

# AVE-CLIP: AudioCLIP-based Multi-window Temporal Transformer for Audio Visual Event Localization

Tanvir Mahmud  
The University of Texas at Austin  
tanvirmahmud@utexas.edu

Diana Marculescu  
The University of Texas at Austin  
dianam@utexas.edu

## Abstract

An audio-visual event (AVE) is denoted by the correspondence of the visual and auditory signals in a video segment. Precise localization of the AVEs is very challenging since it demands effective multi-modal feature correspondence to ground the short and long range temporal interactions. Existing approaches struggle in capturing the different scales of multi-modal interaction due to ineffective multi-modal training strategies. To overcome this limitation, we introduce AVE-CLIP, a novel framework that integrates the AudioCLIP pre-trained on large-scale audio-visual data with a multi-window temporal transformer to effectively operate on different temporal scales of video frames. Our contributions are three-fold: (1) We introduce a multi-stage training framework to incorporate AudioCLIP pre-trained with audio-image pairs into the AVE localization task on video frames through contrastive fine-tuning, effective mean video feature extraction, and multi-scale training phases. (2) We propose a multi-domain attention mechanism that operates on both temporal and feature domains over varying timescales to fuse the local and global feature variations. (3) We introduce a temporal refining scheme with event-guided attention followed by a simple-yet-effective post processing step to handle significant variations of the background over diverse events. Our method achieves state-of-the-art performance on the publicly available AVE dataset with 5.9% mean accuracy improvement which proves its superiority over existing approaches.

## 1. Introduction

Temporal reasoning of multi-modal data plays a significant role in human perception in diverse environmental conditions [10, 37]. Grounding the multi-modal context is critical to current and future tasks of interest, especially those that guide current research efforts in this space, *e.g.*, embodied perception of automated agents [29, 4, 8], human-robot interaction with multi-sensor guidance [25, 6, 2], and active



Figure 1. Example of an audio-visual event (AVE) representing the event of an individual speaking. The person’s voice is audible in all of the frames. Only when the person is visible, an AVE is identified.

sound source localization [34, 22, 18, 27]. Similarly, audio-visual event (AVE) localization demands complex multi-modal correspondence of grounded audio-visual perception [24, 7]. The simultaneous presence of the audio-visual cues over a video frame denotes an audio-visual event. As shown in Fig. 1, the speech of the person is audible in all of the frames. However, the individual speaking is visible in only a few particular frames which represent the AVE. Precise detection of such events greatly depends on the contextual understanding of the multi-modal features over the video frame.

Learning the inter-modal audio-visual feature correspondence over the video frames is one of the major challenges of AVE localization. Effective multi-modal training strategies can significantly improve performance by enhancing the relevant features. Earlier work integrates audio and image encoders pre-trained on large scale unimodal (image/audio) datasets [5, 9] to improve performance [35, 17, 31, 7, 30]. However, such a uni-modal pre-training scheme struggles to extract relevant inter-modal features that are particularly significant for AVEs. Recently, following the wide-spread success of CLIP [19] pre-trained on large-scale vision-language datasets, AudioCLIP [12] has integrated an audio encoder into the vision-language models with large-scale pre-training on audio-image pairs. To enhance the audio-visual feature correspondence for AVEs, we integrate the image and audio encoders from AudioCLIP with effective contrastive fine-tuning that exploits the large-scale pre-trained knowledge from multi-modal datasets instead of uni-modal ones.

Effective audio-visual fusion for multi-modal reasoning over entire video frames is another major challenge for proper utilization of the uni-modal features. Recently, several approaches have focused on using the grounded multi-modal features to generate temporal attention for operating on the intra-modal feature space [36, 17, 31]. Other recent work has applied recursive temporal attention on the aggregated multi-modal features [7, 17, 30]. However, these existing approaches attempt to generalize audio-visual context over the whole video frame and hence struggle to extract local variational patterns that are particularly significant at event transitions. Though generalized multi-modal context over long intervals is of great importance for categorizing diverse events, local changes of multi-modal features are critical for precise event detection at transition edges. To solve this dilemma, we introduce a multi-window temporal transformer based fusion scheme that operates on different timescales to guide attention over sharp local changes with short temporal windows, as well as extract the global context across long temporal windows.

The background class representing uncorrelated audio-visual frames varies a lot over different AVEs for diverse surroundings (Figure 1). In many cases, it becomes difficult to distinguish the background from the event regions due to subtle variations [36]. Xu *et al.* [30] suggests that joint binary classification of the event regions (event/background) along with the multi-class event prediction improves overall performance for better discrimination of the event oriented features. Inspired by this, we introduce a temporal feature refining scheme for guiding temporal attentions over the event regions to introduce sharp contrast with the background. Moreover, we introduce a simple post-processing algorithm that filters out such incorrect predictions in between event transitions by exploiting the high temporal locality of event/background frames in AVEs (Figure 1). By unifying these strategies in the AVE-CLIP framework, we achieve state-of-the-art performance on the AVE dataset which outperforms existing approaches by a considerable margin.

The major contributions of this work are summarized as follows:

- We introduce AVE-CLIP to exploit AudioCLIP pre-trained on large-scale audio-image pairs for improving inter-modal feature correspondence on video AVEs.
- We propose a multi-window temporal transformer based fusion scheme that operates on different timescales of AVE frames to extract local and global variations of multi-modal features.
- We introduce a temporal feature refinement scheme through event guided temporal attention followed by a simple, yet-effective post-processing method to increase contrast with the background.

