

GEXIA: Granularity Expansion and Iterative Approximation for Scalable Multi-grained Video-language Learning

Yicheng Wang^{1,2*}, Zhikang Zhang², Jue Wang², David Fan²,
Zhenlin Xu², Linda Liu², Xiang Hao², Vimal Bhat², and Xinyu Li²

¹Texas A&M University ²Amazon

Abstract

In various video-language learning tasks, the challenge of achieving cross-modality alignment with multi-grained data persists. We propose a method to tackle this challenge from two crucial perspectives: data and modeling. Given the absence of a multi-grained video-text pretraining dataset, we introduce a Granularity Expansion (GEX) method with Integration and Compression operations to expand the granularity of a single-grained dataset. To better model multi-grained data, we introduce an Iterative Approximation Module (IAM), which embeds multi-grained videos and texts into a unified, low-dimensional semantic space while preserving essential information for cross-modal alignment. Furthermore, GEXIA is highly scalable with no restrictions on the number of video-text granularities for alignment. We evaluate our work on three categories of video tasks across seven benchmark datasets, showcasing state-of-the-art or comparable performance. Remarkably, our model excels in tasks involving long-form video understanding, even though the pretraining dataset only contains short video clips.

1. Introduction

Video-language representation learning is crucial for a wide range of video tasks, such as cross-modality retrieval [6, 39, 54], video classification [41, 44, 52], and video question answering [25]. With the advent of CLIP-like models [1, 8, 21, 38], video-language learning has entered the “big data” era, where video-text foundation models [2, 3, 24, 32, 43, 59] are first pretrained on large-scale datasets consisting of diverse video-text pairs aimed at cross-modal representation alignment and then finetuned on small-scale datasets tailored to specific downstream tasks.

Following this paradigm, recent research primarily focuses on two main directions: data and modeling.

On the data front, researchers have introduced larger-scale pretraining datasets [35, 49, 55] with higher-quality video-text pairs using scalable data collection methods. On the modeling side, most approaches [2, 3, 24, 32, 43, 59, 61] focus on single-grained or hierarchical video-text representation alignments. Nevertheless, a frequently overlooked fact should be revisited: video-text data are inherently multi-grained, i.e., videos and texts in the real world naturally have variable durations and text sequence lengths, while the same video can also be paired with texts in different sequence lengths, offering varying levels of detail, and vice versa. This inherent multi-granularity presents a major challenge for modeling video-text alignment. Previous methods [2, 12, 23, 24, 32, 43, 59, 61] often fall short of fundamentally addressing the challenge.

In this work, we propose **GEXIA: Granularity EXpansion and Iterative Approximation** to address this challenge from both data and modeling perspectives. Firstly, existing video-text pre-training datasets [3, 35, 49, 55] follow data collection schemas similar to image-text datasets. As a result, these datasets either provide only seconds-long videos with sentence-long texts [3, 49] or minutes-long videos with audio transcriptions as texts [35, 55]. Constrained by scalability, such datasets are essentially single-grained and cannot reflect the multi-granularity of real-world video-text data, which hinders model performance during pretraining. To this end, we introduce a scalable video-text Granularity EXpansion (GEX) pipeline that generates multi-grained video-text datasets from single-grained ones. Specifically, we propose Integration and Compression operations on videos and texts separately to expand two additional granularities by using only existing samples in the dataset. These operations can be applied sequentially to achieve the desired number of granularities, as shown in Figure 1.

*Work was done as an intern at Amazon.



Figure 1. An overview of the Granularity EXpansion (GEX) pipeline, which expands a single-grained dataset into a multi-grained dataset with video and text integration \oplus_v and \oplus_t and text compression Θ_t operations.

Secondly, most recent works [3, 24, 32, 43, 59] treat input video-text pairs as single-grained regardless of differences in video durations and text lengths, which limits their performance in tasks that involve multi-grained data. Other studies [2, 19, 33, 61] address the issue using hierarchical designs to handle different granularities of data at varying structural levels. However, such approaches face significant limitations in scalability to new granularities or are tailored for specific downstream tasks only. A comprehensive solution for multi-grained video-text alignment across diverse downstream tasks with high scalability has not been thoroughly explored. Aiming at the modeling challenge, we propose an Iterative Approximation Module (IAM) as a general building block for multi-grained video-text alignment presented in Figure 2. We formulate the multi-grained alignment as a three-stage process: dense feature extraction, iterative approximation, and cross-model alignment. First, dense features of the videos and texts are extracted and aggregated separately with an existing image and text encoder. Subsequently, IAMs iteratively and adaptively approximate these dense features into fixed-size low-dimensional embeddings, considering the granularity of inputs. As such, uniform representations of multi-grained videos and texts are embedded in the same low-dimensional space to enable the final cross-modal alignment through conventional contrastive learning.

We take an existing single-grained pretraining dataset to showcase the effectiveness of GEX and conduct extensive evaluations of IAMs across cross-modal retrieval, video classification, and video-question-answering tasks on **seven** widely used datasets of various granularities. Our model, enhanced by the multi-grained pretraining data, consistently achieves state-of-the-art or comparable performance on all tasks compared to baseline models. Notably, our model excels in scenarios requiring long-term visual-text alignment and a deep understanding of long-form videos, highlighting the effectiveness of our approach. In summary, our contributions are as follows.

1. We propose Granularity EXpansion (GEX), a scalable method to automatically expand single-grained video-text datasets into multi-grained ones without additional data collection, addressing the multi-grained pretraining data bottleneck;

2. We introduce the Iterative Approximation Module (IAM), which embeds multi-grained data into a unified, low-dimensional space, enabling efficient and scalable cross-modal alignment;
3. Our approach achieves state-of-the-art or comparable results across multiple benchmarks. We also provide detailed ablations and insights.

2. Related work

2.1. Single-grained Video-Language Learning

Early works [3, 24, 32, 45, 59] have treated input video and text pairs as in single granularity regardless of the video duration and text length. Due to the overlook of inherent multi-granularity, such approaches are not only limited to tasks that involve multi-grained video-text data but also sacrifice performances on general video-text alignments. More specifically, CLIPBert [24] proposes to apply sparse sampling on videos in a universal manner, which leads to poor performance on long-form video understanding tasks. TACo [59] relies on heuristics that nouns and verbs are more likely to be aligned with visual contents, which sacrifices the performance on long, complicated video scenes and texts. Frozen [3] views images as static video clips and pretrain on images and videos simultaneously, whose performance is limited on long videos with large temporal dynamics. Clip4clip [32] empirically studies the approaches of aggregating image-level CLIP features for video retrieval tasks through mean-pooling, LSTM, or single transformer block, which is largely limited to video retrieval tasks and only achieves coarse alignments between videos and text. LF-VILA [43] focuses on single granularity video-text alignment for long-form video understanding. Despite the fact that LF-VILA is limited to long-form videos only, our approach uses only 1/6 of video-text pairs for pretraining while achieving comparable or better performance than LF-VILA on multiple long-form video understanding tasks as shown in Section 4.

