

CLASP: Cross-modal Salient Anchor-based Semantic Propagation for Weakly-supervised Dense Audio-Visual Event Localization

Jinxing Zhou^{1*}, Ziheng Zhou^{2*}, Yanghao Zhou⁴, Yuxin Mao⁵, Zhangling Duan³, Dan Guo^{2,3†}

¹Mohamed Bin Zayed University of Artificial Intelligence

²Hefei University of Technology

³Hefei Comprehensive National Science Center

⁴National University of Singapore

⁵OpenNLP Lab

Abstract

The Dense Audio-Visual Event Localization (DAVEL) task aims to temporally localize events in untrimmed videos that occur simultaneously in both the audio and visual modalities. This paper explores DAVEL under a new and more challenging weakly-supervised setting (W-DAVEL task), where only video-level event labels are provided and the temporal boundaries of each event are unknown. We address W-DAVEL by exploiting *cross-modal salient anchors*, which are defined as reliable timestamps that are well predicted under weak supervision and exhibit highly consistent event semantics across audio and visual modalities. Specifically, we propose a *Mutual Event Agreement Evaluation* module, which generates an agreement score by measuring the discrepancy between the predicted audio and visual event classes. Then, the agreement score is utilized in a *Cross-modal Salient Anchor Identification* module, which identifies the audio and visual anchor features through global-video and local temporal window identification mechanisms. The anchor features after multimodal integration are fed into an *Anchor-based Temporal Propagation* module to enhance event semantic encoding in the original temporal audio and visual features, facilitating better temporal localization under weak supervision. We establish benchmarks for W-DAVEL on both the UnAV-100 and ActivityNet1.3 datasets. Extensive experiments demonstrate that our method achieves state-of-the-art performance.

Introduction

In recent years, audio-visual learning (Shen et al. 2023; Kryklyvets et al. 2025; Zhou et al. 2024b, 2025a; Liu et al. 2025; Li et al. 2025, 2024; Mao et al. 2024; Zhou et al. 2025c; Guo, Huang, and Zhou 2024; Zhao et al. 2025) has received increasing attention due to its potential to jointly model and reason over auditory and visual modalities, facilitating a more comprehensive understanding of complex real-world environments. Various audiovisual learning tasks have been widely explored, such as audio-visual event perception (Tian et al. 2018; Tian, Li, and Xu 2020; Zhou et al. 2024a), sound source localization (Arandjelovic and Zisserman 2017), audio-visual segmentation (Zhou et al. 2022, 2023b, 2025b; Jin et al. 2025;

*These authors contributed equally.

†Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

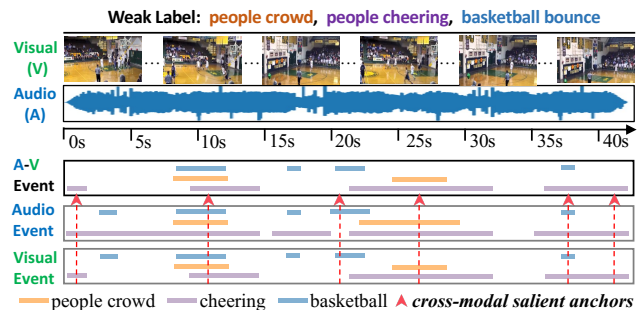


Figure 1: Illustration of the Weakly-supervised Dense Audio-Visual Event Localization (W-DAVEL) task. The W-DAVEL task aims to temporally localize events that co-occur in both audio and visual modalities, using only weak supervision from video-level event labels.

Guo et al. 2023), and more. However, the majority of existing works focus on short-term audio-visual scene understanding.

Recently, Geng *et al.* (Geng et al. 2023) proposed the *Dense Audio-Visual Event Localization (DAVEL)* task, aiming to understand long, untrimmed videos. As shown in Fig. 1, given a long video with its audio track, the DAVEL task aims to temporally localize events occurring in both audio and visual modalities, *i.e.*, the intersection set of the audio events and visual events. Each video usually contains multiple events, and each event may occur at multiple temporal positions, which makes DAVEL a highly challenging task. To date, several studies have attempted to address this problem, with most focusing on more effective audio-visual fusion (Geng et al. 2023; Xing et al. 2024; Zhou et al. 2025d). Although significant progress has been made by these approaches, they are all developed under a *fully-supervised* setting, where the specific start and end timestamp labels of each audio-visual event are provided. However, manually obtaining such fine-grained annotations is time-consuming, labor-intensive, and difficult to scale to large datasets. Therefore, the study of DAVEL under weak supervision emerges as an urgent and practical direction.

In this paper, we introduce the **Weakly-supervised Dense Audio-Visual Event Localization (W-DAVEL)** task, which aims to temporally localize audio-visual events using *only*

video-level labels as supervision. As shown in Fig. 1, given the weak label of ‘people crowd, people cheering, basketball bounce’, the model needs to predict not only the event categories but also their corresponding temporal boundaries. In other words, for each timestamp, the model is expected to predict the dense audio-visual event categories. In the W-DAVEL task, the weak label can provide explicit supervision indicating the presence of events in the video. However, a critical challenge is *how to handle the absence of category supervision along the temporal timeline*. We address this challenge by distributing the event semantics, derived from reliable timestamps guided by video-level supervision, across the temporal dimension through a learnable network. Specifically, even with only video-level event labels, there exist temporal timestamps where the model can confidently infer the audio-visual event categories. We define such timestamps as *cross-modal salient anchors*, which contain the most consistent event semantics across both audio and visual modalities and are distinguishable from the remaining temporal segments. We then model the relationships between temporal features and these anchor features to determine event semantics across the entire timeline.

Based on this idea, we propose a *Cross-modal salient Anchor-based Semantic Propagation (CLASP)* method. Specifically, to identify the salient anchors, we first introduce a *Mutual Event Agreement Evaluation (MEAE)* module. Modality-specific classifiers independently predict audio and visual event probabilities. MEAE measures the discrepancy between the predicted audio and visual event probabilities using the Jensen-Shannon Divergence, from which a mutual event agreement score is obtained. This score reflects the model’s confidence in the consistency of audiovisual predictions at each timestamp. Based on this agreement score, we design a *Cross-modal Salient Anchor Identification (CSAI)* module, which identifies the anchor timestamps through two mechanisms: Global Anchor Identification selects anchors with the top- K agreement scores from a global comparison across the entire video, while Local Anchor Identification follows a similar process but is restricted to multiple local temporal windows. CSAI is applied to both audio and visual modalities. The identified audio and visual anchor features are fused into compact multimodal anchor features, which encode explicit audio-visual event semantics. Then, an *Anchor-based Temporal Propagation (ATP)* module leverages the anchor features to enhance temporal audio and visual features, propagating event semantics from anchors to fine-grained temporal timestamps, facilitating event classification along timeline.