## 2. Related Work

### Audio Visual Event Localization

AVE localization, introduced by Tian *et al.* [24] targets the identification of different types of events (*e.g.*, individual man/woman speaking, crying babies, frying food, musical instruments, *etc.*) at each temporal instance based on audio-visual correspondence. The authors introduced a residual learning method with LSTM guided audio-visual attention relying on simple concatenation and addition fusion. A dual attention matching (DAM) module is introduced by Wu *et al.* [28] for operating on event-relevant features. Zhou *et al.* [36] proposed a positive sample propagation scheme by pruning out the weaker multi-modal interactions. Xuan *et al.* [31, 32] proposed a discriminative multi-modal attention module for sequential learning with an eigen-value based objective function. Duan *et al.* [7] introduced joint co-learning with cyclic attention over the aggregated multi-modal features. Lin and Wang [17] introduced a transformer-based approach that operates on groups of video frames based on audio-visual attention. Xu *et al.* [30] introduced multi-modal relation-aware audio-visual representation learning with an interaction module. Different from existing approaches, AVE-CLIP exploits temporal features from various windows by extracting short and long range multi-modal interactions along with temporal refinement of the event frames.

### Sound Source Localization

The sound source localization task [34] identifies the sounding object in the corresponding video based on the auditory signal. Arda *et al.* [22] introduced an audio-visual classification model that can be adapted for sound source localization without explicit training by utilizing simple multi-modal attention. Wu *et al.* [27] proposed an encoder-decoder based framework to operate on the continuous feature space through likelihood measurements of the sounding sources. Qian *et al.* [18] attempted multiple source localization by exploiting gradient weighted class activation map (Grad-CAM) correspondence on the audio-visual signal. A self-supervised audio-visual matching scheme is introduced by Hu *et al.* [15] with a dictionary learning of the sounding objects. Afouras *et al.* [1] utilized optical flow features along with multimodal attention maps targeting both source localization and audio source separation.

### Large Scale Contrastive Pre-training

To improve the data-efficiency on diverse target tasks, large-scale pre-training of very deep neural networks has been found to be effective for transfer learning [16]. CLIP has introduced vision-language pre-training with self-supervised contrastive learning on large-scale datasets, an approach

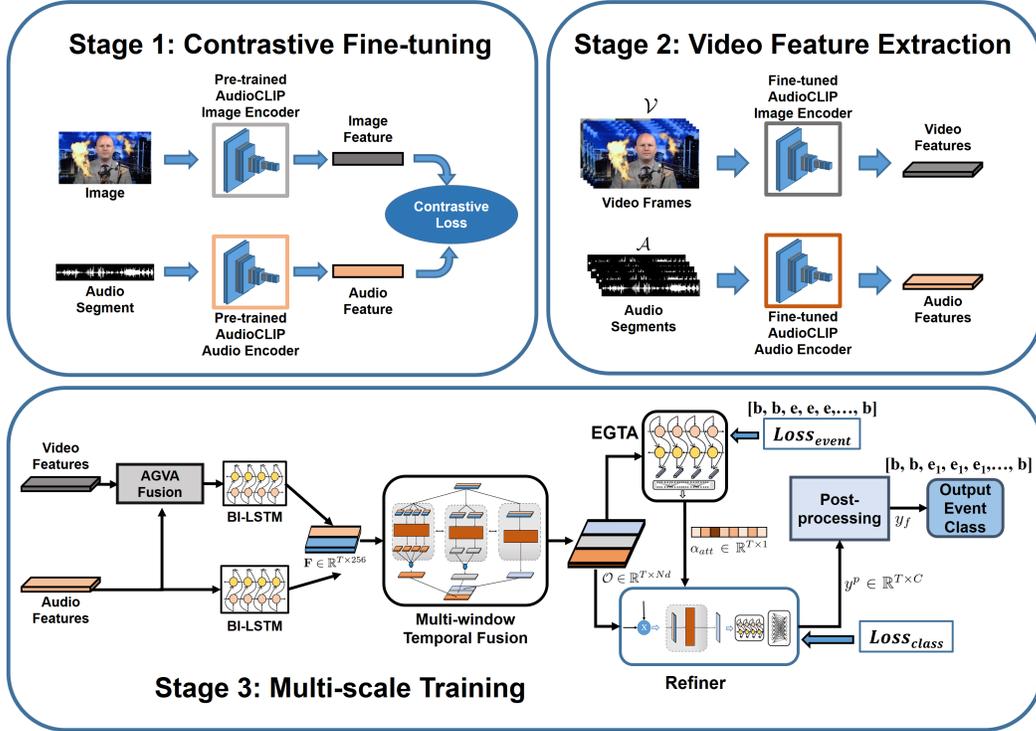


Figure 2. Schematic representation of the proposed method. In stage 1, contrastive fine-tuning is carried on the pre-trained AudioCLIP [12] image and audio encoders with audio-image pairs. In stage 2, video and audio features are extracted with the fine-tuned encoders. In stage 3, multi-scale training is carried at various temporal scales with the proposed multi-window temporal fusion module followed by temporal event refinement and post-processing to enhance event detection.

that received great attention for achieving superior performance on numerous multimodal vision-language tasks [21, 26, 33]. Recently, AudioCLIP [12] has extended the existing CLIP framework by integrating the audio modality with large-scale training utilizing audio-image pairs [9]. Such large-scale pre-training on audio-visual data can be very effective for enhancing multi-modal feature correspondence.

### 3. Proposed Method

In this paper, we introduce AVE-CLIP, a framework that integrates image and audio encoders from AudioCLIP with a multi-window temporal transformer based fusion scheme for AVE localization. Our method comprises three training stages, as presented in Figure 2. Initially, we start with the pre-trained weights of image and audio encoders from AudioCLIP. In stage 1, we extract image and audio segments of corresponding events to initiate fine-tuning of the pre-trained encoders on target AVE-localization frames (Section 3.2). In stage 2, these fine-tuned encoders are deployed to extract the video and audio features from successive video frames and audio segments, respectively (Section 3.3). Later, in stage 3, we introduce the multi-scale training on the extracted audio and video features with the multi-window temporal fusion (MWTF) module that

operates on different temporal windows for generalizing the local and global temporal context (Section 3.4). This is followed by the temporal refinement of the fusion feature through event-guided temporal attention generated with event-label supervision (Section 3.5) along with a hybrid loss function used in training (Section 3.6) and a simple post-processing algorithm that primarily enhances prediction performance during inference by exploiting the temporal locality of the AVEs (Section 3.7).

