

# Beyond Coarse-Grained Matching in Video-Text Retrieval

Aozhu Chen<sup>1</sup>\*, Hazel Doughty<sup>2</sup>, Xirong Li<sup>1</sup>, and Cees G. M. Snoek<sup>3</sup>

<sup>1</sup> Renmin University of China

<sup>2</sup> Leiden University

<sup>3</sup> University of Amsterdam

**Abstract.** Video-text retrieval has seen significant advancements, yet the ability of models to discern subtle differences in captions still requires verification. In this paper, we introduce a new approach for fine-grained evaluation. Our approach can be applied to existing datasets by automatically generating hard negative test captions with subtle single-word variations across nouns, verbs, adjectives, adverbs, and prepositions. We perform comprehensive experiments using four state-of-the-art models across two standard benchmarks (MSR-VTT and VATEX) and two specially curated datasets enriched with detailed descriptions (VLN-UVO and VLN-OOPS), resulting in a number of novel insights: 1) our analyses show that the current evaluation benchmarks fall short in detecting a model’s ability to perceive subtle single-word differences, 2) our fine-grained evaluation highlights the difficulty models face in distinguishing such subtle variations. To enhance fine-grained understanding, we propose a new baseline that can be easily combined with current methods. Experiments on our fine-grained evaluations demonstrate that this approach enhances a model’s ability to understand fine-grained differences.

**Keywords:** Video-Language · Video-Text Retrieval · Fine-grained

## 1 Introduction

In recent years, video-text retrieval has seen huge improvements [4, 9, 14, 17, 22, 29, 31, 34, 44, 50, 53]. This increased success is bolstered by large-scale web-scraped vision-language datasets [16] and the ability to bootstrap training with more readily available image-text pairs [41]. Popular video-text retrieval benchmarks [2, 8, 26, 30, 48, 54] have paved the way for this success, with models often able to recall the correct caption from thousands in the top 10, 5 or even first result. Although impressive, this is only a step towards fully understanding the correspondence between video and text. The collection process of current benchmarks results in a diverse set of videos and corresponding captions which are distinct from each other. This means that current benchmarks only evaluate coarse-grained video retrieval and that we do not know the fine-grained



**a black car is driving down the road**

> **Coarse-grained Negatives**

a person is connecting something to system

this is a video of a live tv show

people are singing on the beach

a little girl does gymnastics

a boy is singing

> **Fine-grained Negatives**

a white car is driving down the road

a black motorcycle is driving down the road

a black car is parked down the road

a black car is driving across the road

a black car is driving down the mountain

**Fig. 1: Coarse vs. Fine-Grained Video Retrieval.**

\* Work was part of an internship at University of Amsterdam

capabilities of current models. As shown in Fig. 1, the ground-truth caption (black) for the given video is distinctly different from the other captions in the test set (coarse-grained negatives in brown). This makes it relatively easy for the model to distinguish between these captions. We refer to matching and retrieving such videos and texts as coarse-grained evaluation. However, it is currently unknown whether a model can discern the subtle single-word differences between the fine-grained negative captions (green) and the original caption, which is challenging and critical for accurate video-text alignment.

In this paper, we propose to go beyond coarse-grained matching in video-text retrieval and focus on fine-grained differences in captions. We create a new fine-grained evaluation approach that can be used with existing datasets by automatically generating new hard negative test captions with subtle single-word variations across different parts-of-speech. We find that current models struggle to distinguish such subtle variations, particularly in prepositions and adverbs. The challenge in addressing this is not only to make a model understand variations in different parts-of-speech but to gain this understanding while maintaining a model’s impressive coarse-grained abilities. To take a first step towards rectifying this, we propose a plug-in method that balances coarse- and fine-grained objectives with phrase-level negatives and fine-grained prompting.

Our contributions can be summarized as: (i) we highlight the coarse-grained nature of current video-retrieval evaluation, (ii) we propose a new fine-grained evaluation approach to analyze models’ sensitivity to individual word variations in different parts-of-speech, (iii) we investigate current models using our proposed evaluation and find they lack fine-grained understanding, and (iv) we propose a simple method which can be easily combined with existing approaches to learn representations suitable for both coarse- and fine-grained retrieval.

## 2 Related Work

**Coarse-Grained Video-Text Retrieval.** Video-text retrieval is an important topic in the vision-language domain that has garnered much attention [1, 10–13, 32, 43, 47, 49, 55, 57, 60]. The aim is to retrieve the correct video given a text query. Thus, it is crucial to align the features of videos to their related texts. Early works in video-text retrieval [45, 57, 58] design different fusion mechanisms for pre-extracted video and text features. Later, Clip-BERT [28] enabled end-to-end training with a sparse sampling strategy on video data in addition to image-text pretraining. Frozen [4] also uses image-text pretraining with a curriculum learning schedule. Several works [5, 6, 19, 24, 33] instead build on CLIP [39], repurposing it for video-text retrieval. However, these works can only understand coarse-grained differences between text and video; they fail to recognize the subtle one-word caption differences investigated in this paper.

**Token-Level Similarity in Video-Text Retrieval.** Several works aim to make the understanding of visual and text correspondence more fine-grained. One category of methods enhances retrieval performance by designing different granularities of cross-modal similarity [20, 27, 34, 36, 50, 56, 64]. For example, TS2Net [31] designs a cross-modal alignment module between caption and frame

features, rather than only the video-level features. X-CLIP [34] calculates similarities between video-sentence, video-word, sentence-frame, and frame-word, using varying levels of similarity to allow the model to perceive multi-granularity alignments. Similarly, UCoFiA [50] uses coarse-to-fine similarity alignment, going a step further with patch-word similarity. While achieving impressive results on current video-to-text benchmarks, we show in this paper that these works still struggle to understand fine-grained single-word differences in video captions.

**Parts-of-Speech in Video-Text Retrieval.** Other methods aim for fine-grained understanding by decomposing the input caption into different parts-of-speech [3,9,15,37,51,53]. For instance, HGR [9] uses hierarchical graph reasoning on verbs, noun phrases, and complete sentences to model a semantic graph of events, actions, entities, and their relationships. JPoSE [53] learns separate embedding spaces for verbs and nouns before combining these into a fine-grained action embedding. ActionInVerb [37] improves verb understanding by leveraging pretrained large language models to generate hard negatives which only change the verbs in sentences. Several works also go beyond verb-noun understanding in video-text retrieval and examine individual parts-of-speech with adverbs [15,16], temporal prepositions [3] and negation [51]. Our work also focuses on fine-grained video-text retrieval with different parts-of-speech, however unlike prior work, it goes beyond investigating verb-noun pairs or single parts-of-speech.