We evaluate our method and establish benchmarks for the W-DAVEL task through comparisons with prior methods from related tasks on the UnAV-100 (Geng et al. 2023) and ActivityNet1.3 (Heilbron et al. 2015) datasets. Experimental results demonstrate that our method outperforms existing approaches by a large margin. For instance, our method surpasses the previous state-of-the-art (SOTA) method CCNet (Zhou et al. 2025d) by 3.1% mAP on UnAV-100 dataset.

In summary, our main contributions are:

- We explore the practical and challenging W-DAVEL task, which aims to perform dense audio-visual event localiza-

tion using only video-level event labels as supervision.

- We propose a salient anchor-based CLASP method, which consists of three core modules: Mutual Event Agreement Evaluation, Cross-modal Salient Anchor Identification, and Anchor-based Temporal Propagation.
- We establish benchmarks for the W-DAVEL task on the UnAV-100 and ActivityNet1.3 datasets, achieving new state-of-the-art results compared to previous methods.

Related Work

Audio-Visual Event Perception (AVEP) aims to temporally localize events in specific categories that occur in audible videos. According to the types of the target events, existing research can be categorized into two branches: Audio-Visual Event Localization (AVEL) (Tian et al. 2018) and Audio-Visual Video Parsing (AVVP) (Tian, Li, and Xu 2020). The AVEL task focuses on identifying the audio-visual events that occur simultaneously in both modalities, while the AVVP task requires further distinguishing events occurring in each individual modality, *i.e.*, the audio events and visual events. To date, many methods for these tasks can operate under weakly supervised settings, where only the video-level event labels are provided. From a model architecture perspective, mainstream methods emphasize audiovisual interaction through various cross-attention fusion strategies (Tian et al. 2018; Xu et al. 2020; Ramaswamy 2020; Mahmud and Marculescu 2023; Yu et al. 2022; Mo and Tian 2022; Jiang et al. 2022), positive segment discovery (Zhou et al. 2021; Zhou, Guo, and Wang 2023; Gao, Chen, and Xu 2023), and background suppression (Xia and Zhao 2022; He et al. 2024). To improve the weak supervision, some methods aim to obtain modality-specific weak labels by denoising the available video label (Wu and Yang 2021; Cheng et al. 2022), or generate temporally specific pseudo-labels (Yung-Hsuan Lai 2023; Zhou et al. 2023a, 2024c; Fan et al. 2023) by utilizing external large foundation models, such as CLIP (Radford et al. 2021) and CLAP (Wu et al. 2023). Although these prior works have made notable progress in addressing weakly supervised AVEP tasks, they are limited to trimmed video scenarios. This study aims to fill the gap in audio-visual event perception for long, untrimmed videos under weak supervision.

Dense Audio-Visual Event Localization (DAVEL) aims to densely identify and localize audio-visual events in untrimmed videos. The pioneering work (Geng et al. 2023) proposes a strong baseline alongside the UnAV-100 benchmark. A cross-modal pyramid transformer encoder is employed to fuse audiovisual features at various temporal scales, and a temporal dependency modeling module is designed to capture correlations between multiple event categories. Subsequently, Xing *et al.* (Xing et al. 2024) leverage the local temporal continuity of audio-visual events to filter out irrelevant cross-modal signals and enhance cross-modal fusion via an adaptive attention mechanism. Zhou *et al.* (Zhou et al. 2025d) introduce CCNet, which improves audiovisual feature fusion via cross-modal consistency and collaborative modeling across multiple temporal granularities. These methods are developed under the fully supervised setting, which requires fine-grained temporal labels for effective model training. In

addition, several approaches, specifically designed for short-term AVEP tasks, such as PPAE (Gao, Chen, and Xu 2025) and NREP (Jiang et al. 2025), also evaluate their models on the UnAV-100 dataset, yet their performances on long-term event localization are far from satisfactory. In contrast, we provide a formal analysis of the weakly supervised DAVEL task and address it by exploring cross-modal salient anchors, achieving superior temporal localization results.

Methodology

Problem Statement

Given a video with T segments at equal intervals, let the t -th visual and audio segments be denoted as V_t and A_t ($t = \{1, \dots, T\}$), respectively. During the training process, only the video-level event label $\mathbf{y} = \{y_i\}_{i=1}^C \in \mathbb{R}^{1 \times C}$ is provided. Here, C is the total number of classes of audio-visual events, and y_i takes a value from $\{0, 1\}$, where $y_i = 1$ indicates that the i -th category exists in the video (but the specific timestamp is unknown). During the evaluation phase, the segment-level ground truth $Y = \{y_n = (t_n^s, t_n^e, c_n)\}_{n=1}^N$ is available for evaluation, where N is the total number of events occurring in the video, t_n^s and t_n^e are the start and end timestamps of the n -th audio-visual event, respectively, and c_n denotes the corresponding event category.

Therefore, a W-DAVEL model aims to predict the start and end timestamps for each identified event. Prior methods (Geng et al. 2023; Zhou et al. 2025d) for fully supervised DAVEL typically formulate the task as a joint category classification and timestamp regression problem. However, the regression subproblem is challenging to solve for the W-DAVEL task due to the lack of explicit timestamp supervision. Therefore, we formulate the W-DAVEL task as an explicit event classification problem, in which the model estimates event category probabilities $\mathbf{p}_{av} \in \mathbb{R}^{T \times C}$.

Framework Overview

An overview of the proposed framework is illustrated in Fig. 2. It consists of five main modules. (1) In the **Unimodal Feature Preparation**, modality-specific pretrained backbones are used to extract preliminary features for both audio and visual modalities. (2) Based on the extracted features, the **Mutual Event Agreement Evaluation (MEAE)** module measures the semantic consistency of event predictions between the audio and visual modalities. An event agreement score is obtained by comparing the predicted audio and visual event probabilities. (3) Given the event agreement score, the **Cross-modal Salient Anchor Identification (CSAI)** module employs two mechanisms to identify salient anchors along the temporal axis. The salient anchors refer to timestamps with the most explicit and consistent event categories across audio and visual modalities, which correspond to large values in the event agreement score. Specifically, Global Anchor Identification directly selects anchors from the entire video, while Local Anchor Identification performs fine-grained anchor selection within multiple, local temporal windows. The anchor audio and visual features identified by both mechanisms are fused to form compact multimodal anchor representations. (4) These fused features subsequently enhance the intra-modal

temporal audio and visual features via the **Anchor-based Temporal Propagation (ATP)**. (5) Finally, the updated audio and visual features are sent to the **Classification** module for event predictions. We elaborate on each module in the next subsections.