2.2. Multi-grained Video-Language Learning

Observing the inherent multi-granularity in video-text alignment, A few recent works [2, 19, 33, 61] address the multi-grained video-text alignment problem either with a

hierarchical design, which limits scalability or target specific downstream tasks only, which limits generalizability. X-CLIP [33] targets video retrieval tasks only and proposes a multi-grained contrastive learning loss and an attention-over-similarity-matrix-module to aggregate multi-grained features, which is limited to retrieval tasks only and unable to extend to video-text granularities beyond seconds-long videos and sentence-long texts. HierVL [2] builds hierarchical video-text embeddings for multi-level video-text alignments and studies two aggregation designs to acquire features of long context windows. This method heavily depends on a high-quality multi-grained video-text pretraining dataset, which is often hard to acquire from the public domain, and can not be well extended to new granularities (HierVL is built on Ego4D, whose videos are specifically collected from non-public domain and texts are written by human narrators). HiTeA [61] proposes to achieve multi-grained video-text alignments by designing two new pretraining tasks, namely cross-modal moment exploration and multi-modal temporal relation exploration. Such tasks extensively extract temporal information from input video sequences, while leading to limitations on long-form videos and more video-text granularities. VideoRecap [19] specifically targets the video captioning problem and also employs a hierarchical design to process multi-grained data.

3. Methodology

3.1. Granularity EXpansion (GEX)

In real-world application scenarios, videos and texts exhibit various levels of granularity. However, due to the scalability constraints of data collection processes, existing large-scale video-text pretraining datasets [3, 35, 49, 55] are mostly single-grained. To overcome the challenge of insufficient multi-grained pretraining data, we introduce a simple and scalable Granularity EXpansion (GEX) pipeline to enhance the granularity of an existing single-grained dataset, as illustrated in Figure 1. The GEX pipeline comprises two main operations: Integration and Compression, which can be further categorized based on the input data type as Video Integration (\oplus_v), Text Integration (\oplus_t), Video Compression (Θ_v), and Text Compression (Θ_t).

Integration (\oplus_v and \oplus_t): The Integration operation combines multiple input video-text pairs of the same granularity into one data sample to create a new granularity. Specifically, we exploit meta-information to identify existing video-text pairs originating from the same source and concatenate (\oplus_v and \oplus_t) them to generate longer video-text pairs. When meta-information is unavailable, random samples can also be concatenated in a similar manner, which still enhances the performance, as shown in Appendix A.

Compression (Θ_v and Θ_t): The Compression, in contrast, condenses individual data samples into more compact forms

to create a new granularity. Specifically, we utilize text summarization via a large language model (LLM) as Θ_t to condense \oplus_t -applied long texts. The summarized texts are much shorter but remain aligned with the corresponding \oplus_v -applied long videos, thereby creating a new granularity. The summarized text is verified to be similar with the original long texts, as demonstrated in Appendix B. We do not perform the video compression Θ_v in the main study, since \oplus_v , \oplus_t , and Θ_t are sufficient to transform a single-grained dataset into a multi-grain dataset for pretraining. Additional exploration of video compression (Θ_v) is included in Appendix F.

GEX is highly scalable and does not rely on human annotation. Both Integration \oplus and Compression Θ operations have no limits to the input data and can be applied sequentially or recursively to expand the granularities of the given dataset to a desired scale. Using GEX, we can generate a multi-grained video-text pretraining dataset (v_{multi}, t_{multi}) from a single-grained dataset (v, t), as shown in Figure 1.

3.2. Multi-grained Video-Text Alignment

Overview: We introduce a three-phase approach for multi-grained video-text alignment, as illustrated in Figure 2. First, we utilize two distinct single-modality encoders for videos and texts to create high-dimensional representations that encapsulate the complete information of the input data, referred to as dense features. The dimensions of these dense features are proportional to the granularity of the input data, whether determined by the video durations or text sequence lengths. In the second phase, given pairs of dense features with variable dimensions, we aim to approximate these features using fixed-size embeddings derived from a universal low-dimensional space. Apparently, higher-dimensional features with richer semantic information are inherently more challenging to approximate with fixed-size embeddings. We employ cross-attention mechanism as the building block to construct an iterative approximation module, thereby allocating more iterations to higher-dimensional dense features for finer approximation. Finally, given multi-grained videos and texts embedded in the same low-dimensional space, we apply conventional contrastive learning to enforce cross-modal alignment.

Dense Feature Extraction: For video-text pairs with varied granularity, the first step is to construct feature representations that capture their complete semantic information. Considering the varying lengths of videos and texts, we propose that the dimensionality of these representations should be proportional to the input size, maintaining consistent information density and minimizing information loss. Additionally, since our goal is to create aligned representations across modalities, it is crucial to keep each data stream separate to prevent information leakage. Therefore, we employ pretrained image and text encoders to extract individual fea-

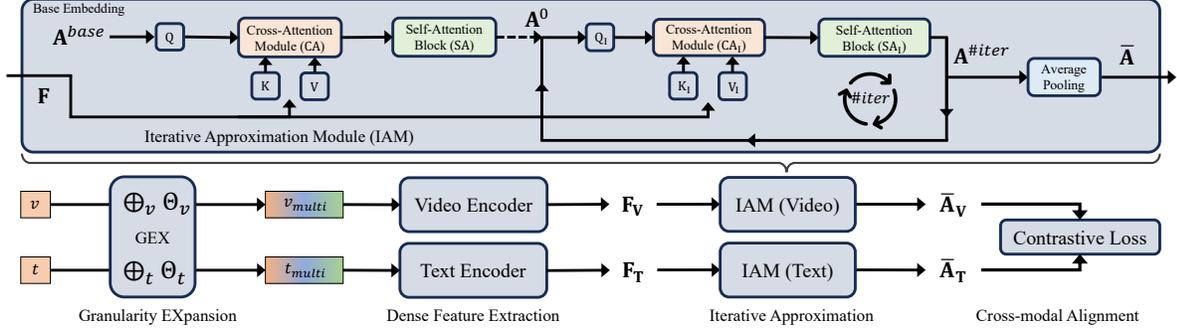


Figure 2. An overview of GEXIA, which consists of Granularity EXpansion (GEX), Dense Feature Extraction, Iterative Approximation Module (IAM), and Cross-modal Alignment. We propose GEX and IAM to address data and modeling challenges, respectively.

tures at the frame and text token levels, respectively, and subsequently aggregate these features along the video temporal and text sequential dimensions to form dense features.

Given d video frames $[f_1, \dots, f_d]$ sampled from a video in v_{multi} , where each frame $f_i \in \mathbb{R}^{H \times W \times 3}$, we first employ a ViT-based image encoder to transform each frame into a tensor in $\mathbb{R}^{P \times C_v}$, where P denotes the patch number and C_v denotes the video feature dimension. Then we concatenate all frame-level features into a tensor $\mathbf{F}_V \in \mathbb{R}^{(d \times P) \times C_v}$. Similarly, given the corresponding text sequence t_{multi} , we tokenize it using byte pair encoding (BPE) and then truncate or pad it into a size of $m \times S$, where m is the maximum number of tokens that the text encoder can process in a single run and S is the size of one token. Subsequently, we apply the text encoder to transform it into token-level dense feature $\mathbf{F}_T \in \mathbb{R}^{m \times C_t}$, where C_t represents the size of the feature.

