

# ECLIPSE: Efficient Long-range Video Retrieval using Sight and Sound

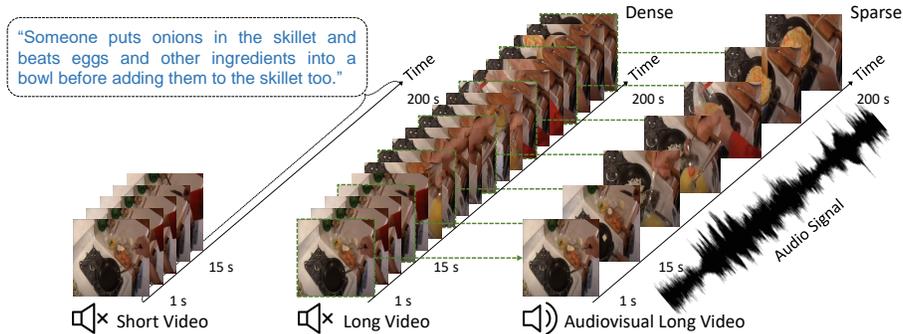
Yan-Bo Lin, Jie Lei, Mohit Bansal, and Gedas Bertasius

Department of Computer Science  
University of North Carolina at Chapel Hill  
{yblin, jielei, mbansal, gedas}@cs.unc.edu

**Abstract.** We introduce an audiovisual method for long-range text-to-video retrieval. Unlike previous approaches designed for short video retrieval (e.g., 5-15 seconds in duration), our approach aims to retrieve minute-long videos that capture complex human actions. One challenge of standard video-only approaches is the large computational cost associated with processing hundreds of densely extracted frames from such long videos. To address this issue, we propose to replace parts of the video with compact audio cues that succinctly summarize dynamic audio events and are cheap to process. Our method, named ECLIPSE (Efficient CLIP with Sound Encoding), adapts the popular CLIP model to an audiovisual video setting, by adding a unified audiovisual transformer block that captures complementary cues from the video and audio streams. In addition to being  $2.92\times$  faster and  $2.34\times$  memory-efficient than long-range video-only approaches, our method also achieves better text-to-video retrieval accuracy on several diverse long-range video datasets such as ActivityNet, QVHighlights, YouCook2, DiDeMo, and Charades. Our code is available at <https://github.com/GenjiB/ECLIPSE>

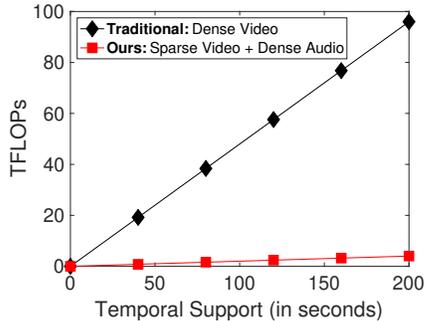
## 1 Introduction

Fueled by the growing availability of video data, the last few years have witnessed remarkable progress in text-to-video retrieval [1,2,3,4,5,6,7,8]. However, modern video retrieval systems are predominantly designed for very short videos (e.g., 5-15 seconds in length). In contrast, the majority of real-world videos often capture complex human actions, which may last several minutes or even hours. For example, consider yourself performing a complex activity of making Japanese Souffle Pancakes, which may take a couple of hours. In a scenario when you forget some of the steps in the recipe, it would be helpful to retrieve a relevant several-minute-long video segment demonstrating how to perform those steps. However, the traditional short-range video retrieval models would struggle with this task due to their inability to analyze longer videos. Combining the strengths of audio and video modalities, we aim to address this limitation by proposing an efficient audiovisual text-to-video retrieval system focused on long-range videos.



**Fig. 2.** Comparison of different high-level frameworks for long-range text-to-video retrieval. Most traditional text-to-video retrieval methods (**Leftmost Column**) are designed for short videos (e.g., 5-15 seconds in duration). Adapting these approaches to several-minute long videos by stacking more input frames (**Middle Column**) is impractical due to excessive computational cost. Instead, our proposed framework operates on sparsely sampled video frames and dense audio cues, which are cheaper to process (**Rightmost Column**). In addition to being more efficient, our framework also achieves higher text-to-video retrieval accuracy than standard video-only approaches.

Among prior vision-and-language methods [2,1,9,10,11,12,13,14], CLIP [15] stands out as one of the most widely adopted models. Several recent approaches extended CLIP to video [16] by independently processing individual video frames and then averaging their predictions across time. However, these approaches are often impractical in the long-range video setting because of the large computational cost required to process hundreds of densely extracted video frames (See Figure 1). Furthermore, while video modality is rich in the information it stores, it also has high informational redundancy (i.e., the video content often changes little in neighboring frames). In contrast, audio can compactly capture information related to human actions [17,18], objects [19,20,21], scenes [22,23] and other complex events [24] while also being cheaper to process [25] than the raw video. For instance, consider a video of a person frying the eggs in a pan. In this example, most of the relevant visual information (e.g., kitchen stove, pan, eggs, etc.) can be captured in just a few video frames, while the temporal dynamics in the scene can be succinctly encoded in the audio stream (e.g., the sounds of the eggs sizzling in a pan, etc.).



**Fig. 1.** Our audiovisual framework scales to long videos more efficiently than dense video-only approaches.

Based on this motivation, we introduce ECLIPSE, an Efficient CLIP with Sound Encoding. Instead of processing many densely-extracted frames from a long video (the middle column in Figure 2), our framework leverages complementary audio and video cues by operating on sparsely sampled video frames accompanied by dense audio (the rightmost column in Figure 2). We demonstrate that compared to dense video-only approaches, our framework is not only more efficient but it is also more accurate.

Our approach adapts CLIP to long-range videos by incorporating a dual-pathway audiovisual attention block into every layer of the Transformer backbone. Such a cross-modal attention mechanism allows our model to (i) incorporate long-range temporal cues from the audio stream into the visual representation, and (ii) conversely inject rich visual features from the video modality into audio representation for improved audio feature expressivity. Such bi-directional exchange of information ensures that both modalities benefit from each other in order to maximize the performance of the downstream application (i.e., long-range text-to-video retrieval). Additionally, we demonstrate that our audiovisual attention block can be easily incorporated into pretrained transformer models such as CLIP [15] without re-training the new model from scratch.

We validate ECLIPSE on several diverse long-range video retrieval benchmarks and show that it achieves state-of-the-art results on ActivityNet [26], QVHighlights [27], DiDeMo [28], YouCook2 [29], and Charades [30] while being  $2.92\times$  faster and  $2.34\times$  memory-efficient than long-range video-only methods.

In summary, our contributions are threefold. First, we propose ECLIPSE, an audiovisual adaptation of CLIP that leverages complementary video and audio cues for long-range video retrieval. Second, we demonstrate that compared to long-range video-only approaches, our audiovisual framework leads to better video retrieval results at a reduced computational cost. Lastly, we provide comprehensive ablation studies investigating the success factors of ECLIPSE .

