

# Open-Vocabulary Video Relation Extraction

Wentao Tian<sup>1</sup>, Zheng Wang<sup>2</sup> <sup>†</sup>, Yuqian Fu<sup>1</sup>, Jingjing Chen<sup>1</sup> <sup>†</sup>, Lechao Cheng<sup>3</sup>

<sup>1</sup> Shanghai Key Lab of Intell. Info. Processing, School of CS, Fudan University

<sup>2</sup> College of Computer Science and Technology, Zhejiang University of Technology

<sup>3</sup> Zhejiang Lab

{wttian22@m., fuyq20@, chenjingjing@fudan.edu.cn, zhengwang@zjut.edu.cn, chenlc@zhejianglab.com}

## Abstract

A comprehensive understanding of videos is inseparable from describing the action with its contextual action-object interactions. However, many current video understanding tasks prioritize general action classification and overlook the actors and relationships that shape the nature of the action, resulting in a superficial understanding of the action. Motivated by this, we introduce Open-vocabulary Video Relation Extraction (OVRE), a novel task that views action understanding through the lens of action-centric relation triplets. OVRE focuses on pairwise relations that take part in the action and describes these relation triplets with natural languages. Moreover, we curate the Moments-OVRE dataset, which comprises 180K videos with action-centric relation triplets, sourced from a multi-label action classification dataset. With Moments-OVRE, we further propose a cross-modal mapping model to generate relation triplets as a sequence. Finally, we benchmark existing cross-modal generation models on the new task of OVRE. Our code and dataset are available at <https://github.com/Iriya99/OVRE>.

## Introduction

Videos contain abundant semantic information, including action, actors (e.g. humans, animals, objects, and other entities), and relationships between actors. To comprehend the dynamic and complex real-world situations depicted in videos, researchers have investigated a wide range of video comprehension tasks. These endeavors allow for the transition of understanding video content from broad semantic concepts to more detailed ones. Despite the variations, all of these tasks converge on a pivotal aspect : extracting semantic information within the videos and constructing a higher-level representation to facilitate comprehension.

Foundational tasks, such as action classification (Kong and Fu 2022) and temporal action localization (Xia and Zhan 2020), primarily center on recognizing broad-level actions within videos, yet they often overlook the specific scenario in which these actions unfold. Consequently, these tasks often struggle to offer a profound understanding of the context and the specific actors that are part of actions. In essence, they concentrate solely on deciphering “what”

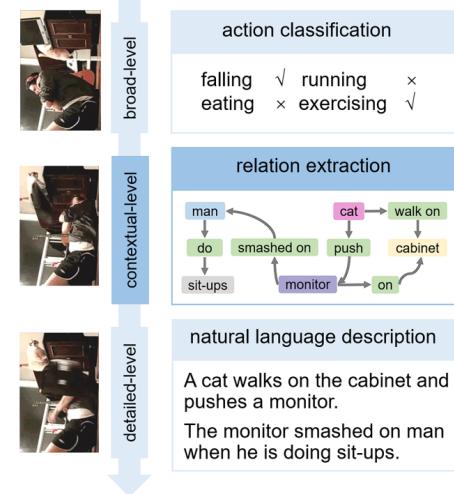


Figure 1: Open-vocabulary Video Relation Extraction enables a contextual-level comprehension of video content, bridging the gap between general action classification and precise language description.

action is transpiring and “when” it takes place, omitting the “who” and “how” aspects. On the other hand, video captioning (Chen, Yao, and Jiang 2019), video grounding (Chen and Jiang 2019; Wang, Chen, and Jiang 2021), and video-text retrieval (Song, Chen, and Jiang 2023) strive to encapsulate the videos’ essence through textual descriptions by mapping them into a joint semantic space. Nevertheless, these textual descriptions often provide a detailed-level overview of the action context, lacking a nuanced comprehension of relations.

To achieve a comprehensive comprehension of visual content, researchers have introduced Video Visual Relation Detection (VidVRD) tasks (Shang et al. 2017, 2019) specifically designed for video analysis. VidVRD tasks are geared toward identifying objects and the relations between them within videos. For instance, VidVRD (Shang et al. 2017) employs available video object detection data to non-exclusively label various simplistic relationships between objects. As depicted in Figure 2 (a), relationships like *touch* and *watch* are inherently *object-centric* relation-

<sup>†</sup>Correspondence to: Zheng Wang, Jingjing Chen.  
Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

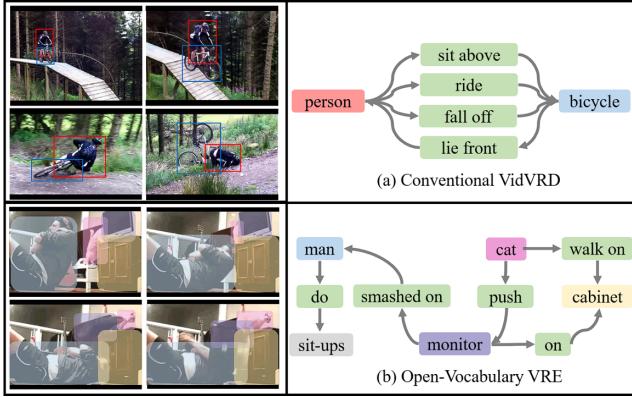


Figure 2: Two falling videos depicted in VidVRD and OVRE diverse in (a) salient objects interconnected throughout frames are exhaustively annotated with diverse relations in VidVRD; (b) OVRE builds a relation graph with various actors and relations that closely related to the falling action.

ships with limited implications for the falling action. Notably, Action Genome (Ji et al. 2019) is distinct in its dedication to an *action-centric* understanding of videos. It dissects actions into spatio-temporal scene graphs. However, it still grapples with a challenge shared by VidVRD – the constraint of a fixed vocabulary range. Consider the scenario shown in Figure 2 (a) involving a bicyclist falling from a steep bridge. Due to the vocabulary limitations, the description fails to encompass the bridge and its relationship with the biker, which is an essential element in understanding the falling action.