Unimodal Feature Preparation

Given the audio and visual segments $\{A_t, V_t\}_{t=1}^T$, we follow prior work (Geng et al. 2023) to extract initial features. Specifically, VGGish (Hershey et al. 2017), pretrained on the AudioSet (Gemmeke et al. 2017) dataset, is employed to extract audio features, denoted as $\mathbf{a} \in \mathbb{R}^{T \times d}$, where d is the channel dimension. The two-stream I3D (Carreira and Zisserman 2017), pretrained on Kinetics (Kay et al. 2017), is employed to extract visual features, denoted as $\mathbf{v} \in \mathbb{R}^{T \times d}$. Here, the audio and visual features are projected into a shared feature space using two convolutional layers with ReLU.

Mutual Event Agreement Evaluation

Our method handles W-DAVEL task by highlighting cross-modal salient anchors, *i.e.*, timestamps that exhibit shared salient event semantics across both audio and visual modalities. To this end, we propose to first measure the mutual agreement on event categories along the temporal dimension.

Given the audio and visual features $\mathbf{a} \in \mathbb{R}^{T \times d}$ and $\mathbf{v} \in \mathbb{R}^{T \times d}$, we first predict event categories independently for each modality using separate classifiers. This is implemented using two linear layers with a LeakyReLU activation in between, formulated as:

$$\mathbf{P}_a = \sigma((\mathbf{a}\mathbf{W}_1^a)\mathbf{W}_2^a), \quad \mathbf{P}_v = \sigma((\mathbf{v}\mathbf{W}_1^v)\mathbf{W}_2^v), \quad (1)$$

where $\mathbf{W}_1^m \in \mathbb{R}^{d \times d}$, $\mathbf{W}_2^m \in \mathbb{R}^{d \times C}$ ($m \in \{a, v\}$), and σ is the sigmoid function to enable multi-label event prediction.

$\mathbf{P}_a, \mathbf{P}_v \in \mathbb{R}^{T \times C}$ represent the event probabilities at each timestamp. A target audio-visual event is considered present only when the predicted class probability is high in both modalities, *i.e.*, both modalities agree that the event exists. We first measure the discrepancy between the predicted distributions \mathbf{P}_a and \mathbf{P}_v using the tool of Jensen-Shannon Divergence (Fuglede and Topsøe 2004), as follows:

$$\begin{aligned} \bar{\mathbf{P}} &= \frac{1}{2}(\mathbf{P}_a + \mathbf{P}_v), \\ \mathbf{d}_{JSD} &= \frac{1}{2}D_{KL}(\bar{\mathbf{P}}, \mathbf{P}_a) + \frac{1}{2}D_{KL}(\bar{\mathbf{P}}, \mathbf{P}_v), \end{aligned} \quad (2)$$

where D_{KL} is the Kullback-Leibler Divergence. After averaging over the category dimension, we have $\mathbf{d}_{JSD} \in \mathbb{R}^{T \times 1}$. Accordingly, the audiovisual mutual agreement score \mathbf{s} can be obtained by:

$$\mathbf{s} = 1 - \mathbf{d}_{JSD}, \quad (3)$$

where $\mathbf{s} \in \mathbb{R}^{T \times 1}$. The higher the value of \mathbf{d}_{JSD} , the more different the distributions of \mathbf{P}_a and \mathbf{P}_v are, indicating lower overlap in event categories, resulting in a lower mutual event agreement score.

Cross-modal Salient Anchor Identification

After obtaining the mutual agreement score $\mathbf{s} \in \mathbb{R}^{T \times 1}$, we employ it as an indicator to identify cross-modal salient anchors along the temporal axis. We develop two mechanisms:

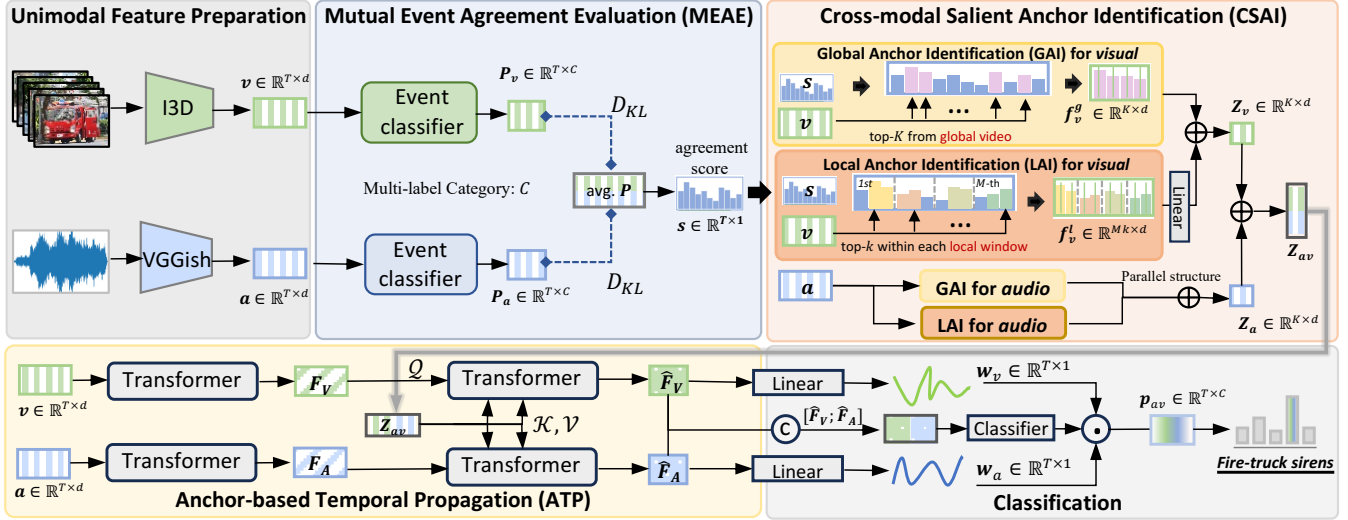


Figure 2: Illustration of our cross-modal salient anchor-based semantic propagation (CLASP) framework for W-DAVEL.

Global Anchor Identification (GAI). The GAI aims to identify the salient anchors by performing a global comparison over the entire video timeline. Specifically, timestamps with the top- K values in s are selected. Let \mathcal{I}^g denote the index of the selected timestamps, where each index \mathcal{I}_i^g belongs to $\{1, \dots, T\}$ for $i = 1, \dots, K$. Using \mathcal{I}^g , we extract the corresponding audio and visual features from a and v , respectively, denoted as $f_a^g \in \mathbb{R}^{K \times d}$ and $f_v^g \in \mathbb{R}^{K \times d}$. f_a^g and f_v^g represent the features that exhibit the most consistent event semantics across both audio and visual modalities.