#### 3.1. Preliminary

Given a video sequence  $\mathbf{S}$  of duration  $T$ , a set of non-overlapping video segments  $\mathcal{V} = \{V^{(1)}, V^{(2)}, \dots, V^{(T)}\}$ , and synchronized audio segments  $\mathcal{A} = \{A^{(1)}, A^{(2)}, \dots, A^{(T)}\}$  of duration  $t$  are extracted. Each video segment  $V^{(i)} = \{v_k^{(i)}\}_{k=1}^P$  consists of  $P$  image frames and the corresponding audio segment  $A^{(i)} = \{a_m^{(i)}\}_{m=1}^Q$  consists of  $Q$  samples, respectively. If the  $i^{th}$  audio-video segment pair represents an event, it is labeled either as event ( $e^{(i)} = 1$ ) or background ( $e^{(i)} = 0$ ). Along with the generic event/background label, each segment of the entire video is labeled with a particular event category. Hence, the set of one-hot

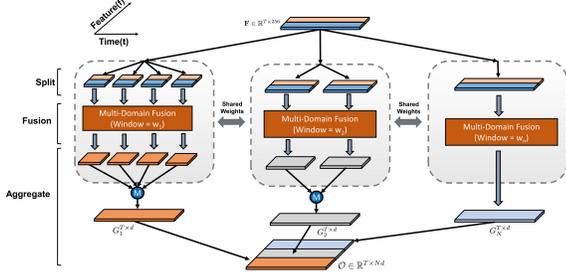


Figure 3. The three phases of the Multi-window Temporal Fusion (MWTF) module. In the **split** phase, the aggregated features are divided into separate temporal blocks based on the window length. In the **fusion** phase, multi-domain fusion is carried out on particular window. In the **aggregation** phase, temporal merging (‘M’) followed by feature concatenation is carried out. The window length can be varied and shared weights can be used in all fusion modules irrespective of window lengths.

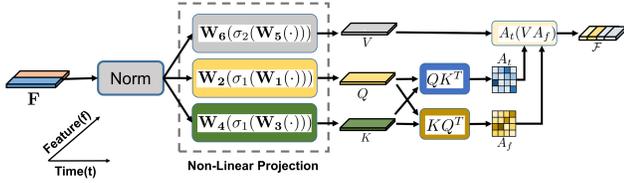


Figure 4. Representations of the **Multi-domain Fusion** process. The attention maps are generated for both of the temporal and feature axis that are applied on the input features through non-linear projection.

encoded labels<sup>1</sup> for the video sequence across  $C$  categories is denoted by,  $\mathbf{Y} = \{y^{(i)}\}_{i=1}^T \in \mathbb{R}^{T \times C}$ . For example, let consider  $\mathbf{Y} = \{b, b, b, e_2, e_2, e_2, e_3, e_3\}$  where  $b$  represents background and  $e_2, e_3$  denote an event of class 2 and 3, respectively. Here, we utilize the class labels ( $y^{(i)}$ ) to generate event label ( $e^{(i)}$ ) as  $\{0, 0, 0, 1, 1, 1, 1\}$  to distinguish between event (1) and background (0).

### 3.2. Contrastive Fine-Tuning on Audio-Image Pairs

We extract positive and negative audio-image pairs from the target dataset where a positive pair corresponds to the same AVE and a negative one represents a mismatch. Initially, we start with the pre-trained Audio and Image encoders from AudioCLIP [12]. Afterwards, we initiate fine-tuning on the extracted audio-image pairs utilizing the InfoNCE loss  $L_{\text{InfoNCE}} = L_I + L_A$ , where  $L_I$  represents the image-to-audio matching loss and  $L_A$  represents the audio-to-image matching loss.  $L_I$  is given by

$$L_I = -\frac{1}{B} \sum_{i=1}^B \sum_{j=1}^B I_{ij} \log \frac{\exp((\mathbf{z}_i^I \mathbf{z}_j^A) / \tau)}{\sum_{k=1}^B \exp((\mathbf{z}_i^I \mathbf{z}_k^A) / \tau)} \quad (1)$$

<sup>1</sup>One-hot encoding changes the dimension over vector space over  $\mathbb{R}$ .

where  $B$  denotes total number of audio-image pairs in a batch,  $\mathbf{z}_i^A, \mathbf{z}_i^I$  represent normalized audio and image features of  $i^{\text{th}}$  pair, respectively,  $I_{ij}$  represents the identity matrix element with  $I_{ii} = 1, \forall i = j$  and  $I_{ij} = 0, \forall i \neq j$ , and  $\tau$  is a trainable temperature. Similarly, we construct the audio-to-image matching loss  $L_A$ .

### 3.3. Video Feature Extraction

The fine-tuned audio and image encoders are deployed to extract the features from the whole video sequence  $S$ . To generate feature map  $\mathbf{v}^i$  from each video segment  $\mathcal{V}^{(i)}$  containing  $K$  image frames, we take the mean of feature maps  $\mathbf{z}_k^I$  for  $k = 1, \dots, T$ . Afterwards, all feature maps from  $T$  video segments are concatenated to generate the video feature  $s_v$  of a particular sequence  $S$ . Similarly, audio feature of each segment  $\mathbf{a}^i$  are concatenated to generate audio feature  $s_a$  of a video sequence, and thus:

$$\mathbf{s}_v = \mathbf{v}^1 \oplus \mathbf{v}^2 \oplus \dots \oplus \mathbf{v}^T \quad (2)$$

$$\mathbf{s}_a = \mathbf{a}^1 \oplus \mathbf{a}^2 \oplus \dots \oplus \mathbf{a}^T \quad (3)$$

where  $\oplus$  denotes feature concatenation and  $T$  denotes the number of segments in a sequence.

