



Advancing Weakly-Supervised Audio-Visual Video Parsing via Segment-Wise Pseudo Labeling

Jinxing Zhou^{1,2} · Dan Guo^{1,3,4} · Yiran Zhong² · Meng Wang^{1,3}

Received: 19 September 2023 / Accepted: 31 May 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

The Audio-Visual Video Parsing task aims to identify and temporally localize the events that occur in either or both the audio and visual streams of audible videos. It often performs in a weakly-supervised manner, where only video event labels are provided, i.e., the modalities and the timestamps of the labels are unknown. Due to the lack of densely annotated labels, recent work attempts to leverage pseudo labels to enrich the supervision. A commonly used strategy is to generate pseudo labels by categorizing the known video event labels for each modality. However, the labels are still confined to the video level, and the temporal boundaries of events remain unlabeled. In this paper, we propose a new pseudo label generation strategy that can explicitly assign labels to each video segment by utilizing prior knowledge learned from the open world. Specifically, we exploit the large-scale pretrained models, namely CLIP and CLAP, to estimate the events in each video segment and generate segment-level visual and audio pseudo labels, respectively. We then propose a new loss function to exploit these pseudo labels by taking into account their category-richness and segment-richness. A label denoising strategy is also adopted to further improve the visual pseudo labels by flipping them whenever abnormally large forward losses occur. We perform extensive experiments on the LLP dataset and demonstrate the effectiveness of each proposed design and we achieve state-of-the-art video parsing performance on all types of event parsing, i.e., audio event, visual event, and audio-visual event. Furthermore, our experiments verify that the high-quality segment-level pseudo labels provided by our method can be flexibly combined with other audio-visual video parsing backbones and consistently improve their performances. We also examine the proposed pseudo label generation strategy on a relevant weakly-supervised audio-visual event localization task and the experimental results again verify the benefits and generalization of our method.

Keywords Audio-visual video parsing · Audio-visual event localization · Pseudo labeling · Label denoising · Audio-visual learning

Communicated by Gunhee Kim.

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. This work is also partially supported by the National Key R&D Program of China (NO.2022ZD0160100).

✉ Dan Guo
guodan@hfut.edu.cn

✉ Meng Wang
eric.mengwang@gmail.com

Jinxing Zhou
zhoujxfut@gmail.com

Yiran Zhong
zhongyiran@gmail.com

1 Introduction

Acoustic and visual signals flood our lives in abundance, and each signal may carry various events. For example, we often see driving cars and pedestrians walking around on the street. Meanwhile, we can hear the beeping of the car horns and the sound of people talking. Humans achieve such a comprehensive understanding of audio-visual events in large part thanks to the simultaneous use of their auditory and visual sensors. To imitate this kind of intelligence for machines,

¹ Hefei University of Technology, Hefei, China

² Shanghai AI Laboratory, Shanghai, China

³ Hefei Comprehensive National Science Center, Hefei, China

⁴ Anhui Zhonghuitong Technology Co., Ltd., Hefei, China

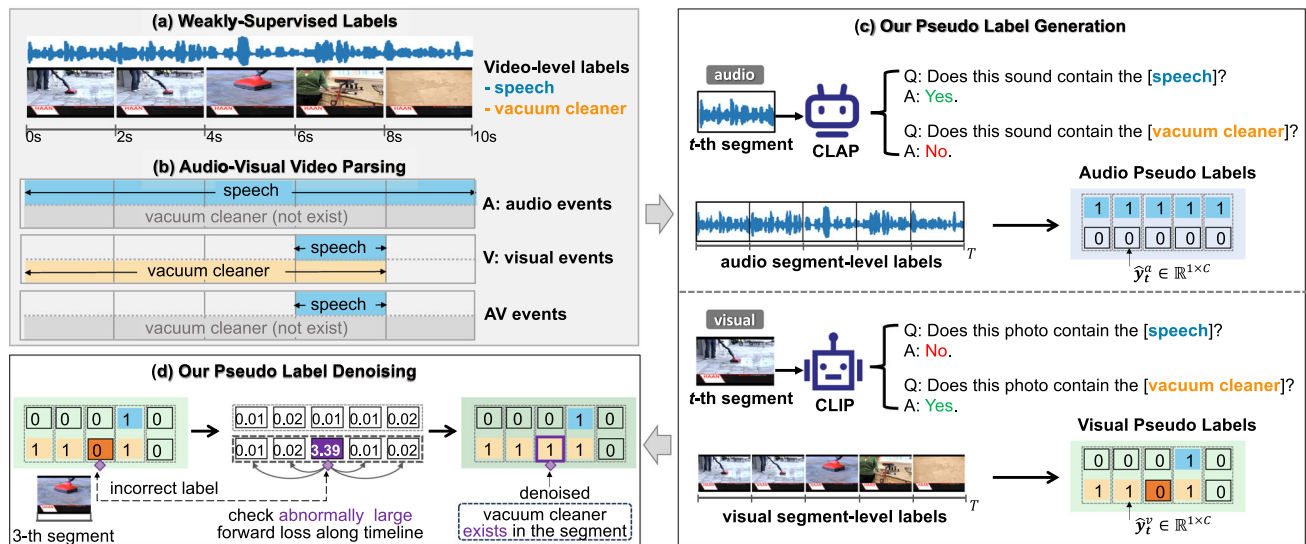


Fig. 1 An illustration of the weakly-supervised audio-visual video parsing (AVVP) task and our pseudo label exploration method. **a** Given a video and its event label (“speech” and “vacuum cleaner”), **b** AVVP task needs to predict and localize the audio events, visual events, and audio-visual events. Note that “vacuum cleaner” only exists in the visual track, while “speech” exists in both audio and visual tracks, resulting in the audio-visual event “speech”. **c** To ease this challenging weakly-supervised task, we aim to explicitly assign reliable segment-level audio and visual pseudo labels. In our pseudo label generation process, the

pretrained CLAP and CLIP models are used to tell what events occur in each audio and visual segment, respectively. **d** We further propose a pseudo label denoising strategy to improve the obtained visual pseudo labels by examining those segments that have abnormally large forward loss values. In the example, visual event *vacuum cleaner* at the third segment is assigned an incorrect pseudo label ‘0’ and gets a large forward loss. Our pseudo-label denoising strategy further amends this, giving the accurate pseudo label ‘1’

many research works started from some fundamental tasks of single modality understanding, such as the audio classification (Hershey et al., 2017; Kong et al., 2018; Kumar et al., 2018; Gong et al., 2021), video classification (Karpathy et al., 2014; Long et al., 2018a, b; Tran et al., 2019), and temporal action localization (Zeng et al., 2019; Chao et al., 2018; Zhu et al., 2021; Gao et al., 2022). The audio classification task focuses on the recognition of the audio modality, while the video classification and temporal action localization tasks focus on the visual modality. With the deepening of research, many works have further explored the multi-modal audio-visual perception (Wei et al., 2022), giving birth to tasks such as sound source localization (Arandjelovic & Zisserman, 2017; Rouditchenko et al., 2019; Arandjelovic & Zisserman, 2018; Senocak et al., 2018; Hu et al., 2020, 2019; Qian et al., 2020; Zhao et al., 2018; Afouras et al., 2020; Zhou et al., 2022b, 2023b; Sun et al., 2023), audio-visual event localization (Tian et al., 2018; Wu et al., 2019; Xu et al., 2020; Zhou et al., 2021; Mahmud & Marculescu, 2022; Rao et al., 2022b; Xia & Zhao, 2022; Wu et al., 2022; Zhou et al., 2023a; Wang et al., 2023), audio-visual video description (Shen et al., 2023) and question answering (Yun et al., 2021; Li et al., 2022; Yang et al., 2022; Song et al., 2022; Li et al., 2023).

Recently, Tian et al. (Tian et al., 2020) proposed a new multi-modal scene understanding task, namely Audio-Visual

Video Parsing (AVVP). Given an audible video, the AVVP task asks to identify what events occur in the audio and visual tracks and in which video segments these events occur. Accordingly, the category and temporal boundary of each event are expected to be predicted for each modality. Note that both the audio and visual tracks may contain multiple distinct events, and these events usually exist in different consecutive segments, it is labor-intensive to provide segment-level event labels for each modality with strong supervision. The fact is that the AVVP is performed in a weakly-supervised setting where only the video label is provided during model training. As the example shown in Fig. 1a, we only know that this video contains the event set of *speech* and *vacuum cleaner*. For each event, the model needs to judge whether it exists in the audio modality (audio event), visual modality (visual event), or both (audio-visual event), and locate the specific temporal segments, respectively. Notably, as illustrated in Fig. 1b, in the AVVP task, the audio-visual event is the intersection of the audio event and visual event, whereas the video label is the union of the audio event and visual event.

In this work, we emphasize there are two main challenges in the AVVP task. **(1) Cross-modal interference from the video label.** As the example shown in Fig. 1b, given the weakly-supervised video label, the audio and the visual track share the same supervision, i.e., {*speech*, *vacuum cleaner*}

together. However, the audio and visual tracks contain distinct events. The *vacuum cleaner* only exists in the visual modality. Thus, during the model training process, the label *vacuum cleaner* will interfere with the audio event parsing. Similarly, the visual event parsing may also be interfered with the audio event label in other samples. **(2) Temporal segment distinction.** Assuming we successfully identify there is an event *vacuum cleaner* in the visual modality, it is still hard to distinguish which segments contain this event (segment level) under the weakly-supervised labels (video level). These two challenges make the AVVP an intractable Multi-modal Multi-Instance Learning (MMIL) problem, namely distinguishing the events from both *modality* and *temporal* perspectives.

In the pioneer work (Tian et al., 2020), a benchmark named Hybrid Attention Network (HAN) is proposed to encode the audio-visual features, which uses attentive pooling to aggregate the audio and visual features to predict events of the video. The weak video label is used as the main supervision. To address this task, they propose to obtain the pseudo labels for separate audio and visual modalities by processing the known video label with label smoothing (Szegedy et al., 2016) technique. Their experimental results indicate that generating pseudo labels for each modality brings significant benefits for supervising event parsing (Tian et al., 2020). The subsequent studies diverge into two branches. Most of them focus on *designing effective networks* to implicitly aggregate the multi-modal features for prediction (Mo & Tian, 2022; Pasi et al., 2022; Lamba et al., 2021; Yu et al., 2022; Lin et al., 2021; Jiang et al., 2022; Gao et al., 2023), while using the video-level pseudo labels generated by HAN (Tian et al., 2020). In contrast, the other new works (Wu & Yang, 2021; Cheng et al., 2022) devote to *generating better pseudo labels* for each modality based on the baseline backbone of HAN. However, the generated pseudo label is denoised from the video label and limited to the video level which only indicates what events exist in each modality (modality perspective). Therefore, it fails to address the second challenge because it remains difficult to distinguish which segments contain the event (temporal perspective).

To deal with the above-mentioned two challenges, our work starts with the intuition that can we explicitly generate pseudo labels for **each segment** of **each modality** to facilitate this MMIL task. This is inspired by two observations: (1) The AVVP models are expected to be well-guided with segment-level labels as such fine-grained labels can provide more explicit supervision information and directly fit the goal of the AVVP task (temporal perspective); (2) The audio and visual signals are processed with independent sensors for humans. We can indeed annotate each modality, specifically for what we hear or see, by leveraging unimodal input (modality perspective). To this end, we propose a **Visual-Audio Pseudo Label exploration (VAPLAN) method** that

aims to generate high-quality segment-level pseudo labels for both visual modality and audio modality and further advances this weakly-supervised AVVP task.

To obtain the visual or audio pseudo labels, a natural idea is to borrow free knowledge from pretrained models for the image or audio classification. However, there is a category misalignment problem between the source and the target datasets using such a strategy. Take generating visual pseudo labels as an example, the models typically pretrained on the ImageNet (Deng et al., 2009) would classify the instance in the AVVP task into *predefined categories of the ImageNet*. However, the predicted category label may not exist in the target LLP dataset of the AVVP task, causing the category misalignment. Different from the traditional image classification models, vision-language pre-training (Alayrac et al., 2022; Jia et al., 2021; Radford et al., 2021) has attracted tremendous attention recently, which can flexibly classify images from an open-category vocabulary and show impressive zero-shot performance. Among those works, Contrastive Language-Image Pretraining (CLIP) (Radford et al., 2021) is a representative one. Given an image, its potential category names are inserted into a predefined text prompt. Then CLIP can score the categories according to the similarity between the encoded texts and the image features. The category with a high similarity score is finally identified as the classification result. Similar to the CLIP, in the audio community, the Contrastive Language-Audio Pretraining (CLAP) (Wu et al., 2023) is trained on a large-scale corpus that incorporates the texts with the semantic-aligned audio. With similar training and inference schemes, CLAP is able to perform audio classification in a zero-shot manner, and satisfactorily identify the category of a given audio from open-vocabulary too.

Inspired by such benefits of large-scale pretraining, we propose a **Pseudo Label Generation (PLG)** module that seeks guidance from the CLIP (Radford et al., 2021) and CLAP (Wu et al., 2023) to generate reliable segment-level visual and audio pseudo labels. A simple illustration of PLG can be seen from Fig. 1c. Given all the potential event labels, CLIP/CLAP acting like an intelligent robot is asked to answer whether the event is contained in the given visual/audio segment. In brief, the queried event categories with high cross-modal similarity scores that exceed the pre-set threshold τ_v/τ_a are finally regarded as the visual/audio pseudo labels. This process can be applied to each video segment, so we can obtain segment-level pseudo labels. We provide more implementation details in Sect. 4.1. The generated pseudo labels are used to provide full supervision for each modality. Going a step further, we consider the generated pseudo labels may contain potential noise since the pseudo labels are non-manually annotated. Especially, some video instances can be challenging even for human annotators due to issues inherent in the collected videos, such as objects in the visual event being too small or obscured. As the example shown in Fig. 1d,

only part of the *vacuum cleaner* is visible in the third segment. PLG only uses the single frame to generate pseudo labels and fails to recognize the visual event *vacuum cleaner* for this segment without contextual information, giving the incorrect pseudo label '0' for this category (denoted by brown box). To alleviate such noise in pseudo labels generated by PLG, we further propose a **Pseudo Label Denoising (PLD)** strategy to re-examine the generated pseudo labels and amend the incorrect ones. Samples with noisy labels are usually hard to learn and often get a large forward propagation loss (Hu et al., 2021a; Kim et al., 2022; Huang et al., 2019). In our work, the large loss comes from those data where the model is unable to give consistent predictions with the pseudo labels. For the video example shown in Fig. 1d, the third segment indeed suffers an abnormally large forward loss whereas the value is 3.39. Note that the values are almost zero for other segments in the same video which are assigned accurate labels. This motivates us to perform a segment-wise denoising by checking the abnormally large forward loss along the timeline. The segments with these controversial pseudo labels will be reassigned, providing a more accurate version. More discussions and implementation details of PLD will be introduced in Sect. 4.3.