**Analysis of Vision-Language Models.** Several studies have investigated the comprehension of fine-grained information in vision-language models [7,18,21,35,42,52,59,61,62]. Foilit [42] creates an erroneous description ('foil'), by replacing nouns in an image caption with similar but incorrect words, testing models with foil classification, foil word detection, and foil word correction. SVO-Probes [21] assess image-text models' understanding of verbs by evaluating images related or unrelated to specific sentences with subject, verb, and object triplets. Madasu *et al.* [35] investigate compositional and semantic understanding in video retrieval by removing, shifting, or replacing, verbs and nouns in the ground-truth captions. Similarly, other works [59,61] also highlight the lack of compositional understanding in image-text models. VL-Checklist [62] instead tests VLMs' understanding of objects, attributes, and relations, while Wray *et al.* [52] focus on models' understanding of different captions with similar semantic meanings. Different to these prior works, this paper explores video-text retrieval models' understanding of nouns, adjectives, verbs, adverbs, and prepositions, focusing specifically on whether models can accurately understand and respond to subtle one-word changes in language.

### 3 Shortcomings of Current Video Retrieval Benchmarks

In this section, we highlight the coarse-grained nature of the current video-text retrieval paradigm by examining current datasets and the captions they contain.

**Definition & Notation.** Video-text retrieval models aim to learn representations to determine the relevance of texts  $T$  to videos  $V$ . Specifically, given a video  $v \in V$ , the aim is to rank the set of texts  $T$  according to their relevance to  $v$  and vice versa. What is learned by a representation thus depends heavily on the

**Table 1: Details of Benchmarks in Text-to-Video Retrieval.** We show the number of videos and captions in the train and test sets, alongside the proportions of captions containing nouns, verbs, adjectives, adverbs, and prepositions.

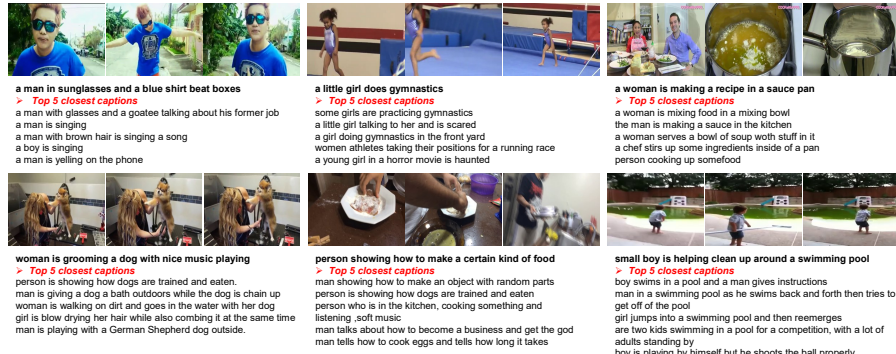
Dataset	Source	Training Set		Test Set		#Vocab	Proportions of Captions					Eval. Metrics
		#Video	#Caption	#Video	#Caption		Noun	Verb	Adj	Adv	Prep	
YouCook2 [63]	YouTube	13.8K	13.8K	3.3K	3.3K	2.3K	1.00	0.95	0.33	0.10	0.75	Recall@k, Av. Rank
DiDeMo [2]	Flicker	10.6K	41.2K	1.0K	4.0K	6.8K	1.00	0.88	0.44	0.24	0.67	Recall@k
ActivityNet [26]	YouTube	10.9K	10.9K	4.9K	4.9K	13.0K	0.99	0.99	0.42	0.33	0.83	Recall@k
VATEX [48]	YouTube	28.9K	289.9K	1.5K	15.0K	29.1K	1.00	0.99	0.50	0.25	0.89	Recall@k, Av. Rank
MSR-VTT [54]	YouTube	9.0K	180.0K	1.0K	20.0K	25.7K	1.00	0.92	0.36	0.09	0.68	Recall@k, Av. Rank
MSVD [8]	YouTube	1.3K	52.9K	670	27.8K	12.9K	1.00	0.96	0.18	0.05	0.46	Recall@k, Av. Rank
TGIF [30]	Tumblr	89.4K	90.4K	11.3K	34.1K	11.4K	1.00	0.99	0.43	0.20	0.78	Recall@k, Av. Rank

content of the captions labeled as relevant and irrelevant to each video. Similarly, the insights gained from model evaluation hinges on the captions present in the test set. If the negative captions correspond to clear visual differences, we do not know if a model can distinguish fine-grained differences. Similarly, if captions can be distinguished by one part-of-speech alone, *e.g.* nouns, we do not know if a model can understand the visual appearance of different parts-of-speech.

**Datasets.** In Tab. 1 we show the statistics of commonly used video-text retrieval benchmarks. We order these by the size of the test set. A larger test set is generally more challenging as a model not only needs to distinguish between more items but there is an increased chance of having fine-grained differences between items. Most datasets originate from YouTube or other websites such as Flickr and Tumblr and are annotated after the fact via crowd-sourcing. During this annotation, an annotator observes a single video and is tasked with creating a relevant caption. This caption is not created by considering the other videos or captions present in the dataset. All datasets are evaluated using Recall@k or Avg. Rank where the relevant caption(s) should be retrieved in the top k results or higher than the irrelevant captions sourced from other videos. The granularity of this evaluation thus relies on the contents of the test set captions which rarely have subtle differences between each other due to the collection process. The proportions of captions with each word type in Tab. 1 also suggest these datasets do not test fine-grained understanding. Subjects, objects (nouns), and actions (verbs) are the most central elements, while the manner of actions (adverbs) is less common in captions.

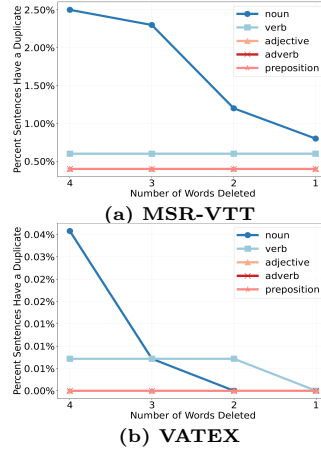
**Qualitative Analysis.** To investigate how fine-grained the captions in these datasets are we take a closer look at example captions in Fig. 2. Here, we show randomly selected videos from MSR-VTT [54] and VATEX [48] alongside the ground-truth caption and the top-5 closest captions in the test set according to CLIP [38] similarity. We see that the closest captions in the test set differ by multiple concepts *e.g.* ‘a little girl does gymnastics’ vs. ‘women athletes taking their position for a running race’. This suggests that these datasets inadequately test the fine-grained capabilities of models. Furthermore, most differences are due to different nouns and verbs, none of the examples share the same noun and verb, differing only in adjectives, adverbs, and prepositions.