To mitigate the aforementioned limitations of VidVRD, we introduce **Open-vocabulary Video Relation Extraction (OVRE)**. OVRE extracts all *action-centric* relation triplets from videos with language descriptions. Take Figure 2 (b) as an example: OVRE dissects the falling action into multiple relation components, such as  $\langle \text{cat}, \text{push}, \text{monitor} \rangle$  and  $\langle \text{monitor}, \text{smashed on}, \text{man} \rangle$ . Departing from VidVRD (Shang et al. 2017) that confined objects and relations to limited categories, we harness the immense potential of large Language Models (LLMs) to articulate action-aware relations using natural language. In essence, we undertake a two-fold process: first, using CLIP visual encoder to encapsulate video semantics; second, mapping these semantics into the linguistic realm for generating relation triplets through a pre-trained LLM. Employing this straightforward baseline, we achieve the ability to articulate relation triplets with unconstrained vocabularies.

To study OVRE, we present Moments-OVRE, an extensive dataset encompassing over 180K diverse videos. These videos are sourced from a subset of Multi-Moments in Time (Monfort et al. 2021)(M-MiT), a multi-label video dataset with an average video duration of three seconds. In Moments-OVRE, we annotate both actors and relations as relation triplets that are most relevant to the action labels without vocabulary limitations. Note that “actor” does not refer only to human individuals but can also refer to objects, animals, or any entities that are actively involved in an ac-

tion. Moments-OVRE confers several advantages: (I) Unrestricted Vocabulary Annotations: Moments-OVRE benefits from an expansive vocabulary encompassing diverse actors and relations, offering a more accurate portrayal of real-world scenarios. (II) Emphasis on Action-Centric Annotations: Our annotations exclusively focus on relation triplets related to actions/events depicted in the video. (III) Ample Dataset Scale: Moments-OVRE contains over 180K videos, establishing itself as the most extensive video relation extraction dataset known to date.

The main contributions are as follows:

- We introduce the novel Open-vocabulary Video Relation Extraction task, aiming to extract action-centric relations within videos.
- To foster ongoing research, we curate Moments-OVRE, an expansive dataset for video action comprehension. Comprising over 180k videos, Moments-OVRE is enriched with comprehensive open-vocabulary relationship annotations.
- We build a simple pipeline designed for generating action-centric relation triplets from raw videos. Extensive experimentation is conducted to validate assorted design considerations employed within the pipeline.

## Related Works

**Video understanding** has long been an active research area in computer vision, encompassing a wide range of diverse aspects. Action recognition (Kay et al. 2017; Goyal et al. 2017; Monfort et al. 2019) categorizes the video content into a predefined set of action classes, temporal action localization (Heilbron et al. 2015; Liu et al. 2022) further emphasizes the specific time points at which these actions occur. Tasks such as video captioning (Chen, Yao, and Jiang 2019), video question answering (Qian et al. 2023), and video-text retrieval (Song et al. 2021) are dedicated to learning a detailed representation of semantic information. In the spectrum of these tasks resides an intermediary domain referred to as relation understanding, dedicated to capturing the diverse relations among objects. This contextual domain provides semantically rich intermediate representations for videos, crucial for comprehending interactions and behaviors depicted within them. VidVRD (Shang et al. 2017) and VidOR (Shang et al. 2019) are representative benchmarks for relation understanding, with finite object and predicate sets. Recently proposed task Open-VidVRD (Gao et al. 2023) manually splits the finite object and predicate sets into base and novel categories and builds the novel prediction setting. However, these existing VRD tasks focus on object-centric relation detection and struggle to bridge the gap between closed-label sets and natural languages, prompting us to propose video relationship extraction under open-vocabulary scenarios. Moreover, we aim to identify informative triplets that portray various actions in the video.

**Contextual-level Visual Comprehension** has been mainly explored in the image domain, including scene graph generation (Ji et al. 2019), visual semantic role labeling (Yatskar, Zettlemoyer, and Farhadi 2016; Sadhu et al. 2021), and human-object interaction (Li et al. 2022; Gkioxari et al. 2018). These tasks have also been extended to the video domain, which requires information aggregation across

Dataset	Task	Annotation	#Subject/Object	#Predicate	#Video
Kinetics400	Action Classification	Action Categories	-	-	306,245
VideoLT			-	-	256,218
MSR-VTT	Captioning, Retrieval	Textual Descriptions	-	-	7,180
VATEX			-	-	41,269
VidVRD	Video Relation Detection	BBoxes, Relation Triplets	35	132	1,000
Action Genome			35	25	10,000
VidSitu	Semantic Role Labeling	Verbs, SRLs, Event Relations	Open	Open	29,220
<b>Moments-OVRE</b>	<b>Video Relation Extraction</b>	<b>Relation Triplets</b>	<b>Open</b>	<b>Open</b>	<b>186,943</b>