PLG and PLD enable the production of high-quality pseudo labels. Furthermore, we find that the obtained segment-level audio and visual pseudo labels contain rich information, indicating *how many categories of events happen in each audio/visual segment* (category-richness) and *how many audio/visual segments a certain category of the event exists in* (segment-richness). Take the visual modality for example, as shown in Fig. 1b, the video-level label indicates that there may be at most *two* events in the visual track, i.e., the *speech* and *vacuum cleaner*. In practice, only the fourth segment contains both *two* events while the first segment contains only *one* event, namely the vacuum cleaner. Therefore, we can denote the visual category richness for the first and the fourth segments as 1/2 and 1, respectively. Similarly, from the perspective of the event categories, the vacuum cleaner event appears in *four* video segments of the entire video which totally contains *five* segments, while the speech event only exists in *one* (the fourth) segment. Thus, we can denote the visual segment richness for events of *vacuum cleaner* and *speech* as 4/5 and 1/5, respectively. Such information about category richness and segment richness can also be observed in the audio track. An AVVP model should be aware of the differences in category richness and segment richness to give correct predictions. Based on this, we propose a **Pseudo Label Exploitation (PLE)** strategy that uses a novel *Richness-aware Loss* to align the richness information contained in model predictions with that contained in pseudo labels. Our experiments verify that the generated pseudo labels combined with the proposed richness-aware loss significantly boost the video parsing performance.

For the challenging audio-visual video parsing task, we conduct a comprehensive study on the exploration of the segment-wise audio and visual pseudo labels, including their generation, exploitation, and denoising. Extensive experimental results demonstrate the effectiveness of our main designs. Besides, our method can also be extended to the related weakly-supervised audio-visual event localization (AVEL) (Tian et al., 2018; Wu et al., 2019; Zhou et al., 2021) task. Overall, our contributions can be summarized as follows:

- We introduce a new approach to explore the pseudo-label strategy for the AVVP task from a more fine-grained level, i.e., the segment level.
- Our proposed pseudo label generation and label denoising strategies successfully provide high-quality segment-wise audio and visual pseudo labels.
- We propose a new richness-aware loss function for superior model optimization, effectively exploiting the segment-richness and category-richness present in the pseudo labels.
- Our method achieves new state-of-the-art in all types of event parsing, including audio event, visual event, and audio-visual event parsing.
- The proposed core designs can be seamlessly integrated into existing frameworks for the AVVP task and AVEL task, leading to enhanced performances.

2 Related Work

Audio-Visual Video Parsing (AVVP). AVVP task needs to recognize what events happen in each modality and localize the corresponding video segments where the events exist. Tian et al. (Tian et al., 2020) first propose this task and design a hybrid attention network to aggregate the intra-modal and inter-modal features. Also, they use the label smoothing (Szegedy et al., 2016) strategy to address the modality label bias from the single video-level label. Some methods focus on network design. Yu et al. (2022) propose a multimodal pyramid attentional network that consists of multiple pyramid units to encode the temporal features. Jiang et al. (2022) use two extra independent visual and audio prediction networks to alleviate the label interference between audio and visual modalities. Mo & Tian (2022) use learnable class-aware tokens to group the semantics from separate audio and visual modalities. To overcome the label interference, Wu & Yang (2021) swap the audio and visual tracks of two event-independent videos to construct new data for model training. The pseudo labels are generated according to the predictions of the reconstructed videos. Cheng et al. (2022) first estimate the noise ratio of the video label and reverse a certain percentage of the label with large forward

losses. Although these methods bring considerable improvements, they can only generate the event label from the video level. Unlikely, we aim to directly obtain high-quality pseudo labels for both audio and visual modalities from the segment level that further helps the video parsing system training.

CLIP/CLAP Pre-Training. Here, we discuss the pre-training technique and elaborate on why we choose the CLIP/CLAP as the base big model for generating pseudo labels in this work. CLIP (Radford et al., 2021) is trained on a dataset with 400 million *image-text* pairs using the contrastive learning technique. This large-scale pretraining enables CLIP to learn efficient representations of the images and texts and demonstrates impressive performance on zero-shot image classification. Its zero-shot transfer ability opens a new scheme to solve many tasks and spawns a large number of research works, such as image caption (Barraco et al., 2022), video caption (Tang et al., 2021), and semantic segmentation (Ma et al., 2022; Ding et al., 2022; Xu et al., 2021; Zhou et al., 2022a; Rao et al., 2022b). Most of the works choose to freeze or fine-tune the image and text encoders of CLIP to extract advanced features for downstream tasks (Tang et al., 2021; Wang et al., 2022; Barraco et al., 2022; Ma et al., 2022; Zhou et al., 2022c). For the zero-shot semantic segmentation, some methods start to use the pretrained CLIP to generate pixel-level pseudo labels which are annotator-free and helpful (Zhou et al., 2022a; Rao et al., 2022b). Similarly to CLIP, CLAP (Wu et al., 2023) is trained using a similar contrastive objective but with 630k *audio-text* pairs and achieves state-of-the-art zero-shot audio classification performance. Recently, some works have started to use CLAP to facilitate downstream tasks, such as audio source separation (Liu et al., 2023b), text-to-audio generation (Liu et al., 2023a), and speech emotion recognition (Pan et al., 2023). In this work, we make a new attempt to borrow the prior knowledge from CLIP/CLAP to ease the challenging weakly-supervised audio-visual video parsing task.

Learning with Pseudo Labels. Deep neural networks achieve remarkable performance in various tasks, largely due to the large amount of labeled data available for training. Recently, some researchers have attempted to generate massive pseudo labels for unlabeled data to further boost model performance. Most methods directly generate and use pseudo labels, which have been proven to be beneficial for various tasks, such as image classification (Yalniz et al., 2019; Xie et al., 2020; Pham et al., 2021; Rizve et al., 2021; Zoph et al., 2020; Hu et al., 2021b), speech recognition (Kahn et al., 2020; Park et al., 2020), and image-based text recognition (Patel et al., 2023). For the studied AVVP task, few works study the impact of pseudo labels and existing several methods focus on disentangling the event pseudo label for each modality from the known video label (Tian et al., 2020; Wu & Yang, 2021; Cheng et al., 2022). However, the obtained pseudo labels are confined to the *video level*. On

the other hand, some new works in other fields notice the potential noise contained in the pseudo labels and propose effective methods to better learn with noisy pseudo labels (Hu et al., 2021a; Kim et al., 2022). Specifically, Hu et al. (2021a) propose to optimize the network by giving much weight to the clean samples while less on the hard-to-learn samples. In the weakly-supervised multi-label classification problem, Kim et al. (2022) propose to correct the false negative labels that are likely to have larger losses. However, these works focus on label refinement for image tasks. Refocusing on our video task, we conduct a comprehensive exploration of pseudo labels, encompassing both their generation and denoising. Specifically, we propose to assign explicit pseudo labels for *each segment of each modality*. We achieve this goal by flexibly sending all the possible event categories to reliable large-scale text-vision/audio models to pick the most likely event categories for each video segment. Furthermore, we propose a new pseudo-label denoising strategy, which performs *segment-wise* denoising to provide pseudo labels with more accurate temporal boundaries within each video. We also provide more in-depth discussions on pseudo-label quality assessment and the denoising effects in different modalities as shown in Sect. 5.2.

3 Preliminary

In this section, we formulate the detail of the AVVP task and briefly introduce the baseline framework HAN (Tian et al., 2020), which is used in both our approach and prior works employing video-level pseudo labels (Wu & Yang, 2021; Cheng et al., 2022) in the AVVP task.

Task Formulation. Given a T -second video sequence $\{V_t, A_t\}_{t=1}^T$, V_t and A_t denote the visual and the audio components at the t -th video segment, respectively. The event label of the video $\mathbf{y}^{v \cup a} \in \mathbb{R}^{1 \times C} = \{y_c^{v \cup a} | y_c^{v \cup a} \in \{0, 1\}, c = 1, 2, \dots, C\}$, where C is the total number of event categories, the superscript ' $v \cup a$ ' denotes the event label of the entire video is the union of the labels of audio and visual modalities, value 1 of $y_c^{v \cup a}$ represents an event with that c -th category happens in the video. Note that $\mathbf{y}^{v \cup a}$ is a weakly-supervised label from the video level, *the label of each individual modality for each video segment is unknown during training. However, the audio events, visual events, and audio-visual events contained in each segment need to be predicted for evaluation.* We denote the probabilities of the video-level visual and audio events as $\{\{\mathbf{p}^v; \mathbf{p}^a\} \in \mathbb{R}^{1 \times C} | p_c^v, p_c^a \in [0, 1]\}$, $\mathbf{p}^{v \cap a} = \mathbf{p}^v * \mathbf{p}^a$ is used to represent the intersection of them. Thus, the probability of the visual events, audio events, and audio-visual events of all video segments can be denoted as $\{\mathbf{P}^v; \mathbf{P}^a; \mathbf{P}^{v \cap a}\} \in \mathbb{R}^{T \times C}$, which need to be predicted.

Baseline Framework. The baseline network HAN (Tian et al., 2020) uses the multi-head attention (MHA) mechanism in Transformer (Vaswani et al., 2017) to encode intra-modal and cross-modal features for audio and visual modalities. We denote the initial audio and visual features extracted by pretrained neural networks (Hershey et al., 2017; He et al., 2016) as $F^a, F^v \in \mathbb{R}^{T \times d}$, where d is the feature dimension. The process of HAN can be summarized as,

$$\begin{cases} \hat{F}^a = F^a + \text{MHA}(F^a, F^a) + \text{MHA}(F^a, F^v), \\ \hat{F}^v = F^v + \text{MHA}(F^v, F^v) + \text{MHA}(F^v, F^a), \end{cases} \quad (1)$$

where $\hat{F}^a, \hat{F}^v \in \mathbb{R}^{T \times d}$ are the updated audio and visual features. The probabilities of segment-wise events for audio and visual modalities are predicted through a fully-connected (FC) layer and a sigmoid function, denoted as $P^a \in \mathbb{R}^{T \times C}$ and $P^v \in \mathbb{R}^{T \times C}$. An attentive pooling layer is further used to transform the segment-level predictions $\{P^a; P^v\}$ to video-level predictions $\{p^a; p^v\} \in \mathbb{R}^{1 \times C}$. By summarizing the audio and visual predictions, p^a and p^v , we obtain the event prediction of the entire video $p^{v \cup a} \in \mathbb{R}^{1 \times C}$. The basic video-level objective for model training is:

$$\mathcal{L} = \mathcal{L}_{\text{bce}}(p^{v \cup a}, y^{v \cup a}) + \mathcal{L}_{\text{bce}}(p^a, \bar{y}^a) + \mathcal{L}_{\text{bce}}(p^v, \bar{y}^v), \quad (2)$$

where \mathcal{L}_{bce} is the binary cross-entropy loss, $y^{v \cup a} \in \mathbb{R}^{1 \times C}$ is the video-level ground truth label and $\{\bar{y}^v; \bar{y}^a\} \in \mathbb{R}^{1 \times C}$ are the video-level visual and audio pseudo labels generated using label smoothing (Szegedy et al., 2016) from $y^{v \cup a}$.

4 Our Method

An overview of our method is shown in Fig. 2. We focus on producing reliable segment-level audio and visual pseudo labels to better supervise the model for audio-visual video parsing. For the backbone, we simply adopt the baseline HAN (Tian et al., 2020) to generate event predictions. Our method provides the following new innovations. (1) We propose a **pseudo label generation** module that uses the pretrained CLIP (Radford et al., 2021) and CLAP (Wu et al., 2023) models to respectively generate reliable visual and audio pseudo labels from the segment level. (2) We then propose a **pseudo label exploitation** strategy to utilize the obtained pseudo labels. Specifically, we design a new *richness-aware loss* to regularize the predictions to be aware of the category richness and segment richness contained in the pseudo labels. This is helpful for model optimization. (3) We also propose a **pseudo label denoising** strategy that further improves the generated visual pseudo labels for those data with abnormally high forward loss values due to being

assigned incorrect pseudo labels. Next, we elaborate on these proposed strategies.

4.1 Pseudo Label Generation (PLG)

PLG aims to generate high-quality visual and audio pseudo labels from the segment level that are expected to alleviate the video-level label interference for single modality and better supervise the model to distinguish video segments. As discussed in Sect. 1, we select the pretrained CLIP (Radford et al., 2021) and CLAP (Wu et al., 2023) to achieve this goal due to their flexible open-vocabulary classification capabilities.

Taking visual modality as an example, we detail the pseudo label generation process. Specifically, each video instance is evenly split into several segments and we sample the middle frame to represent each segment. As shown in Fig. 2-1, for the sampled frame \mathcal{I}_t at the t -th segment, we input it into CLIP image encoder and obtain the visual feature, denoted as $f_t^I \in \mathbb{R}^{1 \times d}$. As for the event category encoding, the default text input of the CLIP text encoder follows the prompt “A photo of a [CLS]” where the [CLS] can be replaced by the potential category names. For the AVVP task, we empirically change the prompt to a more appropriate one, “This photo contains the [CLS]” (An ablation study of prompt in CLIP text encoder will be shown in Sect. 5.2). By replacing the [CLS] in this prompt with each event category and sending the generated texts to the CLIP text encoder, we can obtain the text (with event category) features of all C -class $f_t^T \in \mathbb{R}^{C \times d}$. Then the normalized cosine similarity $s_t \in \mathbb{R}^{1 \times C}$ between the image and event categories can be computed by,

$$s_t = \text{softmax} \left(\frac{f_t^I}{\|f_t^I\|_2} \otimes \left(\frac{f_t^T}{\|f_t^T\|_2} \right)^T \right), \quad (3)$$

where \otimes denotes the matrix multiplication, and \top is the matrix transposition. A high similarity score in s_t indicates that the event category is more likely to appear in the image.