Iterative Approximation: Given variable-sized dense features, we develop a novel Iterative Approximation Module (IAM) to approximate the inputs into fixed-sized and low-dimensional embeddings. An IAM takes a dense feature $\mathbf{F} \in \mathbb{R}^{M \times C}$ and a learnable base embedding $\mathbf{A}^{base} \in \mathbb{R}^{N \times D}$ as inputs, where N is the number of learnable base vectors and D is the dimension. In the video IAM, the notation $\mathbf{F} \in \mathbb{R}^{M \times C}$ corresponds to $\mathbf{F}_V \in \mathbb{R}^{(d \times P) \times C_v}$, with $M = (d \times P)$ and $C = C_v$. For the text IAM, $\mathbf{F} \in \mathbb{R}^{M \times C}$ signifies $\mathbf{F}_T \in \mathbb{R}^{m \times C_t}$, where $M = m$ and $C = C_t$.

The base embedding \mathbf{A}^{base} is designed to capture the essential semantic information of the training data and serves as an optimized initial value of the iterative approximation process during inference. Subsequently, inspired by the design of attention mechanism [46], which models the relationship of features via similarity, we apply the cross-attention module CA_I to execute the approximation taking \mathbf{A}^{base} as the Query and \mathbf{F} as the Key and Value. To align the different dimensions necessary for executing the cross-attention between \mathbf{F} and \mathbf{A}^{base} , and to ensure that the output maintains the same dimensionality as \mathbf{A}^{base} , we implement linear projection layers Q_I , K_I , and V_I , which project all feature dimensions to D . We also add a self-

attention block SA_I following the CA_I to enhance the data approximation. Notably, longer videos and texts, which inherently contain more semantic information, pose greater challenges for effective approximation. Therefore, we repeat the approximation operations with CA_I and SA_I , adjusting the iteration number $\#iter$ in accordance with the content length, with longer videos and texts assigned a greater $\#iter$. Through adaptive approximation iterated for $\#iter$ times, the resultant $\mathbf{A}^{\#iter}$ closely approximates \mathbf{F} , preserving its semantic information in a compact, low-dimensional space.

Meanwhile, inspired by [20, 29, 36, 60, 64] that study the unrolling of the iterative algorithm and indicate that more unrolled iterations lead to higher model performance but lower computation efficiency, we unroll the iterative approximation process one time while assigning remaining iterations to the original iterative block denoted by I to achieve good trade-offs between model performance and training efficiency. For clarity, we represent the cross-attention, self-attention, linear projection layers of the Query, Key, Value, and the output of the unrolled iteration as CA , SA , Q , K , V , and \mathbf{A}^0 , respectively, as illustrated in Figure 2, where the output \mathbf{A}^0 serves as the base embedding of the following iterative approximation.

IAM enables a single model to handle variable-length videos and texts by modulating $\#iter$, thereby maintaining a fixed low-dimensional output space for cross-modal alignment. Additionally, the output embedding $\mathbf{A}^{\#iter} \in \mathbb{R}^{N \times D}$ of the IAM is not only more compact than dense features in size but also retains a high fidelity approximation of the original dense feature $\mathbf{F} \in \mathbb{R}^{M \times C}$, leading to improved model performance on downstream tasks after alignment.

Cross-modal Alignment: Given a batch (batch size B) of video and text pairs, IAMs produce paired embeddings $\mathbf{A}_V^{\#iter} \text{ and } \mathbf{A}_T^{\#iter} \in \mathbb{R}^{B \times N \times D}$ after the iterative approximation. Since $\mathbf{A}_V^{\#iter}$ and $\mathbf{A}_T^{\#iter}$ have the same size regardless of the input granularity, conventional contrastive learning [38] can be directly applied to build alignments. We process paired embeddings with average pooling along dimension N , yielding vectorized features $\bar{\mathbf{A}}_V$ and $\bar{\mathbf{A}}_T \in$

$\mathbb{R}^{B \times D}$ to minimize the Video-to-Text Contrastive (VTC) loss $\mathcal{L}_{\text{VTC}} = \frac{1}{2}(\mathcal{L}_{\text{v2t}} + \mathcal{L}_{\text{t2v}})$, which encompasses both video-to-text and text-to-video losses:

$$\mathcal{L}_{\text{v2t}} = -\mathbb{E}_{\mathbf{v}} \left[\log \frac{f(\mathbf{v}, \mathbf{t})}{\sum_{\mathbf{t}' \in \bar{\mathbf{A}}_{\text{T}}} f(\mathbf{v}, \mathbf{t}')} \right],$$

$$\mathcal{L}_{\text{t2v}} = -\mathbb{E}_{\mathbf{t}} \left[\log \frac{f(\mathbf{v}, \mathbf{t})}{\sum_{\mathbf{v}' \in \bar{\mathbf{A}}_{\text{V}}} f(\mathbf{v}', \mathbf{t})} \right],$$

where $f(\mathbf{v}, \mathbf{t}) = e^{s(\mathbf{v}, \mathbf{t})/\tau}$. s is the cosine similarity and τ is a learnable scaling parameter. This loss function aims to maximize the cosine similarity between matching video-text pairs while minimizing it for non-matching pairs.

4. Experiments

4.1. Experimental Setups

Multi-grained Pretraining Data: We utilize the InternVid-10M-FLT [49] as the base pretraining dataset which consists of ten million short video clips collected from four million YouTube videos. Based on the statistics in [49], we treat video-text pairs in this dataset as single granularity, labeled short-video-short-text (SVST). In alignment with the data distribution of targeted downstream tasks, we choose to expand this base dataset with two additional granularities: long-video-long-text (LVLT) and long-video-short-text (LVST) pairs. More specifically, we perform Video and Text Integration (\oplus_v and \oplus_t) through video concatenation and text concatenation on 4.5 million SVST pairs to build 1 million LVLT video-text pairs. These 4.5 million SVST pairs are selected by their source video IDs and timestamps, excluding those with fewer than four SVST pairs. The concatenation is performed only on videos sharing identical source IDs in time order, along with their corresponding texts. Subsequently, we apply the Text Compression (Θ_t) through text summarization on the 1 million generated long texts using a self-hosted large language model (LLM) (Vicuna 13b-v1.5 [9, 50]) to build 1 million LVST pairs. As demonstrated in Appendix B, the summarized short texts are shown to be semantically similar to the original long texts. Further details on the selection of the LLM and the potential use of Video Compression (Θ_v) are provided in Appendix A and F.

Model Pretraining: During the pretraining phase, we initialize our video and text feature encoder with CLIP weights [38] (ViT-B/32 model). The pretraining is conducted on 32 NVIDIA Tesla-V100 GPUs and takes approximately 2.5 days, covering 5 epochs on our multi-grained dataset. We employ the AdamW optimizer [31], using a batch size of 512. The learning rate for the video and text encoders is set to $1e^{-6}$, while for other components it is set to $1e^{-4}$, accompanied by a cosine learning rate schedule [30]. We process $d = 16$ frames for short videos and $d = 32$ frames for longer ones and empirically set $\#iter = 1$ for short videos/texts and $\#iter = 3$ for long

videos/texts. The ablation on the $\#iter$ setups is in Section 4.4. τ is initially set to 0.07 [38].