**Quantitative Analysis.** To verify our observations from Fig. 2 hold true for the rest of the video-text pairs in these datasets, we conduct statistical analysis



**Fig. 2: Qualitative Analysis on MSR-VTT & VATEX.** We show ground-truth video-text pairs alongside the 5 most similar captions to the ground-truth according to CLIP. We see that captions are coarse-grained, *i.e.* there are multiple difference conceptual differences between captions which mostly relate to the nouns present.

on MSR-VTT and VATEX in Fig. 3. Specifically, we randomly remove a noun from each sentence in the test set and count how many sentences become identical after the removal. We do the same for two nouns, and so forth, until all nouns are removed, thereby testing whether a caption is the same in all other parts-of-speech. We follow the same procedure for verbs, adjectives, adverbs, and prepositions. From Fig. 3 we see that captions rarely differ due to a single word. Even when removing 4 nouns, only 0.25% of captions have a duplicate with the same remaining words. Of course, captions can have the same semantic meaning with different wording, however, the lack of single-word differences in this analysis combined with the qualitative examples highlights the need for fine-grained testing of current video-text retrieval models.

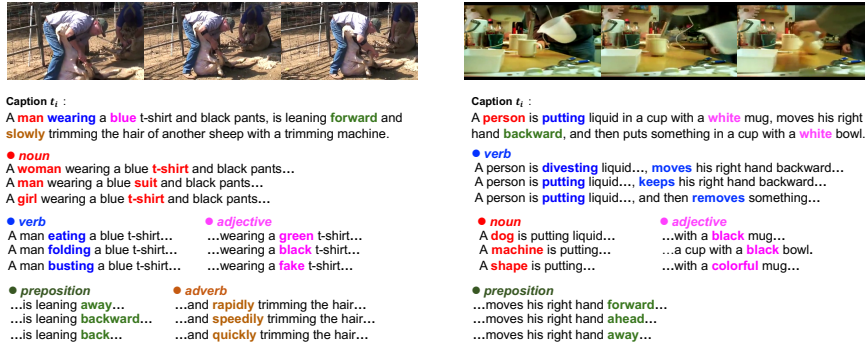


**Fig. 3: Analysis of Sentence Similarity.** Few captions differ by one part-of-speech, particularly outside of nouns.

### 4 PoSRank: A Fine-grained Evaluation Approach

We propose an approach for fine-grained evaluation of video retrieval which complements existing coarse-grained benchmarks. Particularly, our evaluation focuses on changing individual words to create multiple sentences with subtly different meanings. If a model truly understands every part of a sentence, then it should be able to detect these subtle differences between sentences and accurately retrieve the sentence most relevant to the video. We separate our evaluation by different parts-of-speech, *i.e.* nouns, verbs, adjectives, adverbs, and prepositions, to allow a better understanding of the fine-grained capabilities of the model.

**Definition.** We assume we have a set of videos  $V$  with a corresponding set of coarse-grained captions  $T$ . Given a video  $v_i \in V$  and its corresponding caption



**Fig. 4: Examples of Hard Negative Sentences** in our evaluation. By creating subtle single-word variations of the original sentence we test the ability of video-text retrieval models to distinguish fine-grained variations for different parts-of-speech.

$t_i \in T$ , we create a set of negative captions  $T_i^p$  of size  $N$  for video  $v_i$  for each part-of-speech  $p \in P$  we are interested in. We create  $T_i^p$  by modifying single words with the relevant part-of-speech tag from the caption  $t_i$  to generate a fine-grained hard negative sentence. Specifically, for each part-of-speech, we randomly select a word  $w$  from sentence  $t$  with that part-of-speech tag and modify it to word  $w'$  keeping all other words the same. We do this  $K$  times to acquire  $K$  different sentences which make up set  $T_i^p$ . When selecting word  $w'$  we prioritize the antonym of word  $w$ , followed by the antonym of a hypernym or hyponym of the word. Otherwise, we randomly select a word of the same part-of-speech from the dataset vocabulary. By ensuring the part-of-speech tag is the same we ensure that the sentence is grammatically correct and increase the challenge of the evaluation. Prioritizing antonyms also helps the plausibility of the sentence, increases the challenge and avoids false negatives. From inspecting a random subset of 100 captions we found only 4 false negatives and 7 grammatical errors. For each video  $v_i$  we create negative caption sets for nouns  $T_i^{noun}$ , verbs  $T_i^{verb}$ , adjectives  $T_i^{adj}$ , adverbs  $T_i^{adv}$ , and prepositions  $T_i^{prep}$ . Fig. 4 shows examples of the hard negative sentences created in this process. Only changing a single word per caption results in challenging captions where the model has to distinguish fine-grained concepts *e.g.* ‘slow’ vs. ‘fast’, ‘forward’ vs. ‘backward’ as well as color, age, and gender.

**Evaluation Criteria.** We utilize the newly created negative sentences to evaluate the fine-grained performance of video retrieval models. Specifically, for a part-of-speech  $p$  and each video  $v_i$ , we obtain the ranking of the ground-truth sentence  $t_i$  and negatives  $T_i^p$ . Suppose the rank of the ground-truth  $t_i$  within the retrieved results, including both the correct sentence and the negatives, is denoted by  $r_i^p$ . We compute the Mean Reciprocal Rank (MRR) for part-of-speech  $p$ , named  $PoSRank^p$ , as follows:

$$PoSRank^p = \frac{1}{N} \sum_{i=1}^N \frac{1}{r_i^p}, \quad (1)$$

where  $N$  is the number of ground-truth video-text pairs  $(v_i, t_i)$ . In summary,  $PoSRank^p$  measures the effectiveness of the retrieval model in ranking

**Table 2: Statistics of the fine-grained test sets.** We show the number of original captions, their average length, and the proportions of captions containing each part-of-speech. For all datasets, we create a challenging fine-grained test set per part-of-speech.

Dataset	Original Captions			Proportions of Captions					Fine-Grained Test Captions				
	#Train	#Test	Av. Len	noun	verb	adj	adv	prep	noun	verb	adj	adv	prep
MSR-VTT [54]	180,000	20,000	9.3	1.00	0.92	0.36	0.09	0.71	418.13K	367.56K	135.87K	17.09K	8.51K
VATEX [48]	519,820	13,969	14.3	1.00	0.99	0.49	0.24	0.88	293.35K	283.40K	134.59K	32.38K	12.12K
VLN-UVO [46]	12,359	6,561	17.2	1.00	1.00	0.89	0.26	0.86	137.74K	136.48K	120.23K	4.58K	4.96K
VLN-OOPS [46]	26,193	5,335	25.7	1.00	1.00	0.92	0.61	0.98	112.04K	111.62K	99.31K	18.61K	6.38K

the ground-truth sentence as close to the top of the list as possible across the negative sets  $T_i^p$  of each video. If the retrieval model robustly understands the semantics of each sentence part, the original caption should always rank first. Thus, a higher value indicates a better fine-grained understanding of the model with 1.0 being the best score and 0.0 the worst.

## 5 Analysing Fine-grained Ability of Current Models

We evaluate existing video-text retrieval methods using our proposed fine-grained evaluation to answer the following questions: (i) How well do current methods understand fine-grained differences in video retrieval? (ii) Does performance in fine-grained retrieval correlate with existing coarse-grained metrics? (iii) Are certain parts-of-speech more challenging than others for current methods? (iv) Are certain datasets more suitable than others for fine-grained evaluation?