We use a threshold τ_v to select the categories with higher confidence scores in s_t and obtain the score mask m_t . After that, we impose the score mask m_t on the known video-level label $y^{v \cup a}$ with element-wise multiplication \odot to filter out the visual events occurring at t -th segment $\hat{y}_t^v \in \mathbb{R}^{1 \times C}$. This process can be formulated as,

$$\begin{cases} m_t = \mathbb{1}(s_t - \tau_v), \\ \hat{y}_t^v = m_t \odot y^{v \cup a}, \end{cases} \quad (4)$$

where $\mathbb{1}(x_i)$ outputs ‘1’ when the input $x_i \geq 0$ else outputs ‘0’, $i = 1, 2, \dots, C$, and $m_t \in \mathbb{R}^{1 \times C}$.

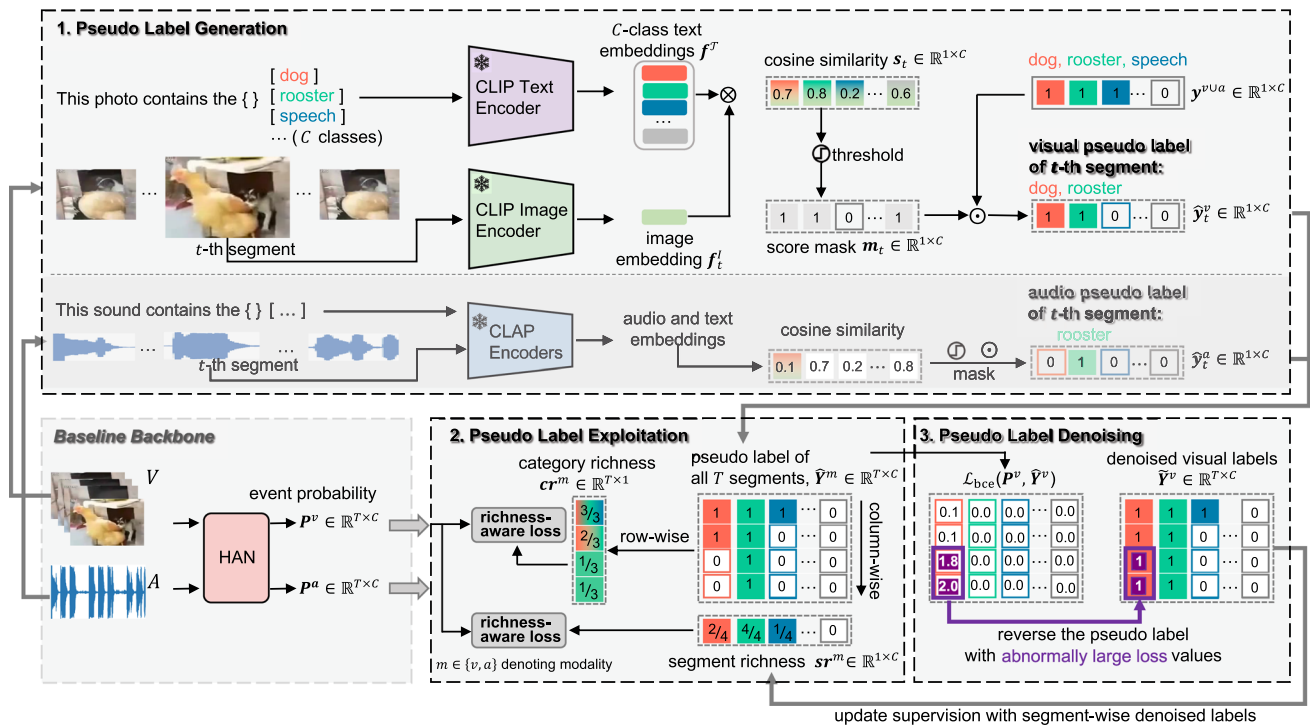


Fig. 2 Overview of our method. As a label refining method, we aim to produce high-quality and fine-grained segment-wise event labels. For the backbone, any existing network for the AVVP task can be used to generate event predictions. Here, we adopt the baseline HAN (Tian et al., 2020). In our solution, we design a *pseudo label generation (PLG)* module, where the pretrained CLIP (Radford et al., 2021) and CLAP (Wu et al., 2023) are used to generate segment-level pseudo labels for the visual and the audio modality, respectively. Notably, the parameters of the CLIP and CLAP are frozen. In the figure, we detail the visual pseudo label generation and simplify that for the audio modality since they share similar pipelines. In brief, the pseudo labels can be identified by thresholding the similarity of visual/audio—(event) text embeddings. For the t -th segment, the video label ‘speech’ is filtered out for the

visual modality and only ‘rooster’ is remained for the audio modality. After that, with the generated pseudo labels, we propose the *pseudo label exploitation (PLE)* by designing a richness-aware loss as a new fully supervised objective to help the model align the category richness and segment richness in the prediction and pseudo label. Lastly, we design a *pseudo label denoising (PLD)* strategy that further refines the pseudo labels by reversing the positions with anomalously large forward loss values. Specifically, we re-examine the pseudo labels along the timeline. Pseudo labels of those segments with abnormal high binary cross-entropy forward loss will be refined (the motivation and implementation detail can be seen in Sect. 4.3). The updated pseudo labels are further used as new supervision for model training. \otimes denotes the matrix multiplication and \odot is the element-wise multiplication

This pseudo label generation process can be applied to all the segments. Therefore, we can obtain the segment-level visual pseudo label for each video, denoted as $\hat{Y}^v = \{\hat{y}_t^v\} \in \mathbb{R}^{T \times C}$. Note that the video-level visual pseudo label $\hat{y}^v \in \mathbb{R}^{1 \times C}$ can be easily obtained from \hat{Y}^v , where $\hat{y}_c^v = \mathbb{1}(\sum_{t=1}^T \hat{Y}_{t,c}^v)$ that means if a category of the event exists in at least one video segment, it is contained in the video-level label.

As for the audio pseudo labels, they can be generated in a similar way but with several adjustments. For brevity, we introduce the main steps here. (1) We use the CLAP model instead of the CLIP for audio pseudo label generation. (2) The audio waveform of the entire video is split into T equal-length segments and each segment is sent to the CLAP audio encoder. (3) We use the prompt “This sound contains the [CLS]” with the event categories as the input of CLAP text encoder. (4) We compute the similarity

score of the text and audio features extracted by CLAP (just like Eq. 3) and use an independent threshold τ_a (replace τ_v in Eq. 4) to select high similarity values. In this way, we obtain the segment-level audio pseudo label $\hat{Y}^a \in \mathbb{R}^{T \times C}$ and the video-level audio pseudo label $\hat{y}^a \in \mathbb{R}^{1 \times C}$ for each video sample.

4.2 Pseudo Label Exploitation (PLE)

The weakly-supervised AVVP task requires predicting for each segment, but only the video-level label is provided. This task would be greatly advanced if segment-level supervision is additionally provided. In this part, we try to exploit the pseudo labels from both the video-level and segment-level since we have obtained pseudo labels of these two levels, namely $\hat{y}^m \in \mathbb{R}^{1 \times C}$ and $\hat{Y}^m \in \mathbb{R}^{T \times C}$, where $m \in \{v, a\}$ denotes the modality type. In particular, for the segment-level

supervision, we propose a new richness-aware optimization objective to help the model align the predictions and pseudo labels. We introduce our pseudo label exploitation strategy in the two aspects below.

Basic video-level loss. Existing methods usually adopt the objective function formulated in Eq. 1 for model training (Wu & Yang, 2021; Yu et al., 2022; Cheng et al., 2022; Mo & Tian, 2022), where $\bar{y}^m \in \mathbb{R}^{1 \times C}$ is the video-level label obtained by label smoothing. Instead, we use the video-level pseudo label $\hat{y}^m \in \mathbb{R}^{1 \times C}$ generated by our PLG module as new supervision. The objective is then updated to,

$$\mathcal{L}_V = \mathcal{L}_{\text{bce}}(\mathbf{p}^{v \cup a}, \mathbf{y}^{v \cup a}) + \mathcal{L}_{\text{bce}}(\mathbf{p}^a, \hat{\mathbf{y}}^a) + \mathcal{L}_{\text{bce}}(\mathbf{p}^v, \hat{\mathbf{y}}^v). \quad (5)$$

New segment-level loss. With the segment-wise pseudo label $\hat{\mathbf{Y}}^m$, we propose a new richness-aware loss that is inspired by the following observations. (1) Each row of the segment-wise pseudo labels, e.g., $\hat{\mathbf{Y}}_t^m \in \mathbb{R}^{1 \times C}$, the t -th row of the pseudo label, indicates whether all the events appear in the t -th segment. For example, we show the visual pseudo label in Fig. 2-2, i.e., $\hat{\mathbf{Y}}^m$ where $m = v$. There are three visual events in the first segment, i.e., the *dog*, *rooster*, and *speech*, $\hat{\mathbf{Y}}_1^v = [1, 1, 1]$, while the last segment only contains one *rooster* event, i.e., $\hat{\mathbf{Y}}_T^v = [0, 1, 0]$. This reflects the richness of the event category in different segments that indicates *how many event categories exist in each segment*. Similarly, the audio pseudo label $\hat{\mathbf{Y}}^a$ tells the category richness of audio events. We define the **category richness** of t -th segment cr_t^m as the ratio of the category number of t -th segment to the total category number of the video, written as,

$$cr_t^m = \frac{\sum_{c=1}^C \hat{\mathbf{Y}}_{t,c}^m}{\sum_{c=1}^C \mathbf{y}_c^{v \cup a}}, \quad (6)$$

where $m \in \{v, a\}$ denotes the visual or audio modality. Therefore, we can obtain the category richness vector of all segments $\mathbf{cr}^m \in \mathbb{R}^{T \times 1}$ for each modality. In the example shown in Fig. 2-2, the visual category richness for the first and last segments, i.e., cr_1^v and cr_T^v , is equal to 1 and 1/3, respectively.

(2) On the other hand, each column of the pseudo labels, e.g., $\hat{\mathbf{Y}}_c^m \in \mathbb{R}^{T \times 1}$, $m \in \{v, a\}$, indicates *how many visual/audio segments contain the event of c -th category*. We denote the **segment richness** of c -th category sr_c^m as the ratio of the number of segments containing that category c to the total segment number of the video, written as below,

$$sr_c^m = \frac{1}{T} \sum_{t=1}^T \hat{\mathbf{Y}}_{t,c}^m. \quad (7)$$

In the example shown in Fig. 2-2, the visual segment richness for the event categories *dog* and *speech*, i.e., sr_1^v and sr_3^v is

equal to 1/2 and 1/4, respectively. Extending to all C event categories, we can obtain the segment richness vector of all the categories $\mathbf{sr}^m \in \mathbb{R}^{1 \times C}$, where $m \in \{v, a\}$ denotes the visual and audio modalities.

So far, regardless of modality $m \in \{v, a\}$, we can obtain the category richness \mathbf{cr}^m and segment richness \mathbf{sr}^m of the pseudo label. With the prediction $\mathbf{P}^m \in \mathbb{R}^{T \times C}$ from the baseline network, we can compute its category richness and segment richness in the same way, denoted as $\mathbf{pcr}^m \in \mathbb{R}^{T \times 1}$ and $\mathbf{psr}^m \in \mathbb{R}^{1 \times C}$. Then, we design the segment-level richness-aware loss \mathcal{L}_S to align the richness of the predictions and the pseudo labels, calculated by,

$$\mathcal{L}_S = \sum_{m \in \{v, a\}} \mathcal{L}_{\text{bce}}(\mathbf{pcr}^m, \mathbf{cr}^m) + \mathcal{L}_{\text{bce}}(\mathbf{psr}^m, \mathbf{sr}^m). \quad (8)$$

The total objective function $\mathcal{L}_{\text{total}}$ for AVVP in this work is the combination of the basic loss \mathcal{L}_V and the richness-aware loss \mathcal{L}_S , i.e.,

$$\mathcal{L}_{\text{total}} = \mathcal{L}_V + \lambda \mathcal{L}_S, \quad (9)$$

where λ is a weight parameter.

4.3 Pseudo Label Denoising (PLD)

In general, PLG can produce trustworthy segment-level pseudo labels, especially when combined with the proposed richness-aware loss, which significantly improves the audio-visual video parsing performance. This can be verified by our experiments shown in Sect. 5.3. Going a step further, we posit that the generated pseudo labels may still encompass some noise. By our observation, the video-level event category pseudo-annotation can be satisfactorily tackled, but the misclassification of specific segments exists along the timeline within each video, particularly when dealing with hard video instances that are difficult to annotate from the segment level. We specifically trace such challenges in the visual modality and observe that without contextual information, separate frames sent to the CLIP may be incorrectly classified, especially in the instances where the visual objects in the images are too diminutive, the images are afflicted by blurriness or inadequate lighting, when portions of the objects are obscured, rendering them arduous to discern, *etc.* As shown in Fig. 2, the *dog* at the last two segments is mostly obscured by the *rooster*, and CLIP fails to recognize the visual event *dog* without contextual information. In this case, the generated pseudo labels do not accurately capture the temporal boundary of the event and would be detrimental to model training. We believe that the segment-level visual pseudo labels can be further refined. As for the audio modality, the audio signal is represented through waveform and it keeps good continuity even if it is split into multiple segments for

pseudo-labeling. This characteristic may help to resist disturbances along the timeline when generating segment-level audio pseudo labels with CLAP. In fact, the quality of audio pseudo labels is indeed better than that of visual pseudo labels. For example, the segment-level F-score metric for audio pseudo labels is ~ 10 points higher than that of visual pseudo labels, as demonstrated in Tables 1, 2. This implies the high quality of audio pseudo labels produced by PLG and highlights the greater difficulty in enhancing the accuracy of visual pseudo labels. We present further discussions with more experimental results in Sect. 5.2.