Local Anchor Identification (LAI). In a long video, audio-visual events often appear at multiple temporal positions. In addition to GAI, we propose LAI to identify salient anchors through a finer-grained temporal comparison. Specifically, as shown in Fig. 2, the video is divided into M temporal windows, each containing $\lfloor \frac{T}{M} \rfloor$ segments. For each temporal window, we extract the corresponding slice of the mutual agreement score s , denoted as s_m ($m = 1, \dots, M$). Then, operations similar to GAI are applied to identify the anchors with the top- k values of s_m . We denote the obtained audio and visual anchor features as $f_a^{l,m}$ and $f_v^{l,m}$, where $f^{l,m} \in \mathbb{R}^{k \times d}$ is the anchor feature set for the m -th temporal window. Therefore, the anchor features across all M windows are aggregated and denoted as $f_a^l = \{f_a^{l,m}\}_{m=1}^M$ and $f_v^l = \{f_v^{l,m}\}_{m=1}^M \in \mathbb{R}^{Mk \times d}$. A linear layer is then applied to transform f_a^l and f_v^l , aligning them with the K anchors obtained from GAI, yielding $f_a^l, f_v^l \in \mathbb{R}^{K \times d}$.

After obtaining the anchor features from both the global video level and the local-temporal window level, we aggregate the two types of anchor features as follows:

$$\begin{aligned} Z_a &= f_a^g + f_a^l, \quad Z_v = f_v^g + f_v^l, \\ Z_{av} &= ([Z_a; Z_v] \mathbf{W}_3) \mathbf{W}_4, \end{aligned} \quad (4)$$

where $\mathbf{W}_3 \in \mathbb{R}^{2d \times 2d}$, $\mathbf{W}_4 \in \mathbb{R}^{2d \times d}$. $[\cdot; \cdot]$ represents the feature concatenation along the channel dimension. $Z_{av} \in \mathbb{R}^{K \times d}$ is the ultimate anchor features, which contain the

salient event category semantics (events mostly likely to occur) in both audio and visual modalities.

Anchor-based Temporal Propagation

The W-DAVEL task requires identifying the event category for each timestamp. Therefore, it is important to capture *temporal* features that accurately reflect event semantics. Our Anchor-based Temporal Propagation (ATP) module enhances temporal interactions using the identified anchors. First, the original audio a and visual v features are passed through separate unimodal two-layer Transformers (Vaswani et al. 2017) to encode the temporal relations within each modality, yielding the updated features $F_A, F_V \in \mathbb{R}^{T \times d}$.

Then, we model the interaction between the modality-specific features (F_A/F_V) and the multimodal anchor feature Z_{av} via a cross-attention mechanism, formulated as:

$$\begin{aligned} \hat{F}_A &= F_A + \text{MHA}(F_A, Z_{av}, Z_{av}), \\ \hat{F}_V &= F_V + \text{MHA}(F_V, Z_{av}, Z_{av}), \end{aligned} \quad (5)$$

$$\text{MHA}(\mathcal{Q}, \mathcal{K}, \mathcal{V}) = \delta \left(\frac{\mathcal{Q} \mathbf{W}^Q (\mathcal{K} \mathbf{W}^K)^\top}{\sqrt{d}} \right) \mathcal{V} \mathbf{W}^V,$$

where $\hat{F}_A, \hat{F}_V \in \mathbb{R}^{T \times d}$ denote the updated temporal features for audio and visual modalities, $\mathbf{W}^Q, \mathbf{W}^K, \mathbf{W}^V \in \mathbb{R}^{d \times d}$ are learnable parameters, and δ is the softmax function.

During this process, the modality-specific features F_A and F_V query the anchor features Z_{av} to assess whether the event categories encoded by Z_{av} are present at each timestamp, reflected by higher cross-attention weights. In this way, salient event semantics are propagated from the anchor features to temporally aligned positions, facilitating fine-grained, timestamp-level prediction.

Classification and Model Training

Inspired by prior work on temporal action localization (Hong et al. 2021; Li et al. 2023), we introduce modality-specific

classifiers to predict the foreground weight for each audio and visual modality, with $w_a, w_v \in \mathbb{R}^{T \times 1}$. Foreground segments with high probability indicate that events, rather than background noise, are more likely to occur. Specifically, the classifier is applied to the final audio or visual features (\hat{F}_A/\hat{F}_V), implemented by a linear layer followed by a sigmoid activation. We use the average of w_a and w_v to represent the overall event foreground score, denoted as $\bar{w} \in \mathbb{R}^{T \times 1}$.

Meanwhile, the audio-visual event probability $p_{av} \in \mathbb{R}^{T \times C}$ can be obtained via an additional classifier applied to the concatenated audio and visual features, *i.e.*, $[\hat{F}_A; \hat{F}_V]$. Then, p_{av} is modulated by the event foreground score \bar{w} to suppress background temporal segments, computed as $p_{av} \odot \bar{w}$, where \odot denotes element-wise multiplication.

The modulated p_{av} can be used for temporal-level prediction during the inference phase. For model training, we compute the video-level audio-visual event prediction $\hat{p} \in \mathbb{R}^{1 \times C}$ by aggregating $p_{av} \in \mathbb{R}^{T \times C}$ over the temporal dimension using MIL pooling, as in (Tian et al. 2018). Then, the cross-entropy loss between \hat{p} and the video-level ground truth label y is computed. Notably, P_a and P_v used in MEAE module (Eq. 1) are regularized in the same way.

Experiments

Experimental Setups

Datasets. We evaluate our method on the UnAV-100 (Geng et al. 2023) and ActivityNet1.3 (Heilbron et al. 2015) datasets. *UnAV-100* serves as the standard benchmark for the fully-supervised DAVEL task. It contains 10,790 untrimmed videos totaling over 126 hours of video. The videos have varying durations, with most exceeding 40 seconds. The video content spans 100 common audio-visual event categories drawn from real-world scenarios, including human and animal activities, musical instrument performances, and various means of transportation, *etc.* Each video contains an average of 2.8 audio-visual events, with a maximum of 23 events per video. According to the standard dataset division scheme, the training, validation, and test sets follow a 3:1:1 split. *ActivityNet1.3* is another large-scale dataset, originally constructed for temporal action localization. We also evaluate on this dataset to assess the robustness of our approach. Specifically, it includes 19,994 long video clips, each with an average duration of approximately 100 seconds and 1.41 activity classes. This dataset covers 203 activity categories, including daily activities, sports competitions, and performing arts.

Evaluation Metric. Following prior work (Geng et al. 2023), we adopt mean Average Precision (mAP) as the primary evaluation metric to assess temporal localization performance. We report mAP values at tIoU thresholds ranging from 0.5 to 0.9 with a step size of 0.1 ([0.5:0.1:0.9]). Meanwhile, the average mAP across threshold range [0.1:0.1:0.9] (‘Avg.’) is reported as a holistic measure of overall model performance.