4.2. Cross-modal Retrieval

Setting: We finetune the pretrained model on ActivityNet Captions (paragraph-to-video) [22], MSR-VTT (9k) [15, 54], and LSMDC [39] datasets for text-to-video and video-to-text retrieval tasks. Following previous work [32, 49], we report both zero-shot and finetuning retrieval performance with recall@K (R@K) and median rank (Mdr). We finetune our pretrained model for 50 epochs on 8 NVIDIA Tesla-V100 GPUs. We adopt an AdamW [31] optimizer [31] and set the learning rate of the video and text encoders to $1e^{-7}$, and the other part $1e^{-5}$ with a cosine annealing scheduler [30]. Based on the average video durations and text length of the datasets, we empirically determine MSR-VTT and LSMDC as SVST, ActivityNet Captions as LVLT and set the $\#iter$ accordingly (see statistics and the $\#iter$ setups in Appendix E).

Retrieval Results: Due to the complexity of training strategies and model backbones design, we present our comparisons from multiple perspectives, as outlined in Table 1.

Compared with models of **using the same ViT-B/32 backbone or having similar FLOPs**, our approach outperforms other baselines [32, 45, 56] on the long video ActivityNet Captions dataset. PerceiverVL [45] uses a similar iterative module, but our approach achieves T2V R@1 gains of +11.4% on ActivityNet, +14.2% on MSR-VTT, and +5.4% on LSMDC. Notably, in zero-shot settings, our approach surpasses the finetuned PerceiverVL on ActivityNet (R@1 +1.4%) and MSR-VTT (R@1 +4.7%), highlighting the superior multi-grained data modeling capability of our IAM module. Compared to CLIP-based models, our model significantly outperforms CLIP4Clip [32] on ActivityNet (T2V/V2T R@1 +4.8%/+2.5%). Although pre-trained on a dataset (10M) which is much smaller than CLIP4Clip [32] and CLIP-ViP [56] (136M), our model performs comparably on the short video MSR-VTT and LSMDC, demonstrating the effectiveness of our approach in expanding data granularity and aligning multi-grained features.

Among **retrieval-specific models**, which incorporate specialized retrieval structures or strategies, our GEXIA as a general pretraining model, also outperforms X-CLIP_{B32} [33] on two datasets (T2V/V2T R@1 +1.0%/+1.1% on ActivityNet and +0.7%/+0.6% on MSR-VTT). Furthermore, it achieves comparable results with PromptSwitch [11] and UCOFIA [51], emphasizing the strengths of our GEX data preprocessing method and IAM module design. Additionally, our method demonstrates low retrieval complexity and high inference efficiency, as detailed in Appendix C.

Although models with **higher computational demands (higher FLOPs)** tend to exhibit better performance, our ap-

Table 1. Text-to-video (T2V) and Video-to-text (V2T) retrieval results using Recall at K (R@K) and Median Rank (MdR) metrics across different setups for ActivityNet Captions, MSR-VTT, and LSMDC datasets. Results for models that are not directly comparable are shown in gray. #PT represents the size of the pretraining dataset. *Best among four setups. **No auxiliary captions.

Text-to-video (T2V) Setup	Method	#PT	Backbone	ActivityNet Captions T2V				MSR-VTT T2V				LSMDC T2V			
				R@1	R@5	R@10	MdR	R@1	R@5	R@10	MdR	R@1	R@5	R@10	MdR
Zero-shot	Frozen [3]	5.5M	TS4mer	-	-	-	-	24.7	46.9	57.2	3.0	13.6	27.9	33.5	32.0
	VIOLET [14]	185.8M	Swin _B	-	-	-	-	25.9	49.5	59.7	-	-	-	-	-
	CLIP4Clip [32]	136M	ViT _{B32}	-	-	-	-	32.0	57.0	66.9	4.0	15.1	28.5	36.4	28.0
	LocVTP [7]	136M	ViT _{B16}	-	-	-	-	32.7	55.7	64.9	4.0	-	-	-	-
	ViCLIP [49]	10M	ViT _{L14}	15.1	-	-	-	24.0	-	-	-	20.1	-	-	-
	Ours	10M	ViT _{B32}	35.3	64.3	76.1	3.0	37.3	63.8	72.5	3.0	12.9	28.3	35.1	35.5
	VideoPrism [65]	618M	ViViT _B	49.6	76.7	-	-	51.4	74.4	-	-	-	-	-	-
Finetuned	CLIPBert [24]	5.6M	Res ₅₀	21.3	49.0	63.5	6.0	-	-	-	-	-	-	-	-
	Perceiver-VL [45]	5.5M	ViT _{B32}	33.9	62.1	76.4	-	32.6	62.1	71.6	-	15.8	37.6	40.1	-
	CLIP4Clip* [32]	136M	ViT _{B32}	40.5	72.4	-	2.0	44.5	71.5	81.6	2.0	22.6	41.8	49.8	11.0
	CLIP-ViP** [56]	136M	ViT _{B32}	-	-	-	-	46.5	72.1	82.5	-	-	-	-	-
Ours	10M	ViT _{B32}	45.3	76.5	86.6	2.0	46.8	72.0	81.5	2.0	21.2	40.9	51.4	9.0	
Finetuned	Frozen [3]	5.5M	TS4mer	28.8	60.9	-	3.0	32.5	61.5	71.2	3.0	15.0	30.8	40.3	20.0
	X-CLIP [33]	-	ViT _{B32}	44.3	74.1	-	-	46.1	73.0	83.1	2.0	23.3	43.0	-	-
	PromptSwitch [11]	-	ViT _{B32}	-	-	-	-	47.8	73.9	82.2	-	23.1	41.7	50.5	-
	UCOFIA [51]	-	ViT _{B32}	45.7	76.0	-	-	49.4	72.1	-	-	-	-	-	-
Retrieval-specific Model	X-CLIP [33]	-	ViT _{B16}	46.2	75.5	-	-	49.3	75.8	84.8	2.0	26.1	48.4	-	-
	CenterCLIP [66]	-	ViT _{B16}	46.2	77.0	87.6	2.0	48.4	73.8	82.0	2.0	24.2	46.2	55.9	8.0
	TaCo [59]	136M	Res ₁₅₂	30.4	61.2	-	3.0	24.5	52.8	65.5	5.0	-	-	-	-
Finetuned	LF-VILA [43]	60M	Swin _B	35.3	65.4	-	3.0	-	-	-	-	-	-	-	-
	VIOLET [14]	185.8M	Swin _B	-	-	-	-	34.5	63.0	73.4	-	16.1	36.6	41.2	-
	LAVENDER [26]	30M	Swin _B	-	-	-	-	40.7	66.9	77.6	-	26.1	46.4	57.3	-
	HiTeA [61]	5M	MViT _B	45.1	73.5	84.2	-	44.4	69.3	78.9	-	27.1	46.2	54.5	-
Model with Higher FLOPs	LocVTP [7]	136M	ViT _{B16}	-	-	-	-	46.3	72.8	82.0	-	-	-	-	-
	ViCLIP [49]	10M	ViT _{L14}	49.8	-	-	-	52.5	-	-	-	33.0	-	-	-
Video-to-text (V2T) Setup				ActivityNet Captions V2T				MSR-VTT V2T				LSMDC V2T			
	Method	#PT	Backbone	R@1	R@5	R@10	MdR	R@1	R@5	R@10	MdR	R@1	R@5	R@10	MdR
Zero-shot	ViCLIP [49]	10M	ViT _{L14}	24.0	-	-	-	41.3	-	-	-	16.9	-	-	-
	Ours	10M	ViT _{B32}	35.0	64.9	77.6	3.0	39.5	64.6	73.0	2.0	13.0	26.8	33.9	43.0
	VideoPrism [65]	618M	ViViT _B	47.9	75.0	-	-	50.2	73.2	-	-	-	-	-	-
Finetuned	CLIP4Clip* [32]	136M	ViT _{B32}	42.5	74.1	85.8	2.0	43.1	70.5	81.2	2.0	20.9	40.7	49.1	11.0
	Ours	10M	ViT _{B32}	45.0	76.4	87.3	2.0	47.4	74.0	81.8	2.0	21.9	40.5	50.2	10.0
Model with Similar FLOPs	X-CLIP [33]	-	ViT _{B32}	43.9	73.9	-	-	46.8	73.3	84.0	2.0	22.5	42.2	-	-
	PromptSwitch [11]	-	ViT _{B32}	-	-	-	-	46.0	74.3	84.8	-	22.0	40.8	50.3	-
	UCOFIA [51]	-	ViT _{B32}	46.3	76.5	-	-	47.1	74.3	-	-	-	-	-	-
	X-CLIP [33]	-	ViT _{B16}	46.4	75.9	-	-	48.9	76.8	84.5	2.0	26.9	46.2	-	-
Retrieval-specific Model	CenterCLIP [66]	-	ViT _{B16}	46.7	77.1	88.0	2.0	47.7	75.0	83.3	2.0	24.5	46.4	55.8	7.0
	ViCLIP [49]	10M	ViT _{L14}	48.1	-	-	-	51.8	-	-	-	32.5	-	-	-