Table 1: Representative video understanding datasets.

frames. Existing works often treat video relation detection as a multi-stage pipeline, including object tracking, tracklet pair generation, and relation classification. These works focus on improving relation classification by leveraging contextual knowledge (Qian et al. 2019) or adding additional modules to enhance predicate representation (Gao et al. 2021) with generated tracklets. Recent trends lean towards one-stage models like VRDFormer (Zheng, Chen, and Jin 2022), employing various queries to integrate spatio-temporal information, facilitating tracklet pair generation and relation classification concurrently. However, their scalability and generalization are hampered by datasets containing a restricted size of objects and relationships. Video captioning, aimed at generating open-ended descriptions, typically uses an encoder-decoder architecture. Over time, this setup has advanced to more sophisticated models (Zhang et al. 2019; Tang et al. 2021; Lin et al. 2022). Recent strides in visual-language representation learning have led to models pre-trained on extensive paired data demonstrating promise in fine-tuning for various multi-modal tasks (Wang et al. 2022; Xu et al. 2023; Chen et al. 2023). Building on the progress in video captioning, we seek to leverage generative models for our OVRE task.

**Datasets for video understanding** have been introduced for multiple tasks over the past years. Kinetics400 (Kay et al. 2017) is a widely adopted category-balanced classification dataset. VideoLT (Zhang et al. 2021), in opposition, follows a naturally long-tailed distribution for action recognition. MSR-VTT (Xu et al. 2016) is a widely adopted dataset for video captioning and retrieval that includes video clips from different domains, each annotated with approximately 20 natural sentences. Recently introduced VATEX (Wang et al. 2019) is a large-scale multilingual video description dataset. VidVRD (Shang et al. 2017) and VidOR (Shang et al. 2019) densely annotate objects and most basic relations. Action Genome (Ji et al. 2019) focuses on action understanding by decomposing them into different relationships and ignoring other relations irrelevant to actions. VidVRD datasets have restricted their objects and relations to a fixed label set and thus are far away from reflecting diverse real-world relations. VidSRL (Sadhu et al. 2021) presents the sole open-vocabulary relation understanding task, while primarily focusing on events. In contrast, OVRE places greater emphasis

on actions, which is a fundamental ability for event comprehension. Detailed statistics are listed in Table 1.

## Video Relation Extraction Benchmark

**Task Formulation** Given a raw video input  $V$ , OVRE targets inferring its corresponding visual relationships  $R$ . Here,  $V \in \mathbb{R}^{T \times H \times W \times C}$  denotes its frame sequence, where  $T$ ,  $H$ ,  $W$ , and  $C$  represent the number of frames, width, height, and channels respectively. The relation triplets are represented as  $R = \{r_1, \dots, r_K\}$ , and each relation  $r_K$  is structured as a triplet:  $\langle \text{subject}, \text{relation}, \text{object} \rangle$ .

Given that actions depicted in videos can encompass a wide spectrum of interactions, and distinct interactions might be common across various actions, our objective is to systematically extract all relationships associated with actions within the video context. Each  $r_i$  consists of a sequence of words  $w = \{w_1, w_2, \dots, w_N\}$ . These individual  $r_i$  are then concatenated to form a token sequence  $R = w_1, w_2, \dots, w_M$ , which serves as a representation of a sequence of relation triplets, collectively portraying all the action-centric relationships. The training objective is to maximize the likelihood of generating  $R$  given the video  $V$ :

$$\max_{\theta} \log p_{\theta}(R|V),$$

where  $\theta$  denotes the model’s trainable parameters. Our key idea is to leverage the powerful capability of generative models to address the OVRE problem. Consequently, the revised training objective can be formulated as follows:

$$\max_{\theta} \sum_{i=1}^M \log p_{\theta}(w_i|V, w_{i-1}).$$

OVRE involves generating relationships using natural language, rendering common metrics wildly used VidVRD metrics such as Precision@K and Recall@K unsuitable for our task. Inspired by the evaluation metrics in image captioning, we advocate for assessing the quality of generated triplets using Bleu, CIDEr, and METEOR scores. For the Bleu metric, we only consider B@1, B@2, and B@3, as each triplet contains at least three words. These metrics face challenges when applied directly to evaluating the overall relation triplets since the generated triplet sequence can be

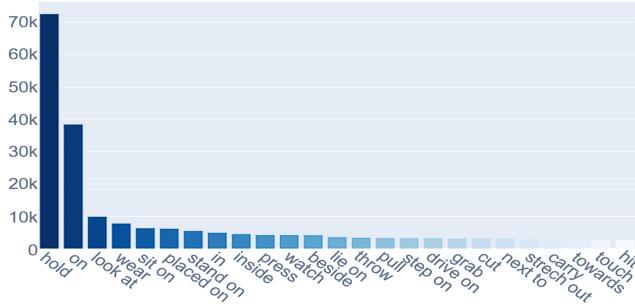


Figure 3: Counts of top 25 relationships in Moments-OVRE.

regarded as an unordered set, whose alignment with the ground truth triplets is unknown.

To match the generated triplet set with the ground truth, we first utilize a sentence embedding model SimCSE (Gao, Yao, and Chen 2021) to quantify the text similarity between the generated and ground truth triplet set. Then we use these embeddings to acquire a similarity matrix  $S$ , where  $S_{ij}$  denotes the cosine similarity score of generated triplet  $\hat{r}_i$  and ground truth triplet  $r_j$ . After obtaining  $S$ , we use the Hungarian algorithm to establish one-to-one matches between the triplets and apply the aforementioned metrics to these paired triplets. Note that the number of generated triplets may differ from that of ground truth triplets. When fewer triplets are generated, unmatched ground truth triplets receive zero scores.

