guitar	36.2	40.2	4.0	42.4	44.4	2.0
Telephone	27.3	41.7	14.4	0.0	50.0	50.0
Baby cry	44.9	48.9	4.0	40.8	47.3	6.5
Blender	89.5	87.0	-2.5	88.0	86.1	-1.9
Clapping	14.3	14.7	0.3	16.4	17.3	0.9
<b>Avg.</b>	<b>52.8</b>	<b>55.5</b>	<b>2.7</b>	<b>54.2</b>	<b>56.9</b>	<b>2.7</b>

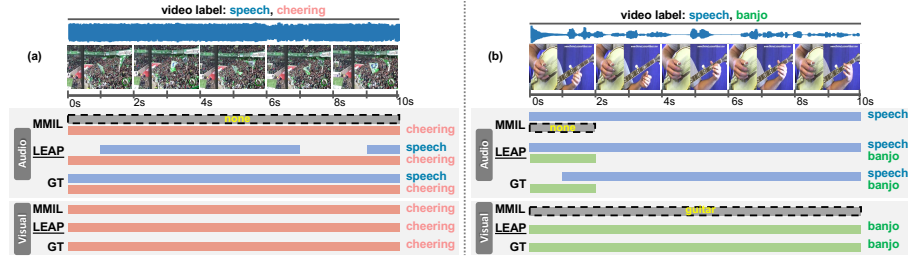
**Fig. 3: Comparison between LEAP and typical MMIL in parsing audio and visual events in each class.**  $\Delta$  denotes the performance improvements of our method compared to MMIL and “Avg.” denotes the average results of all the event classes. MM-Pyr [30] is used as the audio-visual encoder and the event-level metrics are reported.

**Table 3: Comparison between LEAP and typical MMIL in tackling non-overlapping and overlapping events.** The event-level metrics are reported.

Methods		non-overlapping						overlapping					
Encoders	Decoders	A	V	AV	Type@AV	Event@AV	Avg.	A	V	AV	Type@AV	Event@AV	Avg.
HAN [23]	MMIL [23]	66.7	69.6	57.1	64.5	56.6	62.9	49.7	44.1	46.5	46.8	47.5	46.9
	<b>LEAP</b>	<b>68.6</b>	<b>71.9</b>	<b>58.7</b>	<b>66.4</b>	<b>57.6</b>	<b>64.6</b>	<b>50.6</b>	<b>44.4</b>	<b>48.7</b>	<b>47.9</b>	<b>48.5</b>	<b>48.0</b>
MM-Pyr [30]	MMIL [23]	66.5	72.8	58.6	66.0	56.4	64.1	49.7	45.7	49.4	48.3	47.3	48.1
	<b>LEAP</b>	<b>72.1</b>	<b>73.0</b>	<b>60.9</b>	<b>68.7</b>	<b>60.6</b>	<b>67.1</b>	<b>52.4</b>	<b>46.2</b>	<b>50.7</b>	<b>49.7</b>	<b>50.0</b>	<b>49.8</b>

**Comparison on parsing non-overlapping and overlapping events.** We divide the test set of LLP dataset into two subsets: the overlapping set and the non-overlapping set. The former set consists of those videos that contain multiple events in at least one segment, while the remaining videos form the non-overlapping set where one segment only contains a single event with a specific class. As shown in Table 3, the proposed LEAP has better performance than the typical MMIL in parsing both types of events. When employing the MM-Pyr [30] as the audio-visual encoder, our LEAP outperforms MMIL by 3.0% in parsing non-overlapping events. The improvement over MMIL is still significant (1.7%) when dealing with the more challenging overlapping case. These results again verify the superiority of our LEAP in effectively distinguishing different event classes and disentangling overlapping semantics.

**Qualitative comparison on audio-visual video parsing.** As shown in Fig. 4 (a), this video contains two events, *i.e.*, *speech* and *cheering*. Only the *cheering* event exists in the visual track and both the typical MMIL and our LEAP suc-



**Fig. 4: Qualitative examples of audio-visual video parsing.** Compared to MMIL, the proposed LEAP performs better in distinguishing the semantics of non-overlapping and overlapping events.