proach often achieves comparable, and in some cases superior results. For instance, it outperforms previous works such as VIOLET [14] (T2V R@1 +12.3% on MSR-VTT and +5.1% on LSMDC) and LAVENDER [26] (T2V R@1 +6.1% on MSR-VTT), both of which use additional masked language modeling pretraining (MLM). Our model also surpasses LF-VILA [43] (T2V R@1 +10.0% on ActivityNet) and HiTeA [61] (T2V R@1 +0.2% on ActivityNet and +2.4% on MSR-VTT), which are specifically designed for long-form and hierarchical video-text alignment. These achievements highlight the critical role of the IAM in effectively aligning multi-grained video-text pairs.

In the **zero-shot setting**, our method significantly outperforms ViCLIP [49] on the long-video ActivityNet Captions. ViCLIP employs a ViT_{L14} backbone, which demands 17× more computational resources (see Appendix D). Despite this disparity, our model achieves remarkable gains of +20.2%/+11.0% in T2V/V2T R@1, underscoring the effectiveness of the GEX method in enhancing the model during pretraining, especially for long-video datasets. Furthermore, our approach surpasses other methods in T2V retrieval on MSR-VTT, such as CLIP4Clip [32] and LocVTP [7], both of which were pretrained on the large-scale HowTo100M [35] datasets (136M video-text pairs). While CLIP4Clip outperforms our model in

LSMDC zero-shot retrieval due to the closer domain similarity between LSMDC (average clip length: 4.7 seconds) and HowTo100M (average clip length: 4 seconds, compared to our dataset’s 10 seconds per clip), our model still achieves competitive performance after fine-tuning. We also include VideoPrism [65] for reference but do not make direct comparisons, as it leverages a massive (618M) non-public dataset and uses a ViViT backbone, which is 30× more computationally expensive than our approach (see Appendix D).

4.3. Classification

Setting: We evaluate our GEXIA method on video-related classification tasks using the LVU [52], COIN [44], Charades-Ego [41], and How2QA [25] datasets. LVU and COIN are designed for long-form video understanding. Charades-Ego, an action-classification-targeted dataset with egocentric videos, is used to evaluate the domain transfer performance of our approach. How2QA is leveraged to evaluate our approach on a cross-modal classification task.

Following the methodology of LF-VILA [43], we exclusively employ the video branch and add a fully-connected layer for classification on the LVU and COIN datasets reporting the results in terms of accuracy. The *#iter* is set to 3. We select the three content understanding tasks from the

Table 2. Classification results on the LVU, COIN, Charades-Ego, and How2QA datasets. Results from models using larger backbones on the LVU and COIN datasets are shown in gray for reference. #PT denotes the size of the pretraining dataset, ZS indicates zero-shot results, and FT refers to finetuned results.

(a) Accuracy Results on LVU dataset.

Method	Relation	Speaking	Scene
VideoBERT [42]	52.4	37.9	54.9
ObjT4mer [52]	53.1	39.4	56.9
LST [18]	52.4	37.3	62.8
Orthoformer [37]	50.0	38.3	66.3
VIS4mer [18]	57.1	40.8	67.4
Ours	61.9	42.7	70.9
LF-VILA [43]	61.5	41.3	68.0
S5 _{60frame} [47]	61.9	41.8	69.9
S5 _{100frame} [47]	66.7	41.8	73.3
MA-LMM _{100frame} [17]	58.2	44.8	80.3

(c) Results on Charades-Ego dataset.

Method	ZS mAP	FT mAP	Δ_{FTZS}
Actor [40]	-	20.0	-
SSDA [10]	-	23.1	-
I3D [10]	-	25.8	-
Ego-Exo [27]	-	30.1	-
EgoVLP [28]	25.0	32.1	7.1
HierVL-Avg [2]	25.2	32.6	7.4
HierVL-SA [2]	26.0	33.8	7.8
Ours	23.4	33.8	10.4

(b) Results on COIN dataset.

Method	#PT	Accuracy
CLIPBert [24]	5.6M	65.4
MIL-NCE [34]	136M	70.2
SlowFast [13]	136M	71.6
VideoCLIP [53]	136M	72.5
Ours	10M	84.8
TS4mer [4]	0.37M	83.5
TS4mer [4]	136M	85.3
LF-VILA [43]	60M	85.7
S5 _{100frame} [48]	60M	90.4
MA-LMM _{100frame} [17]	-	93.2

(d) Results on How2QA dataset.

Method	#PT	Accuracy
ATP [5]	400M	65.4
Hero [25]	136M	71.6
LF-VILA [43]	60M	76.1
FrozenBiLM [58]	10M	81.5
SeViLA [62]	-	83.7
SiaSamRea [63]	5.6M	84.1
JustAsk [57]	69M	84.4
Ours	10M	84.8

LVU dataset (Relationship, Way of Speaking, and Scene) as outlined in LF-VILA. For the Charades-Ego dataset, we replicate the process used by HierVL [2], treating it as a retrieval task across 157 distinct short action captions, and evaluate using the mean average precision (mAP) metric, with *#iter* configured per LVST settings. How2QA is a multi-choice video question answering (VQA) dataset. It consists of 44K short QA pairs and each question has one correct answer and 3 wrong answers. We concatenate questions and answers as text data. For each downstream task, our pretrained model was finetuned for 100 epochs on 8 NVIDIA Tesla-V100 GPUs. We utilized an AdamW optimizer [31], setting the learning rate for the video and text encoders at $1e^{-7}$, and for other components at $1e^{-6}$ managed by a cosine annealing schedule [30]. Detailed *#iter* settings are provided in Appendix E.