## Dataset Construction

We present the Moments-OVRE dataset tailored for OVRE. Notably, Moments-OVRE boasts distinct attributes such as diverse video content, an extensive compilation of video-triplet pairs, annotations embracing open vocabulary, and a focus on action-centric relationships. In this section, we will detail how we select representative videos, annotate relations, and devise data splits. Additionally, we thoroughly explore and analyze the statistical insights within the Moments-OVRE dataset.

**Video Selection** Since our task focuses on action-centric relationships, several requirements are posed for the selection of videos. First, we prefer videos with multiple actions. Second, videos should contain rich information, involving different scenes, objects, and events in the real world. We choose to annotate videos from Multi Moments in Time (M-MiT) (Monfort et al. 2021) due to its multi-label nature, which allows a more nuanced comprehension of actions within intricate contexts. M-MiT offers several distinct advantages: (I) Unlike commonly utilized action datasets like Kinetics, M-MiT exhibits substantial intra-class diversity. This highlights the necessity of providing more intricate descriptions of relationships to capture variations between videos of the same action labels. (II) A majority of M-MiT videos have a duration of only 3 seconds. This temporal brevity encourages annotators to focus primarily on portraying the action itself, as opposed to the temporal context which is more prevalent in longer videos, such as those in

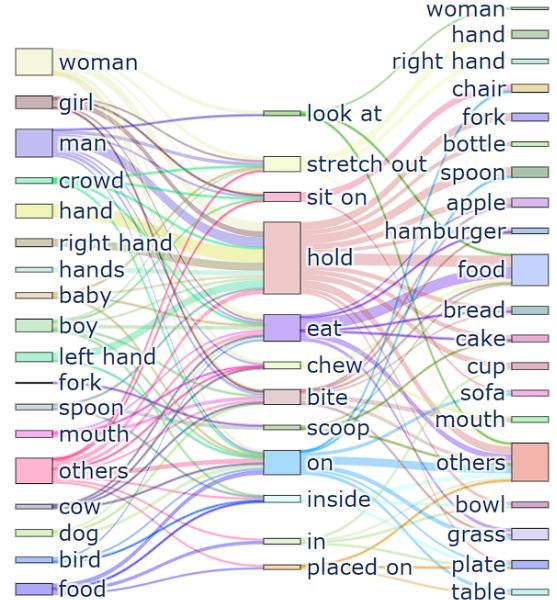


Figure 4: A weighted bipartite mapping of the top 12 most frequent relations in eating videos.

YouCook and Howto100M datasets. Considering the long-tail distribution of action categories in M-MiT, we attempt to relatively balanced sampling videos from all classes. M-MiT has 292 action categories, we sample at least 660 videos per class and additionally select other random videos as some videos may be discarded during annotation.

**Annotation Pipeline** We ask annotators to perform an action-centric relation triplet annotation, which follows these steps: Given the potentially noisy nature of the large-scale M-MiT dataset, annotators first verify the correctness of the action labels. Videos with incorrect labeling are subsequently excluded. To perform action-centric annotation, both the video and the corresponding action labels are presented to the annotators. They first identify and annotate all pertinent objects, and then articulate the associated relationships among these objects. Annotators are instructed to provide descriptions for relation triplets solely when these relationships are relevant to the action labels. To illustrate, refer to Figure 1 (b): the annotation `<monitor, smashed on, man>` is required, whereas the annotation `<guitar, placed on, ground>` is invalid, as the latter constitutes background information rather than being directly tied to the falling action. Besides, we also manually review and correct the low-quality annotations.

**Dataset splits** The data is partitioned into training and testing sets, resulting in 178,480 and 8,463 videos respectively. To ensure maximum consistency with the original splitting, all videos are selected exclusively from their designated sets within the M-MiT collection.