Implementation Details. (1) *Feature extraction.* We preprocess the audio and visual data following prior works (Geng et al. 2023; Heilbron et al. 2015). For the UnAV-100 dataset, we sample the video frames at 25 FPS. Then, the RAFT (Teed and Deng 2020) model is employed to extract the optical flow images. Then, they are combined with 24 consecutive RGB

Methods	Task	0.5	0.6	0.7	0.8	0.9	Avg.
AVEL (Tian et al. 2018)	AVEP	18.7	16.5	14.5	12.3	10.0	19.2
JoMoLD (Cheng et al. 2022)		22.4	19.6	17.3	14.8	12.0	22.9
PPAE (Gao, Chen, and Xu 2025)		23.6	20.5	17.9	15.2	12.4	23.8
NREP (Jiang et al. 2025)		24.9	19.7	17.2	15.6	<u>15.2</u>	26.4
DAVE (Geng et al. 2023)	DAVEL	21.6	19.5	17.1	14.8	11.7	21.6
CCNet (Zhou et al. 2025d)		<u>27.0</u>	<u>24.4</u>	<u>21.3</u>	<u>17.8</u>	13.6	<u>26.9</u>
CLASP (ours)	DAVEL	29.6	26.3	23.0	19.7	15.8	30.0

Table 1: Comparison with prior works on UnAV-100 dataset. All methods are trained and evaluated under weakly supervised settings. The best and second-best results are highlighted in bold and underlined, respectively.

Methods	0.5	0.6	0.7	0.8	0.9	Avg.
ActionFormer (Zhang, Wu, and Li 2022)	23.2	19.7	16.2	11.8	6.2	14.3
DAVE (Geng et al. 2023)	25.5	21.7	17.8	13.0	6.8	15.8
CCNet (Zhou et al. 2025d)	<u>31.7</u>	<u>27.0</u>	<u>22.2</u>	<u>16.3</u>	8.5	<u>19.6</u>
CLASP (ours)	34.6	29.0	23.6	17.6	9.8	21.3

Table 2: Comparison on the ActivityNet1.3 dataset.

frames and sent to a two-stream pretrained I3D (Carreira and Zisserman 2017) model to extract 2048-D visual features. The audio signal is segmented at time intervals of 0.96 seconds using a sliding window of 0.32 seconds. The pretrained VGGish (Hershey et al. 2017) is used to extract 128-D audio features. To facilitate parallel processing of videos with varying duration, a maximum temporal length of $T=224$ is set by clipping or zero-padding. For the ActivityNet1.3 dataset, the sampling rates of video frames and the audio stream are 16 FPS and 16 KHz, respectively. The audio and visual features are extracted using the pretrained VGGish and R(2+1)D-34 model, respectively. T is set to 256. 2) *Model Configuration.* For model training on the UnAV-100 dataset, we train our model for 40 epochs with a batch size of 16. For the training on ActivityNet1.3, the total training epoch and the batch size are 40 and 8, respectively. The Adam optimizer is utilized with a learning rate of 1e-4 for both datasets. The optimal configuration of the key hyperparameters, such as the number of global anchors K and the number of anchors in each local window k , are set to 10 and 4, respectively. The number of local temporal windows M is 14 and 16 for the UnAV-100 and ActivityNet1.3 datasets, respectively. During the inference phase, we utilize the conventional threshold of 0.5 to identify confident event classes from the prediction p_{av} . All experiments are conducted on a NVIDIA A40 (40GB) GPU.

Main Comparison Results

We evaluate our method on the UnAV-100 and ActivityNet1.3 datasets and compare it against prior state-of-the-art (SOTA) methods from related tasks.

Evaluation on the UnAV-100 dataset. Table 1 presents the evaluation results. We compare with SOTA methods designed for the audio-visual event perception (AVEP) task, including AVEL (Tian et al. 2018), JoMoLD (Cheng et al. 2020), PPAE (Gao, Chen, and Xu 2025), and NREP (Jiang et al. 2025). As shown in the table, our method outperforms all these approaches by a significant margin. For instance, our method surpasses the best method, NREP, by 3.6% on the

Methods	Param.(M)	FLOPs(G)	Runtime(s)
DAVEL (Geng et al. 2023)	139.4	52.4	193.0
CCNet (Zhou et al. 2025d)	238.8	72.1	217.6
CLASP (ours)	27.2	5.7	38.7

Table 3: Efficiency comparison with prior methods.

K	0.5	0.6	0.7	0.8	0.9	Avg.
5	29.1	25.6	22.4	19.1	15.2	29.4
10	29.6	26.3	23.0	19.7	15.8	30.0
15	29.3	25.9	22.9	19.4	15.1	29.6
20	28.4	25.4	22.1	19.1	15.8	29.1

Table 4: Impact of the global anchor number K .

‘Avg.’ metric. In addition, our method also outperforms SOTA approaches developed for the DAVEL task. The baseline model DAVE (Geng et al. 2023) achieves only 21.6% on the ‘Avg.’ metric. The recent SOTA method for fully-supervised DAVEL, CCNet (Zhou et al. 2025d), achieves a higher score of 26.9%. However, they still fall significantly short of our method under all tIoU metrics. Notably, we do not compare with LoCo (Xing et al. 2024), which performs comparably to CCNet on the fully-supervised DAVEL task, as its implementation is not publicly available so far. It is worth noting that fully-supervised DAVEL methods can reach around 50% ‘Avg.’. This performance gap highlights the difficulty and potential of the studied W-DAVEL task.

Evaluation on the ActivityNet1.3 dataset. Table 2 presents the evaluation results on ActivityNet1.3 dataset. Our method again outperforms prior SOTA methods from related tasks, surpassing the most competitive method, CCNet (Zhou et al. 2025d), by 1.7% on the ‘Avg.’ metric. These results demonstrate the effectiveness and generalizability of our method.

Efficiency comparison. We compare the model efficiency of the proposed method with previous DAVEL approaches by reporting the number of trainable parameters (M), FLOPs (G), and inference runtime (s). Experiments are conducted on the UnAV-100 dataset. As shown in Table 3, our method is more efficient, with fewer parameters and shorter runtime, striking a good balance between accuracy and efficiency.

Ablation Studies

In this section, we conduct extensive experiments to evaluate the impacts of some key hyperparameters and the effectiveness of each main module design. The ablation experiments are conducted on the standard UnAV-100 dataset.

