



Collaborative Three-Stream Transformers for Video Captioning

Hao Wang^a, Libo Zhang^{b,**}, Heng Fan^c, Tiejian Luo^a

^aSchool of Computer Science and Technology, University of Chinese Academy of Sciences, Beijing, 101408, China

^bInstitute of Software, Chinese Academy of Sciences, Beijing, 100190, China

^cDepartment of Computer Science and Engineering, University of North Texas, Denton 76203, Texas, United States of America

ABSTRACT

As the most critical components in a sentence, subject, predicate and object require special attention in the video captioning task. To implement this idea, we design a novel framework, named Collaborative three-Stream Transformers (COST), to model the three parts separately and complement each other for better representation. Specifically, COST is formed by three branches of transformers to exploit the visual-linguistic interactions of different granularities in spatial-temporal domain between videos and text, detected objects and text, and actions and text. Meanwhile, we propose a cross-granularity attention module to align the interactions modeled by the three branches of transformers, then the three branches of transformers can support each other to exploit the most discriminative semantic information of different granularities for accurate predictions of captions. The whole model is trained in an end-to-end fashion. Extensive experiments conducted on three large-scale challenging datasets, *i.e.*, YouCookII, ActivityNet Captions and MSVD, demonstrate that the proposed method performs favorably against the state-of-the-art methods.

© 2023 Elsevier Ltd. All rights reserved.

1. Introduction

Video captioning aims to generate natural language descriptions of video content, which attracts much attention in recent years along with the rapidly increasing amount of videos recorded in daily life. It can be adopted for a wide range of real-world applications, such as blind people assistance, automatic videos summarization and classification, and intelligent video surveillance. However, as noted in former works (Xiong et al., 2018; Park et al., 2019; Lei et al., 2020), it is very challenging to generate natural paragraph descriptions due to the difficulties of having relevant, less redundant, and semantic coherent sentences.

Recently, researchers attempt to use the transformer model to solve the video captioning task (Vaswani et al., 2017; Dai et al., 2019; Iashin and Rahtu, 2020; Zhu and Yang, 2020; Tang et al., 2021), which relies on the self-attention mechanism to describe the interactions between different modalities of the input data, such as video, audio, and text. In practice, the aforementioned

methods generally concatenate the features extracted from individual modalities, or use self-attention to model the interactions between extracted features. Although they advance the state-of-the-art of video captioning, it is still far from satisfactory in real applications due to the domain gap between different modalities. Thus, a question naturally arises, “*How to reduce the domain gap and capture the interactions among visual and linguistic modalities for video captioning?*”

Before answering this question, let us see the basic grammar rules at first. As pointed out in Krishna et al. (2017a) and Zhou et al. (2018a), a sentence is generally presented as the following form, *e.g.*,

Women wear Arabian skirts on a stage.

where **Subject**, **Object**, and **Predicate** are the three most critical elements, and indicate the objects, the actions of objects, and the interactions among different objects, respectively. We believe that these components correspond to visual representations with different granularities, and modeling multi-modal interactions based on them can effectively reduce the domain gap and help model understand the content of video better. Thus motivated, we propose a novel framework, called Collaborative three-Stream Transformers (COST), for video captioning.

**Corresponding author:

e-mail: libo@iscas.ac.cn (Libo Zhang)

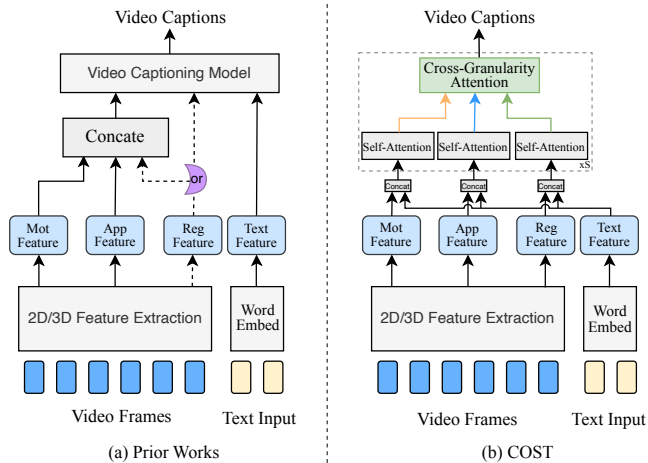


Fig. 1. Comparison between prior works and the proposed COST. Previous works tended to concatenate the visual features, *i.e.*, motion feature (Mot Feature), appearance feature (App Feature) and region feature (Reg Feature), as input to their video captioning model. Differently, we propose a three-branch transformer-based architecture to encode the visual-linguistic interactions of different granularities separately and design the cross-granularity module to complement the interactions with each other.

Different from former methods (Zhou et al., 2018b; Lei et al., 2020; Dai et al., 2019; Wang et al., 2021) which directly concatenate visual features of different granularities and model the correlation between visual and linguistic modality directly (Fig. 1(a)), COST models the visual-linguistic (Nan et al., 2021; Fan and Yang, 2020; Seo et al., 2022; Wei et al., 2020; Luo et al., 2020) interactions of different granularities separately in spatial-temporal domain and then fuses the multi-modal feature in different branches for better video content understanding (Fig. 1(b)). As we argue that these previous works have overlooked the domain gap present in the features extracted from different pre-trained models, and this oversight renders it unreasonable to simply concatenate these features as a unified video representation. Note that, the dashed lines in Fig. 1(a) indicate that there exist some works extracting and combining region features with the concatenated motion and appearance features (Park et al., 2019) together or sending them parallel (Zhang et al., 2020) as the visual input to their models.