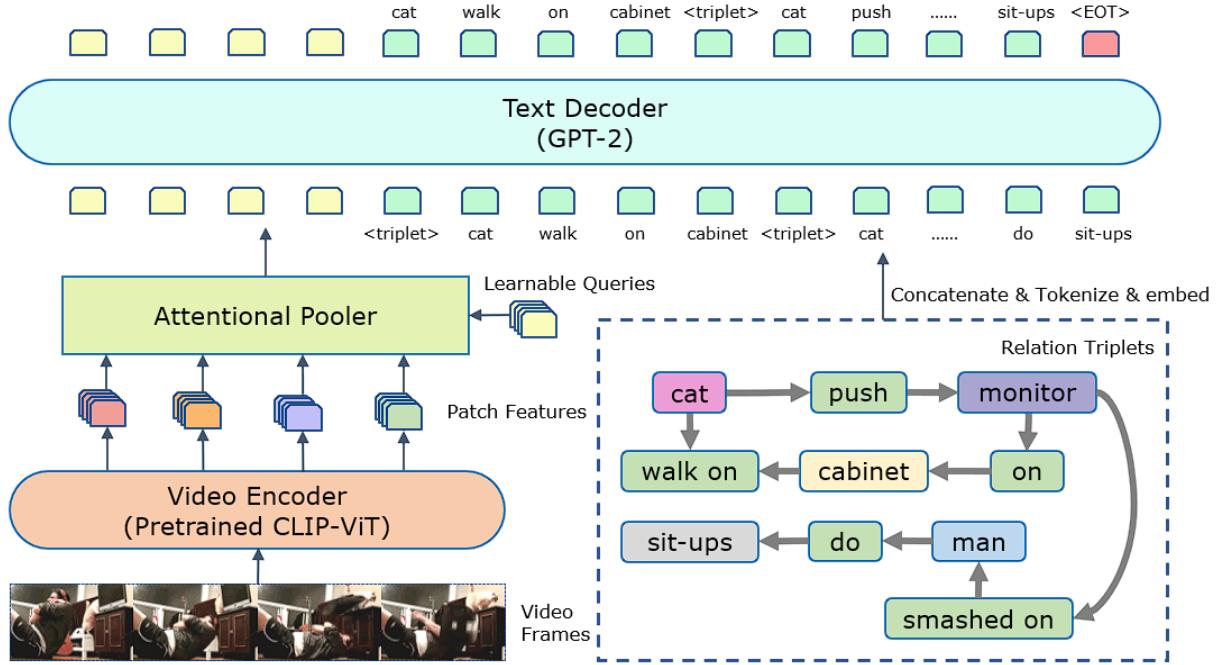


Figure 5: Overview of our model architecture, enabling simple relation generation with the powerful vision-language pre-trained model CLIP-ViT as the video encoder and the large language model GPT-2 as the text decoder.

## Dataset Analysis and Statistics

Overall, Moments-OVRE offers open-vocabulary annotations for 186,943 videos, encompassing a sum of 399,576 relation triplets. As shown in Figure 3, the distribution of the top 25 most frequent relations follows a natural long-tailed distribution. Furthermore, Figure 4 shows frequently described relation triplets of eating videos. Note that *hold* is a prevalent visual relation shared by various kinds of videos, such as painting, and drinking videos. A clue from this observation is that merely relying on a single relation is insufficient to infer the action categories from the videos, and a comprehensive understanding requires recognition of combinations of diverse relations.

## Method

**Overview** The overall framework is illustrated in Figure 5, which mainly includes a Video Encoder, an Attentional Pooler, and a Text Decoder. The video encoder transforms the video into a feature sequence, which is further condensed into prefix information by the attentional pooler and is then fed into the text decoder for relation triplet generation.

**Video Encoder** To extract visual features from a video, we leverage the visual encoder of a pre-trained CLIP (Radford et al. 2021) model. Previous research has demonstrated its exceptional performance across various open-vocabulary recognition tasks (Ni et al. 2022; Tang et al. 2021; Luo et al. 2022). We choose CLIP-ViT (Radford et al. 2021) to extract visual features, which utilizes ViT (Dosovitskiy et al. 2020) to translate video frames into sequences of patches. These visual patches capture finer-grained information, making

them an optimal fit for OVRE. The video encoder translates a video  $V \in \mathbb{R}^{T \times H \times W}$  to  $T \times \frac{HW}{P^2}$  patches.

**Prefix Strategies** With the features from  $V$ , our focus shifts towards transmitting them into the textual domain for relation generation. A simple solution is to directly feed all patches into the decoder, while the excessive number of patches not only leads to additional computational costs but also introduces a lot of redundant information. In light of this, we utilize an attentional pooler denoted as  $F(\cdot)$ , which employs a predetermined number of queries to extract meaningful features from all video patches:  $q_1, \dots, q_m = F(p_1, \dots, p_n, q_1, \dots, q_m)$ , which aggregates spatial-temporal features into a more concise representation.

**Attentional Pooler** Drawing inspiration from VideoCoCa (Yan et al. 2023), we incorporate the Attentional Pooler, which takes learnable queries and patch features as its input, facilitating cross-attention between them. The optimization of both the parameters of the Attentional Pooler and the queries enables the queries to gradually refine their ability to extract significant relationships from the patches. This design proves well-suited to the demands of OVRE. Remarkably, even with a single-layer transformer, our model can generate a diverse set of meaningful visual relationships.

**Text Decoder** We employ GPT-2 (Radford et al. 2019) for text generation. Accordingly, all relation triplets of a video are simply concatenated into a text sequence by a separation token. Then, we map all the words into their corresponding tokens and pad them to the maximum length  $M$  to obtain a sequence of embeddings.

Method	B@1	B@2	B@3	CIDEr	METEOR
ClipCap	29.75	16.32	9.48	125.45	19.25
GIT	35.19	20.12	11.90	155.38	23.06
<b>Ours</b>	<b>37.27</b>	<b>21.92</b>	<b>13.90</b>	<b>174.47</b>	<b>25.07</b>

Table 2: Baseline comparison on Moments-OVRE.

CLIP	GPT2	CIDEr	METEOR
✗	✗	131.85	19.84
✗	✓	165.67	24.12
✓	✓	<b>174.47</b>	<b>25.07</b>

Table 3: Study on vision and language model fine-tuning. The ✓ mark means fine-tuning the corresponding module.

## Experiments

### Experiments Settings

**Video and Text Preprocessing.** For each input video, we first resize it to  $224 \times 224$  and then extract 16 frames using uniform sampling, resulting in 786 patches for each video. We utilize RandAugment (Cubuk et al. 2019) as our data augmentation strategy. For the paired relation triplets, we concatenate the unordered relation triplets into a sequence using the special token `<triplet>`.