Impact of the Global Anchor Number K . In our Global Anchor Identification mechanism, we identify K anchor features from the entire video. We conduct experiments to evaluate its impact by varying the value of K . The results are shown in Table 4. When using quite a few anchors, *e.g.*, $K = 5$, the model has a lower performance at 29.4% on ‘Avg.’ metric. Increasing the global anchor number to 10, the model’s performance improves accordingly. However, continuing to enlarge the number of K introduces a slight decrease in performance. Therefore, we adopt $K = 10$ as the optimal setup.

k	0.5	0.6	0.7	0.8	0.9	Avg.
1	28.3	25.4	22.1	19.2	<u>15.7</u>	29.2
2	30.4	27.1	23.3	<u>19.4</u>	14.3	<u>29.4</u>
4	<u>29.6</u>	<u>26.3</u>	<u>23.0</u>	19.7	15.8	30.0
8	28.9	25.8	22.5	19.1	15.4	29.3

Table 5: Impact of the local anchor number k .

M	0.5	0.6	0.7	0.8	0.9	Avg.
4	29.1	25.8	22.5	19.3	15.3	29.4
8	29.4	26.0	22.6	19.2	15.1	29.7
14	29.6	26.3	23.0	19.7	15.8	30.0
28	29.4	25.9	22.5	19.2	15.1	29.6

Table 6: Impact of the local temporal window number M .

Impact of the Local Anchor Number k . In our Local Anchor Identification mechanism, we identify k anchor features from each local temporal window. We conduct experiments to evaluate its impact. As shown in Table 5, the optimal performance is achieved when $k = 4$. Similar to the trend observed for the aforementioned K , the model performance gradually improves with increasing k before declining when k exceeds 4. Therefore, we set $k = 4$ as the default setting.

Impact of the Local Temporal Window Number M . In our Local Anchor Identification mechanism, we divide the entire video into M temporal windows and identify the k salient anchors in each window. We fix k to 4 and conduct experiments to explore the impact of M . As shown in Table 6, the optimal performance is achieved when $M = 14$. Using fewer or more temporal windows, the model’s performance slightly decreases. Therefore, we employ $M = 14$ as the optimal setup. It is noteworthy that our model still maintains better performance compared to prior approaches under varied hyperparameters (K , k , and M), demonstrating the superiority and robustness of our framework.

Ablation Study on the MEAE Module. In our Mutual Event Agreement Evaluation (MEAE) module, we utilize the initial audio and visual features (\mathbf{a} , \mathbf{v}) extracted by pretrained backbones to compute the agreement score. We also explore a variant that employs the features after going through the unimodal transformers (\mathbf{F}_A , \mathbf{F}_V). The experimental results are shown in Table 7. Using the vanilla audiovisual features \mathbf{a} and \mathbf{v} has superior performance. This indicates that it is not necessary to perform unimodal temporal relation modeling in the mutual event agreement evaluation phase.

Ablation Study on the CSAI Module. Our Cross-modal Salient Anchor Identification (CSAI) module introduces two mechanisms: Global Anchor Identification (GAI) and Local Anchor Identification (LAI). We investigate their contributions by ablating each mechanism. As shown in Table 8, removing either the GAI or LAI mechanism leads to a performance drop. These results highlight the effectiveness of both mechanisms and suggest their complementarity. We also observe that using only LAI yields slightly better performance than using only GAI. This indicates that anchor identification

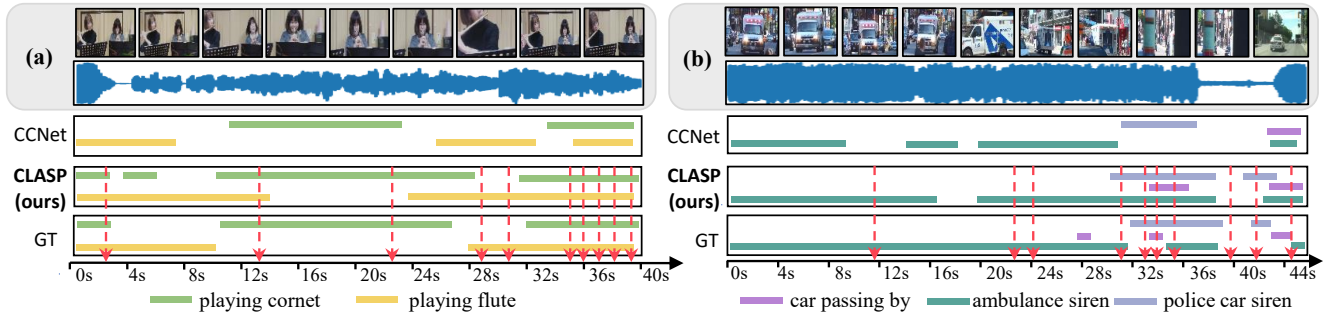


Figure 3: Visualization examples of the weakly-supervised dense audio-visual event localization. We compare our method with CCNet (Zhou et al. 2025d). ‘GT’ is the ground truth. The red arrows mark the identified global ten salient anchor timestamps.

Setups	0.5	0.6	0.7	0.8	0.9	Avg.
F_A, F_V	29.0	25.9	22.8	19.5	15.3	29.6
a, v	29.6	26.3	23.0	19.7	15.8	30.0

Table 7: Ablation study on MEAE module. We explore the impacts of different audio and visual features used in mutual event agreement score computation.

GAI	LAI	0.5	0.6	0.7	0.8	0.9	Avg.
✗	✓	29.4	26.2	22.3	19.4	15.5	29.8
✓	✗	29.2	25.9	22.8	19.3	15.1	29.5
✓	✓	29.6	26.3	23.0	19.7	15.8	30.0

Table 8: Ablation study on the two anchor identification mechanisms in CSAI module. ‘GAI’ and ‘LAI’ are abbreviations for Global Anchor Identification and Local Anchor Identification mechanisms in CSAI.

within multiple local temporal windows is more beneficial, as it facilitates finer-grained semantic comparisons over time. **Ablation Study on the ATP Module.** Our Anchor-based Temporal Propagation (ATP) module utilizes the identified multimodal anchor features to interact with the temporal audio and visual features. We conduct ablation experiments to explore the impact of anchor features. As shown in Table 9, without using the anchors features in ATP, *i.e.*, only using the modality-specific audio and visual features obtained by unimodal transformers, the model’s performance drops in all metrics, resulting in an overall decrease of 0.7% on ‘Avg.’ metric. These results demonstrate the benefits of anchor features. The contained salient semantics of the audio-visual events in the anchor features can be propagated to enhance the temporal features, facilitating the event localization under the weakly supervised setting.