### 3.4. Multi-scale Training with Multi-window Temporal Fusion (MWTF) Transformer

For better discrimination of the local feature variations particularly at the event transition edges, it is required to fuse multi-modal features over short temporal windows. However, the general context of the entire video is essential for better event classification. The proposed Multi-Window Temporal Fusion (MWTF) module effectively solves this issue by incorporating multi-domain attention over various temporal scales of the entire video, an approach that addresses the impact of both local and global variations (Figure 3).

Initially, to re-scale the video feature representation ( $s_v$ ) in accordance with the audio representation ( $s_a$ ) for early fusion, we adopt the Audio-Guided Video Attention (AGVA) module presented before [24, 36, 7]. The re-scaled video feature  $\mathbf{s}_{v,a} \in \mathbb{R}^{T \times 1024}$  and corresponding audio feature  $\mathbf{s}_a \in \mathbb{R}^{T \times 1024}$  are processed with separate bidirectional long short term memory (BiLSTM) layers for generating  $\mathbf{v} \in \mathbb{R}^{T \times 256}$  and  $\mathbf{a} \in \mathbb{R}^{T \times 256}$ , respectively. Afterwards, temporal aggregation of audio-visual features is carried out to generate  $\mathbf{F} \in \mathbb{R}^{T \times 512}$ .

In the MWTF module, we incorporate  $N$  sub-modules that operate on different timescales depending on the window length,  $\{w_i \in \mathbb{R}\}_{i=1}^N$ . The basic operations in each sub-module are divided into three stages: split, fusion, and aggregation. In the **split** phase of the  $i^{\text{th}}$  sub-module, the aggregated feature  $\mathbf{F}$  is segmented into  $c_i$  blocks based on window length  $w_i (w_i \times c_i = T)$  that generates  $\{F_i^{(k)} \in \mathbb{R}^{w_i \times 512}\}_{k=1}^{c_i}$ , and thus,

$$F_i^{(k)} = \mathbf{F}[w_i(k-1) : w_i k]; \forall k \in \{1, 2, \dots, c_i\} \quad (4)$$

In addition, it is possible to use varying window lengths in a sub-module totaling the number of time steps  $T$ .

Following the split action, the multi-domain attention guided **fusion** operation is carried out on each  $k^{th}$  feature block  $F_i^{(k)}$  of the  $i^{th}$  sub-module. The multi-domain attention operation is illustrated in Figure 4. Considering the two-domain distribution of each block  $F_i^{(k)}$ , we introduce joint temporal (TA) and feature attention (FA) mechanisms by reusing the weights of similar transformations.

Firstly, each block of  $F_i^{(k)} \in \mathbb{R}^{w_i \times 512}$  features is transformed to generate query vector  $Q_i^{(k)} \in \mathbb{R}^{w_i \times d}$ , key vector  $K_i^{(k)} \in \mathbb{R}^{w_i \times d}$ , and value vector  $V_i^{(k)} \in \mathbb{V}^{w_i \times d}$ , such that,

$$Q_i^{(k)} = \mathbf{W}_2(\sigma_1(\mathbf{W}_1(Norm(F_i^{(k)}))) \quad (5)$$

$$K_i^{(k)} = \mathbf{W}_4(\sigma_1(\mathbf{W}_3(Norm(F_i^{(k)}))) \quad (6)$$

$$V_i^{(k)} = \mathbf{W}_6(\sigma_2(\mathbf{W}_5(Norm(F_i^{(k)}))) \quad (7)$$

$$\forall i \in \{1, \dots, N\}, k \in \{1, \dots, c_i\}$$

where  $\mathbf{W}_1 \in \mathbb{R}^{d \times 256}$ ,  $\mathbf{W}_2 \in \mathbb{R}^{d \times d}$ ,  $\mathbf{W}_3 \in \mathbb{R}^{d \times 256}$ ,  $\mathbf{W}_4 \in \mathbb{R}^{d \times d}$ ,  $\mathbf{W}_5 \in \mathbb{R}^{d \times 256}$ ,  $\mathbf{W}_6 \in \mathbb{R}^{d \times d}$ , and  $\sigma_1(\cdot)$ ,  $\sigma_2(\cdot)$  denote the *tanh* and *Relu* activation, respectively.

Afterwards, we process each query  $Q_i^{(k)}$  and key  $K_i^{(k)}$  vectors on the temporal and feature domains to generate  $\beta_{t,i}^{(k)} \in \mathbb{R}^{w_i \times w_i}$  and  $\beta_{f,i}^{(k)} \in \mathbb{R}^{d \times d}$ , respectively, such that

$$\beta_{t,i}^{(k)} = (Q_i^{(k)} K_i^{(k)T}) / \sqrt{d} \quad (8)$$

$$\beta_{f,i}^{(k)} = (K_i^{(k)} Q_i^{(k)T}) / \sqrt{w_i} \quad (9)$$

$$\forall i \in \{1, \dots, N\}, k \in \{1, \dots, c_i\}$$

Then, we generate the temporal attention map  $A_{t,i}^{(k)} \in \mathbb{R}^{w_i \times w_i}$  by applying *softmax* over the rows (*axis* = 1) of  $\beta_{t,i}^{(k)} \in \mathbb{R}^{w_i \times w_i}$ . In addition, the feature attention map  $A_{f,i}^{(k)} \in \mathbb{R}^{d \times d}$  is generated by applying *softmax* over the columns (*axis* = 0) of  $\beta_{f,i}^{(k)} \in \mathbb{R}^{d \times d}$ .