**Training Setting.** We train the generation model using cross-entropy loss and employ teacher forcing to accelerate the training process. All models are optimized using AdamW optimizer, with  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , a batch size of 16, and weight decay of  $1e-3$ . The initial learning rate is set to  $1e-6$  for CLIP,  $2e-5$  for GPT-2, and  $1e-3$  for Attention-Pooler. We applied learning rate warm-up during the early 5% training steps followed by cosine decay. We trained the networks for 50 epochs on 8 Nvidia V100 GPUs and chose the model with the highest CIDEr score as the final model.

### Main Results

**Baselines models.** Previous VidVRD methods focused on predicting relationships over detected objects, which is essentially a classification task and thus cannot be applied to OVRE. Therefore, we introduce several generative models as baseline models.

- ClipCap (Mokady, Hertz, and Bermano 2021) is an image captioning model that utilizes the same visual encoder and text decoder as we do. It uses a mapping network to convert CLIP embeddings into GPT-2 prefixes. To apply ClipCap for videos, we follow the most common strategy that treats each video frame as an individual image and then performs a mean pooling layer along the temporal dimension to obtain a global video representation.

- GIT (Wang et al. 2022) stands as a vision-language generative model, demonstrating strong performance across numerous generation tasks. This achievement is attributed to its effective optimization of the language model loss during

Features	CIDEr	METEOR
Region	77.38	13.14
Frame	115.85	18.11
Patch	<b>165.67</b>	<b>24.12</b>

Table 4: Comparisons with different visual features (w/o fine-tuning visual encoder).

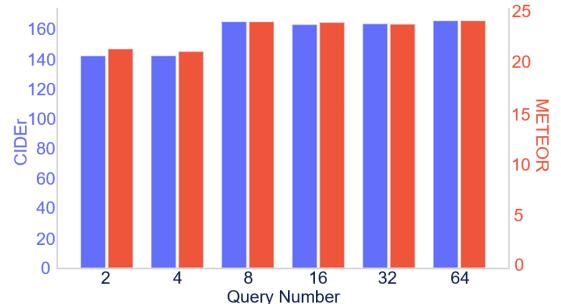


Figure 6: Impact of the query numbers on the performance. For each query number, we report CIDEr (blue) and METEOR scores (red) over the Moments-OVRE test set.

pre-training, involving a substantial collection of image-text pairs. We directly fine-tune GIT<sub>B</sub> to generate relation triplets without making further modifications.

**Result and Analysis.** We present our results on Moments-OVRE in Table 2 and compare our approach with baseline methods trained under the same training settings. Our approach outperforms baseline generative methods, achieving a higher METEOR score (+6.22) than ClipCap and (+2.01) than GIT. We find that although GIT was pre-trained on 0.8B image-text pairs and achieved impressive performance on video captioning datasets, it did not perform as well as our approach on the OVRE task. This could be attributed to the fact that the image-text generative pre-training does not directly facilitate the understanding of fine-grained information such as relationships in videos.

### Ablation Study

**The Impact of Query Numbers.** We first investigate the impact of query numbers. Specifically, we experiment with query number=2, 4, 8, 16, 32, and 64. Figure 6 shows that an extremely small number of queries will yield inferior generation results since the limited queries are insufficient to extract the content in video patches. As the number grows from 2 to 8, the generated results show significant improvement. Subsequently, when the number of queries continues to double, the performance of the model gradually becomes saturated. We choose query number=64 as it demonstrated the best performance across all metrics.

**Exploration of Vision and Language Model Fine-tuning.** Recent research (Rasheed et al. 2023) has shown that fully fine-tuned CLIP can effectively bridge the modality gap in the video domain. As shown in Table 3, though the train-



Figure 7: Comparisons of generated triplets across diverse OVRE methods. The illustration highlights accurately described triplets in light green, triplets with semantic correlation in light blue, and irrelevant triplets in light red.

able parameters increased, fine-tuning the CLIP-ViT can improve the CIDEr score by 8.8. Additionally, we delve into the outcomes of fine-tuning the text decoder. Experiments reveal that freezing GPT-2’s parameters results in a notable reduction in the CIDEr score. We attribute this decline to that GPT-2 is primarily good at generating natural language rather than triplets. Therefore, fine-tuning becomes imperative to improve its ability to generate rare triplet sequences.

**The Effect of Different Visual Features.** Our experimental investigations involve ablations on different granularity of visual features. Within our proposed framework, we employ patch features extracted from videos as prefixes for the text decoder. Furthermore, we explore two alternative representations as inputs to the attentional pooler: (I) Region features: Following the common VidVRD practice, we extract a sequence of objects and subsequently employ a tracking algorithm to obtain 5 tracklet features per video. These features replace patch features as input to the model. Specifically, we utilize RegionCLIP (Zhong et al. 2021) pre-trained from LVIS to crop bounding boxes and seqNMS (Han et al. 2016) for object tracking. (II) Frame features: We directly utilize features extracted from individual frames using CLIP, concatenating them to form a representation of frame-level features. As depicted in Table 4, both frame features and region features exhibit poor performance. Notably, frame features capture the overall visual content of an image but overlook finer details such as objects and relationships. Surprisingly, region features perform worse than frame features, which we attribute to the limited generalization capability of existing object detectors.

## Conclusion

In this paper, we introduce a new task named OVRE, where the model is required to generate all relationship triplets associated with the video actions. Concurrently, we present the corresponding Moments-OVRE dataset, which encompasses a diverse set of videos along with annotated relationships. We conduct extensive experiments on Moments-OVRE and demonstrated the superiority of our proposed approach over other baseline methods. We hope that our task and dataset will inspire more intricate and generalizable research in the realm of video understanding.