### 5.1 Experiment Setup

**Datasets.** According to the approach described in Sec. 4, we create fine-grained evaluation sets for the commonly used MSR-VTT and VATEX. Additionally, we use VLN-UVO and VLN-OOPS [46] originally proposed for video narrative grounding, as these contain longer, more detailed descriptions and thus have a greater number of different part-of-speech with which to create new fine-grained captions, *e.g.* MSR-VTT has 9.3 words per sentence vs. 17.2 from VLN-UVO. VLN-UVO and VLN-OOPS also have higher percentages of adjectives, adverbs, and prepositions. For instance, in VLN-OOPS 61% of sentences contain at least one adverb, while in MSR-VTT only 9% of sentences have an adverb. More detailed statistics can be seen in Tab. 2.

**Baselines.** We use four recent, publicly available video retrieval models covering a variety of different backbones and training strategies. **Frozen** [4] is a joint embedding model for end-to-end video retrieval. It utilizes a ViT backbone pre-trained on CC3M [41] (image-text pairs) and WebVid-2M [4] (video-text pairs) with a contrastive loss. **TS2Net** [31] initializes its visual and text encoders with CLIP. It uses a token shift module which aims to capture subtle local movements between frames as well as a token selection module to minimize spatial and temporal redundancy. **X-CLIP** [34] also bases its model on CLIP and focuses on computing fine-grained cross-modal similarities. Specifically, it obtains the final video-text similarity by fusing video-sentence, video-word, sentence-frame, and frame-word similarities. **UCoFiA** [50] aligns text and video at both coarse

**Table 3: Fine-Grained Evaluation with Average PoSRank.** All four methods struggle to distinguish fine-grained differences in all datasets.

Method	MSR-VTT	VATEX	VLN-UVO	VLN-OOPS	Mean
Frozen [4]	0.285	0.243	0.252	0.249	0.257
X-CLIP [34]	0.343	0.278	0.301	0.282	0.301
UCoFiA [50]	0.351	0.268	0.308	0.299	0.306
TS2-Net [31]	0.351	0.283	0.310	0.293	0.309

**Fig. 5: Correlation Between Coarse & Fine-Grained Performance.** There is a weak correlation between performance on coarse and fine-grained metrics, however good coarse-grained performance does not indicate good fine-grained performance.

and fine-grained granularities. This is achieved by three types of matching: video-sentence, frame-sentence, and patch-word. This model is also based on CLIP.

**Implementation Details.** For each original caption, we create  $K=20$  negative captions for each of the five parts-of-speech using SpaCy for parsing.

## 5.2 Results


### How well do current methods understand fine-grained differences?

We show the performance of the four tested video-text retrieval models with our *PoSRank* metric, averaged over the five parts-of-speech, in Tab. 3. For all methods and datasets, the *PoSRank* scores range from 0.24 to 0.36 far from the best score of 1. This demonstrates that current methods have poor fine-grained understanding and are not sensitive to subtle single-word changes in the caption. CLIP-based methods X-CLIP, UCoFiA and TS2-Net all perform slightly better than Frozen. This may be due to the CLIP-based methods having a better understanding of different parts-of-speech or due to these approaches learning better video-text representations. We analyze both possibilities in the following questions. Regardless, all methods evaluated struggle with our proposed fine-grained evaluation and have much room for improvement.

### Does performance in fine-grained retrieval correlate with existing coarse-grained metrics?

Fig. 5 shows mean recall@{1,5,10} of video-to-text and text-to-video retrieval alongside the mean *PoSRank*. There is some correlation between performance on coarse and fine-grained metrics. For instance, Frozen is the worst in both coarse-grained and fine-grained retrieval. However, at the higher end the correlation is weak with the best-performing model in coarse-grained achieving only 3rd place on fine-grained for VLN-UVO and VLN-OOPS. This demonstrates that good coarse-grained performance does not indicate good fine-grained performance and further justifies the need for our fine-grained *PoSRank* evaluation. Furthermore, the differences between methods in coarse-grained met-





**Fig. 6: Fine-Grained Evaluation Per Part-of-Speech.** Models find fine-grained differences in adverbs and prepositions the most difficult to distinguish.

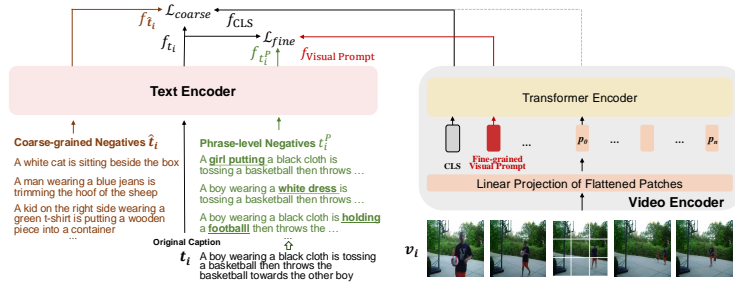
rics can be large, *e.g.* 71 for Frozen vs. 85 for UCoFiA on VATEX, while performance is more similar in our fine-grained *PoSRank*, *e.g.* 0.24 vs. 0.27. This demonstrates that all current methods, regardless of their success on current coarse-grained metrics, have a long way to go in fine-grained understanding.

**Are certain parts-of-speech more challenging than others?** Fig. 6 shows *PoSRank* performance for individual parts-of-speech for all methods on MSR-VTT and VLN-UVO. From these results, we see that current models can better distinguish differences in adjectives, with nouns being the second easiest. This is likely due to the reliance of video-text retrieval models on image pretraining where nouns and adjectives are common; three of the four methods make use of CLIP models while Frozen also pretrains on images and single video frames. Adverbs and prepositions are the hardest parts-of-speech on both datasets for almost all models. This is likely due to these parts-of-speech requiring more understanding of the spatial and temporal relationships between different objects and frames, unlike nouns and adjectives which can often be understood from the appearance of a single object in a single frame. Results for VATEX and VLN-OOPS follow similar trends and can be found in the supplementary material.

**Are certain datasets more suitable than others for fine-grained evaluation?** The fine-grained sets we create for VATEX, VLN-UVO and VLN-OOPS are all more challenging for current methods than the one created for MSR-VTT. For instance, the average *PoSRank* value for XCLIP is 0.278 on VATEX and 0.282 on VLN-OOPS, while on MSR-VTT, it is 0.343 (Tab. 3). This is due to longer and more complex original captions offering more potential for challenging and subtle variations that better test the fine-grained capabilities of a model. VLN-UVO and VLN-OOPS also provide more challenges for the most difficult parts-of-speech. For instance, on VLN-UVO models achieves  $\sim 0.2$  adverb *PoSRank*, compared to  $\sim 0.3$  on MSR-VTT (Fig. 6).