In this section, we propose a pseudo label denoising (PLD) strategy that aims to recheck the pseudo labels generated by PLG and further refine the inaccurate ones (noisy pseudo labels). Our PLD is inspired by the works that conduct label denoising with the help of the forward propagation loss for image tasks (Kim et al., 2022; Hu et al., 2021a). In general, a large forward loss means that the trained model does not give the same prediction as the labels for a sample. There are two main reasons for this: (1) the provided label is correct but the video data is hard to learn and the model does not learn an effective representation for it; (2) the label itself is incorrect. In this work, our PLD aims to leverage the forward loss to check the temporal continuity of segment-level pseudo labels in each video and amend the abnormal segments when they belong to the second case.

Specifically, we first use the objective function shown in Eq. 9 to train a baseline model. Then, we compute the element-wise forward loss matrix by measuring the binary cross entropy between the prediction \mathbf{P}^m and the pseudo label $\hat{\mathbf{Y}}^m$, denoted as $\mathcal{M}^m = \mathcal{L}_{\text{bce}}(\mathbf{P}^m, \hat{\mathbf{Y}}^m) \in \mathbb{R}^{T \times C}$, where $m \in \{v, a\}$ denotes the visual and the audio modality. Denote the j -th column of \mathcal{M}^m as $\mathcal{M}_{:,j}^m \in \mathbb{R}^{T \times 1}$, it indicates the loss value of all segments for the specific j -th event category. In the example shown in Fig. 2-3, we display the forward loss matrix for the visual modality and find that the last two video segments have much larger forward losses than other segments for the *dog* category; they actually contain this event like other segments. The abnormally large loss value is caused by the fact that the last two segments are assigned incorrect visual pseudo labels. Therefore, the matrix \mathcal{M}^m can reflect those segments whose pseudo labels contain potential noise and require refinement.

Note that the pseudo label $\hat{\mathbf{y}}^m \in \mathbb{R}^{1 \times C}$ indicates the predicted event categories that appear in each modality. We trust the event category $\hat{\mathbf{y}}^m$ and use it to mask the matrix \mathcal{M}^m . There are two steps for the matrix \mathcal{M}^m masking. **Step I:** For other event categories that do not occur in the video sample, their pseudo labels will be eased by setting zeros in \mathcal{M}^m . For the example shown in Fig. 2-2, we only need to denoise the pseudo labels for the three columns of $\hat{\mathbf{y}}^m$ that corresponds to the predicted event categories of *dog*, *rooster* and *speech*.

The calculation of the masked matrix \mathcal{M}^m can be computed by,

$$\mathcal{M}^m = f_{\text{rpt-T}}(\hat{\mathbf{y}}^m) \odot \mathcal{M}^m, \quad (10)$$

where $\mathcal{M}^m \in \mathbb{R}^{T \times C}$, and $f_{\text{rpt-T}}(\hat{\mathbf{y}}^m)$ denotes the operation of repeating $\hat{\mathbf{y}}^m$ along the temporal dimension for T times, and $f_{\text{rpt-T}}(\hat{\mathbf{y}}^m) \in \mathbb{R}^{T \times C}$.

Step II: Returning to the masked forward loss of all video segments of the j -th category $\mathcal{M}_{:,j}^m \in \mathbb{R}^{T \times 1}$, we treat the average of the top- K smallest loss values of $\mathcal{M}_{:,j}^m$ as the threshold μ_j^m . μ_j^m is the tolerable forward loss within a video sample. If the loss of some segments is abnormally larger than μ_j^m , they may have incorrect pseudo labels. Comparing the forward loss of each segment with μ_j^m , we can obtain a binary mask vector $\phi_j^m \in \mathbb{R}^{T \times 1}$, where ‘1’ reflects that the segment has a larger loss than μ_j^m . This process can be written as,

$$\begin{cases} \mu_j^m = f_{\text{avg}}(f_{\text{k}}(\mathcal{M}_{:,j}^m)), \\ \phi_j^m = \mathbb{1}(\mathcal{M}_{:,j}^m - \alpha \cdot \mu_j^m), \end{cases} \quad (11)$$

where f_{k} and f_{avg} denotes the top- K minimum loss selection and the average operation, respectively. Note that we set a scaling factor α to magnify the averaged loss. It is used to better ensure that anomalous loss is caused by incorrect pseudo labels rather than the data not being well learned.

Extending Eq. 11 to all the event categories, we obtain the binary mask matrix of the video $\Phi^m = \{\phi_j^m\} \in \mathbb{R}^{T \times C}$. Afterwards, the segment-level pseudo label $\hat{\mathbf{Y}}^m$ produced by PLG can be refined by reversing the positions that have unusually large loss values reflected by Φ^m , denoted as $\tilde{\mathbf{Y}}^m = f_{\sim}(\hat{\mathbf{Y}}^m, \Phi^m)$. As shown in Fig. 2-2, for the event *dog* again, the visual pseudo labels generated by PLG are ‘0’ for the last two segments (indicating that there is no *dog*) and get a large forward loss (marked by the purple box in Fig. 2-3). This indicates that the visual pseudo labels of these two segments are incorrect (actually containing *dog*) and are thus reversed during the denoising process. We display more examples in Fig. 7 to illustrate the pseudo label denoising process. Finally, the pseudo labels refined by PLD can be taken as new supervision for the model training.

5 Experiments

5.1 Experimental Setup

Dataset. Experiments for the AVVP task are conducted on the publicly available *Look, Listen, and Parse (LLP)* (Tian et al., 2020) dataset. It contains 11,849 videos spanning over 25 common audio-visual categories, involving scenes such

as humans, animals, vehicles, musical instruments, *etc.* Each video is 10 s long and around 61% of the videos contain more than one event category. Videos of the LLP dataset are split into 10,000 for training, 649 for validation, and 1200 for testing. The training set is provided with only the video-level labels, i.e., the label union of the audio events and visual events. For validation and test sets, the segment-wise event labels for each audio and visual modality are additionally provided.

Evaluation metrics. Following existing works (Tian et al., 2020; Cheng et al., 2022; Wu & Yang, 2021; Yu et al., 2022), we evaluate our method by measuring the parsing results of all the types of events, namely audio events (**A**), visual events (**V**), and audio-visual events (**AV**, both audible and visible). The average parsing result of the three types is denoted as the “**Type@AV**” metric. Different from **Type@AV** metric, “**Event@AV**” metric calculates the F-score considering the predictions of the audio and the visual events together. For the above event types, both the segment-level and event-level F-scores are used as evaluation metrics. The segment-level metric measures the quality of the predicted events by comparing them with the ground truth for each video segment. And the event-level metric treats consecutive segments containing the same event category as a whole event, and computes the F-score based on $mIoU = 0.5$ as the threshold. Therefore, the event-level F-score metric is more difficult because it requires the model to predict a satisfactory temporal boundary of the event.

Implementation details. (1) *Feature extraction.* For the LLP dataset, each video is divided into 10 consecutive 1-second segments. For a fair comparison, we adopt the same feature extractors to extract the audio and visual features. Specifically, the VGGish (Hershey et al., 2017) network pretrained on AudioSet (Gemmeke et al., 2017) dataset is used to extract the 128-dim audio features. The pretrained ResNet152 (He et al., 2016) and R(2+1)D (Tran et al., 2018) are used to extract the 2D and 3D visual features, respectively. The low-level visual feature is the concatenation of 2D and 3D visual features. (2) *Pseudo label preparation.* For each video in the training set of the LLP dataset, we first offline generate the segment-wise visual and audio pseudo labels using our PLG module. We use the ViT-B/32-based CLIP (Vaswani et al., 2017) and HTSAT-RoBERTa-based CLAP (Wu et al., 2023) to conduct the pseudo label generation, and their parameters are frozen. (3) *Training procedure.* The objective function \mathcal{L}_{total} shown in Eq. 9 is used to train the baseline model HAN (Tian et al., 2020). The hyperparameter λ in Eq. 9 for balancing the video-level and the segment-level losses is empirically set to 0.5. This pretrained model is then used in our PLD to further refine the pseudo labels. The refined pseudo labels are used to supervise the baseline model training again. For all the training processes, we adopt the Adam optimizer to train the model with a mini-batch size

Table 1 Parameter study of the threshold τ_v and prompt used in the VISUAL pseudo label generation

Parameter setup		Precision	Segment. (V)	Event. (V)
τ_v	prompt			
–	–	66.96	58.65	53.48
0.040	VP1	85.31	70.29	64.68
0.041		86.88	71.08	64.82
0.042		72.19	51.51	43.13
0.041	VP1	86.88	71.08	64.82
	VP2	85.64	68.96	61.83
	VP3	84.69	67.60	60.98
	VP4	86.75	70.29	63.78

Different setups are used to generate segment-level pseudo labels; consequently, we can obtain the corresponding video-level pseudo labels. Here, we report the precision between the visual pseudo label and the ground truth from the video level. Also, we report the segment-level and event-level F-scores. ‘–’ denotes the result of directly assigning video labels as the visual event labels and each event happens at all the visual segments. The specific expressions of the prompts are introduced in our main text. This experiment is conducted on the validation set of the LLP dataset

of 32 and the learning rate of 3×10^{-4} . The total training epoch is set to 30. All experiments are conducted with PyTorch (Paszke et al., 2019) on a single NVIDIA GeForce RTX-2080-Ti GPU. The codes, pseudo labels, and pretrained models will be released.

5.2 Parameter Studies

We perform parameter studies of essential parameters used in our method, namely the score threshold τ_v/τ_a and the text prompt for CLIP/CLAP used in the PLG module, and the top- K and scaling factor α used in the PLD strategy. Experiments in this section are conducted on the validation set of the LLP dataset of which the segment-level event labels are accessible. Thus, we also verify the quality of pseudo labels through *correctness measurements* in this part.

Study of the thresholds and prompts in PLG. τ_v/τ_a is the threshold to select high scores of the cosine similarity between the event category and the visual/audio segment in the mask calculation (Eq. 4). We first explore the impact of τ_v on the **visual pseudo label generation**. As shown in the upper part of Table 1, we used the default prompt **VP1**—“This photo contains the [CLS]” and test several values of τ_v to generate visual pseudo labels. Then, we report the category precision between the pseudo labels and the ground truth at the video level, and the segment-level and event-level F-scores to measure the quality of the generated pseudo labels. As shown in the Table, the pseudo label with the best quality is obtained when $\tau_v = 0.041$. And all the evaluation metrics drop significantly when τ_v changes from 0.041 to 0.042. We argue such sensitivity is related to

Table 2 Parameter study of the threshold τ_a and prompt used in the AUDIO pseudo label generation

Parameter setup		Segment. (A)	Event. (A)	Average
τ_a	prompt			
–	–	77.07	63.84	70.45
0.037	AP1	79.79	70.77	75.28
0.038		80.01	70.87	75.28
0.039		80.23	71.27	75.75
0.040		80.18	71.70	75.44
0.037	AP2	80.06	70.74	75.40
0.038		80.32	71.54	75.93
0.039		80.20	71.00	75.60
0.040		80.03	69.91	74.97

Different setups are used to generate segment-level audio pseudo labels. Here, we report the segment-level and event-level F-scores between the audio pseudo label and the ground truth. The last column shows the average value of these two evaluation metrics, which is used to select the best setup. ‘–’ denotes the result of directly treating the video labels as the audio event labels and each event happens at all the audio segments. The specific expressions of the prompts are introduced in our main text. This experiment is conducted on the validation set of the LLP dataset

the *softmax* operation in Eq. 3 that squeezes the similarity score into small logits. The metrics for visual modality are acceptable up to the threshold of $\tau_v = 0.041$. Using the same experimental strategy, we explore the impacts of threshold τ_a in **audio pseudo label generation**. The experimental results are shown in Table 2 and we find that the optimal audio pseudo labels are obtained when τ_a is equal to 0.038.

Furthermore, we explore the impact of prompts used in the PLG. The prompts are combined with the event categories and sent as text inputs to the CLIP or CLAP text encoder. **For the visual pseudo label generation**, specifically, we test four types of prompts, i.e., our default **VP1**—“This photo contains the [CLS]”, **VP2**—“This photo contains the scene of [CLS]”, **VP3**—“This photo contains the visual scene of [CLS]” and **VP4**—“This is a photo of the [CLS]”. We use these different prompts to generate pseudo labels and compare them with the ground truth. As shown in the lower part of Table 1, visual pseudo labels generated using these different prompts remain relatively consistent. The pseudo label has the highest F-score using the **VP1** prompt. Therefore, we use the prompt **VP1** as the default setup for visual pseudo label generation in our following experiments. Notably, the precision of the video-level visual pseudo label reaches about 87% under the optimal setup, whereas the precision of directly assigning video labels as the visual event labels (i.e., without prompt) is only $\sim 67\%$. This reveals that PLG can satisfactorily disentangle visual events from weak video labels. **For the audio pseudo label generation**, we test two prompts,

Table 3 Parameter study of the K and scaling factor α used in the VISUAL pseudo label denoising

Parameter setup		Segment. (V)	Event. (V)	Average
K	α			
–	–	71.08	64.82	67.95
4	30	72.45	67.82	70.13
5		72.99	68.28	70.63
6		72.17	66.90	69.53
5	20	72.85	68.10	70.47
	30	72.99	68.28	70.63
	40	72.82	68.12	70.47

Different values of K and α are tested for the segment-wise visual pseudo label denoising. The segment-level and event-level F-scores of the denoised visual pseudo labels are reported. The last column is the average result. ‘–’ denotes the result of the visual pseudo label generated by PLG without label denoising. This experiment is conducted on the validation set of the LLP dataset

i.e., the **AP1**—“This is a sound of [CLS]” and **AP2**—“This sound contains the [CLS]”, to generate segment level audio pseudo labels. Then, we report the segment-level and event-level F-scores of the audio events under different setups and use their average value to select the best one. As shown in the Table 2, performances moderately change under different setups, and the best performance is obtained when using the **AP2** prompt and τ_a equals 0.038. We thereby use this optimal setup as the default for audio pseudo label generation. It is noteworthy that the event-level F-score is only around 64% if simply assigning the video labels to all the audio segments (without prompt). In contrast, this metric is around 72% for our generated audio pseudo labels. This reveals the vital role of segment-level event identification.