**Limitations:** (I) This version of Moment-OVRE has currently omitted BBox annotation due to the high cost of annotation. We are committed to progressively enhancing this dataset and intend to introduce BBox annotations in upcoming versions of Moments-OVRE. (II) For extracting action-centric relations, leveraging commonsense among action categories and relations (Yang et al. 2018) or implicit knowledge-driven representation learning methods (Li et al. 2023; Li, Tang, and Mei 2018) have shown promise. We will consider these knowledge-driven methods in future work.

## Acknowledgements

Jingjing Chen is supported partly by the National Natural Science Foundation of China (NSFC) project (No. 62072116). Zheng Wang is supported by the NSFC project (No. 62302453). Lechao Cheng is supported partly by the NSFC project (No. 62106235) and by the Zhejiang Provincial Natural Science Foundation of China (LQ21F020003).

## References

Chen, S.; He, X.; Guo, L.; Zhu, X.; Wang, W.; Tang, J.; and Liu, J. 2023. Valor: Vision-audio-language omni-perception pretraining model and dataset. *arXiv preprint arXiv:2304.08345*.

Chen, S.; and Jiang, Y.-G. 2019. Semantic proposal for activity localization in videos via sentence query. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 8199–8206.

Chen, S.; Yao, T.; and Jiang, Y.-G. 2019. Deep Learning for Video Captioning: A Review. In *IJCAI*, volume 1, 2.

Cubuk, E. D.; Zoph, B.; Shlens, J.; and Le, Q. V. 2019. RandAugment: Practical data augmentation with no separate search. *CoRR*, abs/1909.13719.

Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.

Gao, K.; Chen, L.; Huang, Y.; and Xiao, J. 2021. Video relation detection via tracklet based visual transformer. In *Proceedings of the 29th ACM international conference on multimedia*, 4833–4837.

Gao, K.; Chen, L.; Zhang, H.; Xiao, J.; and Sun, Q. 2023. Compositional Prompt Tuning with Motion Cues for Open-vocabulary Video Relation Detection. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Gao, T.; Yao, X.; and Chen, D. 2021. Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.

Gkioxari, G.; Girshick, R.; Dollár, P.; and He, K. 2018. Detecting and recognizing human-object interactions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8359–8367.

Goyal, R.; Kahou, S. E.; Michalski, V.; Materzynska, J.; Westphal, S.; Kim, H.; Haenel, V.; Fruend, I.; Yianilos, P.; Mueller-Freitag, M.; Hoppe, F.; Thurau, C.; Bax, I.; and Memisevic, R. 2017. The “Something Something” Video Database for Learning and Evaluating Visual Common Sense. In *2017 IEEE International Conference on Computer Vision (ICCV)*, 5843–5851.

Han, W.; Khorrami, P.; Paine, T. L.; Ramachandran, P.; Babaeizadeh, M.; Shi, H.; Li, J.; Yan, S.; and Huang, T. S. 2016. Seq-NMS for Video Object Detection. *arXiv:1602.08465*.

Heilbron, F. C.; Escorcia, V.; Ghanem, B.; and Niebles, J. C. 2015. ActivityNet: A large-scale video benchmark for human activity understanding. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 961–970.

Ji, J.; Krishna, R.; Fei-Fei, L.; and Niebles, J. C. 2019. Action Genome: Actions as Composition of Spatio-temporal Scene Graphs. *CoRR*, abs/1912.06992.

Kay, W.; Carreira, J.; Simonyan, K.; Zhang, B.; Hillier, C.; Vijayanarasimhan, S.; Viola, F.; Green, T.; Back, T.; Natsev, A.; Suleyman, M.; and Zisserman, A. 2017. The Kinetics Human Action Video Dataset. *ArXiv*, abs/1705.06950.

Kong, Y.; and Fu, Y. 2022. Human action recognition and prediction: A survey. *International Journal of Computer Vision*, 130(5): 1366–1401.

Li, Y.-L.; Liu, X.; Wu, X.; Li, Y.; Qiu, Z.; Xu, L.; Xu, Y.; Fang, H.-S.; and Lu, C. 2022. Hake: a knowledge engine foundation for human activity understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Li, Z.; Tang, H.; Peng, Z.; Qi, G.-J.; and Tang, J. 2023. Knowledge-guided semantic transfer network for few-shot image recognition. *IEEE Transactions on Neural Networks and Learning Systems*.

Li, Z.; Tang, J.; and Mei, T. 2018. Deep collaborative embedding for social image understanding. *IEEE transactions on pattern analysis and machine intelligence*, 41(9): 2070–2083.

Lin, K.; Li, L.; Lin, C.-C.; Ahmed, F.; Gan, Z.; Liu, Z.; Lu, Y.; and Wang, L. 2022. Swinbert: End-to-end transformers with sparse attention for video captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17949–17958.

Liu, Y.; Wang, L.; Wang, Y.; Ma, X.; and Qiao, Y. 2022. FineAction: A Fine-Grained Video Dataset for Temporal Action Localization. *arXiv:2105.11107*.

Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2022. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508: 293–304.

Mokady, R.; Hertz, A.; and Bermano, A. H. 2021. Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*.