Long-form Video Understanding: Our GEXIA method achieves state-of-the-art performance in video classification tasks **on the LVU dataset**, as shown in Table 2a, even outperforming some models that may have been pretrained on larger datasets or utilize backbones with higher FLOPs. Notably, while the recent work S5_{60frame} [48] is specifically designed for long-form videos with state compression, our model still demonstrates superior performance (Accuracy +0.9% for Speaking, +1.0% for Scene), validating the success of our proposed IAM module. Furthermore, our model surpasses LF-VILA [43], which is also designed for long-form video pretraining, by a considerable margin (Accuracy +0.4% for Relationship, +1.4% for Speaking, +1.9% for Scene). Although the MA-LMM method [17] uses a pretrained large multimodal model and achieves higher results on the Speaking and Scene labels, our model still outperforms it on the Relationship label. Our model also attains the highest score **on the COIN dataset** among comparable methods [13, 24, 34, 42]. The results, shown in Table 2b, are

Table 3. Retrieval results of finetuned model on ActivityNet Captions dataset across different pretraining data setups.

Data Setup	Text-to-video			Video-to-text		
	R@1	R@5	R@10	R@1	R@5	R@10
short-video-short-text (SVST)	41.6	72.8	83.7	41.2	72.6	84.0
+long-video-long-text (LVLT)	42.8	72.8	84.0	41.8	73.0	84.3
+long-video-short-text (LVST)	43.1	73.2	84.5	42.6	73.4	84.8

competitive with LF-VILA [43] and TimeSformer [4], even though those models are pretrained on more data or utilize backbones with higher FLOPs.

These findings highlight the effectiveness of our video branch model in capturing long-form temporal information. Our model’s significant advantage stems from its ability to handle the varying information density across the temporal dimension in long-form videos. This adaptability to different granularities is a key factor in our model’s success.

Domain Transfer to Egocentric Video: We present the results from the Charades-Ego dataset in Table 2c to demonstrate the robustness of our model. While our model shows relatively lower zero-shot results compared to baseline methods due to the domain gap between our pretrained data and the egocentric evaluation data, it significantly outperforms other baselines after finetuning. Notably, it achieves performance comparable to the HierVL [2] method, which is pretrained on the in-domain Ego4D dataset [16]. The key observation is the substantial disparity between our model’s zero-shot and finetuned mAP results, highlighting its ability to adapt effectively to different video content after finetuning, even when pretrained on an out-of-domain dataset.

Cross-modal Classification: We finally show results on the How2QA dataset in Table 2d. Our model outperforms the baseline methods, achieving the highest accuracy. Our result significantly surpasses the video-text pretraining models [25, 43], and also outperforms the method leveraging LLMs [62], as well as those specifically designed for VQA tasks [57, 63]. The superior performance demonstrates that the low-dimensional embeddings of multi-grained data generated by IAM approximate dense features well without losing critical details in videos and texts.

4.4. Ablation Studies

We conduct ablation studies to investigate different settings of GEX and IAM, especially pretraining data compositions and iteration numbers (*#iter*). We also study the setups of the encoders, the random Integration in GEX, and the selections of LLM in GEX detailed in Appendix A.

Pretraining Data Compositions: To validate the performance improvement from introducing multi-grained data with GEX, we pretrain our model with different compositions of pretraining data and finetune the model on ActivityNet for retrieval. We use *#iter*=1 to establish a fair comparison. As shown in Table 3, the data composition with three different granularities works best while the training without multi-grained data is the worst. This result demon-

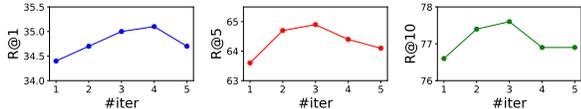


Figure 3. Zero-shot T2V retrieval results with different $\#iter$ for long video/text data on the ActivityNet Captions dataset.

strates the effectiveness of the proposed GEX pipeline.

Iteration Number Selection for Long Videos and Texts:

We empirically study the best option of the $\#iter$ for long videos and texts using ActivityNet Captions T2V retrieval task as a benchmark. As shown in Figure 3, it is clear that, ranging from 1 to 5, $\#iter = 3$ strikes the best balance between effectiveness and efficiency for long videos and texts. Therefore, we set $\#iter = 3$ for long video/text data in the training and inference stage.

Iteration Number in the Pretraining Stage: To fully understand the impact of iteration numbers in IAMs on multi-grained information extraction and alignments, we test different setups of video-text $\#iter$ in the pretraining stage. Given the optimal setup of $\#iter$ during pretraining, i.e., SVST: 1-1; LVLT: 3-3; LVST: 3-1, we introduce two sub-optimal setups: fixed $\#iter$ to 1-1 or 3-3 across SVST, LVLT, and LVST regardless of their actual granularity. Subsequently, we conduct zero-shot retrieval experiments on ActivityNet Captions [6] and MSR-VTT [54] datasets, aiming to evaluate the performance of the pre-trained models while isolating the results from any influences of finetuning.

As shown in Table 4a, using a 3-3 only setup yields results that are comparable to the optimal setup on the short-video-short-text (SVST) MSR-VTT dataset while using a 1-1 only setup results in a large performance drop. This reveals the fact that assigning smaller $\#iter$ to short inputs is sufficient while assigning larger $\#iter$ to long inputs is necessary during pretraining to achieve better performance. For the long-video-long-text (LVLT) ActivityNet Captions dataset, the optimal setup outperforms both the 1-1 only and 3-3 only setups. This underscores the importance of proper $\#iter$ counts to the length of the input during pretraining, as this strategy is crucial for achieving the best results in information extraction and alignment with multi-grained data. Overall, we set the $\#iter = 1$ for short videos/texts and $\#iter = 3$ for long videos/texts.

Iteration Number in the Inference Stage: We explore the different iteration numbers and their impact in the inference stage after pretraining the model with multi-grained data pairs. As reported in Table 4b, it becomes evident that the number of iterations should be aligned with the form of the video and text content. Benchmarked against InternVid [49], which comprises short videos (average 10 sec) with short texts (9.4 words), the ActivityNet Captions dataset — characterized by an average video length of 180 sec and an average text length of 49.2 words (see Appendix E) — falls into the long-video-long-text (LVLT)

Table 4. Ablation study results on different setups of $\#iter$.

(a) Zero-shot retrieval results across different **pretraining** $\#iter$ setups.