Qualitative Results

We present qualitative visualizations for dense audio-visual event localization and compare our method with the prior SOTA method CCNet (Zhou et al. 2025d). As illustrated in Fig. 3(a), the video contains two audio-visual events in

Setups	0.5	0.6	0.7	0.8	0.9	Avg.
w/o. anchor	28.9	25.6	22.5	19.3	15.2	29.3
w. anchor	29.6	26.3	23.0	19.7	15.8	30.0

Table 9: Ablation study on ATP module. ‘w/o. anchor’ represents that the identified anchor features are not used.

a musical performance scene: *playing cornet* and *playing flute*. Both CCNet and our method correctly identify these two events. However, CCNet fails to localize the *playing cornet* during the initial two seconds. Moreover, our method provides more accurate temporal boundaries for both events between 28s and 40s. Similarly, in Fig. 3(b), CCNet misses the *car passing by* and *ambulance siren* events occurring between 32s and 40s, while our method successfully identifies both. We also mark the positions of the identified global salient anchors ($K=10$) using red arrows. These anchors are unevenly distributed, often aligning with the moments where major events occur. The identified anchors play a crucial role in guiding event semantic perception and propagating it to the remaining temporal timestamps through our anchor-based mechanism, thereby enhancing temporal localization under weak supervision.

Conclusion

We explore the W-DAVEL task, aiming to achieve dense audio-visual event localization under a weakly-supervised setting where only the video-level event labels are provided. Our method, CLASP, addresses this task by exploiting the cross-modal salient anchors, *i.e.*, temporal timestamps that exhibit the most consistent event semantics across audio and visual modalities. A mutual event agreement evaluation module is used to assess the consistency of audio and visual event predictions. The resulting agreement score serves as the indicator for identifying cross-modal salient anchors through both global video-level and local temporal window-level comparisons. The identified anchor features are then leveraged to enhance temporal audio and visual features, which are used for final event classification. Extensive experiments verify the effectiveness of our method. We hope our work encourages further research into this challenging yet practical task.

Acknowledgments

We would like to express our sincere gratitude to the anonymous reviewers for their invaluable comments and insightful suggestions. This work is supported by National Key R&D Program of China (NO.2024YFB3311602), Natural Science Foundation of China (62272144), the Anhui Provincial Natural Science Foundation (2408085J040), and the Major Project of Anhui Provincial Science and Technology Breakthrough Program (202423k09020001), and the Fundamental Research Funds for the Central Universities (JZ2024HGTG0309, JZ2024AHST0337).