Study of the K and α in PLD. For each predicted event category, the top- K smallest forward loss along the temporal dimension is magnified by α and used as the threshold to determine which segments’ pseudo labels should be refined (Eq. 11). *The segment-level and event-level F-scores of the events are used to evaluate the quality of the denoised pseudo labels.* For the visual pseudo label denoising, the results in Table 3 indicate that denoised visual pseudo labels ensure significantly better results than the original labels generated by PLG. In particular, the event-level F-score is improved by 3.46%. Observing Table 3, the optimal setup are $K = 5$ and $\alpha = 30$. Under this setup, the segment-level and event-level F-scores of the visual pseudo labels of the validation set achieve 72.99% and 68.28%, respectively. For the audio pseudo label denoising, as shown in Table 4, the denoised audio pseudo labels are slightly better than the pseudo labels generated by PLG under the optimal setup ($K = 6$, $\alpha = 400$). As discussed in Sect. 4.3, PLD is proposed to alleviate the potentially discontinuous pseudo-event labels that happened during PLG and provide better temporal boundaries of

Table 4 Parameter study of the K and scaling factor α used in the AUDIO pseudo label denoising

Parameter setup		Segment. (A)	Event. (A)	Average
K	α			
–	–	80.32	71.54	75.93
5	400	79.63	70.88	75.25
6		80.43	71.68	76.06
7		80.15	71.33	75.74
6	300	80.16	71.27	75.72
	400	80.43	71.68	76.06
	500	80.40	71.27	75.72

Different values of K and α are tested for the segment-wise audio pseudo label denoising. The segment-level and event-level F-scores of the denoised audio pseudo labels are reported. The last column is the average result. ‘–’ denotes the result of the audio pseudo label generated by PLG without label denoising. This experiment is conducted on the validation set of the LLP dataset

the events. We argue that the discontinuity of pseudo labels of audio events rarely occurs due to the temporal characteristics of audio data, thus leading to a slight improvement for the audio modality as shown in Table 4. Besides, from Tables 3 and 4, we observe an interesting phenomenon that the segment-level and event-level F-scores of audio pseudo labels without PLD (80.32% and 71.54%) remain superior to those of the denoised visual pseudo labels (72.99% and 68.28%). This suggests the high quality of audio pseudo labels generated by PLG and underscores the greater difficulty in denoising visual pseudo labels. We ultimately strike a balance between the second computational costs and denoising improvements and refrain from applying PLD to the audio modality in our experiment setup.

5.3 Ablation Studies

In this section, we provide some ablation studies to explore the impact of each module in our method. The experimental results are shown in Table 5. The row with id-① denotes the performance of the baseline HAN (Tian et al., 2020).

Impact of the PLG. To further verify the benefits of PLG, we use the generated pseudo labels to supervise the model training. Note that the vanilla HAN (id-① in Table 5) is trained with the video-level pseudo label obtained by using label smoothing on the given weak label (Eq. 2). For a fair comparison, we only use the video-level pseudo labels generated by PLG as the model supervision (Eq. 5). As shown in row-② of Table 5, utilizing the video-level pseudo label generated by our PLG significantly improves the visual event parsing performances. The visual metric (V) increases from 52.9 to 64.1% at the segment level while from 48.9 to 60.2% at the event level. These improvements reflect that our PLG generates more accurate video-level pseudo labels for the

visual modality, better distinguishing the event categories and guiding the model training. The improvement in audio event parsing is not pronounced in this situation. We anticipate that the temporally continuous audio segments are more challenging to distinguish under weak video-level supervision. Additionally, the visual features can encapsulate more distinct event semantics, thereby promoting model optimization that is more beneficial to the visual modality. Even so, the utilization of more fine-grained, segment-level pseudo labels generated by our PLG (see ids ③ and ④ in Table 5) significantly enhances both the audio and visual event parsing performances.


Our PLG is able to generate high-quality pseudo labels at the segment level, which can be verified by the results shown in Tables 1 and 2. In Fig. 3a, we further display the event-level F-scores of the generated audio and visual pseudo labels of each event category and provide more discussions. As seen, the audio and visual pseudo labels have satisfactory F-scores for most of the categories. The highest F-score is 93.5% for audio event *Accordion* and 91.7% for visual event *Blender*, respectively. Besides, we also find that each modality faces some intractable event categories, such as the *speech* for visual modality and *cat* for audio modality. We argue this is caused by the unbalanced data distribution and some categories are particularly difficult for visual recognition, such as *speech*, *cheering*, and *clapping*. Nevertheless, our PLG generally provides reliable audio and visual pseudo labels from both the video level and segment level, ensuring better model learning.

Impact of the PLE. Our PLE uses the proposed richness-aware loss \mathcal{L}_S in Eq. 8 to exploit the pseudo labels from segment-level, which is taken as a complement to the video-level supervision. At first, we make an ablation study to explore the effect of the respective richness component. As shown in Table 6, “SR” and “CR” denote the segment richness loss and category richness loss between the predictions and pseudo labels, respectively. From Table 6, we can find that each of them can effectively improve the model performance since the studied AVVP task requires distinguishing both the video segments and the event categories. When both types of richness information are used, the pseudo labels fully demonstrate the capability for model optimization. To further validate its superiority, we compare it with a native variant that directly computes the binary cross entropy loss between the predictions and the pseudo labels, denoted as $\mathcal{L}'_S = \mathcal{L}_{\text{bce}}(\mathbf{P}^v, \hat{\mathbf{Y}}^v) + \mathcal{L}_{\text{bce}}(\mathbf{P}^a, \hat{\mathbf{Y}}^a)$. As shown in the row-③ and ④ of Table 5, both \mathcal{L}'_S and the proposed \mathcal{L}_S are beneficial for the audible video parsing since they all provide segment-level supervision. Nevertheless, the proposed RL loss is more helpful. The conventional cross-entropy loss relies on ‘hard’ segment-wise alignments between predictions and pseudo labels. In contrast, our proposed richness-aware loss exploits the pseudo labels


Table 5 Ablation study of the main modules

Id	Main modules			Segment-level					Event-level				
	PLG	PLE	PLD	A	V	AV	Type@AV	Event@AV	A	V	AV	Type@AV	Event@AV
①	×	×	×	60.1	52.9	48.9	54.0	55.4	51.3	48.9	43.0	47.7	48.0
②	✓	×	×	59.8	64.1	57.5	60.5	58.3	50.8	60.2	50.7	53.9	49.3
③	✓	✓- \mathcal{L}'_S	×	61.5	64.7	58.6	61.6	60.0	54.5	61.0	52.4	55.9	52.7
④	✓	✓- \mathcal{L}_S	×	61.2	65.8	59.1	62.0	60.2	54.8	62.4	52.6	56.6	53.3
⑤	✓	✓	✓	62.4	66.7	60.3	63.1	61.4	55.7	63.3	53.7	57.6	54.3

Id-① denotes the performance of the baseline backbone HAN (Tian et al., 2020). \mathcal{L}_S is the proposed richness-aware loss (Eq. 8). \mathcal{L}'_S is a native loss that simply computes the binary cross entropy loss of the prediction and pseudo label. We report the results on the test set of the LLP dataset



(a) Event-level F-scores of the audio and visual pseudo labels generated by PLG for each event category



(b) Event-level F-scores of the visual pseudo labels obtained by PLG and PLD for each event category

Fig. 3 Event-level F-scores of pseudo labels for each event category. **a** We display the event-level F-scores of audio and visual pseudo labels generated by PLG. **b** Compared to PLG, PLD further improves the

event-level F-scores for most categories, providing more accurate visual pseudo labels. All the results are reported on the validation set of the LLP dataset

by aligning predictions from two *independent* dimensions: category-richness (*cr*) and segment-richness (*sr*). According to the definitions of *cr* (Eq. 6) and *sr* (Eq. 7), their values are expressed as percentages (*‘soft’ ratios*) and are independent. This design makes the model trained with our richness-aware loss automatically balance and utilize the soft supervisions from category-richness and segment-richness. Experimental results shown in Table 6 indicate the superiority of our flexible design of richness-aware loss.

Impact of the PLD. The impact of PLD can be observed from two aspects. On one hand, PLD provides more accurate pseudo labels than PLG. As the quality measurement of visual pseudo labels shown in Table 3 on the validation set, the average F-score is 67.95% for PLG while it is 70.63% for PLD. In Fig. 3b, we show event-level F-scores for the visual pseudo labels obtained by PLG and PLD of each event category. PLD further improves the F-scores for most categories (18/25), e.g., the metrics for events *Fire alarm* and *Blender*

increase substantially by 18.2% and 6.8%, respectively. On the other hand, visual pseudo labels generated by PLD are more helpful than PLG for model training. We update the visual pseudo labels as the new supervision to train the HAN model. As shown in row-⑤ of Table 5, the model has superior

Table 6 Richness-aware loss \mathcal{L}_S under different configurations

Loss \mathcal{L}_S		Segment-level		Event-level	
\mathcal{SR}	\mathcal{CR}	Type@AV	Event@AV	Type@AV	Event@AV
×	×	60.5	58.3	53.9	49.3
×	✓	61.8	60.2	56.4	52.9
✓	×	61.3	59.6	56.1	52.6
✓	✓	62.0	60.2	56.6	53.3

\mathcal{SR} and \mathcal{CR} denote that we only compute \mathcal{L}_S with the segment richness and category richness alignment, respectively

Table 7 Comparison with the state-of-the-arts

Method	Segment-level					Event-level				
	A	V	AV	Type@AV	Event@AV	A	V	AV	Type@AV	Event@AV
AVE (Tian et al., 2018)	47.2	37.1	35.4	39.9	41.6	40.4	34.7	31.6	35.5	36.5
AVSDN (Lin et al., 2019)	47.8	52.0	37.1	45.7	50.8	34.1	46.3	26.5	35.6	37.7
HAN (Tian et al., 2020)	60.1	52.9	48.9	54.0	55.4	51.3	48.9	43.0	47.7	48.0
MM-Pyramid (Yu et al., 2022)	60.9	54.4	50.0	55.1	57.6	52.7	51.8	44.4	49.9	50.5
MGN (Mo & Tian, 2022)	60.8	55.4	50.4	55.5	57.2	51.1	52.4	44.4	49.3	49.1
CVCMS (Lin et al., 2021)	59.2	59.9	53.4	57.5	58.1	51.3	55.5	46.2	51.0	49.7
DHHN (Jiang et al., 2022)	61.3	58.3	52.9	57.5	58.1	54.0	55.1	47.3	51.5	51.5
▲MA (Wu & Yang, 2021)	60.3	60.0	55.1	58.9	57.9	53.6	56.4	49.0	53.0	50.6
▲JoMoLD (Cheng et al., 2022)	61.3	63.8	57.2	60.8	59.9	53.9	59.9	49.6	54.5	52.5
▲VAPLAN (ours)	62.4	66.7	60.3	63.1	61.4	55.7	63.3	53.7	57.6	54.3
★VAPLAN (ours)	69.0	70.2	63.5	67.6	67.9	61.9	66.4	56.9	61.7	60.1

▲ represents these methods are all focused on generating better pseudo labels for the AVVP task and are all developed on the baseline HAN (Tian et al., 2020) backbone. ★ denotes we further implement our method with the more advanced visual and audio features extracted by CLIP and CLAP, respectively. Results are reported on the test set of the LLP dataset

Table 8 Generalization of our method on other audio-visual video parsing backbones

Method	Segment-level					Event-level				
	A	V	AV	Type@AV	Event@AV	A	V	AV	Type@AV	Event@AV
MGN (Mo & Tian, 2022)	<u>60.8</u>	55.4	50.4	55.5	57.2	51.1	52.4	44.4	49.3	49.1
MGN + MA	60.2	61.9	55.5	59.2	58.7	<u>50.9</u>	59.7	49.6	53.4	<u>49.9</u>
MGN + PLG	60.1	<u>63.3</u>	<u>56.5</u>	<u>60.0</u>	<u>58.9</u>	50.3	<u>60.9</u>	<u>50.2</u>	<u>53.8</u>	49.4
MGN + PLD	61.0	64.3	57.1	60.8	60.1	51.1	61.9	50.6	54.5	50.4
MM-Pyramid (Yu et al., 2022)	60.9	54.4	50.0	55.1	57.6	52.7	51.8	44.4	49.9	50.5
MM-Pyramid + MA	61.1	60.3	55.8	59.7	59.1	53.8	56.7	<u>49.4</u>	54.1	51.2
MM-Pyramid + PLG	60.2	<u>65.4</u>	<u>58.3</u>	<u>61.3</u>	<u>60.1</u>	<u>54.5</u>	<u>62.0</u>	52.8	<u>56.4</u>	<u>53.0</u>
MM-Pyramid + PLD	<u>61.0</u>	66.4	58.5	62.0	60.9	55.0	63.0	52.8	56.9	53.4

Our method can generate reliable segment-level audio and visual pseudo labels which can be directly used for other methods in the AVVP task too. We evaluate two representative backbones, namely the MGN (Mo & Tian, 2022) and MM-Pyramid (Yu et al., 2022). The pseudo labels generated by our PLG and refined by our PLD consistently boost these models. Both PLG and PLD are also superior to the existing method MA (Wu & Yang, 2021) that provides video-level pseudo labels. The best and second-best results of each evaluation metric are **bold** and underlined, respectively

Table 9 Generalization of our method on the weakly-supervised audio-visual event localization task

Method	label objective	
	video-level	segment-level (ours)
AVEL	67.1	69.2 _(+2.1)
PSP	72.1	74.3 _(+2.2)
CMBS	72.2	74.4 _(+2.2)

Given the only video-level event label, this task needs to localize the temporal video segments that contain the audio-visual event, i.e., the audio and visual segments simultaneously describe the same event. We extend our pseudo label generation strategy to this task and generate segment-level event labels. We test several SOTA models on this task, namely AVEL (Tian et al., 2018), PSP (Zhou et al., 2021), and CMBS (Xia & Zhao, 2022). All of them can be further improved using our segment-level pseudo labels as the objective. This experiment is conducted on the AVE (Tian et al., 2018) dataset

performance on all types of event parsing. This again reveals that the visual pseudo labels obtained by PLD are more accurate than by PLG and can better supervise the multi-modal parsing model. These results verify the effectiveness of the label denoising strategy in PLD.