These multi-domain attention maps are sequentially applied on each axis of  $V_i^{(k)}$  to generate the modified feature map  $\mathcal{F}_i^{(k)} \in \mathbb{R}^{w_i \times d}$  by,

$$\mathcal{F}_i^{(k)} = A_{t,i}^{(k)}(V_i^{(k)} A_{f,i}^{(k)}) \quad (10)$$

$$\forall i \in \{1, \dots, N\}, k \in \{1, \dots, c_i\}$$

Finally, in the **aggregate** phase, the modified feature maps of each  $i^{th}$  block in a sub-module are temporally concatenated to generate  $G_i \in \mathbb{R}^{T \times d}$  such that,

$$G_i = (\mathcal{F}_i^{(1)} || \mathcal{F}_i^{(2)} || \dots || \mathcal{F}_i^{(c_i)}); \forall i \in \{1, \dots, N\} \quad (11)$$

where ‘||’ denotes the feature concatenation along temporal axis.

Afterwards, all the modified feature maps from each sub-module are concatenated along channel axis maintaining temporal relation to generate  $\mathcal{O} \in \mathbb{R}^{T \times Nd}$  as,

$$\mathcal{O} = G_1 \oplus G_2 \oplus \dots \oplus G_{c_i} \quad (12)$$

where  $\oplus$  denotes feature concatenation along channel axis.

### 3.5. Event-Guided Temporal Refinement

As the background class represents miss-aligned audio-visual pairs from all other classes, it often becomes difficult to distinguish them from event classes in case of subtle variations. To enhance the contrast between the event-segments and backgrounds, we introduce supervised event-guided temporal attention (EGTA) over the event region. Following EGTA, we refine the contrasted event segments with an additional stage of single window fusion for better discrimination of the event categories.

To generate EGTA  $\alpha_{att} \in \mathbb{R}^{T \times 1}$ , the fusion vector  $\mathcal{O}$  is passed through a BiLSTM module as,

$$\gamma = BiLSTM(Norm(\mathcal{O})) \quad (13)$$

$$\alpha_{att} = \sigma_3(\mathbf{W}_7(\gamma)) \quad (14)$$

where  $\gamma \in \mathbb{R}^{T \times d'}$ ,  $\mathbf{W}_7 \in \mathbb{R}^{1 \times d'}$ , and  $\sigma_3(\cdot)$  represents the *sigmoid* activation function.

Afterwards, in the refining phase, we apply the EGTA mask  $\alpha_{att}$  over the fusion vector  $\mathcal{O}$  to generate  $\mathcal{O}' \in \mathbb{R}^{T \times Nd}$  by,

$$\mathcal{O}' = \mathcal{O} \odot (\alpha_{att} \mathbf{I}) \quad (15)$$

where  $\mathbf{I} \in \mathbb{R}^{1 \times Nd}$  denotes a broadcasting vector with all ones, and  $\odot$  represents the element-wise multiplication.

As a last step, we incorporate a single window fusion with window length  $w = T$  to generate  $\mathcal{O}_w \in \mathbb{R}^{T \times d}$  by refining event-concentrated vector  $\mathcal{O}'$ . Finally, we obtain the final event category prediction  $y^p \in \mathbb{R}^{T \times C}$  after applying another sequential BiLSTM layer as,

$$y^p = softmax(\mathbf{W}_8(BiLSTM(Norm(\mathcal{O}_w)))) \quad (16)$$

where  $\mathbf{W}_8 \in \mathbb{R}^{C \times d}$  for  $C$  categories of AVE.

### 3.6. Loss Function

To guide the event-attention ( $\alpha_{att}$ ) for refinement, we use an event label loss  $\mathcal{L}_e$ . Also, for the multi-class prediction  $y^p$ , an event category loss  $\mathcal{L}_c$  is incorporated. The total loss  $L$  is obtained by combining  $\mathcal{L}_e$  and  $\mathcal{L}_c$  as,

$$L = \lambda_1 \overbrace{\mathcal{L}_{BCE}(\alpha_{att}, E)}^{\mathcal{L}_e} + \lambda_2 \overbrace{\mathcal{L}_{CE}(y^p, y)}^{\mathcal{L}_c} \quad (17)$$

where  $\lambda_1, \lambda_2$  denote weighting factors,  $E \in \mathbb{R}^{T \times 1}$  denotes the binary event label, and  $y \in \mathbb{R}^{T \times C}$  denotes the one-hot encoded multi-class event categories over the time-frame.

### 3.7. Post-Processing Algorithm during Inference

Due to the sequential nature of AVEs, it is expected that they have high locality over the entire video frame. Therefore, AVEs are typically clustered together and isolated non-AVEs can be viewed as anomalies. We exploit this property to filter the generated event prediction  $y^p$  during inference for obtaining the final prediction  $y_f$ . Here, we consider a window length  $W$  to represent the minimum number of consecutive predictions required for considering any change as an anomaly. As such, all non-matching  $y^p$  values are corrected according to the prevailing one.

## 4. Experiments and Analysis

### 4.1. Experimental Setup

#### Audio-Visual Event Dataset

The Audio-Visual Event Dataset, introduced by Tian *et al.* [24], is widely used for the audio-visual event localization task. The dataset contains 4,143 video clips along with the audio containing 28 different events including daily human activities, instrument performances, animal actions, and vehicle activities. Each video clip is 10 seconds long with temporal start/end annotations for all events. According to existing work [24, 36, 17], training/validation/test splits of 3,309/402/402 are considered for evaluation of the proposed method.

#### Evaluation Metrics

Following existing work [31, 36, 17, 7], we consider the final classification accuracy of multi-class events over the entire video as the evaluation metric. Along with the background, 29 event classes are considered for per-second prediction over the 10s video duration where the video sampling rate varies from 16 to 48. The background category includes all the misaligned audio-visual segments that don't belong to any of the 28 main categories.