## 2 Related Work

**Text-to-Video Retrieval.** The association of text descriptions and videos provides rich supervisory signals for developing robust text-to-video retrieval systems. Self-supervised learning approaches in this area achieve impressive results using contrastive loss [3,31,32,33,34,35,36], masked language modeling [37,38,39,40,41], or masked feature prediction [42]. Additionally, several prior methods propose to incorporate rich audio/speech information for video-and-text representations learning, either by fusing cross-modal signals [6,43,44] or masking inputs from different modalities during training [7,45]. Furthermore, with large-scale pre-training on millions of image and text pairs, CLIP [15] has achieved impressive results on a wide array of vision-and-language tasks. Recently, CLIP-based approaches [46,47,16,48,49,13,14,50,51] have also been used in video by aggregating image-level outputs across different time steps.

Unlike these prior methods, which are designed for short-range videos (e.g., 5-15s), we aim to design an audiovisual framework for retrieving long videos (e.g.,

several minutes in length). Compared to the existing CLIP-based approaches, which are difficult to adapt to long videos due to the large computational cost of processing many densely-extracted video frames, we propose to leverage compact audio cues in order to reduce the need for the costly video modality. This enables efficient adaptation of CLIP to long video retrieval.

**Audiovisual Learning.** Audio and video synchronization is commonly used for self-supervised audio-visual learning [52,53,54,19,22,55,23,56,57,58,59]. Aside from self-supervised learning, many recent methods were proposed for audiovisual event classification [60,25,17,24,61]. Furthermore, the recent popularity of Transformers [62,63,64,65] have enabled a wide array of architectures for jointly modeling audio and video data [66,67,68,69,70,71,72]. Compared to these prior approaches, our approach focuses on efficient long-range text-to-video retrieval. Specifically, we aim to leverage audio cues in order to reduce the computational cost of processing long videos.

**Long Sequence Modeling.** Recent work in the natural language processing (NLP) domain [73,74,75] proposed to approximate the self-attention operator for long sequence modeling. While these approaches are effective in NLP, they are still very costly in the video domain due to the high dimensionality of video inputs. Furthermore, as demonstrated by recent work in the video domain [75], such approximation techniques lead to a substantial accuracy drop while producing limited efficiency gains for video recognition tasks. Additionally, we note that these approximation mechanisms are often incompatible with pretrained vision-and-language models such as CLIP (due to different network architectures).

### 3 ECLIPSE: Efficient CLIP with Sound Encoding

Our goal is to design an efficient framework that leverages audiovisual cues for long-range text-to-video retrieval. Instead of processing many densely-extracted frames from a long video, which is costly, our framework operates on sparsely sampled video frames accompanied by dense audio. We adapt CLIP to long-range videos by adding a dual-pathway audiovisual attention block into every layer of the Transformer backbone. Our video retrieval framework consists of three high-level components: (i) multimodal input embeddings, (ii) an audiovisual backbone for processing video and audio modalities, and (iii) a contrastive video-to-text matching objective. Below we provide more details behind each of these components. We also illustrate our framework in Figure 3.

#### 3.1 Obtaining Multimodal Input Embeddings

**Video, Audio and Text Inputs.** Our framework takes audio, video, and text modalities as its inputs. For video modality, we consider video clips  $X \in \mathbb{R}^{T \times H \times W \times 3}$  consisting of  $T$  RGB frames of size  $H \times W$ , sampled uniformly from the whole input video. For audio, we use  $T$  audio spectrograms  $Z \in \mathbb{R}^{T \times M \times C}$ , each spanning  $t$  seconds and centered around each of  $T$  video frames. Here,  $M$  and  $C$  depict spatial spectrogram dimensions. Lastly, the text is represented as a

sequence  $y = (y_1, \dots, y_L)$  where  $y_i$  represents a distinct word in the textual video description and  $L$  is the length of the description (i.e., the number of words).

**Video Patch Decomposition.** Following the ViT [63], we decompose each frame into  $N$  non-overlapping patches, each of size  $P \times P$ , and flatten these patches into vectors  $\mathbf{x}_{(p,t)} \in \mathbb{R}^{3P^2}$  where  $p = 1, \dots, N$  denotes spatial locations and  $t = 1, \dots, T$  indicates a frame index.

**Video Patch Embeddings.** Video patches from each frame  $\mathbf{x}_{(p,t)}$  are linearly mapped into vectors  $\mathbf{v}_{(p,t)}^{(0)} \in \mathbb{R}^d$ , for  $p = 1 \dots N$ , and  $t = 1 \dots T$ . Afterward, we also augment each visual token with spatiotemporal position information as is done in [76]. A specialized CLS token  $\mathbf{v}_{cls}^{(0)}$  is prepended to the visual sequence of each frame. Finally, the embeddings  $\mathbf{V}^{(0)} \in \mathbb{R}^{T \times (N+1) \times d}$  are used as visual inputs to our ECLIPSE model.

**Audio Embeddings.** Given an audio spectrogram  $Z_t \in \mathbb{R}^{M \times C}$ , an audio encoder maps it into audio embeddings  $\mathbf{A}_t^{(0)} \in \mathbb{R}^d$  for each timestep  $t = 1 \dots T$  where as before,  $T$  denotes the number of video frames. We note that the audio encoder can be either a CNN [77,78,79] or a Transformer [80,81]. Afterward, the audio embeddings  $\mathbf{A}^{(0)} \in \mathbb{R}^{T \times d}$  are fed into ECLIPSE together with the visual tokens  $\mathbf{V}^{(0)} \in \mathbb{R}^{T \times (N+1) \times d}$ .

**Text Embeddings.** We use a pretrained CLIP [15] text encoder to embed a textual video description  $y = (y_1, \dots, y_L)$  into a textual embedding  $\mathbf{g} \in \mathbb{R}^d$  where  $\mathbf{g}$  corresponds to the CLS token of a given text sequence.