## 6 Training for Fine-grained Understanding

Our fine-grained evaluation shows there is still much room for improvement in current methods. To improve fine-grained understanding, we introduce a new baseline (Fig. 7) that can be easily combined with existing approaches. Beginning from an existing coarse-grained video-text retrieval model (Sec. 6.1), we train with fine-grained ‘word-level’ negatives (Sec. 6.2). While this improves the model’s ability to distinguish fine-grained text differences, it hurts the ability to



**Fig. 7: Fine-grained Training for Video-Text Retrieval** incorporates phrase-level negatives and fine-grained prompting to identify coarse and fine-grained differences.

distinguish coarse-grained differences. We strike a balance between coarse and fine with ‘phrase-level negatives’ (Sec. 6.3) and better allow the model to learn from both types of negatives with fine-grained prompts (Sec. 6.4).

### 6.1 Coarse-Grained Training

A typical retrieval model learns from a set of videos  $V$  and a corresponding set of coarse-grained captions  $T$ . As in coarse-grained evaluation each caption  $t_i \in T$  is paired with a single video  $v_i \in V$ , although a video may have multiple corresponding captions. The goal is to give matching video text pairs  $(v_i, t_i)$  high similarity. This similarity should be higher than the same video with a less relevant text  $(v_i, t'_i)$  or the same text with a less relevant video  $(v'_i, t_i)$ . Assuming all unpaired videos and texts are less relevant to each other than the pairs, a model can be trained with the following  $\mathcal{L}_{coarse}$  loss:

$$\begin{aligned} \mathcal{L}_{v2t} &= -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(s(v_i, t_i))}{\sum_{j=1}^B \exp(s(v_i, t_j))}, \\ \mathcal{L}_{t2v} &= -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(s(v_i, t_i))}{\sum_{j=1}^B \exp(s(v_j, t_i))}, \\ \mathcal{L}_{coarse} &= \mathcal{L}_{v2t} + \mathcal{L}_{t2v}, \end{aligned} \quad (2)$$

where  $s(\cdot)$  is a similarity function and  $B$  is the batch size. While this loss is effective for training models for existing coarse-grained video retrieval metrics, models trained in this way do not learn to distinguish fine-grained word-level differences. This is because just like the test set (see Sec. 3) the set of captions  $T$  in training does not contain fine-grained differences.

### 6.2 Training with Word-Level Negatives

In addition to using our sets of fine-grained negative captions  $T_i^p$  for evaluation, we take inspiration from prior works on hard negative [25, 40] mining and also incorporate our fine-grained negatives during training to enhance the model’s sensitivity to all parts of a sentence. For each video-text pair  $(v_i, t_i)$  in a batch, we generate  $\mathbb{N}$  word-level negative sentences for each of the five parts-of-speech we are interested in, using the method described in Section Sec. 4. We incorporate these word-level negatives into training using the following contrastive loss:

$$\mathcal{L}_{fine} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(s(v_i, t_i))}{\sum_{p \in P} \sum_{t \in T_i^p} \exp(s(v_i, t))}, \quad (3)$$

where  $P$  is the set of POS-tags of interest and  $T_i^p$  is the set of word-level negative captions created from caption  $t_i$  for POS-tag  $p$ .

Since we want the model to be both suitable for coarse and fine-grained evaluation we combine the two losses as follows:

$$\mathcal{L} = \mathcal{L}_{coarse} + \lambda \mathcal{L}_{fine}, \quad (4)$$

where  $\lambda$  is the weight to balance the coarse and fine loss. We use  $\lambda \ll 1$  as the similarity between a video and its word-level hard negatives should be greater than the similarity between a video and the coarse-grained negatives.

### 6.3 Training with Phrase-Level Negatives

Training with the loss in Eq. (4) enables the model to make fine-grained word-level distinctions, particularly for nouns, adjectives and verbs which are plentiful in the training data enabling the creation of a large variety of negatives. However, even with this loss models struggle to balance coarse and fine-grained, with results decreasing on coarse-grained metrics as fine-grained performance increases. Due to the vastly different objectives of coarse and fine-grained video retrieval current models struggle to learn both tasks simultaneously.


To address this we propose using ‘phrase-level’ negatives which balance coarse and fine granularities. Our phrase-level negatives modify two consecutive words in a sentence. As with the word-level negatives, for a video  $v_i$  with corresponding caption  $t_i$  we randomly select a word  $w_1$  from caption  $t_i$  to modify which is one of the five parts-of-speech of interest. Additionally, with 50:50 chance we either modify the preceding or following word which is also one of the five parts-of-speech. For example, if *black* is the word selected from *a boy wearing a black t-shirt* then we could obtain the following phrase-level negatives: *a boy removing a white t-shirt* or *a boy wearing a black coat*. As in the word-level negatives and fine-grained evaluation, words are replaced with other words from the dataset with the same POS-tag, prioritizing antonyms. All remaining words are unmodified.

### 6.4 Fine-Grained Prompting

Our phrase-level negatives allow the model to better balance between coarse and fine-grained distinctions. However, models still struggle to optimize for both objectives and underperform a model specialized for either coarse or fine-grained. To remedy this we incorporate prompts to inform the model which task it should focus on. Specifically, we use visual prompts [23] to learn fine-grained similarity. This gives us additional output tokens of the transformer so the same model can predict both fine and coarse-grained similarity. We train the output of the CLS token with the coarse-grained loss  $\mathcal{L}_{coarse}$  (Eq. (2)), which establishes a general understanding of the content. Concurrently, we train the output of the visual prompt with  $\mathcal{L}_{fine}$  (Eq. (3)), enhancing the model’s capability to detect subtle nuances. This allows for a comprehensive and detailed understanding of video content, aligning video with the text queries more precisely.

**Table 4: Model Components Ablation.** Phrase-level negative with fine-grained prompting balances coarse (mean recall) and fine-grained (PoSRank) understanding.

Training-strategy	Coarse-Grained ( $\uparrow$ )		Fine-Grained ( $\uparrow$ )				
	V2T	T2V	noun	adj	verb	adv	prep
Coarse-Grained Training	56.7	41.0	0.367	0.440	0.315	0.201	0.184
Word-Level Negatives	47.5	40.0	0.894	0.864	0.969	0.468	0.701
Phrase-Level Negatives	52.5	40.2	0.839	0.723	0.913	0.382	0.527
Fine-Grained Prompting	52.9	39.8	0.860	0.781	0.954	0.439	0.555

**Fig. 8: Visual Attention.** Coarse-grained training focuses on the most salient regions, while word-level negatives focuses on subtle details necessary for distinguishing fine-grained elements of the caption. Fine-grained prompting balances the two.

## 7 Experiments

We first describe the implementation details and then analyze the contribution of each method component and the impact of introduced hyperparameters. Finally, we demonstrate that our method can be plugged into various existing video-text retrieval models and is effective on multiple different datasets.