Monfort, M.; Andonian, A.; Zhou, B.; Ramakrishnan, K.; Bargal, S. A.; Yan, T.; Brown, L.; Fan, Q.; Gutfreund, D.; Vondrick, C.; et al. 2019. Moments in Time Dataset: one million videos for event understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–8.

Monfort, M.; Pan, B.; Ramakrishnan, K.; Andonian, A.; McNamara, B. A.; Lascelles, A.; Fan, Q.; Gutfreund, D.; Feris, R.; and Oliva, A. 2021. Multi-Moments in Time: Learning and Interpreting Models for Multi-Action Video Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1.

Ni, B.; Peng, H.; Chen, M.; Zhang, S.; Meng, G.; Fu, J.; Xiang, S.; and Ling, H. 2022. Expanding Language-Image Pretrained Models for General Video Recognition. *arXiv:2208.02816*.

Qian, T.; Cui, R.; Chen, J.; Peng, P.; Guo, X.; and Jiang, Y.-G. 2023. Locate before answering: Answer guided question localization for video question answering. *IEEE Transactions on Multimedia*.

Qian, X.; Zhuang, Y.; Li, Y.; Xiao, S.; Pu, S.; and Xiao, J. 2019. Video Relation Detection with Spatio-Temporal Graph. In *Proceedings of the 27th ACM International Conference on Multimedia, MM '19*, 84–93. New York, NY, USA: Association for Computing Machinery. ISBN 9781450368896.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.

Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I.; et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.

Rasheed, H.; khattak, M. U.; Maaz, M.; Khan, S.; and Khan, F. S. 2023. Finetuned CLIP models are efficient video learners. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*.

Sadhu, A.; Gupta, T.; Yatskar, M.; Nevatia, R.; and Kembhavi, A. 2021. Visual semantic role labeling for video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5589–5600.

Shang, X.; Di, D.; Xiao, J.; Cao, Y.; Yang, X.; and Chua, T.-S. 2019. Annotating Objects and Relations in User-Generated Videos. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval*, 279–287. ACM.

Shang, X.; Ren, T.; Guo, J.; Zhang, H.; and Chua, T.-S. 2017. Video Visual Relation Detection. *Proceedings of the 25th ACM international conference on Multimedia*.

Song, X.; Chen, J.; and Jiang, Y.-G. 2023. Relation Triplet Construction for Cross-modal Text-to-Video Retrieval. In *Proceedings of the 31st ACM International Conference on Multimedia*, 4759–4767.

Song, X.; Chen, J.; Wu, Z.; and Jiang, Y.-G. 2021. Spatial-temporal graphs for cross-modal text2video retrieval. *IEEE Transactions on Multimedia*, 24: 2914–2923.

Tang, M.; Wang, Z.; Liu, Z.; Rao, F.; Li, D.; and Li, X. 2021. Clip4caption: Clip for video caption. In *Proceedings of the 29th ACM International Conference on Multimedia*, 4858–4862.

Wang, J.; Yang, Z.; Hu, X.; Li, L.; Lin, K.; Gan, Z.; Liu, Z.; Liu, C.; and Wang, L. 2022. GIT: A Generative Image-to-text Transformer for Vision and Language. arXiv:2205.14100.

Wang, X.; Wu, J.; Chen, J.; Li, L.; Wang, Y.-F.; and Wang, W. Y. 2019. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4581–4591.

Wang, Z.; Chen, J.; and Jiang, Y.-G. 2021. Visual co-occurrence alignment learning for weakly-supervised video moment retrieval. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1459–1468.

Xia, H.; and Zhan, Y. 2020. A Survey on Temporal Action Localization. *IEEE Access*, 8: 70477–70487.

Xu, H.; Ye, Q.; Yan, M.; Shi, Y.; Ye, J.; Xu, Y.; Li, C.; Bi, B.; Qian, Q.; Wang, W.; Xu, G.; Zhang, J.; Huang, S.; Huang, F.; and Zhou, J. 2023. mPLUG-2: A Modularized Multi-modal Foundation Model Across Text, Image and Video. arXiv:2302.00402.

Xu, J.; Mei, T.; Yao, T.; and Rui, Y. 2016. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5288–5296.

Yan, S.; Zhu, T.; Wang, Z.; Cao, Y.; Zhang, M.; Ghosh, S.; Wu, Y.; and Yu, J. 2023. VideoCoCa: Video-Text Modeling with Zero-Shot Transfer from Contrastive Captioners. arXiv:2212.04979.

Yang, S.; Gao, Q.; Saba-Sadiya, S.; and Chai, J. 2018. Commonsense justification for action explanation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2627–2637.

Yatskar, M.; Zettlemoyer, L.; and Farhadi, A. 2016. Situation recognition: Visual semantic role labeling for image understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5534–5542.

Zhang, W.; Wang, B.; Ma, L.; and Liu, W. 2019. Reconstruct and represent video contents for captioning via reinforcement learning. *IEEE transactions on pattern analysis and machine intelligence*, 42(12): 3088–3101.

Zhang, X.; Wu, Z.; Weng, Z.; Fu, H.; Chen, J.; Jiang, Y.-G.; and Davis, L. S. 2021. Videolt: Large-scale long-tailed video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7960–7969.

Zheng, S.; Chen, S.; and Jin, Q. 2022. VRDFormer: End-to-End Video Visual Relation Detection with Transformers. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 18814–18824.

Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; and Gao, J. 2021. RegionCLIP: Region-based Language-Image Pre-training. arXiv:2112.09106.