### 3.2 Audiovisual Attention Block

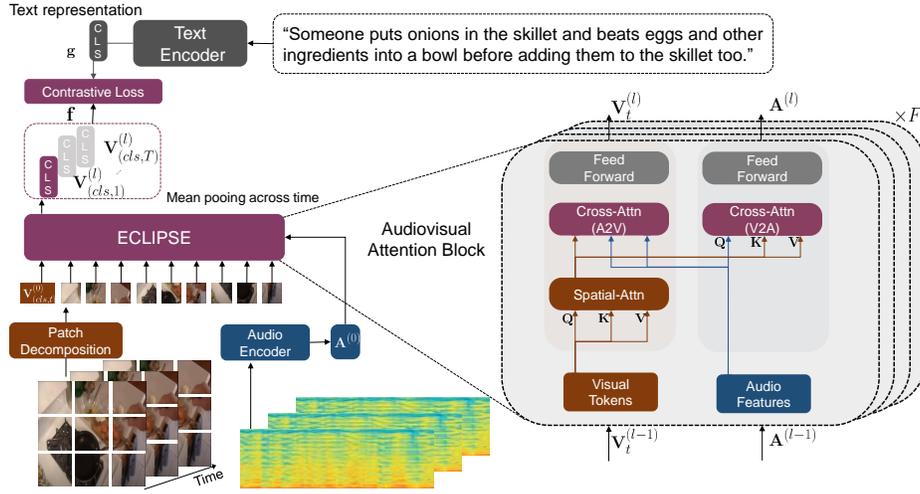
Although videos contain rich information, they are also redundant and costly to process. In contrast, audio is more compact and cheaper. Thus, we propose an audiovisual attention block that gradually incorporates relevant audio cues into the visual representation. Our audiovisual attention block consists of three distinct attention schemes: (i) spatial visual attention, (ii) audio-to-video attention, and (iii) video-to-audio attention. We next describe each of these attention schemes in more detail.

**Multi-Head Self-Attention.** All of our three attention schemes are implemented using a standard multi-head self-attention:

$$\text{MHA}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax} \left( \frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}} \right) \mathbf{V}, \quad (1)$$

where  $\mathbf{Q}, \mathbf{K}, \mathbf{V}$  are the query, key and value matrices obtained using learnable projection weights  $\mathbf{W}^Q, \mathbf{W}^K, \mathbf{W}^V \in \mathbb{R}^{d \times d}$  respectively. With this formal description of the MHA function, we can now proceed to the definitions of the three attention schemes in our audiovisual attention block.

**Spatial Attention.** In order to preserve the pretrained network structure of CLIP, we use an identical spatial attention scheme as in CLIP. Intuitively, spatial attention enables our model to obtain discriminative frame-level representation by aggregating relevant information from the visual tokens in the individual



**Fig. 3.** We adapt CLIP [15] to long-range text-to-video retrieval by adding an efficient audiovisual attention block into the Transformer architecture. First, we obtain fixed dimensional text, audio, and visual feature embeddings. Afterward, the visual and audio embeddings are fed into our ECLIPSE audiovisual backbone, which injects relevant audio information to video and vice-versa. This is accomplished using a dual-pathway audiovisual attention block (illustrated on the right), which is stacked on top of each other  $F$  times. Afterward, the audiovisual video segments are aggregated using temporal pooling, and the model is optimized by maximizing the similarity between audiovisual and textual embeddings using a contrastive loss function.

video frames. We can implement this scheme using our previously defined MHA function as:

$$\mathbf{S}_t^{(\ell)} = \text{MHA}(\mathbf{V}_t^{(\ell-1)}, \mathbf{V}_t^{(\ell-1)}, \mathbf{V}_t^{(\ell-1)}) + \mathbf{V}_t^{(\ell-1)}. \quad (2)$$

Here,  $\mathbf{S}_t^{(\ell)} \in \mathbb{R}^{(N+1) \times d}$  is our newly computed spatial self-attention representation for frame  $t$ , and  $\mathbf{V}_t^{(\ell-1)}$  is a visual patch representation for frame  $t$  from the previous transformer layer  $l-1$ , which is used as input to the transformer layer  $l$ . Note that in the spatial self-attention, the multi-head self-attention is applied independently for each of  $T$  video frames. As discussed above, this enables us to preserve the network structure of the original CLIP model, which is essential for good text-to-video retrieval performance. For brevity, we omit the layer normalization operation, which is applied to  $\mathbf{V}_t^{(\ell)}$  before feeding it to the spatial attention block. The right part of Figure 3 provides a visual illustration of where spatial attention fits within our audiovisual attention block.

**Audio-to-Video Attention (A2V).** To efficiently incorporate temporal audio cues into static video frame representation, we use an audio-to-video (A2V) attention mechanism, which is also illustrated in the right part of Figure 3 (labeled as Cross-Attn A2V module). This operation can be written as:

$$\mathbf{V}_t^{(\ell)} = \text{MHA}(\mathbf{S}_t^{(\ell-1)}, \mathbf{A}^{(\ell-1)}, \mathbf{A}^{(\ell-1)}) + \mathbf{S}_t^{(\ell-1)}. \quad (3)$$

Here,  $\mathbf{A}^{(\ell-1)} \in \mathbb{R}^{T \times d}$  depicts our previously defined audio representation at layer  $l - 1$ , and  $\mathbf{S}_t^{(\ell-1)} \in \mathbb{R}^{(N+1) \times d}$  denotes a spatial video representation at timestep  $t$  computed using our previously defined spatial attention block. Intuitively, the new visual representation  $\mathbf{V}_t^{(\ell)}$  is computed as a weighted summation of the audio features, which enables the model to incorporate long-range audio cues into the visual features. Furthermore, because the audio representation is compact, the operation above can be implemented efficiently.

**Video-to-Audio Attention (V2A).** Conversely, to inject rich visual information into compact audio features, we use a video-to-audio (V2A) attention mechanism (illustrated in Figure 3 as Cross-Attn V2A module). We implement this attention scheme as:

$$\mathbf{A}_t^{(\ell)} = \text{MHA}(\mathbf{A}_t^{(\ell-1)}, \mathbf{S}_t^{(\ell-1)}, \mathbf{S}_t^{(\ell-1)}) + \mathbf{A}_t^{(\ell-1)}. \quad (4)$$

At a high level, the operation above computes a new audio feature representation for each timestep  $t$  as a weighted combination of all the visual token features at timestep  $t$ . This allows us to improve the richness of the audio representation.

**Final Audiovisual Representation.** Following CLIP4Clip [16], we stack our audiovisual attention block  $F$  times ( $F$  typically being set to 12). Afterward, we perform temporal pooling over the CLS tokens across all video frames, to obtain the final audiovisual representation  $\mathbf{f} \in \mathbb{R}^d$ .

### 3.3 Loss Function

We use the same contrastive video-to-text matching loss as in [16]. Specifically, we compute the similarity between text and video using a normalized dot product between the two embeddings  $\mathbf{f}$  and  $\mathbf{g}$ . We consider the matching text-video pairs in a given batch as positive samples and all the other pairs in that same batch as negative samples. To train our model, we minimize the sum of the video-to-text and text-to-video matching losses [16].