**Table 4: Comparison with state-of-the-arts.** <sup>▲</sup> denotes those methods are developed on the baseline HAN [23]. <sup>▼</sup> denotes those methods that focus on designing stronger audio-visual encoders. The best and second-best results are **bolded** and underlined, respectively.

Methods	Venue	Segment-level					Event-level				
		A	V	AV	Type@AV	Event@AV	A	V	AV	Type@AV	Event@AV
HAN [23]	ECCV'20	60.1	52.9	48.9	54.0	55.4	51.3	48.9	43.0	47.7	48.0
<sup>▲</sup> CVCMS [16]	NeurIPS'21	59.2	59.9	53.4	57.5	58.1	51.3	55.5	46.2	51.0	49.7
<sup>▲</sup> MA [27]	CVPR'21	60.3	60.0	55.1	58.9	57.9	53.6	56.4	49.0	53.0	50.6
<sup>▲</sup> JoMoLD [2]	ECCV'22	61.3	63.8	57.2	60.8	59.9	53.9	59.9	49.6	54.5	52.5
<sup>▲</sup> BPS [20]	ICCV'23	<b>63.1</b>	63.5	57.7	61.4	60.6	54.1	60.3	51.5	55.2	52.3
<sup>▲</sup> VALOR [31]	NeurIPS'23	61.8	<b>65.9</b>	<u>58.4</u>	<u>62.0</u>	<u>61.5</u>	<u>55.4</u>	<u>62.6</u>	<u>52.2</u>	<u>56.7</u>	<u>54.2</u>
HAN [23] + <b>LEAP (ours)</b>	-	<u>62.7</u>	<u>65.6</u>	<b>59.3</b>	<b>62.5</b>	<b>61.8</b>	<b>56.4</b>	<b>63.1</b>	<b>54.1</b>	<b>57.8</b>	<b>55.0</b>
MM-Pyr [30]	MM'22	60.9	54.4	50.0	55.1	57.6	52.7	51.8	44.4	49.9	50.5
<sup>▼</sup> MGN [18]	NeurIPS'22	60.8	55.4	50.4	55.5	57.2	51.1	52.4	44.4	49.3	49.1
<sup>▼</sup> DHHN [11]	MM'22	61.3	58.3	52.9	57.5	58.1	54.0	55.1	47.3	51.5	51.5
<sup>▼</sup> CMPAE [5]	CVPR'23	<u>64.2</u>	<u>66.4</u>	<u>59.2</u>	<u>63.3</u>	<u>62.8</u>	<u>56.6</u>	<u>63.7</u>	<u>51.8</u>	<u>57.4</u>	<u>55.7</u>
MM-Pyr [30] + <b>LEAP (ours)</b>	-	<b>64.8</b>	<b>67.7</b>	<b>61.8</b>	<b>64.8</b>	<b>63.6</b>	<b>59.2</b>	<b>64.9</b>	<b>56.5</b>	<b>60.2</b>	<b>57.4</b>

cessfully recognize this visual event. However, when there are overlapping events in audio modality, MMIL totally misses the audio event *speech*. In contrast, our LEAP correctly identifies this event and gives satisfactory segment-level predictions. Similarly, in Fig. 4(b), MMIL fails to recognize the audio event *banjo* in the initial two segments, whereas our LEAP successfully disentangles *banjo* semantic, even though it overlaps with *speech*. Besides, MMIL incorrectly identifies the non-overlapping visual event *banjo* as the similar event *guitar*, while our LEAP predicts the correct category. These results demonstrate the superiority of our method which disentangles different semantics into separate label embeddings, benefiting the various category recognition and overlapping event distinction.

#### 4.4 Comparison with the State-of-the-Arts

We compare our method with prior works. As shown in the upper part of Table 4, our LEAP-based model is superior to those methods developed based on HAN [23]. It is noteworthy that the most competitive work VALOR [31] also uses the segment-level pseudo labels as supervision but adopts the typical

**Table 5: Generalization of our LEAP to the audio-visual event localization (AVEL) task.** “DCH” denotes the default event decoding paradigm in this task that directly classifies audio-visual events by transforming hidden features.