References

- Arandjelovic, R.; and Zisserman, A. 2017. Look, listen and learn. In *ICCV*, 609–617.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 6299–6308.
- Cheng, H.; Liu, Z.; Zhou, H.; Qian, C.; Wu, W.; and Wang, L. 2022. Joint-Modal Label Denoising for Weakly-Supervised Audio-Visual Video Parsing. In *ECCV*, 431–448.
- Cheng, Y.; Wang, R.; Pan, Z.; Feng, R.; and Zhang, Y. 2020. Look, listen, and attend: Co-attention network for self-supervised audio-visual representation learning. In *ACM MM*, 3884–3892.
- Fan, Y.; Wu, Y.; Du, B.; and Lin, Y. 2023. Revisit weakly-supervised audio-visual video parsing from the language perspective. *NeurIPS*, 36: 40610–40622.
- Fuglede, B.; and Topsoe, F. 2004. Jensen-Shannon divergence and Hilbert space embedding. In *International Symposium on Information Theory*, 31.
- Gao, J.; Chen, M.; and Xu, C. 2023. Collecting Cross-Modal Presence-Absence Evidence for Weakly-Supervised Audio-Visual Event Perception. In *CVPR*, 18827–18836.
- Gao, J.; Chen, M.; and Xu, C. 2025. Learning probabilistic presence-absence evidence for weakly-supervised audio-visual event perception. *TPAMI*, 4787–4802.
- Gemmeke, J. F.; Ellis, D. P.; Freedman, D.; Jansen, A.; Lawrence, W.; Moore, R. C.; Plakal, M.; and Ritter, M. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *ICASSP*, 776–780.
- Geng, T.; Wang, T.; Duan, J.; Cong, R.; and Zheng, F. 2023. Dense-localizing audio-visual events in untrimmed videos: A large-scale benchmark and baseline. In *CVPR*, 22942–22951.
- Guo, C.; Huang, H.; and Zhou, Y. 2024. Enhance audio-visual segmentation with hierarchical encoder and audio guidance. *Neurocomputing*, 594: 127885.
- Guo, R.; Ying, X.; Chen, Y.; Niu, D.; Li, G.; Qu, L.; Qi, Y.; Zhou, J.; Xing, B.; Yue, W.; Shi, J.; Wang, Q.; Zhang, P.; and Liang, B. 2023. Audio-Visual Instance Segmentation. *arXiv preprint arXiv:2310.18709*.
- He, X.; Liu, X.; Li, Y.; Zhao, D.; Shen, G.; Kong, Q.; Yang, X.; and Zeng, Y. 2024. Cace-net: Co-guidance attention and contrastive enhancement for effective audio-visual event localization. In *ACM MM*, 985–993.
- Heilbron, F. C.; Escorcia, V.; Ghanem, B.; and Niebles, J. C. 2015. ActivityNet: A large-scale video benchmark for human activity understanding. In *CVPR*, 961–970.
- 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 *ICASSP*, 131–135.
- Hong, F.-T.; Feng, J.-C.; Xu, D.; Shan, Y.; and Zheng, W.-S. 2021. Cross-modal Consensus Network for Weakly Supervised Temporal Action Localization. In *ACM MM*, 1591–1599.
- Jiang, X.; Xu, X.; Chen, Z.; Zhang, J.; Song, J.; Shen, F.; Lu, H.; and Shen, H. T. 2022. DHHN: Dual Hierarchical Hybrid Network for Weakly-Supervised Audio-Visual Video Parsing. In *ACM MM*, 719–727.
- Jiang, X.; Xu, X.; Zhu, L.; Sun, Z.; Cichocki, A.; and Shen, H. T. 2025. Resisting noise in pseudo labels: Audible video event parsing with evidential learning. *TNNLS*, 10874–10888.
- Jin, D.; Zhou, Y.; Zhou, J.; Ma, J.; Guo, R.; and Guo, D. 2025. SimToken: A Simple Baseline for Referring Audio-Visual Segmentation. *arXiv preprint arXiv:2509.17537*.
- Kay, W.; Carreira, J.; Simonyan, K.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; Natsev, P.; et al. 2017. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 1–22.
- Kryklyvets, Y.; Kurpath, M. I.; Mullappilly, S. S.; Zhou, J.; Khan, F. S.; Anwer, R. M.; Khan, S.; and Cholakkal, H. 2025. MAViS: A Multimodal Conversational Assistant For Avian Species. In *EMNLP*, 28601–28627.
- Li, G.; Cheng, D.; Ding, X.; Wang, N.; Wang, X.; and Gao, X. 2023. Boosting Weakly-Supervised Temporal Action Localization With Text Information. In *CVPR*, 10648–10657.
- Li, Z.; Guo, D.; Zhou, J.; Zhang, J.; and Wang, M. 2024. Object-aware adaptive-positivity learning for audio-visual question answering. In *AAAI*, 3306–3314.
- Li, Z.; Zhou, J.; Zhang, J.; Tang, S.; Li, K.; and Guo, D. 2025. Patch-level sounding object tracking for audio-visual question answering. In *AAAI*, 5075–5083.
- Liu, X.; Xia, N.; Zhou, J.; Li, Z.; and Guo, D. 2025. Towards energy-efficient audio-visual classification via multimodal interactive spiking neural network. *ACM Transactions on Multimedia Computing, Communications and Applications*, 21(5): 1–24.
- Mahmud, T.; and Marculescu, D. 2023. Ave-clip: Audioclip-based multi-window temporal transformer for audio visual event localization. In *WACV*, 5158–5167.
- Mao, Y.; Shen, X.; Zhang, J.; Qin, Z.; Zhou, J.; Xiang, M.; Zhong, Y.; and Dai, Y. 2024. TAVGBench: Benchmarking text to audible-video generation. In *ACM MM*, 6607–6616.
- Mo, S.; and Tian, Y. 2022. Multi-modal Grouping Network for Weakly-Supervised Audio-Visual Video Parsing. In *NeurIPS*.
- 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 *ICML*, 8748–8763.
- Ramaswamy, J. 2020. What makes the sound?: A dual-modality interacting network for audio-visual event localization. In *ICASSP*, 4372–4376.
- 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 *CVPR*, 10585–10596.
- Teed, Z.; and Deng, J. 2020. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 402–419.
- Tian, Y.; Li, D.; and Xu, C. 2020. Unified multisensory perception: Weakly-supervised audio-visual video parsing. In *ECCV*, 436–454.
- Tian, Y.; Shi, J.; Li, B.; Duan, Z.; and Xu, C. 2018. Audio-visual event localization in unconstrained videos. In *ECCV*, 247–263.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS*, 1–15.
- Wu, Y.; Chen, K.; Zhang, T.; Hui, Y.; Berg-Kirkpatrick, T.; and Dubnov, S. 2023. Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *ICASSP*, 1–5. IEEE.
- Wu, Y.; and Yang, Y. 2021. Exploring heterogeneous clues for weakly-supervised audio-visual video parsing. In *CVPR*, 1326–1335.
- Xia, Y.; and Zhao, Z. 2022. Cross-modal Background Suppression for Audio-Visual Event Localization. In *CVPR*, 19957–19966.
- Xing, L.; Qu, H.; Yan, R.; Shu, X.; and Tang, J. 2024. Locality-aware cross-modal correspondence learning for dense audio-visual events localization. *arXiv preprint arXiv:2409.07967*.
- Xu, H.; Zeng, R.; Wu, Q.; Tan, M.; and Gan, C. 2020. Cross-modal relation-aware networks for audio-visual event localization. In *ACM MM*, 3893–3901.
- Yu, J.; Cheng, Y.; Zhao, R.-W.; Feng, R.; and Zhang, Y. 2022. MM-Pyramid: Multimodal Pyramid Attentional Network for Audio-Visual Event Localization and Video Parsing. In *ACM MM*, 6241–6249.
- Yung-Hsuan Lai, Y.-C. F. W., Yen-Chun Chen. 2023. Modality-Independent Teachers Meet Weakly-Supervised Audio-Visual Event Parser. In *NeurIPS*, 1–19.
- Zhang, C.-L.; Wu, J.; and Li, Y. 2022. ActionFormer: Localizing Moments of Actions with Transformers. In *ECCV*, 492–510.
- Zhao, P.; Zhou, J.; Zhao, Y.; Guo, D.; and Chen, Y. 2025. Multimodal class-aware semantic enhancement network for audio-visual video parsing. In *AAAI*, 10448–10456.
- Zhou, J.; Guo, D.; Guo, R.; Mao, Y.; Hu, J.; Zhong, Y.; Chang, X.; and Wang, M. 2024a. Towards Open-Vocabulary Audio-Visual Event Localization. *arXiv preprint arXiv:2411.11278*.
- Zhou, J.; Guo, D.; Mao, Y.; Zhong, Y.; Chang, X.; and Wang, M. 2024b. Label-anticipated Event Disentanglement for Audio-Visual Video Parsing. In *ECCV*, 1–22.
- Zhou, J.; Guo, D.; and Wang, M. 2023. Contrastive Positive Sample Propagation along the Audio-Visual Event Line. *TPAMI*, 7239–7257.
- Zhou, J.; Guo, D.; Zhong, Y.; and Wang, M. 2023a. Improving audio-visual video parsing with pseudo visual labels. *arXiv preprint arXiv:2303.02344*.
- Zhou, J.; Guo, D.; Zhong, Y.; and Wang, M. 2024c. Advancing Weakly-Supervised Audio-Visual Video Parsing via Segment-Wise Pseudo Labeling. *IJCV*, 5308–5329.
- Zhou, J.; Li, Z.; Yu, Y.; Zhou, Y.; Guo, R.; Li, G.; Mao, Y.; Han, M.; Chang, X.; and Wang, M. 2025a. Mettle: Meta-Token Learning for Memory-Efficient Audio-Visual Adaptation. *arXiv preprint arXiv:2506.23271*.
- 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 preprint arXiv:2301.13190*.
- Zhou, J.; Wang, J.; Zhang, J.; Sun, W.; Zhang, J.; Birchfield, S.; Guo, D.; Kong, L.; Wang, M.; and Zhong, Y. 2022. Audio-visual segmentation. In *ECCV*, 386–403.
- Zhou, J.; Zheng, L.; Zhong, Y.; Hao, S.; and Wang, M. 2021. Positive sample propagation along the audio-visual event line. In *CVPR*, 8436–8444.
- Zhou, J.; Zhou, Y.; Han, M.; Wang, T.; Chang, X.; Cholakkal, H.; and Anwer, R. M. 2025b. Think before you segment: An object-aware reasoning agent for referring audio-visual segmentation. *arXiv preprint arXiv:2508.04418*.
- Zhou, Y.-H.; Huang, H.; Guo, C.; Tu, R.-C.; Xiao, Z.; Wang, B.; and Mao, X.-L. 2025c. ALOHA: Adapting Local Spatio-Temporal Context to Enhance the Audio-Visual Semantic Segmentation. *ACM TOMM*.
- Zhou, Z.; Zhou, J.; Qian, W.; Tang, S.; Chang, X.; and Guo, D. 2025d. Dense audio-visual event localization under cross-modal consistency and multi-temporal granularity collaboration. In *AAAI*, 10905–10913.