### 3.4 Implementation Details

Our ECLIPSE follows CLIP4Clip [16] setting where text encoder and visual encoder are initialized with CLIP weights [15]. Specifically, we initialize the spatial attention weights with the weights from CLIP. We also use CLIP weights to initialize both of our cross-modal attention blocks. We attach zero-initialized linear projection layers to the outputs of both cross-modal attention blocks so that the initial outputs of these blocks would be set to zero. Unless otherwise noted, for all of our experiments, we use a ViT-B/32 with uniformly sampled 32-frame inputs spanning the whole input video. The visual frames are extracted at 3 fps. We implement ECLIPSE using Pytorch [82] and conduct the training on four NVIDIA A6000 GPUs. For a fair comparison with the baselines, we set

the batch size to 64. For audio encoder, we use ResNet-18 [83] pre-trained on VGGSound [79]. We sample 10-second audio clips in the neighborhood around the sampled video frame and process the raw audio into a spectrogram as is done in [79]. We train our model with Adam optimizer [84] and set the learning rate to  $1e-7$  for text encoder and spatial attention in Eq. 2. The frame-level *CLS* tokens are averaged to obtain the final video embedding.

Furthermore, the maximum text input is set to 64 tokens for DiDeMo and QVHighlights, and 128 for ActivityNet Captions and YouCook2.

## 4 Experimental Setup

### 4.1 Downstream Datasets

We evaluate ECLIPSE on five diverse long-range datasets: ActivityNet Captions [26], QVHighlights [27], DiDeMo [28], YouCook2 [29], and Charades [30].

**ActivityNet Captions** [26] consists of 20,000 YouTube human activity videos, each annotated with temporally localized sentence descriptions, with a total of 100,000 sentences. The average video length is 180 seconds, which makes this dataset well suited for verifying our model’s ability to retrieve long-range videos. We follow [16,85,6,2] to evaluate paragraph-to-video retrieval, where we concatenate all the sentence descriptions to form a paragraph. Since there is no test set provided, we evaluate the video retrieval results on the *val1* split.

**QVHighlights** [27] contains 3,164 videos (10,148 clips) from YouTube, covering a wide range of topics, including everyday activities in lifestyle vlog videos to social and political activities in news videos. Each video is temporally annotated with multiple text queries describing distinct spans of the video. The average video length is around 8 minutes. The original dataset is created for moment localization and highlight detection. Here we re-purpose it for text-to-video retrieval by evaluating it in paragraph-to-video retrieval setup as ActivityNet Captions. We use the standard splits for training, validation, and testing.

**DiDeMo** [28] contains 10,464 Flickr videos with 40,543 temporally localized sentences. The average video length is 30 seconds. Similar to ActivityNet Caption, we evaluate paragraph-to-video retrieval on DiDeMo. We use the standard splits for training, validation, and testing.

**YouCook2** [29] consists of 2,000 videos capturing 89 complex recipes with total duration of 176 hours. The average video length is 5.26 minutes. Each video is annotated with multiple temporally localized captions. Similar to ActivityNet Captions, we evaluate all methods in the paragraph-to-video retrieval setting. We use standard splits for training, validation, and testing.

**Charades** [30] contains 9,848 videos with the corresponding textual descriptions. The average video length is about 28 seconds. We use standard train and test splits for training and testing.

### 4.2 Evaluation Metrics

We use standard video retrieval evaluation metrics [9,16] such as text-to-video  $R@1$ ,  $R@5$ ,  $R@10$ , and mean rank (MnR) to validate the effectiveness of our

**Table 1. ActivityNet Captions.** We compare ECLIPSE with previous video retrieval methods. In the column Pretrain, C,G,H,W,CW,V denote COCO Captions [86], Visual Genome Captions [87], HowTo100M [88], WIT [15], CC3M [89]+WebVid2M [3] and VGGSound [79] datasets respectively. The performance is evaluated using text-to-video retrieval R@1, R@5, R@10 and MnR metrics. ECLIPSE achieves the best reported accuracy on this benchmark. We also note that using a stronger visual backbone (i.e., ViT-B/16 vs. ViT-B/32) also leads to better video retrieval performance.

Method	Pretrain	Frames	R@1 ↑	R@5 ↑	R@10 ↑	MnR ↓
CE [43]	-	-	18.2	47.7	-	23.1
ClipBERT [9]	C+G	40	21.3	49.0	-	-
TT-CE [1]	-	64	23.4	57.2	-	-
MMT [6]	H	-	28.7	61.4	-	16.0
FiT [3]	CW	32	28.9	57.8	71.2	-
SSB [31]	H	-	29.2	61.6	-	-
HiT [2]	H	-	29.6	60.7	-	-
CLIP4Clip [16] (ViT-B/32)	W	64	40.7	71.8	83.4	8.2
<b>ECLIPSE</b> (ViT-B/32)	W+V	32	<b>42.3</b>	<b>73.2</b>	<b>83.8</b>	<b>8.2</b>
<b>ECLIPSE</b> (ViT-B/16)	W+V	32	<b>45.3</b>	<b>75.7</b>	<b>86.2</b>	<b>6.2</b>

ECLIPSE model. Since our model is built on CLIP, which is pretrained on a large-scale image-and-text dataset [15], the comparisons with some of the previous methods are not directly applicable. Therefore, in all of our evaluations, we use a publicly available state-of-the-art CLIP4Clip [16] video retrieval system as our primary baseline.

## 5 Results and Analysis

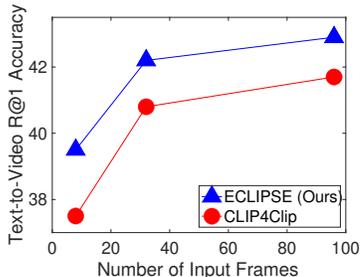
### 5.1 ActivityNet Captions

**Comparison to the State-of-the-Art.** In Table 1, we report the results of our method on ActivityNet Captions. These results indicate several interesting findings. First, we notice that the gap between CLIP-based methods (i.e., CLIP4Clip, ECLIPSE) and other previous approaches is significant ( $> 10\%$  in R@1 metric). This result justifies our motivation to build on the powerful CLIP model. Second, our results indicate that ECLIPSE outperforms CLIP4Clip by a substantial margin (1.6% in R@1), which suggests the usefulness of temporal audio cues. Third, we also note that unlike CLIP4Clip, which operates on 64 frame inputs, ECLIPSE achieves higher accuracy while processing fewer frames (i.e., 32). Lastly, we show that using a stronger visual backbone (i.e., ViT-B/16 vs. ViT-B/32) leads to improved video retrieval performance.

**Accuracy vs. Number of Frames.** We next investigate the trade-off between video retrieval accuracy and the number of input frames. In Figure 4, we plot the long-range text-to-video retrieval accuracy (i.e., R@1) as a function of the number of input frames. Based on these results, we observe that

ECLIPSE consistently outperforms CLIP4Clip in 8, 32 and 96-frame regimes. Furthermore, we notice that ECLIPSE achieves higher accuracy than CLIP4Clip even when operating on a much smaller number of video frames (e.g., 32 vs 96).