**Implementation Details.** We use the same datasets and methods as described in Sec. 5. For all models we use the same training settings as the original paper with the addition of our fine-grained training, see the supplementary material for further details. Unless otherwise specified, we use loss weight  $\lambda=0.2$ , number of hard negatives  $\mathbb{N}=16$ , and number of visual prompts  $\mathbb{T}=1$ . All experiments were performed on an NVIDIA RTX A6000 GPU.

**Evaluation Criteria.** For coarse-grained evaluation, we report average recall over R@1, R@5 and R@10. For fine-grained evaluation, we report the proposed *PosRank* for nouns, adjectives, verbs, adverbs and prepositions.

### 7.1 Ablation Study

To validate the effectiveness our proposed method, we first carry out several ablation studies using the X-CLIP model on the VLN-UVO dataset.

**Model Components.** Tab. 4 demonstrates the effect of each component of our proposed method. Initially, we train X-CLIP [34] as in the original paper, *i.e.* coarse-grained training. This performs well on the standard coarse-grained evaluation metrics but struggles on our fine-grained evaluation, particularly for adverbs and prepositions. When adding fine-grained training with word-level negatives the fine-grained evaluation improves for all parts-of-speech, but this comes at the cost of decreased video-to-text coarse-grained retrieval. This is because the focus is shifted from the most salient regions to fine-grained details (see Figure 8). Replacing the word-level negatives with phrase-level negatives better balances the coarse and fine-grained objectives, for instance, coarse-grained video-to-text retrieval (V2T) increases from 47.5 to 52.5. Using fine-grained prompting with phrase-level negatives further improves the fine-grained evaluation results as it allows the model to separately specialize in coarse- and fine-grained differences.

**Table 5: Loss weight  $\lambda$ .**  $\lambda=0.2$  is the best balance between fine-grained (PoS-Rank) and coarse-grained (mean recall).

$\lambda$	Coarse-Grained ( $\uparrow$ )		Fine-Grained ( $\uparrow$ )				
	V2T	T2V	noun	adj	verb	adv	prep
0.1	48.2	39.8	0.871	0.814	0.959	0.367	0.627
0.2	47.5	40.0	0.894	0.864	0.969	0.468	0.701
0.3	42.7	40.0	0.903	0.875	0.976	0.459	0.711
0.4	41.4	39.5	0.906	0.880	0.977	0.482	0.730
0.5	43.3	39.4	0.906	0.881	0.976	0.523	0.728

**Table 6: Num. of hard negatives  $\mathbb{N}$ .**  $\mathbb{N} = 16$  provides a reasonable balance between performance and training efficiency.

$\mathbb{N}$	Coarse-Grained ( $\uparrow$ )		Fine-Grained ( $\uparrow$ )				
	V2T	T2V	noun	adj	verb	adv	prep
2	53.3	39.9	0.778	0.692	0.882	0.282	0.475
4	48.6	39.8	0.875	0.808	0.957	0.344	0.655
8	48.7	39.8	0.876	0.834	0.962	0.420	0.656
16	47.5	40.0	0.894	0.864	0.969	0.468	0.701
32	44.0	40.0	0.901	0.869	0.975	0.449	0.701

**Loss weighting  $\lambda$ .** We evaluate the impact of  $\lambda$ , which balances coarse and fine-grained loss functions, in Tab. 5. As  $\lambda$  increases, the model focuses on the fine-grained loss, leading to improved fine-grained performance.  $\lambda = 0.2$  represents the best balance between fine-grained and coarse-grained performance.

**Number of hard negative sentences  $\mathbb{N}$ .** We explore the number of hard negative sentences  $\mathbb{N}$  in Tab. 6. As  $\mathbb{N}$  increases fine-grained retrieval improves but coarse-grained performance declines.  $\mathbb{N}=16$  represents a compromise between fine-grained and coarse-grained performance.

## 7.2 Effectiveness with Different Datasets.

We demonstrate our model’s effectiveness for different datasets with X-CLIP in Tab. 7. On all datasets, our approach provides much better fine-grained retrieval performance than standard coarse-grained training, for instance on VLN-OOPS we obtain 0.941 verb and 0.527 preposition *PoSRank* vs. 0.236 and 0.142 for coarse-grained training. Our approach is also more effective at coarse-grained metrics than the naive fine-grained training with word-level negatives, *e.g.* we obtain 67.8 V2T on VLN-OOPS compared to 62.9 for word-level negatives. However, there is still much potential for improvement in both coarse and fine-grained metrics, particularly for adverbs and prepositions in the fine-grained evaluation.

**Table 7: Effectiveness with Different Datasets.** Our approach recognizes fine-grained differences (PoSRank) on different datasets while maintaining coarse-grained ability (mean recall).

Dataset	Training-strategy	Coarse-Grained ( $\uparrow$ )		Fine-Grained ( $\uparrow$ )				
		V2T	T2V	noun	adj	verb	adv	prep
MSR-VTT [54]	Coarse-Grained Training	66.7	67.0	0.430	0.445	0.314	0.304	0.225
	Word-Level Negatives	64.5	67.0	0.854	0.818	0.871	0.675	0.817
	Phrase-Level Negatives	65.6	67.7	0.841	0.795	0.881	0.512	0.713
	Fine-Grained Prompting	66.2	67.4	0.887	0.897	0.861	0.723	0.846
VATEX [48]	Coarse-Grained Training	91.1	82.3	0.312	0.366	0.276	0.270	0.166
	Word-Level Negatives	85.3	78.2	0.875	0.876	0.911	0.811	0.720
	Phrase-Level Negatives	86.6	79.3	0.794	0.757	0.763	0.672	0.494
	Fine-Grained Prompting	90.3	81.5	0.839	0.774	0.870	0.738	0.633
VLN-OOPS [46]	Coarse-Grained Training	70.1	52.3	0.409	0.446	0.236	0.176	0.142
	Word-Level Negatives	62.9	51.1	0.877	0.854	0.965	0.570	0.675
	Phrase-Level Negatives	67.1	51.5	0.834	0.736	0.823	0.456	0.449
	Fine-Grained Prompting	67.8	53.4	0.844	0.783	0.941	0.540	0.527

**Table 8: Effectiveness with Different Models.** Our approach can be combined with existing models for a large improvement in distinguishing fine-grained differences (PoSRank) and maintains coarse-grained performance (mean recall).