#### Implementation Details

We incorporate the pre-trained audio and image encoders from the AudioCLIP [12] framework that fine-tunes the pre-trained CLIP [19] framework with audio-image pairs extracted from the large-scale AudioSet [9] dataset. The

	Method	Accuracy(%)
uni-modal	Audio-based [14]	59.5
	Video-based [23]	55.3
multi-modal (with Co-Attention Fusion)	AVEL [24]	74.7
	DAM [28]	74.5
	PSP [36]	77.8
	AVIN [20]	75.2
	RFJC [7]	76.2
multi-modal (with Transformer Fusion)	AV-Transformer [17]	76.8
	<b>AVE-CLIP (ours)</b>	<b>83.7</b>

Table 1. Performance comparison of the state-of-the-art methods on AVE classification. Various uni-modal methods and multi-modal fusion strategies are compared.

Image/Audio Encoders	Accuracy(%)
with AudioClip-Encoders (w/o Fine-tuning)	81.1
with AudioClip-Encoders (with Fine-tuning)	<b>83.7</b>
without AudioCLIP-Encoders	79.3

Table 2. Impact of the pre-trained AudioCLIP image and audio encoders with contrastive fine-tuning on the AVE-CLIP framework.

audio encoder is the ESResNeXt model [11] that is based on the ResNeXt-50 [3] architecture and the image encoder is a ResNet-50 [13] model. We used combination of four MWTF modules for experiments that is defined empirically. The weights of the hybrid loss function are empirically chosen to be ( $\lambda_1 = 0.3, \lambda_2 = 0.7$ ). For evaluation, each network combination is trained for 300 epochs on 256 AMD EPYC 7742 CPUs with 2 Quadro GV100 and 2 A100-PCIE-40GB GPUs.

### 4.2. Comparison of State-of-the-Art methods

AVE-CLIP is compared against several state-of-the-art methods in Table 1. The multi-modal approaches outperform the uni-modal ones by a great margin. This is expected given the richer context provided by the multi-modal analysis over the uni-modal counterpart. The multi-modal audio-visual fusion strategies play a very critical role in the AVE localization performance. Traditionally, various audio-visual co-attention fusion schemes have been explored to enhance the temporal event features that provide comparable performances. Recently, Lin *et al.* [17] introduced a transformer based multi-modal fusion approach that incorporates instance-level attention to follow visual context over consecutive frames. However, the proposed AVE-CLIP architecture achieves the best performance **with an accuracy of 83.7%** that outperforms the corresponding transformer based approach by **6.9%**. Moreover, the AVE-CLIP provides **5.9%** higher accuracy compared to the best-performing co-attention approach proposed by Zhou *et al.* [36].

Strategies		Accuracy(%)
multi-modal Fusion	with PSP [36]	77.8
	with AV-Transformer [17]	76.8
	with only MWTF and AudioCLIP encoders (ours)	82.0
Temporal Refining	with MWTF + Refiner (ours)	82.5
	with MWTF + EGTA + Refiner (ours)	83.2
<b>Complete Model</b>	<b>with MWTF + EGTA + Refiner + Post Processing (ours)</b>	<b>83.7</b>

Table 3. Impact of various building blocks on the performance of the proposed AVE-CLIP architecture.

Attention Axis	Window Length, w			
	2s	5s	10s	Variable
<b>Temporal</b>	78.1	79.3	81.2	79.4
<b>Feature</b>	78.5	79.1	81.6	79.2
<b>Multi-domain</b>	79.0	79.8	<b>81.6</b>	79.7

Table 4. Accuracy (%) obtained with a **single** fusion block of MWTF module for different attention domains with various window lengths. The variable window length denotes a combination of (3s, 3s, 4s) windows.

Window Combination (with Two Att.)	Accuracy(%)	
	Shared Weights	Independent Weights
10s	-	81.6
10s + 5s	82.4	82.7
10s + 3s*	82.1	82.5
5s + 2s	80.6	81.2
10s + 5s + 3s*	83.2	82.8
10s + 5s + 2s	82.9	83.0
10s + 5s + 3s* + 2s	<b>83.7</b>	83.3

Table 5. Impact of using different combinations of window for the Multi-window Temporal Fusion (MWTF) module. ‘\*’ indicates the use of variable window lengths for 3s.

### 4.3. Ablation Study

To analyze the effect of individual modules in proposed AVE-CLIP, we carried out a detailed ablation study over the baseline approach. The final AVE-CLIP architecture integrates the best performing structures of the building blocks.

#### Effect of Contrastive Fine-tuning in Encoders

We incorporate pre-trained image and audio encoders into the AVE-CLIP framework from AudioCLIP that are subjected to contrastive fine-tuning in training stage 1 (Section 3.2). The effect of these encoders on the final performance of AVE-CLIP is summarized in Table 2. For the baseline comparisons, we adopted the similar VGG-19 backbone pretrained on ImageNet [5] to initially extract the video features, and another VGG-like network pre-trained on AudioSet [9] to extract audio features following existing work [31, 17, 36]. We observe that the best performance of 83.7% is achieved by the AudioCLIP encoders with contrastive fine-tuning which improves accuracy by 4.4% over the uni-modal encoders. Moreover, the contrastive fine-

tuning phase improves the accuracy of AVE-CLIP by 2.6% which shows its effectiveness on AVE localization.

#### Effect of Multi-window Temporal Fusion (MWTF)

For analyzing the effect of the MWTF module in AVE-CLIP, the rest of the modules (temporal refining, post-processing) are replaced with simple fully connected layers followed by a *softmax* classifier. Moreover, to compare with other multi-modal fusion schemes, the PSP based fusion [36] and the AV-Transformer based fusion from [17] are considered. The resulting accuracy is summarized in Table 3. The proposed MWTF module with AudioCLIP encoders provides significant improvements over existing counterparts which shows its effectiveness.