AVEL Paradigms	AVE [24]	PSP [38]	CMBS [29]
DCH (default)	68.2	74.3	74.5
<b>LEAP (ours)</b>	<b>68.8</b> (+0.6)	<b>76.6</b> (+2.3)	<b>77.9</b> (+3.4)

MMIL [23] for event decoding. In contrast, we combine HAN with the proposed LEAP which has better performance. Methods listed in the lower part of Table 4 primarily focus on designing stronger audio-visual encoders and we report their optimal performance. CMPAE [5] is most competitive because it additionally selects thresholds for each event class during event inference while we directly use the threshold of 0.5 as in baselines [23, 30]. Without bells and whistles, we show that the proposed LEAP equipped with the baseline encoder MM-Pyr [30] has achieved new state-of-the-art performance in all types of event parsing.

#### 4.5 Generalization to AVEL Task

We finally extend our label semantic-based projection (LEAP) decoding paradigm to one related audio-visual event localization (AVEL) task, which aims to localize video segments containing events both audible and visible. We evaluate three typical audio-visual encoders in this task, including AVE [24], PSP [38], and CMBS [29]. We combine them with our decoding paradigm, LEAP, based on the official codes. As shown in Table 5, our LEAP is also superior to the default paradigm in this task, consistently boosting the vanilla models. The improvement further increases when using stronger audio-visual encoders. This indicates the generalization of our method and also verifies the benefits of introducing semantically independent label embeddings for the distinctions of different events.

## 5 Conclusion

Addressing the audio-visual video parsing task, this paper presents a straightforward yet highly effective label semantic-based projection (LEAP) method to enhance the event decoding phase. LEAP disentangles the potentially overlapping semantics by iteratively projecting the latent audio/visual features into separate label embeddings associated with distinct event classes. To facilitate the projection, we propose a semantic-aware optimization strategy, which adopts a novel audio-visual semantic similarity loss to enhance feature encoding. Extensive experimental results demonstrate that our method outperforms the typical video decoder MMIL in parsing all types of events and in handling overlapping events. Our method is not only compatible with existing representative audio-visual encoders for AVVP but also benefits the AVEL task. We anticipate our approach to serve as a new video parsing paradigm for the relevant community.

**Acknowledgement** We sincerely appreciate the anonymous reviewers for their positive feedback. This work was supported by the National Key R&D Program of China (NO.2022YFB4500601), the National Natural Science Foundation of China (72188101, 62272144, 62020106007, and U20A20183), the Major Project of Anhui Province (202203a05020011), and the Fundamental Research Funds for the Central Universities.

## References

1. Chen, H., Zhu, D., Zhang, G., Shi, W., Zhang, X., Li, J.: Cm-cs: Cross-modal common-specific feature learning for audio-visual video parsing. In: ICASSP. pp. 1–5 (2023)
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: ECCV. pp. 431–448 (2022)
3. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR. pp. 248–255 (2009)
4. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 pp. 1–16 (2018)
5. Gao, J., Chen, M., Xu, C.: Collecting cross-modal presence-absence evidence for weakly-supervised audio-visual event perception. In: CVPR. pp. 18827–18836 (2023)
6. Gemmeke, J.F., Ellis, D.P., Freedman, D., Jansen, A., Lawrence, W., Moore, R.C., Plakal, M., Ritter, M.: Audio set: An ontology and human-labeled dataset for audio events. In: ICASSP. pp. 776–780 (2017)
7. Guo, D., Li, K., Hu, B., Zhang, Y., Wang, M.: Benchmarking micro-action recognition: Dataset, methods, and applications. IEEE TCSVT pp. 6238–6252 (2024)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
9. 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: ICASSP. pp. 131–135 (2017)
10. Hu, D., Nie, F., Li, X.: Deep multimodal clustering for unsupervised audiovisual learning. In: CVPR. pp. 9248–9257 (2019)
11. 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: ACM MM. pp. 719–727 (2022)
12. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 pp. 1–22 (2017)
13. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR. pp. 1–15 (2014)
14. Li, G., Wei, Y., Tian, Y., Xu, C., Wen, J.R., Hu, D.: Learning to answer questions in dynamic audio-visual scenarios. In: CVPR. pp. 19108–19118 (2022)
15. Li, Z., Guo, D., Zhou, J., Zhang, J., Wang, M.: Object-aware adaptive-positivity learning for audio-visual question answering. In: AAAI. pp. 3306–3314 (2024)
16. 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. In: NeurIPS. pp. 1–13 (2021)