5.4 Comparison with the State-of-the-Arts

We report the performance of our VAPLAN on the test set of the LLP dataset. The comparison results with existing methods are shown in Table 7. Our method achieves superior performance on all types of event parsing. **First**, compared to the baseline HAN (Tian et al., 2020) on which our method is developed, our method significantly improves the performance. Especially for the visual event parsing (V in the table), the segment-level metric is lifted from 52.9% to 66.7% (\uparrow 13.8%), and the event-level metric is improved from 48.9% to 63.3% (\uparrow 14.4%). **Second**, our method outperforms other competitors on the track of generating pseudo labels for the AVVP task. As shown in the low part of Table 7, our method generally exceeds the previous state-of-the-art JoMoLD (Cheng et al., 2022) by about 1.5 points for the audio event parsing, and around 3 points for the visual event and audio-visual event parsing. Both JoMoLD (Cheng et al., 2022) and MA (Wu & Yang, 2021) generate audio-visual pseudo labels from the video level, while our method can provide audio-visual pseudo labels from a more fine-grained segment level. Our video parsing model can be better supervised and optimized, resulting in better performance. **Furthermore**, we report the result of our method using the visual and audio features respectively extracted by CLIP and CLAP. As shown in the last row of Table 7, all types of event parsing performance can be further significantly improved. In particular, the audio event parsing benefits more from such advanced feature representations. As shown, its performance improves by 6.6% and 6.2% for the segment-level and event-level F-scores, respectively. These improvements demonstrate the effectiveness and superiority of our method.

5.5 Generalization of Our Method

Generalization on other AVVP backbones. A core contribution of our method is that it can provide high-quality segment-level audio and visual pseudo labels, which then better guide the model optimization. Our method can also be applied to other existing backbones in the AVVP task. To explore its impact, we examine two recently proposed networks, i.e., MGN (Mo & Tian, 2022) and MM-Pyramid (Yu et al., 2022). Specifically, we train the models using the pseudo labels generated by our PLG and refined by our PLD, respectively. The experimental results are shown in Table 8. Both PLG and PLD significantly boost the vanilla models, especially in the visual event and audio-visual event parsing.

Take the MM-Pyramid (Yu et al., 2022) method for example, the segment-level visual event parsing performance is improved from 54.4% to 65.4% and 66.4% by using our PLG and PLD, respectively. PLD is superior due to the additional label denoising strategy. Such improvements can also be observed for MGN (Mo & Tian, 2022). Besides, it is worth noting that these two backbones perform better when combined with our (segment-level) pseudo labels than the (video-level) pseudo labels generated by the previous method MA (Wu & Yang, 2021). These results again indicate that our method is able to provide better fine-grained pseudo labels and demonstrate the superiority and generalizability of our method.

Generalization on the AVEL task. We also extend our pseudo label generation strategy to a related audio-visual event localization (AVEL) task. We explore the challenging weakly-supervised setting where the model needs to localize those video segments containing the audio-visual events (an event is both audible and visible) given only the video-level event category label. Previous AVEL methods merely use the known video-level labels as the objective for model training. Here we try to generate segment-level pseudo labels for this task as we did for the weakly-supervised AVVP task. Similarly, we use the pretrained CLIP and CLAP models to generate segment-level visual and audio pseudo labels, respectively. The audio-visual event pseudo labels are the intersection of them. In this way, we know if there is an audio-visual event in each video segment. Then such segment-level pseudo labels can be used as a new objective to supervise the model training. We test three representative audio-visual event localization methods whose official codes are available, namely the AVEL (Tian et al., 2018), PSP (Zhou et al., 2021) and CMBS (Xia & Zhao, 2022). We conduct experiments on the corresponding AVE (Tian et al., 2018) dataset and the results are shown in Table 9. The second column shows the performance of vanilla models with only the video-level supervision. The last column shows that these models can be significantly improved by around 2 points when using our segment-level pseudo labels.

We also present some qualitative examples for a more intuitive comparison. As shown in Fig. 4a, the audio-visual event *church bell* occurs exclusively in the last five video segments. The previous state-of-the-art method, CMBS, incorrectly assumes this event to be present in the first five segments as well. In contrast, our method yields accurate localization results. The reason is that vanilla CMBS relies solely on the known weak event label (video-level) to supervise model training, while our method is capable of generating high-quality pseudo labels at the segment level. In the lower part of Fig. 4a, we illustrate our pseudo label generation process. Our method accurately identifies that the *church bell* event exists in all the visual segments but is present only in the last five audio segments, which results in the precise audio-visual

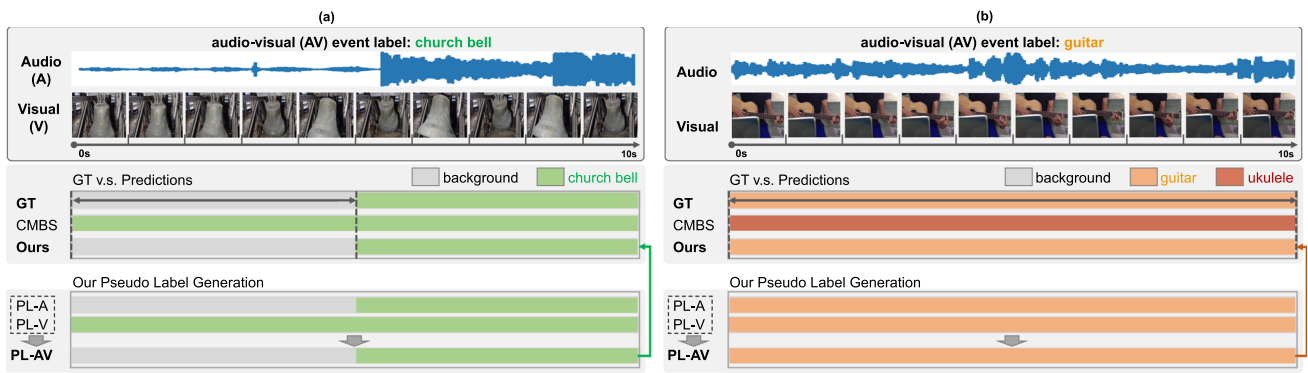


Fig. 4 Qualitative examples for the weakly-supervised audio-visual event localization task. This task aims to temporally locate those segments containing events that are both audible and visible. The previous state-of-the-art method, CMBS (Xia & Zhao, 2022), utilizes only the video-level weak labels for model training and predictions. In contrast, our method can generate high-quality segment-level pseudo labels, offering fine-grained supervision during training and producing more

accurate localization results. “GT” denotes the ground truth. “PL-A” and “PL-V” represent our segment-level pseudo labels for the audio and visual modalities, respectively. The audio-visual event pseudo labels (“PL-AV”) result from the intersection of “PL-A” and “PL-V”. Our method surpasses the vanilla CMBS model in distinguishing between the background and audio-visual events (a) as well as among different audio-visual event categories (b)

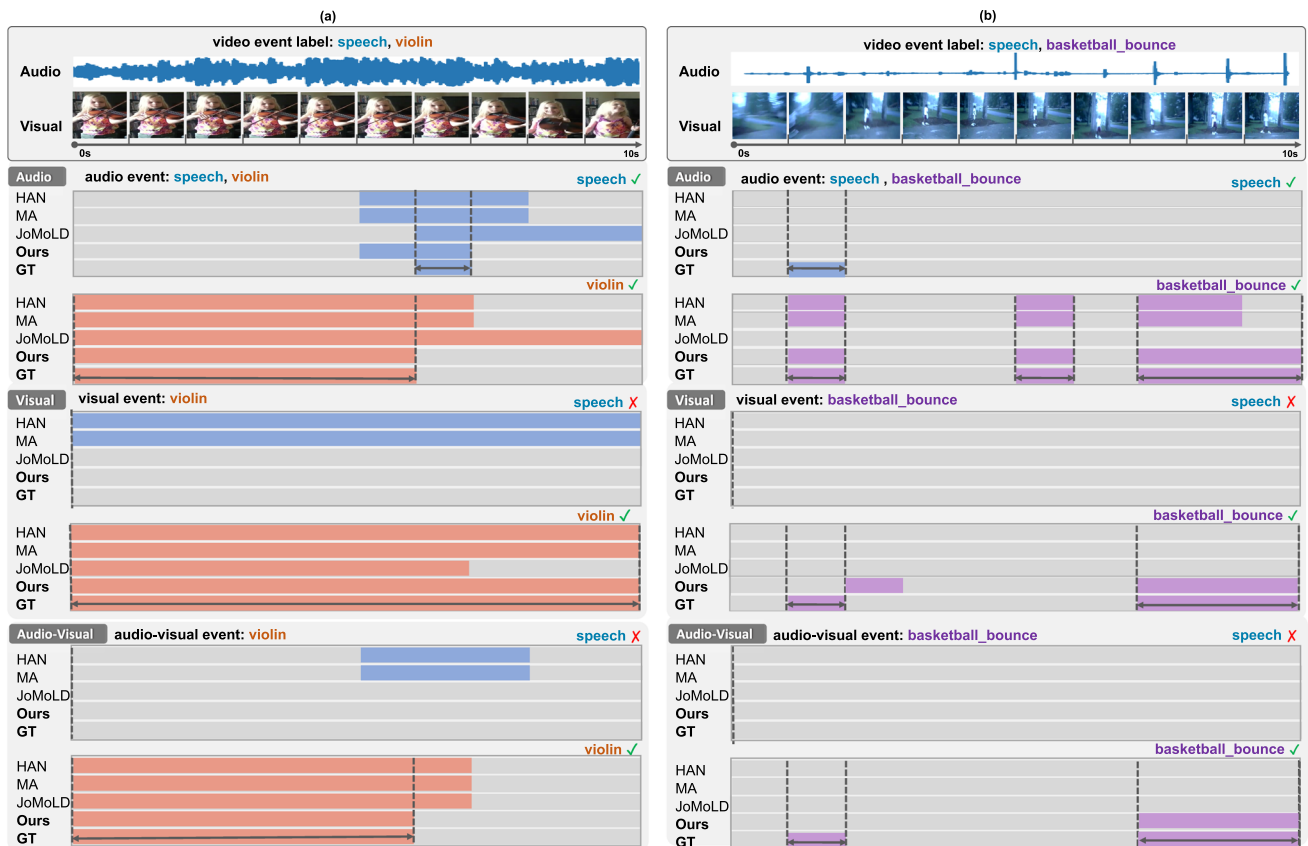


Fig. 5 Qualitative examples of the audio-visual video parsing using different methods. We compare our method with the HAN (Tian et al., 2020), MA (Wu & Yang, 2021) and JoMoLD (Cheng et al., 2022). “GT” denotes the ground truth. Our method successfully recognizes

that there is only one visual event *violin* in (a) or *basketball bounce* in (b). Our method is also more accurate in parsing the audio events and audio-visual events, providing better temporal boundaries of the events