The MWTF module provides a generic framework to operate on various temporal resolutions with shared weights for effectively guiding the multi-modal attention. To analyze the effect of various temporal window lengths, the performance when using a single fusion block in MWTF module is explored in Table 4. The model performs better with increasing window length while achieving the best performance with 10s window length. Moreover, we can observe the consistent performance improvements with multi-domain attention over their single domain counterparts. Although the fusion scheme with a smaller attention window achieves more discriminating features emphasizing high frequency local variations, it misses the global context which is particularly significant for differentiating event categories.

The performance for combinations of varying window lengths is provided in Table 5. By incorporating these fusion windows into the 10s window, performance increases considerably when compared with the baseline. Despite the lower performance for smaller window lengths, these configurations are better at extracting discriminatory features which are particularly critical to determining the sharp edges of event transitions. Hence, the combination of larger and smaller window features is effective for generalizing global low frequency features as well as local high frequency variations at the transition edges. Moreover, we observe that independent weights over different fusion module perform better compared to their shared counterparts with fewer windows. However, when the number of windows increases, such advantages appear to shrink due to the increased complexity of the fusion mechanism.

#### Effect of Event-Guided Temporal Feature Refining

Downstream from the fusion module, AVE-CLIP includes temporal feature refining which consists of two phases: the event-guided temporal attention (EGTA) mask generation, and the corresponding feature refinement. From Table 3, it can be seen the effect of temporal refinement with

	Method	Accuracy(%)
EGTA	with supervision	<b>83.7</b>
	without supervision	83.1
Refiner	window = 10s	<b>83.7</b>
	window = 5s	83.2
	window = (10s +5s)	83.6

Table 6. The effect of event label supervision on the Event-Guided Temporal Attention (EGTA) module and the effect of various fusion window lengths on the Refiner module.

window length	1s	2s	3s	4s	5s
Accuracy(%)	83.2	83.5	<b>83.7</b>	82.9	82.3

Table 7. The effect of different window lengths of the post-processing module during Inference. The best performing modules are considered in AVE-CLIP architecture for the baseline.

EGTA which produces an improvement of 0.7% in accuracy. Moreover, the performance of different combinations for feature refining is provided in Table 6. It is possible to generate the EGTA module without event guided supervision which, as a result, simplifies the loss function to the simple cross-entropy loss. However, with event-label supervision, the model distinguishes the event frames better against backgrounds which in turn provides better performance. For the refiner, the single-window based fusion with  $w = 10s$  generates the best performance since multi-window fusion becomes gradually saturated in this phase.

### Effect of Post-Processing Algorithm

Considering the sequential nature of the events, the proposed post-processing method is found to be very effective for achieving better prediction during inference. As the prediction of event category is generated on a per-second basis, incorrect predictions can be reduced by considering a window of consecutive predictions. The effects of different window lengths on the post-processing method are summarized in Table 7. The best performance is achieved for a 3s window length. With a smaller window length, the effect of filtering is reduced over longer events whereas the larger windows reduces the performance in shorter events.

### 4.4. Qualitative Analysis

The qualitative performance of the proposed AVE-CLIP is demonstrated in Figure 5 for two audio-visual events. For comparative analysis, we have shown the performance of the PSP model ([36]) as well. In the first event, the AVE represents a moving helicopter. Though the helicopter is visible in the first frame, it is a background event due to the absence of the flying helicopter sound. Only the middle three frames capture the AVE through audio-visual correspondence. Our proposed method perfectly distinguishes the helicopter event whereas PSP ([36]) fails at the chal-

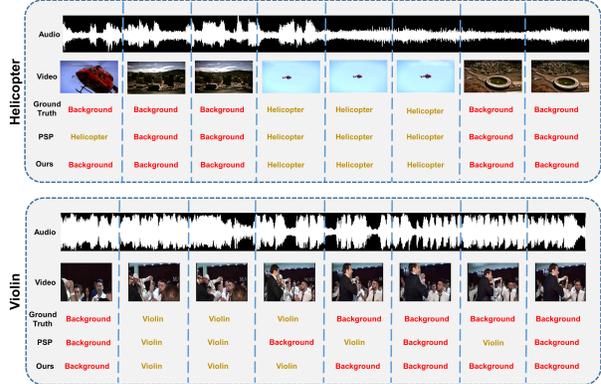


Figure 5. Visual representation of performances of PSP ([36]) and AVE-CLIP on two AVE events (helicopter and violin). The AVE-CLIP performs better to localize event transitions.

lenging first frame. The second event representing a person playing the violin is very challenging given that the violin is hardly visible. Though the sound of violin is present throughout, the image of violin is visible in only few frames that represents the AVE. The PSP ([36]) method generates some incorrect predictions at event transitions. However, the proposed AVE-CLIP perfectly distinguishes the event frames, which demonstrates its effectiveness for generalizing local variations. Furthermore, AVE-CLIP achieves better performance in many challenging cases that demands different scales of temporal reasoning throughout video.

## 5. Conclusion

In this paper, we introduced AVE-CLIP that uses Audio-CLIP encoders in conjunction with a multi-scale temporal fusion based transformer architecture for improving AVE localization performance. We show that the effect of Audio-CLIP encoders with contrastive fine-tuning is significant in AVE-localization for generating improved multi-modal representation. Our results show that local feature variations are essential for event transition detection while global variations are critical for identifying different event classes. The proposed multi-window fusion module exploits both local and global variations with multi-domain attention thereby significantly improving performance. The temporal refining of the event frames simplifies the event classification task which improves multi-class AVE localization performance. Finally, by exploiting the sequential nature of AVEs with a simple post-processing scheme, we were able to achieve state-of-the-art performance on the AVE dataset.

## Acknowledgements

This research was supported in part by the Office of Naval Research, Minerva Program, and a UT Cockrell School of Engineering Doctoral Fellowship.