**Computational Cost Analysis.** We note that compared to the video-only approaches, our proposed ECLIPSE uses an additional audio modality. However, we would also like to emphasize that we use audio to improve the efficiency of the costly video-only approaches rather than merely improving the absolute video retrieval accuracy. In Table 2, we compare the computational cost of a 96-frame CLIP4Clip with our 32-frame ECLIPSE. Based on these results, we observe that ECLIPSE uses  $2.3\times$  less GPU memory, runs  $2.92\times$  faster, and achieves better accuracy (42.3 vs. 41.7) than CLIP4Clip. This suggests that replacing the costly video modality with the audio makes our retrieval framework more efficient and also improves its accuracy.



**Fig. 4.** We compare ECLIPSE with CLIP4Clip with a varying number of frames. Our method outperforms CLIP4Clip while using the same number or even fewer frames.

## 5.2 Results on Other Long-range Datasets

Next, we validate our approach on four other long-range video datasets: QVHighlights [27] (QVH), DiDeMo [28], YouCook2 [29] (YC2), and Charades [30]. Since long-range video retrieval is a relatively unexplored subarea of research, we note that QVHighlights, YouCook2, and Charades are not formally used for the long-range video retrieval task. However, all three of these datasets contain (i) long videos and (ii) multiple annotated text descriptions of short-term segments within each long video. Thus, to re-purpose these datasets for long-range text-to-video retrieval, we follow the protocol of ActivityNet Captions [26]. Specifically, we concatenate the textual descriptions of all short-term segments in a given

**Table 2.** We compare the computational cost of a 32-frame ECLIPSE with a 96-frame CLIP4Clip [16] on ActivityNet Captions. Both methods are built using a ViT-B/32 architecture. Despite using fewer frames, ECLIPSE outperforms CLIP4Clip. Additionally, our method uses  $2.3\times$  less GPU memory, runs  $2.92\times$  faster and is generally more efficient as indicated by the number of GFLOPs (i.e., 827 vs 1251).

Method	Num. Frames	Inference GFLOPs ↓	GPU Mem. (in MB) ↓	Samples per Sec. ↑	T2V R@1 ↑
CLIP4Clip	96	1251	24,802	17.39	41.7
<b>ECLIPSE</b>	32	<b>827</b>	<b>10,637</b>	<b>50.93</b>	<b>42.3</b>

**Table 3.** Our results on QVHighlights [27] (QVH), YouCook2 [29] (YC2), Charades [30] and DiDeMo [28] using the  $R@1$  T2V metric. A 32-frame ECLIPSE with a ViT-B/32 backbone outperforms prior approaches while also being more efficient.

Method	Pretrain	Frames	QVH	DiDeMo	YC2	Charades	GFLOPs
ClipBERT [9]	C+G	32	43.2	20.4	29.8	6.7	-
FiT [3]	CW	32	55.0	35.8	32.2	11.9	1426
CLIP4Clip [16]	W	96	70.2	42.5	37.6	13.9	1251
<b>ECLIPSE</b>	W+V	32	<b>70.8</b>	<b>44.2</b>	<b>38.5</b>	<b>15.7</b>	<b>827</b>

long video and treat it as a paragraph-to-video retrieval task similar to [26]. In our comparisons, we also include other recent video retrieval methods such as ClipBERT [9], Frozen in Time (FiT) [3], and CLIP4Clip [16].

In Table 3, we show that a 32-frame ECLIPSE with a ViT-B/32 backbone outperforms prior methods on all four datasets. Additionally, we point out that our method is more efficient than both FiT [3] and CLIP4Clip [16] (827 vs. 1251 vs. 1426 in GFLOPs).

**Potential Overlap between Audio and Video Datasets.** Next, we want to verify that the videos used to pretrain our audio encoders were not present in the test sets of the video retrieval benchmarks. Upon our investigation, we found that the overlap between VGGSound, which was used to pretrain our best audio model, and ActivityNet Captions was small, i.e., 42 out of 4,926 videos (0.8%). Furthermore, there were no overlaps between the VGGSound and all of the other datasets. To validate that our original conclusions on ActivityNet still hold, we conducted additional experiments on the deduplicated ActivityNet dataset where the overlapping test videos were removed. We used the same CLIP4Clip and ECLIPSE methods as in Table 1. We report that ECLIPSE achieved 42.3% T2V  $R@1$  accuracy while CLIP4Clip obtained 40.8%. Both of these results are almost identical to the results in Table 1 (i.e., 42.3% and 40.7% respectively).

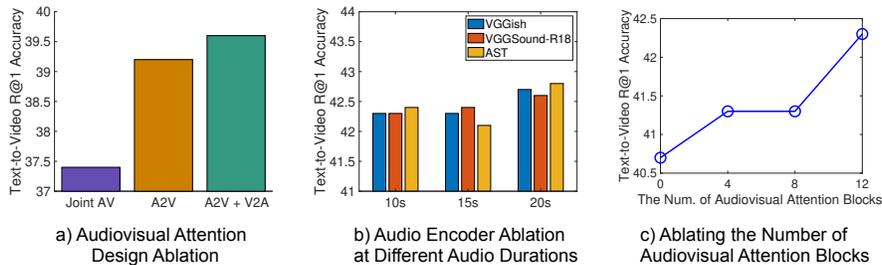
### 5.3 Ablation Studies

Next, we investigate how different design choices of our model affect the long-range video retrieval accuracy on the ActivityNet Captions dataset [26].

**Audiovisual Block Design.** First, we validate the effectiveness of our audiovisual attention block by comparing it to (i) a joint audiovisual attention that processes concatenated video and audio tokens, (ii) the variant of our model that only uses audio-to-video (A2V) attention (Eq. 3) and (iii) our final model that uses both audio-to-video (A2V) and video-to-audio (V2A) attentions (Eq. 3 and Eq. 4). For efficiency, all models are trained using 8-frame inputs.

From the results in Figure 5a, we observe that using a bi-directional audiovisual attention (i.e., both audio-to-video (A2V) and video-to-audio (V2A)) leads to the best  $R@1$  text-to-video retrieval accuracy on ActivityNet.