Method	Training-strategy	Coarse-Grained ( $\uparrow$ )		Fine-Grained ( $\uparrow$ )				
		V2T	T2V	noun	adj	verb	adv	prep
Frozen [4]	Coarse-Grained Training	38.8	38.6	0.332	0.346	0.207	0.215	0.160
	Word-Level Negatives	37.0	37.2	0.718	0.642	0.888	0.411	0.574
	Phrase-Level Negatives	37.1	36.5	0.770	0.701	0.887	0.487	0.627
	Fine-Grained Prompting	39.9	39.5	0.786	0.733	0.927	0.490	0.631
UCoFiA [50]	Coarse-Grained Training	54.4	39.7	0.391	0.423	0.317	0.227	0.184
	Word-Level Negatives	43.9	39.0	0.894	0.859	0.973	0.493	0.696
	Phrase-Level Negatives	50.0	39.1	0.853	0.723	0.925	0.378	0.544
	Fine-Grained Prompting	52.5	38.9	0.840	0.828	0.946	0.467	0.695
TS2-Net [31]	Coarse-Grained Training	54.1	39.2	0.375	0.453	0.316	0.220	0.184
	Word-Level Negatives	47.4	39.1	0.885	0.852	0.968	0.461	0.665
	Phrase-Level Negatives	46.6	39.2	0.889	0.856	0.968	0.476	0.660
	Fine-Grained Prompting	54.6	39.7	0.871	0.805	0.966	0.475	0.539
X-CLIP [34]	Coarse-Grained Training	56.7	41.0	0.367	0.440	0.315	0.201	0.184
	Word-Level Negatives	47.5	40.0	0.894	0.864	0.969	0.468	0.701
	Phrase-Level Negatives	52.5	40.2	0.839	0.723	0.913	0.382	0.527
	Fine-Grained Prompting	52.9	39.8	0.860	0.781	0.954	0.439	0.555

### 7.3 Effectiveness with Different Models

A benefit of our approach is that it can be easily combined with existing models to enable fine-grained video-text retrieval. We demonstrate this in Tab. 8 with results on VLN-UVO for Frozen [4], X-CLIP [34], and UCoFiA [50]. Our model is easy to plug into each of these methods achieving much better fine-grained performance than coarse-grained training while also maintaining or beating the original coarse-grained performance. For instance, our method with TS2-Net obtains 54.6 V2T and 0.966 verb *PoSRank* vs. 54.1 and 0.316 from coarse-grained training. TS2-Net and UCoFiA are the best-performing models on the fine-grained evaluation, however, there is no clear winner with TS2-Net performing best on nouns, verbs, and adverbs and UCoFiA performing best on adjectives and prepositions. While our approach successfully balances performance on coarse and fine-grained metrics there is still much potential for improvement in fine-grained understanding, particularly for adverbs and prepositions.

## 8 Conclusion

In this paper, we conducted an in-depth analysis of current video retrieval evaluations finding that current benchmarks fall short in detecting a model’s ability to perceive subtle single concept differences. We introduce a fine-grained approach to analyze models’ sensitivity to individual word variations across different parts-of-speech. Testing this metric on current methods reveals that these models are not sensitive to such variations, especially for adverbs and prepositions. To bridge this gap, we introduce a simple, but effective baseline which can be easily combined with existing models. While experiments show that the proposed method effectively enhances the models’ ability to understand fine-grained differences, there is much potential for future work.

**Acknowledgements:** We thank Yunhua Zhang for the helpful discussions, and Dennis Koelma for his invaluable technical support.

## References

1. Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., et al.: Flamingo: a visual language model for few-shot learning. In: *NeurIPS* (2022)
2. Anne Hendricks, L., Wang, O., Shechtman, E., Sivic, J., Darrell, T., Russell, B.: Localizing moments in video with natural language. In: *ICCV* (2017)
3. Bagad, P., Tapaswi, M., Snoek, C.G.: Test of time: Instilling video-language models with a sense of time. In: *CVPR* (2023)
4. Bain, M., Nagrani, A., Varol, G., Zisserman, A.: Frozen in time: A joint video and image encoder for end-to-end retrieval. In: *ICCV* (2021)
5. Bain, M., Nagrani, A., Varol, G., Zisserman, A.: A clip-hitchhiker’s guide to long video retrieval. *arXiv preprint arXiv:2205.08508* (2022)
6. Buch, S., Eyzaguirre, C., Gaidon, A., Wu, J., Fei-Fei, L., Niebles, J.C.: Revisiting the " video" in video-language understanding. In: *CVPR* (2022)
7. Chan, D.M., Myers, A., Vijayanarasimhan, S., Ross, D.A., Seybold, B., Canny, J.F.: What’s in a caption? dataset-specific linguistic diversity and its effect on visual description models and metrics. In: *CVPR* (2022)
8. Chen, D., Dolan, W.: Collecting highly parallel data for paraphrase evaluation. In: *CVPR* (2011)
9. Chen, S., Zhao, Y., Jin, Q., Wu, Q.: Fine-grained video-text retrieval with hierarchical graph reasoning. In: *CVPR* (2020)
10. Chen, S., Zhao, Y., Jin, Q., Wu, Q.: Fine-grained video-text retrieval with hierarchical graph reasoning. In: *CVPR* (2020)
11. Chen, Y.C., Li, L., Yu, L., El Kholy, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: *ECCV* (2020)
12. Cheng, Y., Wei, F., Bao, J., Chen, D., Zhang, W.: Cico: Domain-aware sign language retrieval via cross-lingual contrastive learning. In: *CVPR* (2023)
13. Croitoru, I., Bogolin, S.V., Leordeanu, M., Jin, H., Zisserman, A., Albanie, S., Liu, Y.: Teactext: Crossmodal generalized distillation for text-video retrieval. In: *ICCV* (2021)
14. Dong, J., Li, X., Xu, C., Yang, X., Yang, G., Wang, X.: Dual encoding for video retrieval by text. *TPAMI* (2021)
15. Doughty, H., Laptev, I., Mayol-Cuevas, W., Damen, D.: Action modifiers: Learning from adverbs in instructional videos. In: *CVPR* (2020)
16. Doughty, H., Snoek, C.G.M.: How Do You Do It? Fine-Grained Action Understanding with Pseudo-Adverbs. In: *CVPR* (2022)
17. Faghri, F., Fleet, D.J., Kiros, J.R., Fidler, S.: VSE++: improving visual-semantic embeddings with hard negatives. In: *BMVC* (2018)
18. Fan, Z., Wei, Z., Li, Z., Wang, S., Fan, J.: Negative sample is negative in its own way: Tailoring negative sentences for image-text retrieval. In: *NAACL* (2022)
19. Fang, H., Xiong, P., Xu, L., Chen, Y.: Clip2video: Mastering video-text retrieval via image clip. *arXiv preprint arXiv:2106.11097* (2021)
20. Gou, Y., Ko, T., Yang, H., Kwok, J., Zhang, Y., Wang, M.: Leveraging per image-token consistency for vision-language pre-training. In: *CVPR* (2023)
21. Hendricks, L.A., Nematzadeh, A.: Probing image-language transformers for verb understanding. In: *ACL* (2021)
22. Hu, F., Chen, A., Wang, Z., Zhou, F., Dong, J., Li, X.: Lightweight attentional feature fusion: A new baseline for text-to-video retrieval. In: *ECCV* (2022)

23. Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: ECCV (2022)
24. Jiang, H., Zhang, J., Huang, R., Ge, C., Ni, Z., Lu, J., Zhou, J., Song, S., Huang, G.: Cross-modal adapter for text-video retrieval. arXiv preprint arXiv:2211.09623 (2022)
25. Kalantidis, Y., Sariyildiz, M.B., Pion, N., Weinzaepfel, P., Larlus, D.: Hard negative mixing for contrastive learning. In: NeurIPS (2020)
26. Krishna, R., Hata, K., Ren, F., Fei-Fei, L., Carlos Niebles, J.: Dense-captioning events in videos. In: ICCV (2017)
27. Lee, K.H., Chen, X., Hua, G., Hu, H., He, X.: Stacked cross attention for image-text matching. In: ECCV (2018)
28. Lei, J., Li, L., Zhou, L., Gan, Z., Berg, T.L., Bansal, M., Liu, J.: Less is more: Clipbert for video-and-language learning via sparse sampling. In: CVPR (2021)
29. Li, X., Xu, C., Yang, G., Chen, Z., Dong, J.: W2VV++: Fully deep learning for ad-hoc video search. In: ACM MM (2019)
30. Li, Y., Song, Y., Cao, L., Tetreault, J.R., Goldberg, L., Jaimes, A., Luo, J.: TGIF: A new dataset and benchmark on animated gif description. In: CVPR (2015)
31. Liu, Y., Xiong, P., Xu, L., Cao, S., Jin, Q.: Ts2-net: Token shift and selection transformer for text-video retrieval. In: ECCV (2022)
32. Lu, H., Huo, Y., Yang, G., Lu, Z., Zhan, W., Tomizuka, M., Ding, M.: Uniadapter: Unified parameter-efficient transfer learning for cross-modal modeling. In: ICLR (2024)
33. Luo, H., Ji, L., Zhong, M., Chen, Y., Lei, W., Duan, N., Li, T.: Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neuro-computing* **508**, 293–304 (2022)
34. Ma, Y., Xu, G., Sun, X., Yan, M., Zhang, J., Ji, R.: X-clip: End-to-end multi-grained contrastive learning for video-text retrieval. In: ACM MM (2022)
35. Madasu, A., Lal, V.: ICSV: Investigating compositional and syntactic understanding in video retrieval models. In: CVPR (2024)
36. Messina, N., Amato, G., Esuli, A., Falchi, F., Gennaro, C., Marchand-Maillet, S.: Fine-grained visual textual alignment for cross-modal retrieval using transformer encoders. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* **17**(4), 1–23 (2021)
37. Momeni, L., Caron, M., Nagrani, A., Zisserman, A., Schmid, C.: Verbs in action: Improving verb understanding in video-language models. In: ICCV (2023)
38. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: ICML (2021)
39. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
40. Robinson, J., Chuang, C.Y., Sra, S., Jegelka, S.: Contrastive learning with hard negative samples. In: ICLR (2020)
41. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In: ACL (2018)
42. Shekhar, R., Pezzelle, S., Klimovich, Y., Herbelot, A., Nabi, M., Sangineto, E., Bernardi, R.: Foil it! find one mismatch between image and language caption. In: ACL (2017)
43. Shu, F., Chen, B., Liao, Y., Wang, J., Liu, S.: Mac: Masked contrastive pre-training for efficient video-text retrieval. *IEEE Transactions on Multimedia* (2024)



44. Tian, K., Zhao, R., Xin, Z., Lan, B., Li, X.: Holistic features are almost sufficient for text-to-video retrieval. In: CVPR (2024)
45. Torabi, A., Tandon, N., Sigal, L.: Learning language-visual embedding for movie understanding with natural-language. arXiv preprint arXiv:1609.08124 (2016)
46. Voigtlaender, P., Changpinyo, S., Pont-Tuset, J., Soricut, R., Ferrari, V.: Connecting Vision and Language with Video Localized Narratives. In: CVPR (2023)
47. Wang, X., Zhu, L., Yang, Y.: T2vlad: global-local sequence alignment for text-video retrieval. In: CVPR (2021)
48. Wang, X., Wu, J., Chen, J., Li, L., Wang, Y.F., Wang, W.Y.: VATEX: A large-scale, high-quality multilingual dataset for video-and-language research. In: ICCV (2019)
49. Wang, Y., Shi, P.: Video-text retrieval by supervised multi-space multi-grained alignment. CoRR (2023)
50. Wang, Z., Sung, Y.L., Cheng, F., Bertasius, G., Bansal, M.: Unified coarse-to-fine alignment for video-text retrieval. In: ICCV (2023)
51. Wang, Z., Chen, A., Hu, F., Li, X.: Learn to understand negation in video retrieval. In: ACM MM (2022)
52. Wray, M., Doughty, H., Damen, D.: On semantic similarity in video retrieval. In: CVPR (2021)
53. Wray, M., Larlus, D., Csorika, G., Damen, D.: Fine-grained action retrieval through multiple parts-of-speech embeddings. In: ICCV (2019)
54. Xu, J., Mei, T., Yao, T., Rui, Y.: MSR-VTT: A large video description dataset for bridging video and language. In: CVPR (2016)
55. Yang, J., Bisk, Y., Gao, J.: Taco: Token-aware cascade contrastive learning for video-text alignment. In: ICCV (2021)
56. Yao, L., Huang, R., Hou, L., Lu, G., Niu, M., Xu, H., Liang, X., Li, Z., Jiang, X., Xu, C.: Filip: Fine-grained interactive language-image pre-training. In: ICLR (2022)
57. Yu, Y., Kim, J., Kim, G.: A joint sequence fusion model for video question answering and retrieval. In: ECCV (2018)
58. Yu, Y., Ko, H., Choi, J., Kim, G.: Video captioning and retrieval models with semantic attention. arXiv preprint arXiv:1610.02947 (2016)
59. Yuksekgonul, M., Bianchi, F., Kalluri, P., Jurafsky, D., Zou, J.: When and why vision-language models behave like bags-of-words, and what to do about it? In: ICLR (2023)
60. Zellers, R., Lu, J., Lu, X., Yu, Y., Zhao, Y., Salehi, M., Kusupati, A., Hessel, J., Farhadi, A., Choi, Y.: Merlot reserve: Neural script knowledge through vision and language and sound. In: CVPR (2022)
61. Zhang, B., Zhang, P., Dong, X., Zang, Y., Wang, J.: Long-clip: Unlocking the long-text capability of clip. In: ECCV (2024)
62. Zhao, T., Zhang, T., Zhu, M., Shen, H., Lee, K., Lu, X., Yin, J.: Vl-checklist: Evaluating pre-trained vision-language models with objects, attributes and relations. arXiv preprint arXiv:2207.00221 (2022)
63. Zhou, L., Xu, C., Corso, J.J.: Procnets: Learning to segment procedures in untrimmed and unconstrained videos. In: AAAI (2018)
64. Zou, X., Wu, C., Cheng, L., Wang, Z.: Tokenflow: Rethinking fine-grained cross-modal alignment in vision-language retrieval. arXiv preprint arXiv:2209.13822 (2022)