## References

- [1] Triantafyllos Afouras, Andrew Owens, Joon Son Chung, and Andrew Zisserman. Self-supervised learning of audio-visual objects from video. In *ECCV*, pages 208–224. Springer, 2020.
- [2] Partha Chakraborty, Sabbir Ahmed, Mohammad Abu Yousuf, Akm Azad, Salem A Alyami, and Mohammad Ali Moni. A human-robot interaction system calculating visual focus of human’s attention level. *IEEE Access*, 9:93409–93421, 2021.
- [3] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017.
- [4] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *CVPR*, pages 1–10, 2018.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. IEEE, 2009.
- [6] Guanglong Du, Mingxuan Chen, Caibing Liu, Bo Zhang, and Ping Zhang. Online robot teaching with natural human-robot interaction. *IEEE Transactions on Industrial Electronics*, 65(12):9571–9581, 2018.
- [7] Bin Duan, Hao Tang, Wei Wang, Ziliang Zong, Guowei Yang, and Yan Yan. Audio-visual event localization via recursive fusion by joint co-attention. In *WACV*, pages 4013–4022, 2021.
- [8] Chuang Gan, Yiwei Zhang, Jiajun Wu, Boqing Gong, and Joshua B Tenenbaum. Look, listen, and act: Towards audio-visual embodied navigation. In *ICRA*, pages 9701–9707. IEEE, 2020.
- [9] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In *ICASSP*, pages 776–780. IEEE, 2017.
- [10] Esam Ghaleb, Mirela Popa, and Stylianos Asteriadis. Multimodal and temporal perception of audio-visual cues for emotion recognition. In *ACII*, pages 552–558. IEEE, 2019.
- [11] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. Esresne (x) t-fbsp: Learning robust time-frequency transformation of audio. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2021.
- [12] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. Audioclip: Extending clip to image, text and audio. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 976–980. IEEE, 2022.
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. Cnn architectures for large-scale audio classification. In *ICASSP*, pages 131–135. IEEE, 2017.
- [15] Di Hu, Rui Qian, Minyue Jiang, Xiao Tan, Shilei Wen, Errui Ding, Weiyao Lin, and Dejing Dou. Discriminative sounding objects localization via self-supervised audiovisual matching. *NeurIPS*, 33:10077–10087, 2020.
- [16] Yanghao Li, Saining Xie, Xinlei Chen, Piotr Dollar, Kaiming He, and Ross Girshick. Benchmarking detection transfer learning with vision transformers. *arXiv preprint arXiv:2111.11429*, 2021.
- [17] Yan-Bo Lin and Yu-Chiang Frank Wang. Audiovisual transformer with instance attention for audio-visual event localization. In *ACCV*, 2020.
- [18] Rui Qian, Di Hu, Heinrich Dinkel, Mengyue Wu, Ning Xu, and Weiyao Lin. Multiple sound sources localization from coarse to fine. In *ECCV*, pages 292–308. Springer, 2020.
- [19] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [20] Janani Ramaswamy. What makes the sound?: A dual-modality interacting network for audio-visual event localization. In *ICASSP*, pages 4372–4376. IEEE, 2020.
- [21] Aditya Sanghi, Hang Chu, Joseph G Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshah. Clip-forge: Towards zero-shot text-to-shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18603–18613, 2022.
- [22] Arda Senocak, Hyeonggon Ryu, Junsik Kim, and In So Kweon. Less can be more: Sound source localization with a classification model. In *WACV*, pages 3308–3317, 2022.
- [23] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [24] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. In *ECCV*, pages 247–263, 2018.
- [25] Antigoni Tsiami, Panagiotis Paraskevas Filntisis, Niki Efthymiou, Petros Koutras, Gerasimos Potamianos, and Petros Maragos. Far-field audio-visual scene perception of multi-party human-robot interaction for children and adults. In *ICAASP*, pages 6568–6572. IEEE, 2018.
- [26] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3835–3844, 2022.
- [27] Yifan Wu, Roshan Ayyalasomayajula, Michael J Bianco, Dinsh Bharadia, and Peter Gerstoft. Sslide: Sound source localization for indoors based on deep learning. In *ICASSP*, pages 4680–4684. IEEE, 2021.
- [28] Yu Wu, Linchao Zhu, Yan Yan, and Yi Yang. Dual attention matching for audio-visual event localization. In *ICCV*, pages 6292–6300, 2019.

- [29] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *CVPR*, pages 9068–9079, 2018.
- [30] Haoming Xu, Runhao Zeng, Qingyao Wu, Minghui Tan, and Chuang Gan. Cross-modal relation-aware networks for audio-visual event localization. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 3893–3901, 2020.
- [31] Hanyu Xuan, Lei Luo, Zhenyu Zhang, Jian Yang, and Yan Yan. Discriminative cross-modality attention network for temporal inconsistent audio-visual event localization. *IEEE Transactions on Image Processing*, 30:7878–7888, 2021.
- [32] Hanyu Xuan, Zhenyu Zhang, Shuo Chen, Jian Yang, and Yan Yan. Cross-modal attention network for temporal inconsistent audio-visual event localization. In *AAAI*, volume 34, pages 279–286, 2020.
- [33] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022.
- [34] Hang Zhao, Chuang Gan, Andrew Rouditchenko, Carl Vondrick, Josh McDermott, and Antonio Torralba. The sound of pixels. In *ECCV*, pages 570–586, 2018.
- [35] Hang Zhou, Xudong Xu, Dahua Lin, Xiaogang Wang, and Ziwei Liu. Sep-stereo: Visually guided stereophonic audio generation by associating source separation. In *ECCV*, pages 52–69. Springer, 2020.
- [36] Jinxing Zhou, Liang Zheng, Yiran Zhong, Shijie Hao, and Meng Wang. Positive sample propagation along the audio-visual event line. In *CVPR*, pages 8436–8444, 2021.
- [37] Zheng Zhu, Wei Wu, Wei Zou, and Junjie Yan. End-to-end flow correlation tracking with spatial-temporal attention. In *CVPR*, pages 548–557, 2018.