**Different Audio Encoders.** Next, we study how different audio encoders affect the video retrieval performance of our model. Specifically, we experiment



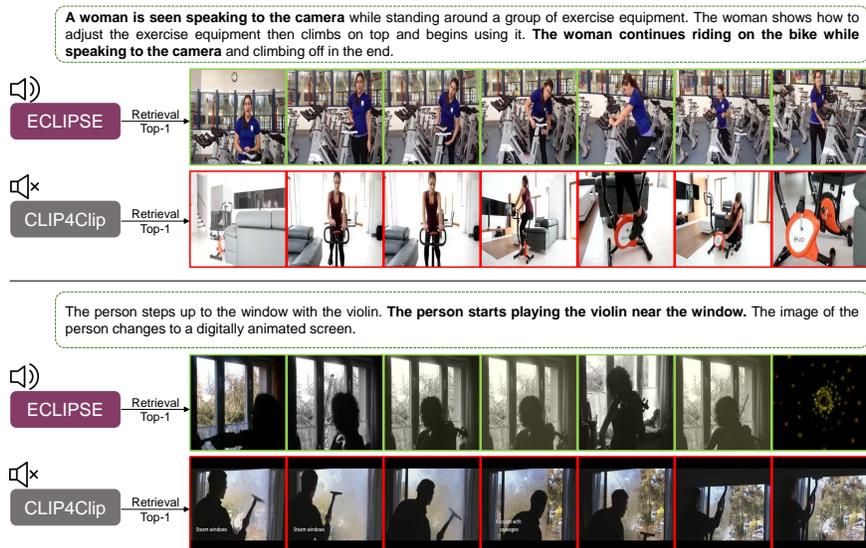
**Fig. 5.** (a) In the left subfigure, we study different audiovisual block design. Joint AV refers to standard self-attention applied to concatenated audio and video tokens. A2V refers to a single cross-modal audio-to-video attention block (Eq. 3). Lastly, A2V+V2A depicts our dual-pathway attention block design (Eq. 3 and Eq. 4). Based on these results, we observe that dual-pathway attention achieves the best performance. For efficiency, we use 8 frame inputs for these experiments. (b) In the middle subfigure, we also investigate different audio encoders applied to different duration audio segments. These results indicate that (i) longer audio typically improves the performance, (ii) ECLIPSE is robust to different audio encoders. (c) In the right subfigure, we study video retrieval accuracy as a function of the number of audiovisual attention blocks. Based on these results, we observe that injecting our proposed audiovisual attention block into every layer of our 12-layer ECLIPSE model leads to the best performance.

with CNN-based audio encoders such as VGGish [78] and VGGSound [79], and also a transformer-based audio encoder AST [81]. Our results in Figure 5b suggest that our framework is robust to the choice of an audio encoder, as all three audio encoders produce a similar performance.

**Audio Duration.** Additionally, we investigate how audio duration affects the accuracy of a long-range video retrieval task. In Figure 5b, we experiment with audio spectrograms of 10, 20, and 30 seconds. Our results indicate that longer audio duration leads to better performance. However, the performance gain is relatively small (i.e., 0.5% in R@1), suggesting that 10s audio spectrograms are typically sufficient to capture relevant audio cues.

**The Number of Audiovisual Attention Blocks.** In Figure 5c, we also study the video retrieval performance (using R@1) as a function of the number of audiovisual attention blocks in our 12-layer ECLIPSE model. Using  $k$  audiovisual attention blocks implies that these  $k$  audiovisual blocks are injected into the first  $k$  layers of the network while the remaining  $12 - k$  layers only consider visual information. Our results indicate that video retrieval performance decreases when we use fewer audiovisual attention blocks. In other words, our method achieves the best video retrieval accuracy when the audiovisual attention block is inserted into every layer of our ECLIPSE architecture.

**The Importance of CLIP Pretraining.** To highlight the importance of CLIP pretraining, we compare CLIP pretraining with the ImageNet-21k pretraining. We use the ViT-B/32 architecture for these experiments. We report that compared to the ImageNet-21k pretraining, CLIP pretraining leads to 27.1%,



**Fig. 6.** Here, we illustrate our qualitative retrieval results on ActivityNet Captions [26]. We compare our audiovisual ECLIPSE model with a video-only CLIP4Clip [16]. For a given a textual query (depicted in a green block), we visualize each method’s top-1 retrieved video. Our results indicate that the video-only CLIP4Clip struggles with retrieval when textual queries include audio event descriptions, e.g., “a woman speaking to the camera”, “a person playing the violin,” etc. (see bolded text). In these cases, CLIP4Clip fails to retrieve the correct video instances, whereas ECLIPSE effectively leverages audiovisual cues for a successful retrieval.

34.6%, 24.2%, 42.5% better T2V R@1 retrieval accuracy on ActivityNet, DiDeMo, YouCook2, and QVHighlights respectively. These results suggest that CLIP pre-training is essential for good downstream video retrieval performance.

**Single Modality Baselines.** We also report the results of (i) 180-second audio-only, (ii) 64-frame video-only, and (iii) our 32-frame audiovisual methods. On ActivityNet Captions, the three approaches achieve 2.7%, 40.7%, and 42.3% R1 T2V retrieval accuracy respectively. These results indicate that jointly modeling audio and video achieves the best accuracy. We also note that while audio alone obtains poor accuracy, audio effectively complements video in our audiovisual approach. We observe similar trends on all other datasets too.

## 5.4 Qualitative Results

**Video Retrieval Results.** In Figure 6, we also illustrate some of the qualitative video retrieval results on ActivityNet Captions [26]. Specifically, for a given textual query (illustrated in the green blocks in Figure 6), we visualize the top-1 retrieved video by our audiovisual ECLIPSE and the video-only CLIP4Clip baseline. Based on these results, we observe that the video-only CLIP4Clip method



**Fig. 7.** Here, we illustrate qualitative sound localization results of our method. Note that our ECLIPSE is not explicitly trained for the sound localization task. In other words, ECLIPSE learns implicit associations between objects and sounds while being optimized with respect to the video retrieval task.

struggles to retrieve videos for textual queries that include audio-based event descriptions. For instance, in the first example of Figure 6, the textual query mentions an audio-based event of “a woman speaking to the camera” (see bolded text). Furthermore, the textual query in the second example also involves a sound event of “a person playing the violin.”

Since CLIP4Clip, does not have any audiovisual modeling capabilities, it fails to retrieve the correct video in these cases. In contrast, ECLIPSE retrieves the correct videos in all three illustrated cases, thus, highlighting the importance of incorporating video and audio cues for effective long-range video retrieval.

**Sound Localization Results.** In Figure 7, we also demonstrate qualitative sound localization results of our method. By computing the similarity between audio features and visual patches, we can obtain saliency maps that are indicative of sound sources in the video. Note that our method does not require any additional sound localization training objective. In other words, ECLIPSE successfully learns associations between sound sources and objects (e.g., a woman talking, a man playing the violin, a man using a chainsaw) as a byproduct of being trained for the video retrieval task.

## 6 Conclusions

In this paper, we present a novel audiovisual framework, ECLIPSE, for long-range video retrieval. By replacing costly and redundant parts of the video, with compact audio cues, ECLIPSE efficiently processes long-range videos while also obtaining better performance than standard video-only methods. Our audiovisual framework is (i) flexible, (ii) fast, (iii) memory-efficient, and (iv) it achieves state-of-the-art results on five diverse long-range video benchmarks. In the future, we plan to extend our method to other multimodal video understanding tasks such as video question answering and video captioning.