17. Mao, Y., Zhang, J., Xiang, M., Zhong, Y., Dai, Y.: Multimodal variational auto-encoder based audio-visual segmentation. In: ICCV. pp. 954–965 (2023)
18. Mo, S., Tian, Y.: Multi-modal grouping network for weakly-supervised audio-visual video parsing. In: NeurIPS. pp. 1–12 (2022)
19. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: EMNLP. pp. 1532–1543 (2014)
20. Rachavarapu, K., A. N., R.: Boosting positive segments for weakly-supervised audio-visual video parsing. In: ICCV. pp. 10192–10202 (2023)
21. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763 (2021)
22. Shen, X., Li, D., Zhou, J., Qin, Z., He, B., Han, X., Li, A., Dai, Y., Kong, L., Wang, M., et al.: Fine-grained audible video description. In: CVPR. pp. 10585–10596 (2023)
23. Tian, Y., Li, D., Xu, C.: Unified multisensory perception: Weakly-supervised audio-visual video parsing. In: ECCV. pp. 436–454 (2020)
24. Tian, Y., Shi, J., Li, B., Duan, Z., Xu, C.: Audio-visual event localization in unconstrained videos. In: ECCV. pp. 247–263 (2018)
25. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: CVPR. pp. 6450–6459 (2018)
26. Wei, Y., Hu, D., Tian, Y., Li, X.: Learning in audio-visual context: A review, analysis, and new perspective. arXiv preprint arXiv:2208.09579 (2022)
27. Wu, Y., Yang, Y.: Exploring heterogeneous clues for weakly-supervised audio-visual video parsing. In: CVPR. pp. 1326–1335 (2021)
28. Wu, Y., Zhu, L., Yan, Y., Yang, Y.: Dual attention matching for audio-visual event localization. In: ICCV. pp. 6292–6300 (2019)
29. Xia, Y., Zhao, Z.: Cross-modal background suppression for audio-visual event localization. In: CVPR. pp. 19989–19998 (2022)
30. 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: ACM MM. pp. 6241–6249 (2022)
31. Yung-Hsuan Lai, Yen-Chun Chen, Y.C.F.W.: Modality-independent teachers meet weakly-supervised audio-visual event parser. In: NeurIPS. pp. 1–19 (2023)
32. Zhang, J., Li, W.: Multi-modal and multi-scale temporal fusion architecture search for audio-visual video parsing. In: ACM MM. pp. 3328–3336 (2023)
33. Zhou, J., Guo, D., Wang, M.: Contrastive positive sample propagation along the audio-visual event line. TPAMI pp. 7239–7257 (2023)
34. Zhou, J., Guo, D., Zhong, Y., Wang, M.: Improving audio-visual video parsing with pseudo visual labels. arXiv preprint arXiv:2303.02344 (2023)
35. Zhou, J., Guo, D., Zhong, Y., Wang, M.: Advancing weakly-supervised audio-visual video parsing via segment-wise pseudo labeling. IJCV pp. 1–22 (2024)
36. Zhou, J., Shen, X., Wang, J., Zhang, J., Sun, W., Zhang, J., Birchfield, S., Guo, D., Kong, L., Wang, M., et al.: Audio-visual segmentation with semantics. arXiv preprint arXiv:2301.13190 (2023)
37. Zhou, J., Wang, J., Zhang, J., Sun, W., Zhang, J., Birchfield, S., Guo, D., Kong, L., Wang, M., Zhong, Y.: Audio-visual segmentation. In: ECCV. pp. 386–403 (2022)
38. Zhou, J., Zheng, L., Zhong, Y., Hao, S., Wang, M.: Positive sample propagation along the audio-visual event line. In: CVPR. pp. 8436–8444 (2021)