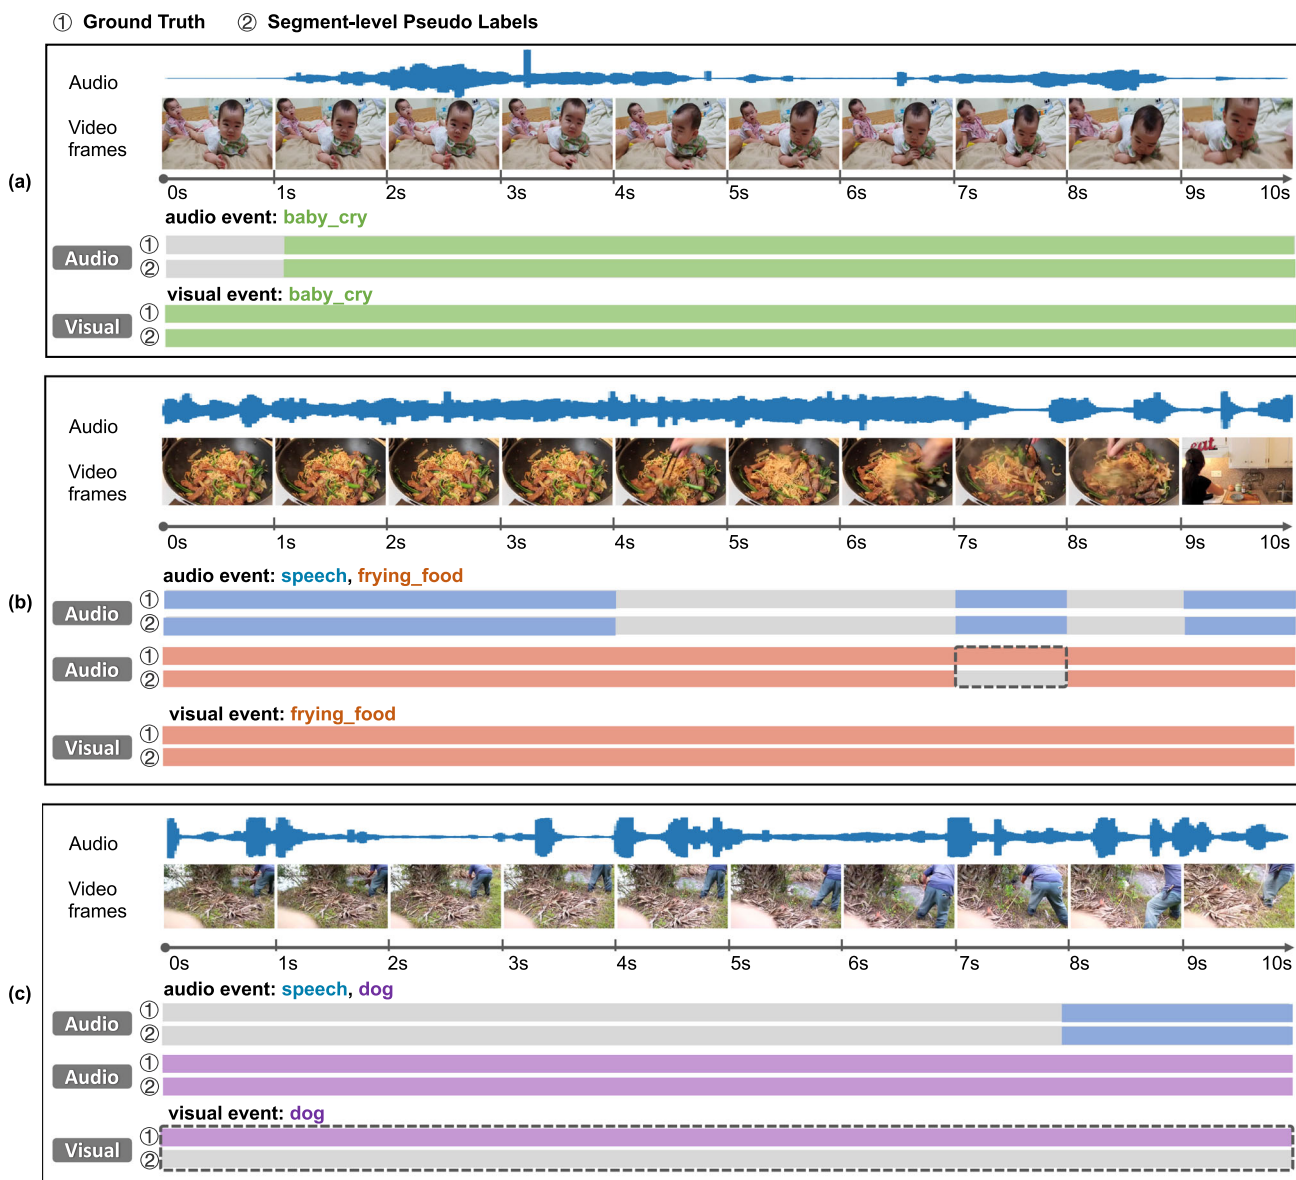


Fig. 6 Typical and challenging visualization examples of the generated audio and visual pseudo labels. “①” and “②” denote the ground truth and the obtained pseudo labels, respectively. **a** In these typical cases where the events are clearly represented in the audio and visual signals, our method can generate accurate segment-level pseudo labels.

event pseudo label and then better supervises the model training and predictions. Similar benefits can also be observed from Fig. 4b, the vanilla CMBS incorrectly classifies the audio-visual event *guitar* to be the *ukulele*. In contrast, our method can generate accurate segment-level pseudo labels, thereby ensuring superior predictions. These results again verify the generalization of our method and we believe our method can also help to address other related audio-visual tasks lacking fine-grained supervision.

We also display some challenging examples: the audio event is mixed with other sounds (b) or the visual event is hard to perceive (c). In general, our method can provide satisfactory audio and visual pseudo labels

5.6 Qualitative Results on the AVVP Task

Visualization examples of the audio-visual video parsing.

We first display some qualitative video parsing examples in Fig. 5. We compare our method with HAN (Tian et al., 2020), MA (Wu & Yang, 2021), and JoMoLD (Cheng et al., 2022). Both MA and JoMoLD are developed on the HAN and try to generate video-level pseudo labels for better model training. As shown in Fig. 5a, two events exist in the video, i.e., *speech* and *violin*, while the visual event only contains the

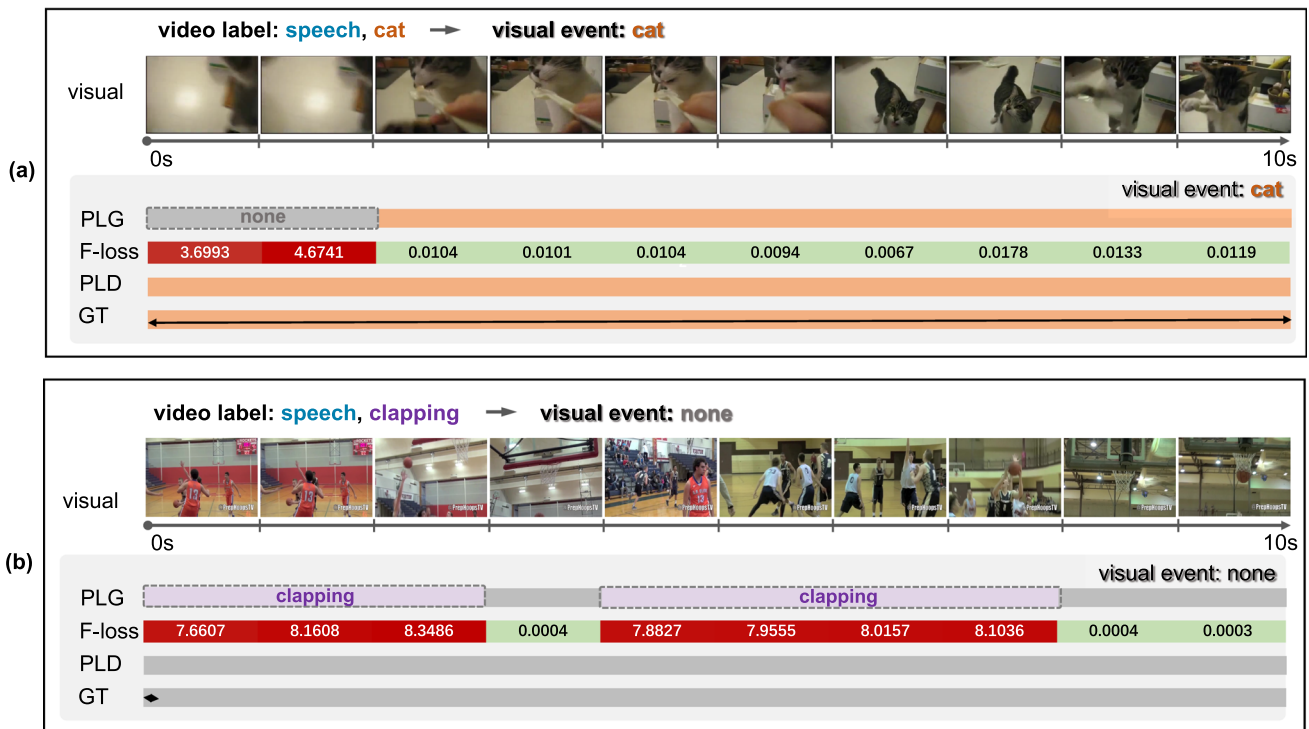


Fig. 7 Qualitative visualization examples of the pseudo label denoising. Here, we take the visual modality as an example since it faces more challenges in both pseudo label generation and denoising processes. “GT” denotes the ground truth. “F-loss” represents the forward loss between the model predictions and the pseudo labels generated by PLG (Eq. 10). PLG basically disentangles the visual event(s) from

the weak video label, yielding well-defined segment-wise event categories. Additionally, PLD helps alleviate potential label noise for those segments along the timeline in the same video whose pseudo labels generated by PLG suffer abnormally large loss values. The improved labels are highlighted by the dotted box

violin. For audio event parsing, although all methods correctly recognize the two events occurring in the audio track, our method locates more exact temporal segments. Also, our method accurately recognizes the visual event *violin* and provides superior audio-visual event parsing. In Fig. 5b, both the events *speech* and *basketball bounce* exist in the video. All methods miss the audio event *speech*. The reason may be that the *speech* event only happens in the second segment and the audio signal contains some noise from outdoors. It is hard to distinguish them. For visual and audio-visual event parsing, only our method provides satisfactory prediction for the audio event *basketball bounce*. Although our method incorrectly identifies that the third segment contains this event, we argue that there may be an annotation mistake. The basketball player in this segment is clearer than in the second segment. If true, our result is more correct. These video samples demonstrate the superiority of our method, which leverages high-quality segment-level pseudo labels to better supervise model training.

Visualization examples of the obtained pseudo labels.

In this part, we display the pseudo labels of some typical and challenging video samples. Our method is able to provide high-quality segment-level audio and visual pseudo labels.

As shown in Fig. 6a, the *baby cry* event is clearly represented in the video and our method successfully recognizes it in both audio and visual tracks. The temporal boundaries of the generated pseudo labels highly match the ground truth. Our method performs well in handling similar cases with explicit audio and visual event signals. Turning to Fig. 6b, our method generates accurate pseudo labels for the visual event *frying food* and audio event *speech*. The audio event *frying food* in the eighth segment is not identified. The difficulty is that the sound of *frying food* is mixed with the louder sound of *speech*, which causes the *frying food* event to be missed. The compound audio classification is still a challenging task in the community. In Fig. 6c, our method satisfactorily generates segment-level pseudo labels for all the audio events but fails to recognize the visual event *dog*. The *dog* in the visual frames is too small (located around the man’s feet in the figure) to be identified. This situation is hard to judge even for a human annotator. The pseudo labels can be further explored in the future if considering more specific techniques for these challenging cases. Nevertheless, our method can generally provide reliable segment-level pseudo labels.

Visualization of the pseudo label denoising. As shown in Fig. 7, we show two visualization examples to reflect the

impact of pseudo label denoising. Here, we take the more challenging visual pseudo label denoising as an example. As shown in Fig. 7a, the video-level label contains the events of *speech* and *cat*, where *speech* does not exist in the visual modality. PLG successfully recognizes that only *cat* event happens in the visual track. However, since the object is too blurry in the first two segments, the event *cat* is incorrectly recognized. As a result, the forward loss values for these two segments are significantly greater, possibly 300 to 400 times larger than the other segments, as shown in the Fig. 7a. Contributing to the proposed label denoising (PLD) strategy, we make the correction. Observing Fig. 7b, there are no visual events. PLG mistakenly classifies a few segments as the event *clapping* because the player's movements are complex in these segments. This inaccuracy is once again evident through the abnormally high forward losses. PLD also rectifies these erroneous pseudo labels. By analysis, the pseudo labels generated by PLG rely on the prior knowledge of event categories from the pretrained CLIP, while PLD benefits from an additional revision process (—the joint exploration of the predictions and pseudo labels through the forward loss calculation in each video) to possibly correct inaccurate segment-level pseudo labels in PLG.

6 Conclusion

We propose a Visual-Audio Pseudo Label exploration (VAPLAN) method for the weakly-supervised audio-visual video parsing task. VAPLAN is a new attempt to generate segment-level pseudo labels in this field, which starts with a pseudo label generation module that uses the reliable CLIP and CLAP models to determine the visual events and audio events occurring in each modality (at the segment level) as pseudo labels. We then exploit the category richness and segment richness contained in the pseudo labels and propose a new richness-aware loss as fine-grained supervision for the AVVP task. Furthermore, we propose a pseudo label denoising strategy to refine the visual pseudo labels and better guide the predictions. Qualitative and quantitative experimental results on the LLP dataset corroborate that our method can effectively generate and exploit high-quality segment-level pseudo labels. All these proposed techniques can be directly used in the community. We also extend our method to a related weakly-supervised audio-visual event localization task and the experimental results verify the effectiveness and generalization of our method. We believe this work will not only facilitate future research on the studied audio-visual video parsing task but also inspire other related audio-visual topics seeking better supervision.

Acknowledgements We would like to thank Dr. Liang Zheng for his constructive suggestions. We also sincerely appreciate the anonymous reviewers for their positive feedback and professional comments.

Data Availability The LLP dataset for the studied audio-visual video parsing is publicly available from the official website <https://github.com/YapengTian/AVVP-ECCV20>. The AVE dataset for the audio-visual event localization task can be accessed at <https://github.com/YapengTian/AVE-ECCV18>. Tables 1–9 and figures 3–7 were generated with our source codes, which will be released at our GitHub repository <https://github.com/jasongief/VPLAN>.

References

- Afouras, T., Owens, A., Chung, J. S., & Zisserman, A. (2020). Self-supervised learning of audio-visual objects from video. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 208–224).
- Alayrac, J. B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al. (2022). Flamingo: A visual language model for few-shot learning. [arXiv:2204.14198](https://arxiv.org/abs/2204.14198)
- Arandjelovic, R., & Zisserman, A. (2017). Look, listen and learn. In *Proceedings of the IEEE international conference on computer vision (ICCV)* (pp. 609–617).
- Arandjelovic, R., & Zisserman, A. (2018). Objects that sound. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 435–451).
- Barraco, M., Cornia, M., Cascianelli, S., Baraldi, L., & Cucchiara, R. (2022). The unreasonable effectiveness of clip features for image captioning: An experimental analysis. In *Workshops of proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 4662–4670).
- Chao, Y. W., Vijayanarasimhan, S., Seybold, B., Ross, D. A., Deng, J., & Sukthankar, R. (2018). Rethinking the faster r-cnn architecture for temporal action localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 1130–1139).
- Cheng, H., Liu, Z., Zhou, H., Qian, C., Wu, W., & Wang, L. (2022). Joint-modal label denoising for weakly-supervised audio-visual video parsing. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 431–448).
- Deng, J., Dong, W., Socher, R., Li, L. J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 248–255).
- Ding, J., Xue, N., Xia, G. S., & Dai, D. (2022). Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 11583–11592).
- Gao, J., Chen, M., & Xu, C. (2022). Fine-grained temporal contrastive learning for weakly-supervised temporal action localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 19999–20009).
- Gao, J., Chen, M., & Xu, C. (2023). 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 (CVPR)* (pp. 18827–18836).
- Gemmeke, J. F., Ellis, D. P., Freedman, D., Jansen, A., Lawrence, W., Moore, R. C., Plakal, M., & Ritter, M. (2017). Audio set: An ontology and human-labeled dataset for audio events. In *IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 776–780).
- Gong, Y., Chung, Y. A., & Glass, J. (2021). Ast: Audio spectrogram transformer. [arXiv:2104.01778](https://arxiv.org/abs/2104.01778)

- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 770–778).
- 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. (2017). Cnn architectures for large-scale audio classification. In *IEEE International conference on acoustics, speech and signal processing (ICASSP)* (pp. 131–135).
- Hu, D., Nie, F., & Li, X. (2019). Deep multimodal clustering for unsupervised audiovisual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 9248–9257).
- Hu, D., Qian, R., Jiang, M., Tan, X., Wen, S., Ding, E., Lin, W., & Dou, D. (2020). Discriminative sounding objects localization via self-supervised audiovisual matching. In *Advances in neural information processing systems (NeurIPS)* (pp. 10077–10087).
- Hu, P., Peng, X., Zhu, H., Zhen, L., Lin, J. (2021a). Learning cross-modal retrieval with noisy labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 5403–5413).
- Hu, Z., Yang, Z., Hu, X., Nevatia, R. (2021b). Simple: Similar pseudo label exploitation for semi-supervised classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 15099–15108).
- Huang, J., Qu, L., Jia, R., & Zhao, B. (2019). O2u-net: A simple noisy label detection approach for deep neural networks. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)* (pp. 3326–3334).
- Jia, C., Yang, Y., Xia, Y., Chen, Y. T., Parekh, Z., Pham, H., Le, Q., Sung, Y. H., Li, Z., & Duerig, T. (2021). Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning (ICML)* (pp. 4904–4916).
- Jiang, X., Xu, X., Chen, Z., Zhang, J., Song, J., Shen, F., Lu, H., & Shen, H. T. (2022). Dhhn: Dual hierarchical hybrid network for weakly-supervised audio-visual video parsing. In *Proceedings of the ACM international conference on multimedia (ACM MM)* (pp. 719–727).
- Kahn, J., Lee, A., & Hannun, A. (2020). Self-training for end-to-end speech recognition. In *IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 7084–7088).
- Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., & Fei-Fei, L. (2014). Large-scale video classification with convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 1725–1732).
- Kim, Y., Kim, J. M., Akata, Z., & Lee, J. (2022). Large loss matters in weakly supervised multi-label classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 14156–14165).
- Kong, Q., Xu, Y., Wang, W., & Plumbley, M. D. (2018). Audio set classification with attention model: A probabilistic perspective. In *IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 316–320).
- Kumar, A., Khadkevich, M., & Fügen, C. (2018). Knowledge transfer from weakly labeled audio using convolutional neural network for sound events and scenes. In *IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 326–330).
- Lamba, J., Akula, J., Dabral, R., Jyothi, P., Ramakrishnan, G., et al. (2021). Cross-modal learning for audio-visual video parsing. [arXiv:2104.04598](https://arxiv.org/abs/2104.04598)
- Li, G., Wei, Y., Tian, Y., Xu, C., Wen, J. R., & Hu, D. (2022). Learning to answer questions in dynamic audio-visual scenarios. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 19108–19118).
- Li, Z., Guo, D., Zhou, J., Zhang, J., & Wang, M. (2023). Object-aware adaptive-positivity learning for audio-visual question answering. In *Proceedings of the AAAI conference on artificial intelligence (AAAI)* (pp. 1–10).
- Lin, Y. B., Li, Y. J., & Wang, Y. C. F. (2019). Dual-modality seq2seq network for audio-visual event localization. In *IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 2002–2006).
- Lin, Y. B., Tseng, H. Y., Lee, H. Y., Lin, Y. Y., & Yang, M. H. (2021). Exploring cross-video and cross-modality signals for weakly-supervised audio-visual video parsing. In *Advances in neural information processing systems (NeurIPS)* (pp. 11449–11461).
- Liu, H., Chen, Z., Yuan, Y., Mei, X., Liu, X., Mandic, D., Wang, W., & Plumbley, M. D. (2023a). Audioldm: Text-to-audio generation with latent diffusion models. [arXiv:2301.12503](https://arxiv.org/abs/2301.12503)
- Liu, X., Kong, Q., Zhao, Y., Liu, H., Yuan, Y., Liu, Y., Xia, R., Wang, Y., Plumbley, M. D., & Wang, W. (2023b). Separate anything you describe. [arXiv:2308.05037](https://arxiv.org/abs/2308.05037)
- Long, X., Gan, C., De Melo, G., Wu, J., Liu, X., & Wen, S. (2018a). Attention clusters: Purely attention based local feature integration for video classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 7834–7843).
- Long, X., Gan, C., Melo, G., Liu, X., Li, Y., Li, F., & Wen, S. (2018b). Multimodal keyless attention fusion for video classification. In *Proceedings of the AAAI conference on artificial intelligence (AAAI)* (pp. 1–8).
- Ma, C., Yang, Y., Wang, Y., Zhang, Y., & Xie, W. (2022). Open-vocabulary semantic segmentation with frozen vision-language models (pp. 1–21). [arXiv:2210.15138](https://arxiv.org/abs/2210.15138)
- Mahmud, T., & Marculescu, D. (2022). Ave-clip: Audioclip-based multi-window temporal transformer for audio visual event localization. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision (WACV)* (pp. 1–10).
- Mo, S., & Tian, Y. (2022). Multi-modal grouping network for weakly-supervised audio-visual video parsing. In *Advances in neural information processing systems (NeurIPS)*
- Pan, Y., Hu, Y., Yang, Y., Yao, J., Fei, W., Ma, L., & Lu, H. (2023). Gemo-clap: Gender-attribute-enhanced contrastive language-audio pretraining for speech emotion recognition. [arXiv:2306.07848](https://arxiv.org/abs/2306.07848)
- Park, D. S., Zhang, Y., Jia, Y., Han, W., Chiu, C. C., Li, B., Wu, Y., & Le, Q. V. (2020). Improved noisy student training for automatic speech recognition (pp. 1–5). [arXiv:2005.09629](https://arxiv.org/abs/2005.09629)
- Pasi, P. S., Nemani, S., Jyothi, P., & Ramakrishnan, G. (2022). Investigating modality bias in audio visual video parsing. [arXiv:2203.16860](https://arxiv.org/abs/2203.16860)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems (NeurIPS)* (pp. 1–12).
- Patel, G., Allebach, J. P., & Qiu, Q. (2023). Seq-ups: Sequential uncertainty-aware pseudo-label selection for semi-supervised text recognition. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision (WACV)* (pp. 6180–6190).
- Pham, H., Dai, Z., Xie, Q., & Le, Q. V. (2021). Meta pseudo labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 11557–11568).
- Qian, R., Hu, D., Dinkel, H., Wu, M., Xu, N., & Lin, W. (2020). Multiple sound sources localization from coarse to fine. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 292–308).
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. (2021). Learning transferable visual models from natural language supervision. In *International conference on machine learning (ICML)* (pp. 8748–8763).

- Rao, V., Khalil, M. I., Li, H., Dai, P., & Lu, J. (2022a). Dual perspective network for audio-visual event localization. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 689–704).
- Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., & Lu, J. (2022b). Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 18082–18091).
- Rizve, M. N., Duarte, K., Rawat, Y. S., & Shah, M. (2021). In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning (pp. 1–20). [arXiv:2101.06329](https://arxiv.org/abs/2101.06329)
- Rouditchenko, A., Zhao, H., Gan, C., McDermott, J., & Torralba, A. (2019). Self-supervised audio-visual co-segmentation. In *IEEE international conference on acoustics* (pp. 2357–2361). IEEE: Speech and Signal Processing (ICASSP).
- Senocak, A., Oh, T. H., Kim, J., Yang, M. H., & Kweon, I. S. (2018). Learning to localize sound source in visual scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)* (pp. 4358–4366).
- Shen, X., Li, D., Zhou, J., Qin, Z., He, B., Han, X., Li, A., Dai, Y., Kong, L., Wang, M., et al. (2023). Fine-grained audible video description. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 10585–10596).
- Song, P., Guo, D., Zhou, J., Xu, M., & Wang, M. (2022). Memorial gan with joint semantic optimization for unpaired image captioning. *IEEE Transactions on Cybernetics*, 4388–4399.
- Sun, W., Zhang, J., Wang, J., Liu, Z., Zhong, Y., Feng, T., Guo, Y., Zhang, Y., Barnes, N. (2023). Learning audio-visual source localization via false negative aware contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 6420–6429).
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 2818–2826).
- Tang, M., Wang, Z., Liu, Z., Rao, F., Li, D., & Li, X. (2021). Clip4caption: Clip for video caption. In *Proceedings of the ACM international conference on multimedia (ACM MM)* (pp. 4858–4862).
- Tian, Y., Li, D., & Xu, C. (2020). Unified multisensory perception: Weakly-supervised audio-visual video parsing. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 436–454).
- Tian, Y., Shi, J., Li, B., Duan, Z., & Xu, C. (2018). Audio-visual event localization in unconstrained videos. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 247–263).
- Tran, D., Wang, H., Torresani, L., Feiszli, M. (2019). Video classification with channel-separated convolutional networks. In *Proceedings of the IEEE international conference on computer vision (ICCV)* (pp. 5552–5561).
- Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., & Paluri, M. (2018). A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 6450–6459).
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems (NeurIPS)* (pp. 1–11).
- Wang, H., Zha, Z. J., Li, L., Chen, X., & Luo, J. (2023). Context-aware proposal-boundary network with structural consistency for audio-visual event localization. In *IEEE transactions on neural networks and learning systems* (pp. 1–11).
- Wang, Z., Lu, Y., Li, Q., Tao, X., Guo, Y., Gong, M., Liu, T. (2022). Cris: Clip-driven referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 11686–11695).
- Wei, Y., Hu, D., Tian, Y., & Li, X. (2022). Learning in audio-visual context: A review, analysis, and new perspective. [arXiv:2208.09579](https://arxiv.org/abs/2208.09579)
- Wu, Y., & Yang, Y. (2021). Exploring heterogeneous clues for weakly-supervised audio-visual video parsing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 1326–1335).
- Wu, Y., Chen, K., Zhang, T., Hui, Y., Berg-Kirkpatrick, T., & Dubnov, S. (2023). Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 1–5).
- Wu, Y., Zhu, L., Yan, Y., & Yang, Y. (2019). Dual attention matching for audio-visual event localization. In *Proceedings of the IEEE international conference on computer vision (ICCV)* (pp. 6292–6300).
- Wu, Y., Zhang, X., Wang, Y., & Huang, Q. (2022). Span-based audio-visual localization. In *Proceedings of the ACM international conference on multimedia (ACM MM)* (pp. 1252–1260).
- Xia, Y., & Zhao, Z. (2022). Cross-modal background suppression for audio-visual event localization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 19989–19998).
- Xie, Q., Luong, M. T., Hovy, E., & Le, Q. V. (2020). Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 10687–10698).
- Xu, H., Zeng, R., Wu, Q., Tan, M., & Gan, C. (2020). Cross-modal relation-aware networks for audio-visual event localization. In *Proceedings of the ACM international conference on multimedia (ACM MM)* (pp. 3893–3901).
- Xu, M., Zhang, Z., Wei, F., Lin, Y., Cao, Y., Hu, H., & Bai, X. (2021). A simple baseline for zero-shot semantic segmentation with pre-trained vision-language model. [arXiv:2112.14757](https://arxiv.org/abs/2112.14757)
- Yalniz, I. Z., Jégou, H., Chen, K., Paluri, M., & Mahajan, D. (2019). Billion-scale semi-supervised learning for image classification. [arXiv:1905.00546](https://arxiv.org/abs/1905.00546)
- Yang, P., Wang, X., Duan, X., Chen, H., Hou, R., Jin, C., & Zhu, W. (2022). Avqa: A dataset for audio-visual question answering on videos. In *Proceedings of the 30th ACM international conference on multimedia (ACM MM)* (pp. 3480–3491).
- Yu, J., Cheng, Y., Zhao, R. W., Feng, R., & Zhang, Y. (2022). Mm-pyramid: Multimodal pyramid attentional network for audio-visual event localization and video parsing. In *Proceedings of the ACM international conference on multimedia (ACM MM)* (pp. 6241–6249).
- Yun, H., Yu, Y., Yang, W., Lee, K., & Kim, G. (2021). Pano-avqa: Grounded audio-visual question answering on 360deg videos. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)* (pp. 2031–2041).
- Zeng, R., Huang, W., Tan, M., Rong, Y., Zhao, P., Huang, J., & Gan, C. (2019). Graph convolutional networks for temporal action localization. In *Proceedings of the IEEE international conference on computer vision (ICCV)* (pp. 7094–7103).
- Zhou, C., Loy, C. C., Dai, B. (2022a). Extract free dense labels from clip. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 696–712).
- Zhao, H., Gan, C., Rouditchenko, A., Vondrick, C., McDermott, J., & Torralba, A. (2018). The sound of pixels. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 570–586).
- Zhou, J., Guo, D., Wang, M. (2023a). Contrastive positive sample propagation along the audio-visual event line. In *IEEE transactions on pattern analysis and machine intelligence (TPAMI)*.

- Zhou, J., Shen, X., Wang, J., Zhang, J., Sun, W., Zhang, J., Birchfield, S., Guo, D., Kong, L., Wang, M., et al. (2023b). Audio-visual segmentation with semantics. [arXiv:2301.13190](https://arxiv.org/abs/2301.13190)
- Zhou, J., Wang, J., Zhang, J., Sun, W., Zhang, J., Birchfield, S., Guo, D., Kong, L., Wang, M., & Zhong, Y. (2022b). Audio-visual segmentation. In *Proceedings of the European conference on computer vision (ECCV)* (pp. 386–403).
- Zhou, J., Zheng, L., Zhong, Y., Hao, S., & Wang, M. (2021). Positive sample propagation along the audio-visual event line. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)* (pp. 8436–8444).
- Zhou, Z., Zhang, B., Lei, Y., Liu, L., & Liu, Y. (2022c). Zeg-clip: Towards adapting clip for zero-shot semantic segmentation. [arXiv:2212.03588](https://arxiv.org/abs/2212.03588)
- Zhu, Z., Tang, W., Wang, L., Zheng, N., & Hua, G. (2021). Enriching local and global contexts for temporal action localization. In *Proceedings of the IEEE international conference on computer vision (ICCV)* (pp. 13516–13525).
- Zoph, B., Ghiasi, G., Lin, T. Y., Cui, Y., Liu, H., Cubuk, E. D., & Le, Q. (2020). Rethinking pre-training and self-training. In *Advances in neural information processing systems (NeurIPS)* (pp. 3833–3845).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.