Pretraining Video-Text $\#iter$ setup	ActivityNet Captions T2V			ActivityNet Captions V2T		
	R@1	R@5	R@10	R@1	R@5	R@10
Only 1-1	35.2	63.9	75.6	34.4	63.6	76.6
Only 3-3	35.0	64.0	76.2	34.0	64.1	76.4
1-1&3-1&3-3	35.3	64.3	76.3	35.0	64.9	77.6

Pretraining Video-Text $\#iter$ setup	MSRVTT T2V			MSRVTT V2T		
	R@1	R@5	R@10	R@1	R@5	R@10
Only 1-1	36.6	63.5	72.3	39.8	63.5	72.3
Only 3-3	36.7	64.8	72.3	40.0	64.3	73.4
1-1&3-1&3-3	37.3	63.8	72.5	39.5	64.6	73.0

(b) Experiment results across different **inference** $\#iter$ setups. FT: Finetuned results; ZS: Zero-shot results.

Inference Video-Text $\#iter$ setup	FT ActivityNet Captions T2V			ZS MSRVT T2V		
	R@1	R@5	R@10	R@1	R@5	R@10
1-1	43.1	73.2	84.5	37.3	63.8	72.5
3-1	44.8	75.5	86.0	36.3	62.5	72.1
3-3	45.3	76.5	86.6	36.3	62.8	71.7

Inference Video-Text $\#iter$ setup	ZS LSMDC T2V			ZS Charades-Ego
	R@1	R@5	R@10	mAP
1-1	12.9	28.3	35.5	22.9
3-1	13.8	27.7	35.3	23.4
3-3	14.1	26.6	33.5	23.3

category. For the ActivityNet, the optimal iteration configuration is found to be video-text $\#iter=3-3$. In contrast, the MSR-VTT and LSMDC, which have average video lengths of 14.8 sec and 4.7 sec, and average text lengths of 9.3 and 9.7 words, respectively, fit into the short-video-short-text (SVST) category. For these datasets, the video-text $\#iter=1-1$ configuration yields the best results. The Charades-Ego, with an average video length of 31.5 sec and an average text length of 3.9 words, is categorized as a long-video-short-text (LVST) dataset. Here, the video-text $\#iter=3-1$ configuration is most effective.

We also visualize the video-text alignment with different setups in Appendix G. These findings not only affirm the versatility of our model in processing multi-grained data but also highlight its ability to adapt during the inference stage with just the modification of the $\#iter$ parameter. This adaptability is crucial for handling various types of video-text pairs, ensuring optimal performance across diverse downstream datasets.

5. Conclusion and Future Work

In this paper, we propose GEXIA for scalable multi-grained video-text learning. We first introduce GEX, a video-text data granularity expansion method that transforms single-grained video-text datasets into a multi-grained format. We also present a multi-grained video-text alignment framework, which features the IAM specifically designed to embed and align multi-grained information effectively. Remarkably, our model achieves state-of-the-art or comparable results across seven downstream datasets, underscoring the effectiveness of both the GEX and IAM. In the future, we aim to develop a multi-grained video-text benchmark to evaluate our model and other baseline approaches comprehensively. We hope that the GEXIA framework will prove valuable to the broader research community.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning. *ArXiv*, abs/2204.14198, 2022. [1](#)
- [2] Kumar Ashutosh, Rohit Girdhar, Lorenzo Torresani, and Kristen Grauman. Hiervl: Learning hierarchical video-language embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23066–23078, 2023. [1](#), [2](#), [3](#), [7](#)
- [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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1728–1738, 2021. [1](#), [2](#), [3](#), [6](#)
- [4] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, volume 2, page 4, 2021. [7](#)
- [5] Shyamal Buch, Cristóbal Eyzaguirre, Adrien Gaidon, Jiajun Wu, Li Fei-Fei, and Juan Carlos Niebles. Revisiting the “video” in video-language understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2917–2927, 2022. [7](#)
- [6] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. [1](#), [8](#)
- [7] Meng Cao, Tianyu Yang, Junwu Weng, Can Zhang, Jue Wang, and Yuexian Zou. Locvtp: Video-text pre-training for temporal localization. In *European Conference on Computer Vision*, 2022. [6](#)
- [8] Yuxiao Chen, Jianbo Yuan, Yu Tian, Shijie Geng, Xinyu Li, Ding Zhou, Dimitris N. Metaxas, and Hongxia Yang. Revisiting multimodal representation in contrastive learning: From patch and token embeddings to finite discrete tokens. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15095–15104, June 2023. [1](#)
- [9] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. <https://lmsys.org/blog/2023-03-30-vicuna/>, March 2023. [5](#)
- [10] Jinwoo Choi, Gaurav Sharma, Manmohan Chandraker, and Jia-Bin Huang. Unsupervised and semi-supervised domain adaptation for action recognition from drones. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1717–1726, 2020. [7](#)
- [11] Chaorui Deng, Qi Chen, Pengda Qin, Da Chen, and Qi Wu. Prompt switch: Efficient clip adaptation for text-video retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15648–15658, 2023. [5](#), [6](#)
- [12] David Fan, Jue Wang, Shuai Liao, Zhikang Zhang, Vimal Bhat, and Xinyu Li. Text-guided video masked autoencoder. In *European Conference on Computer Vision*, pages 282–298. Springer, 2025. [1](#)
- [13] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. [7](#)
- [14] Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, and Zicheng Liu. Violet: End-to-end video-language transformers with masked visual-token modeling. *arXiv preprint arXiv:2111.12681*, 2021. [6](#)
- [15] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 214–229. Springer, 2020. [5](#)
- [16] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022. [7](#)
- [17] Bo He, Hengduo Li, Young Kyun Jang, Menglin Jia, Xuefei Cao, Ashish Shah, Abhinav Shrivastava, and Ser-Nam Lim. Ma-Imm: Memory-augmented large multimodal model for long-term video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. [7](#)
- [18] Md Mohaiminul Islam and Gedas Bertasius. Long movie clip classification with state-space video models. In *European Conference on Computer Vision*, pages 87–104. Springer, 2022. [7](#)
- [19] Md Mohaiminul Islam, Ngan Ho, Xitong Yang, Tushar Nagarajan, Lorenzo Torresani, and Gedas Bertasius. Video recap: Recursive captioning of hour-long videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18198–18208, 2024. [2](#), [3](#)
- [20] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *International conference on machine learning*, pages 4651–4664. PMLR, 2021. [4](#)
- [21] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, 2021. [1](#)
- [22] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pages 706–715, 2017. [5](#)