## References

1. Ioana Croitoru, Simion-Vlad Bogolin, Marius Leordeanu, Hailin Jin, Andrew Zisserman, Samuel Albanie, and Yang Liu. Teachtext: Crossmodal generalized distillation for text-video retrieval. In *ICCV*, 2021. [1](#), [2](#), [9](#)
2. Song Liu, Haoqi Fan, Shengsheng Qian, Yiru Chen, Wenkui Ding, and Zhongyuan Wang. Hit: Hierarchical transformer with momentum contrast for video-text retrieval. In *ICCV*, 2021. [1](#), [2](#), [8](#), [9](#)
3. Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. [1](#), [3](#), [9](#), [11](#)
4. Xiaohan Wang, Linchao Zhu, and Yi Yang. T2vlad: Global-local sequence alignment for text-video retrieval. In *CVPR*, 2021. [1](#)
5. Michael Wray, Hazel Doughty, and Dima Damen. On semantic similarity in video retrieval. In *CVPR*, 2021. [1](#)
6. Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, 2020. [1](#), [3](#), [8](#), [9](#)
7. Valentin Gabeur, Arsha Nagrani, Chen Sun, Karteek Alahari, and Cordelia Schmid. Masking modalities for cross-modal video retrieval. In *WACV*, 2022. [1](#), [3](#)
8. Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. VideoCLIP: Contrastive pre-training for zero-shot video-text understanding. In *EMNLP*, 2021. [1](#)
9. Jie Lei, Linjie Li, Luwei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *CVPR*, 2021. [2](#), [8](#), [9](#), [11](#)
10. Jianfeng Dong, Xirong Li, and Cees GM Snoek. Word2visualvec: Image and video to sentence matching by visual feature prediction. *arXiv Preprint*, 2016. [2](#)
11. Ran Xu, Caiming Xiong, Wei Chen, and Jason Corso. Jointly modeling deep video and compositional text to bridge vision and language in a unified framework. In *AAAI*, 2015. [2](#)
12. Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv Preprint*, 2014. [2](#)
13. Han Fang, Pengfei Xiong, Luhui Xu, and Yu Chen. Clip2video: Mastering video-text retrieval via image clip. *arXiv Preprint*, 2021. [2](#), [3](#)
14. Zijian Gao, Jingyu Liu, Sheng Chen, Dedan Chang, Hao Zhang, and Jinwei Yuan. Clip2tv: An empirical study on transformer-based methods for video-text retrieval. *arXiv Preprint*, 2021. [2](#), [3](#)
15. 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 *ICML*, 2021. [2](#), [3](#), [5](#), [6](#), [7](#), [9](#)
16. Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. CLIP4Clip: An empirical study of clip for end to end video clip retrieval. *arXiv Preprint*, 2021. [2](#), [3](#), [7](#), [8](#), [9](#), [10](#), [11](#), [13](#)
17. Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. Attention bottlenecks for multimodal fusion. In *NeurIPS*, 2021. [2](#), [4](#)
18. Evangelos Kazakos, Jaesung Huh, Arsha Nagrani, Andrew Zisserman, and Dima Damen. With a little help from my temporal context: Multimodal egocentric action recognition. In *BMVC*, 2021. [2](#)

19. Relja Arandjelović and Andrew Zisserman. Objects that sound. In *ECCV*, 2018. [2](#), [4](#)
20. Arun Balajee Vasudevan, Dengxin Dai, and Luc Van Gool. Sound and visual representation learning with multiple pretraining tasks. In *CVPR*, 2022. [2](#)
21. Triantafyllos Afouras, Yuki M Asano, Francois Fagan, Andrea Vedaldi, and Florian Metze. Self-supervised object detection from audio-visual correspondence. In *CVPR*, 2022. [2](#)
22. Yusuf Aytar, Carl Vondrick, and Antonio Torralba. Soundnet: Learning sound representations from unlabeled video. In *NeurIPS*, 2016. [2](#), [4](#)
23. Humam Alwassel, Dhruv Mahajan, Lorenzo Torresani, Bernard Ghanem, and Du Tran. Self-supervised learning by cross-modal audio-video clustering. In *NeurIPS*, 2020. [2](#), [4](#)
24. Yan-Bo Lin, Hung-Yu Tseng, Hsin-Ying Lee, Yen-Yu Lin, and Ming-Hsuan Yang. Exploring cross-video and cross-modality signals for weakly-supervised audio-visual video parsing. In *NeurIPS*, 2021. [2](#), [4](#)
25. Ruohan Gao, Tae-Hyun Oh, Kristen Grauman, and Lorenzo Torresani. Listen to look: Action recognition by previewing audio. In *CVPR*, 2020. [2](#), [4](#)
26. Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *ICCV*, 2017. [3](#), [8](#), [10](#), [11](#), [13](#)
27. Jie Lei, Tamara L Berg, and Mohit Bansal. Qvhighlights: Detecting moments and highlights in videos via natural language queries. In *NeurIPS*, 2021. [3](#), [8](#), [10](#), [11](#)
28. Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, 2017. [3](#), [8](#), [10](#), [11](#)
29. Luwei Zhou, Chenliang Xu, and Jason J Corso. Towards automatic learning of procedures from web instructional videos. In *AAAI*, 2018. [3](#), [8](#), [10](#), [11](#)
30. Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, 2016. [3](#), [8](#), [10](#), [11](#)
31. Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metze, Alexander Hauptmann, Joao Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. In *ICLR*, 2021. [3](#), [9](#)
32. Elad Amrani, Rami Ben-Ari, Daniel Rotman, and Alex Bronstein. Noise estimation using density estimation for self-supervised multimodal learning. In *AAAI*, 2020. [3](#)
33. Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, 2020. [3](#)
34. Michael Wray, Diane Larlus, Gabriela Csurka, and Dima Damen. Fine-grained action retrieval through multiple parts-of-speech embeddings. In *ICCV*, 2019. [3](#)
35. Yuying Ge, Yixiao Ge, Xihui Liu, Dian Li, Ying Shan, Xiaohu Qie, and Ping Luo. Bridging video-text retrieval with multiple choice questions. In *CVPR*, 2022. [3](#)
36. Haoyu Lu, Nanyi Fei, Yuqi Huo, Yizhao Gao, Zhiwu Lu, and Ji-Rong Wen. Cots: Collaborative two-stream vision-language pre-training model for cross-modal retrieval. In *CVPR*, 2022. [3](#)
37. Jinpeng Wang, Yixiao Ge, Guanyu Cai, Rui Yan, Xudong Lin, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Object-aware video-language pre-training for retrieval. In *CVPR*, 2022. [3](#)
38. Linchao Zhu and Yi Yang. Actbert: Learning global-local video-text representations. In *CVPR*, 2020. [3](#)

39. Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. HERO: Hierarchical encoder for Video+Language omni-representation pre-training. In *EMNLP*, 2020. 3
40. Youngjae Yu, Jongseok Kim, and Gunhee Kim. A joint sequence fusion model for video question answering and retrieval. In *ECCV*, 2018. 3
41. Youngjae Yu, Hyungjin Ko, Jongwook Choi, and Gunhee Kim. End-to-end concept word detection for video captioning, retrieval, and question answering. In *CVPR*, 2017. 3
42. Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Jason Li, Taroon Bharti, and Ming Zhou. Univl: A unified video and language pre-training model for multimodal understanding and generation. *arXiv Preprint*, 2020. 3
43. Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. Use what you have: Video retrieval using representations from collaborative experts. In *BMVC*, 2019. 3, 9
44. Niluthpol Chowdhury Mithun, Juncheng Li, Florian Metze, and Amit K Roy-Chowdhury. Learning joint embedding with multimodal cues for cross-modal video-text retrieval. In *ICMR*, 2018. 3
45. Antoine Miech, Ivan Laptev, and Josef Sivic. Learning a text-video embedding from incomplete and heterogeneous data. *arXiv Preprint*, 2018. 3
46. Xing Cheng, Hezheng Lin, Xiangyu Wu, Fan Yang, and Dong Shen. Improving video-text retrieval by multi-stream corpus alignment and dual softmax loss. *arXiv Preprint*, 2021. 3
47. Zeyu Wang, Yu Wu, Karthik Narasimhan, and Olga Russakovsky. Multi-query video retrieval. *arXiv Preprint*, 2022. 3
48. Maksim Dzabraev, Maksim Kalashnikov, Stepan Komkov, and Aleksandr Petiushko. Mdmmt: Multidomain multimodal transformer for video retrieval. In *CVPRW*, 2021. 3
49. Jesús Andrés Portillo-Quintero, José Carlos Ortiz-Bayliss, and Hugo Terashima-Marín. A straightforward framework for video retrieval using clip. In *MCP*, 2021. 3
50. Satya Krishna Gorti, Noël Vouitsis, Junwei Ma, Keyvan Golestan, Maksims Volkovs, Animesh Garg, and Guangwei Yu. X-pool: Cross-modal language-video attention for text-video retrieval. In *CVPR*, 2022. 3
51. Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. A clip-hitchhiker’s guide to long video retrieval. *arXiv Preprint*, 2022. 3
52. Andrew Owens and Alexei A. Efros. Audio-visual scene analysis with self-supervised multisensory features. In *ECCV*, 2018. 4
53. Bruno Korbar, Du Tran, and Lorenzo Torresani. Cooperative learning of audio and video models from self-supervised synchronization. In *NeurIPS*, 2018. 4
54. Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *ICCV*, 2017. 4
55. Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. Ambient sound provides supervision for visual learning. In *ECCV*, 2016. 4
56. Yuki M Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. In *NeurIPS*, 2020. 4
57. Shuang Ma, Zhaoyang Zeng, Daniel McDuff, and Yale Song. Active contrastive learning of audio-visual video representations. In *ICLR*, 2021. 4
58. Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. In *CVPR*, 2021. 4

59. Pedro Morgado, Ishan Misra, and Nuno Vasconcelos. Robust audio-visual instance discrimination. In *CVPR*, 2021. 4
60. Yan-Bo Lin, Yu-Jhe Li, and Yu-Chiang Frank Wang. Dual-modality seq2seq network for audio-visual event localization. In *ICASSP*, 2019. 4
61. Shuang Ma, Zhaoyang Zeng, Daniel McDuff, and Yale Song. Contrastive learning of global and local video representations. In *NeurIPS*, 2021. 4
62. Yunhua Zhang, Hazel Doughty, Ling Shao, and Cees GM Snoek. Audio-adaptive activity recognition across video domains. In *CVPR*, 2022. 4
63. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 4, 5
64. Yan-Bo Lin and Yu-Chiang Frank Wang. Audiovisual transformer with instance attention for audio-visual event localization. In *ACCV*, 2020. 4
65. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 4
66. Nina Shvetsova, Brian Chen, Andrew Rouditchenko, Samuel Thomas, Brian Kingsbury, Rogerio Feris, David Harwath, James Glass, and Hilde Kuehne. Everything at once—multi-modal fusion transformer for video retrieval. page CVPR, 2022. 4
67. Rowan Zellers, Jiasen Lu, Ximing Lu, Youngjae Yu, Yanpeng Zhao, Mohammadreza Salehi, Aditya Kusupati, Jack Hessel, Ali Farhadi, and Yejin Choi. Merlot reserve: Neural script knowledge through vision and language and sound. In *CVPR*, 2022. 4
68. Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and text. In *NeurIPS*, 2021. 4
69. Yanpeng Zhao, Jack Hessel, Youngjae Yu, Ximing Lu, Rowan Zellers, and Yejin Choi. Connecting the dots between audio and text without parallel data through visual knowledge transfer. *arXiv Preprint*, 2021. 4
70. Jean-Baptiste Alayrac, Adrià Recasens, Rosalia Schneider, Relja Arandjelović, Jason Ramapuram, Jeffrey De Fauw, Lucas Smaira, Sander Dieleman, and Andrew Zisserman. Self-supervised multimodal versatile networks. In *NeurIPS*, 2020. 4
71. Brian Chen, Andrew Rouditchenko, Kevin Duarte, Hilde Kuehne, Samuel Thomas, Angie Boggust, Rameswar Panda, Brian Kingsbury, Rogerio Feris, David Harwath, James Glass, and Micha Picheny. Multimodal clustering networks for self-supervised learning from unlabeled videos. In *ICCV*, 2021. 4
72. Yan-Bo Lin and Yu-Chiang Frank Wang. Exploiting audio-visual consistency with partial supervision for spatial audio generation. In *AAAI*, 2021. 4
73. Sinong Wang, Belinda Z Li, Madian Khabza, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv Preprint*, 2020. 4
74. Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamás Sarlós, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J. Colwell, and Adrian Weller. Rethinking attention with performers. In *ICLR*, 2021. 4
75. Mandela Patrick, Dylan Campbell, Yuki Asano, Ishan Misra, Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and João F Henriques. Keeping your eye on the ball: Trajectory attention in video transformers. In *NeurIPS*, 2021. 4
76. Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 5

77. Jort F. Gemmeke, Daniel P. W. 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*, 2017. 5
78. Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, Channing Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron Weiss, and Kevin Wilson. Cnn architectures for large-scale audio classification. In *ICASSP*, 2017. 5, 12
79. Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audio-visual dataset. In *ICASSP*, 2020. 5, 8, 9, 12
80. Yuan Gong, Yu-An Chung, and James Glass. AST: Audio Spectrogram Transformer. In *INTEERSPEECH*, 2021. 5
81. Yuan Gong, Yu-An Chung, and James Glass. Psla: Improving audio tagging with pretraining, sampling, labeling, and aggregation. *TASLP*, 2021. 5, 12
82. Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019. 7
83. Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 8
84. Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 8
85. Bowen Zhang, Hexiang Hu, and Fei Sha. Cross-modal and hierarchical modeling of video and text. In *ECCV*, 2018. 8
86. Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv Preprint*, 2015. 9
87. Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017. 9
88. Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, 2019. 9
89. Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, 2018. 9