- [23] Seon-Ho Lee, Jue Wang, Zhikang Zhang, David Fan, and Xinyu Li. Video token merging for long-form video understanding. *arXiv preprint arXiv:2410.23782*, 2024. [1](#)
- [24] Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7331–7341, 2021. [1](#), [2](#), [6](#), [7](#)
- [25] 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 *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2046–2065, 2020. [1](#), [6](#), [7](#)
- [26] Linjie Li, Zhe Gan, Kevin Lin, Chung-Ching Lin, Zicheng Liu, Ce Liu, and Lijuan Wang. Lavender: Unifying video-language understanding as masked language modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [6](#)
- [27] Yanghao Li, Tushar Nagarajan, Bo Xiong, and Kristen Grauman. Ego-exo: Transferring visual representations from third-person to first-person videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6943–6953, 2021. [7](#)
- [28] Kevin Qinghong Lin, Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Z XU, Difei Gao, Rong-Cheng Tu, Wenzhe Zhao, Weijie Kong, et al. Egocentric video-language pretraining. *Advances in Neural Information Processing Systems*, 35:7575–7586, 2022. [7](#)
- [29] Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10556–10565, 2020. [4](#)
- [30] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2016. [5](#), [7](#)
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. [5](#), [7](#)
- [32] 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 and captioning. *Neurocomputing*, 508:293–304, 2022. [1](#), [2](#), [5](#), [6](#)
- [33] Yiwei Ma, Guohai Xu, Xiaoshuai Sun, Ming Yan, Ji Zhang, and Rongrong Ji. X-clip: End-to-end multi-grained contrastive learning for video-text retrieval. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 638–647, 2022. [2](#), [3](#), [5](#), [6](#)
- [34] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9879–9889, 2020. [7](#)
- [35] 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. [1](#), [3](#), [6](#)
- [36] Vishal Monga, Yuelong Li, and Yonina C. Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38:18–44, 2019. [4](#)
- [37] Mandela Patrick, Dylan Campbell, Yuki Asano, Ishan Misra, Florian Metze, Christoph Feichtenhofer, Andrea Vedaldi, and Joao F Henriques. Keeping your eye on the ball: Trajectory attention in video transformers. *Advances in neural information processing systems*, 34:12493–12506, 2021. [7](#)
- [38] 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. [1](#), [4](#), [5](#)
- [39] Anna Rohrbach, Atousa Torabi, Marcus Rohrbach, Niket Tandon, Christopher Pal, Hugo Larochelle, Aaron Courville, and Bernt Schiele. Movie description. *International Journal of Computer Vision*, 123:94–120, 2017. [1](#), [5](#)
- [40] Gunnar A Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Actor and observer: Joint modeling of first and third-person videos. In *proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7396–7404, 2018. [7](#)
- [41] Gunnar A Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Charades-ego: A large-scale dataset of paired third and first person videos. *arXiv preprint arXiv:1804.09626*, 2018. [1](#), [6](#)
- [42] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7464–7473, 2019. [7](#)
- [43] Yuchong Sun, Hongwei Xue, Ruihua Song, Bei Liu, Huan Yang, and Jianlong Fu. Long-form video-language pre-training with multimodal temporal contrastive learning. *Advances in neural information processing systems*, 35:38032–38045, 2022. [1](#), [2](#), [6](#), [7](#)
- [44] Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. Coin: A large-scale dataset for comprehensive instructional video analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1207–1216, 2019. [1](#), [6](#)
- [45] Zineng Tang, Jaemin Cho, Jie Lei, and Mohit Bansal. Perceiver-vl: Efficient vision-and-language modeling with iterative latent attention. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4410–4420, 2023. [2](#), [5](#), [6](#)
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [4](#)
- [47] Jue Wang, Wentao Zhu, Pichao Wang, Xiang Yu, Linda Liu, Mohamed Omar, and Raffay Hamid. Selective structured

- state-spaces for long-form video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6387–6397, 2023. 7
- [48] Jue Wang, Wentao Zhu, Pichao Wang, Xiang Yu, Linda Liu, Mohamed Omar, and Raffay Hamid. Selective structured state-spaces for long-form video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6387–6397, 2023. 7
- [49] Yi Wang, Yanan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Ma, Xinyuan Chen, Yaohui Wang, Ping Luo, Ziwei Liu, et al. Internvid: A large-scale video-text dataset for multimodal understanding and generation. *arXiv preprint arXiv:2307.06942*, 2023. 1, 3, 5, 6, 8
- [50] Yicheng Wang, Jiayi Yuan, Yu-Neng Chuang, Zhuoer Wang, Yingchi Liu, Mark Cusick, Param Kulkarni, Zhengping Ji, Yasser Ibrahim, and Xia Hu. Dhp benchmark: Are llms good nlg evaluators? *arXiv preprint arXiv:2408.13704*, 2024. 5
- [51] Ziyang Wang, Yi-Lin Sung, Feng Cheng, Gedas Bertasius, and Mohit Bansal. Unified coarse-to-fine alignment for video-text retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2816–2827, 2023. 5, 6
- [52] Chao-Yuan Wu and Philipp Krahenbuhl. Towards long-form video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1884–1894, 2021. 1, 6, 7
- [53] Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metzger, Luke Zettlemoyer, and Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *arXiv preprint arXiv:2109.14084*, 2021. 7
- [54] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296, 2016. 1, 5, 8
- [55] Hongwei Xue, Tiankai Hang, Yanhong Zeng, Yuchong Sun, Bei Liu, Huan Yang, Jianlong Fu, and Baining Guo. Advancing high-resolution video-language representation with large-scale video transcriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5036–5045, 2022. 1, 3
- [56] Hongwei Xue, Yuchong Sun, Bei Liu, Jianlong Fu, Ruihua Song, Houqiang Li, and Jiebo Luo. Clip-vip: Adapting pre-trained image-text model to video-language alignment. In *The Eleventh International Conference on Learning Representations*, 2023. 5, 6
- [57] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1686–1697, 2021. 7
- [58] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Zero-shot video question answering via frozen bidirectional language models. *Advances in Neural Information Processing Systems*, 35:124–141, 2022. 7
- [59] Jianwei Yang, Yonatan Bisk, and Jianfeng Gao. Taco: Token-aware cascade contrastive learning for video-text alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11562–11572, 2021. 1, 2, 6
- [60] Jinrui Yang, Xianhang Li, Druv Pai, Yuyin Zhou, Yi Ma, Yaodong Yu, and Cihang Xie. Scaling white-box transformers for vision. *arXiv preprint arXiv:2405.20299*, 2024. 4
- [61] Qinghao Ye, Guohai Xu, Ming Yan, Haiyang Xu, Qi Qian, Ji Zhang, and Fei Huang. Hitea: Hierarchical temporal-aware video-language pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15405–15416, 2023. 1, 2, 3, 6
- [62] Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. Self-chained image-language model for video localization and question answering. *Advances in Neural Information Processing Systems*, 36, 2024. 7
- [63] Weijiang Yu, Haoteng Zheng, Mengfei Li, Lei Ji, Lijun Wu, Nong Xiao, and Nan Duan. Learning from inside: Self-driven siamese sampling and reasoning for video question answering. *Advances in Neural Information Processing Systems*, 34:26462–26474, 2021. 7
- [64] Yaodong Yu, Sam Buchanan, Druv Pai, Tianzhe Chu, Ziyang Wu, Shengbang Tong, Benjamin Haeffele, and Yi Ma. White-box transformers via sparse rate reduction. *Advances in Neural Information Processing Systems*, 36:9422–9457, 2023. 4
- [65] Long Zhao, Nitesh Bharadwaj Gundavarapu, Liangzhe Yuan, Hao Zhou, Shen Yan, Jennifer J Sun, Luke Friedman, Rui Qian, Tobias Weyand, Yue Zhao, et al. Videoprism: A foundational visual encoder for video understanding. In *Forty-first International Conference on Machine Learning*. 6
- [66] Shuai Zhao, Linchao Zhu, Xiaohan Wang, and Yi Yang. Centerclip: Token clustering for efficient text-video retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 970–981, 2022. 6