Specifically, the proposed COST consists of three transformer branches, including the Video-Text, Detection-Text, and Action-Text transformers. The Video-Text transformer is used to model the interactions between the global video appearances and linguistic texts, which makes the model perceive the general content of the video. The Detection-Text transformer aims at accurately locating objects in individual video frames, which enforces the model to focus on the objects being aligned in the visual and linguistic modalities, *i.e.*, indicating the Subjects and Objects in caption sentences. The Action-Text transformer is designed to model the actions/relations of objects between the visual and linguistic modalities, *i.e.*, indicating the Predicate in caption sentences. Meanwhile, to align the interactions modeled by the three branches of transformers, we introduce a cross-granularity attention model in COST. In particular, the similarity between interactions from different branches is com-

puted to represent the relevance among visual modalities and help inject the information from other interactions. Note that we use the terminology “granularity” to describe the different levels of detail or scales at which the same input data is processed. For example, our motion features and region features (detection features) are both extracted from images, yet they provide distinct information. In contrast, “modality” is a related term that refers to the various types of inputs that a machine learning system processes, such as text, image, and audio. Each modality requires specific processing techniques and algorithms to extract valuable information and patterns effectively. In addition, we introduce additional training objectives for Detection-Text and Action-Text streams separately to align the semantics of embeddings to the underlying video information which is supposed to be conveyed at the time of feature extraction, instead of only setting cross-entropy loss between generated captions and ground-truth descriptions like most former methods. As we believe that introducing appropriate guidance can ensure the uniqueness of information in each stream and provide complementary information for each other. In this way, different branches of transformers support each other to exploit more discriminative semantic information in different modalities and granularities, and enforce the model to pay more attention on generating the accurate Subject, Object and Predicate predictions. The whole model is trained in an end-to-end fashion using Adam algorithm (Kingma and Ba, 2015).

Extensive experiments are conducted on three publicly challenging datasets, *i.e.*, YouCookII (Zhou et al., 2018a), ActivityNet Captions (Krishna et al., 2017a) and MSVD (Chen and Dolan, 2011), to demonstrate the superior performance of the proposed method compared to the state-of-the-art methods (Zhou et al., 2018b; Dai et al., 2019; Park et al., 2019; Lei et al., 2020; Wang et al., 2021; Seo et al., 2022; Lin et al., 2021). Specifically, our method achieves very competitive CIDEr scores with TSN (Wang et al., 2019) features as input, *i.e.*, 45.54% and 24.77%, on the YouCookII *val* set and the ActivityNet *ae-test* set, surpasses or approaches the state-of-the-arts.

The main contributions of this paper are summarized as follows:

1. We propose a novel framework, *i.e.*, Collaborative three-Stream Transformers (COST), that leverages multiple transformer branches to explore various components in a sentence for video captioning.
2. We design a simple but effective cross-granularity attention module to align the interactions modeled by different transformer branches, which supports each other to exploit discriminative semantic cue from different granularities for more accurate predictions of captions.
3. We specially introduce a new training objective for COST that constrains the semantics of embeddings in Detection-Text branch and Action-Text branch to supply for caption generation, further enhancing performance.
4. Extensive experiments conducted on three challenging datasets show that our method performs favorably against the state-of-the-art methods.

2. Related Work

Video Captioning. Over the past few years, video captioning has received increasing attention from both computer vision and natural language processing community (Xu et al., 2021), which aims to generate linguistic description for the video content. Early works mainly focus on template-based approaches. For example, Guadarrama et al. (2013) and Das et al. (2013) first detect the visual objects in a video with human-crafted features and then use them to fill the pre-defined templates with slots. However, such methods are restricted in generating semantically-rich sentences due to the high dependence on fixed templates and language rules.

Motivated by the success of neural network in translation task (Sutskever et al., 2014), the methods of taking captioning task as translation task became popular (Venugopalan et al., 2015b; Yu et al., 2016; Pan et al., 2017; Zhang et al., 2018). The nature of these methods is performing sequence-to-sequence learning in an encoder-decoder paradigm (Chen et al., 2019), where the convolutional neural networks (CNNs) (Zheng et al., 2020) and Long-Short Term Memory (LSTM) networks (Pei et al., 2019; Park et al., 2019) are generally adopted to extract discriminative feature embeddings from input video and generate accurate captions separately. To alleviate the heavy computational burden of applying 2D-CNN, especially 3D-CNN to dense frame inputs for visual feature extraction, the methods in this paradigm mostly operate on the pre-extracted features and the subsequent works mainly concentrate on the improvements in feature extraction or utilization, including 1) multi-modal feature extraction (Wang et al., 2018; Hao et al., 2018; Hori et al., 2018; Xu et al., 2019; Zhang et al., 2020) and 2) feature utilization optimization (Venugopalan et al., 2015b,a; Yao et al., 2015; Li et al., 2017; Gao et al., 2020; Chen and Jiang, 2019). In the former one, HACA (Wang et al., 2018) proposes a hierarchically LSTM-based network to learn and align the attentive representations of both visual and audio features at different granularities. DS-RNN (Xu et al., 2019) proposes a dual-stream framework to model visual and semantic features independently and decode the hidden states from both modalities jointly. ORG-TRL (Zhang et al., 2020) not only proposes an object relational graph to connect each object in video and do relational reasoning by graph convolutional network, but also designs a teacher-recommended learning method to utilize the external language model to improve the generation of caption model by integrating the linguistic knowledge. In the later one, LSTM-YT (Venugopalan et al., 2015b) proposes that mean pooling features across all frames is a rational representation for generating simple descriptions to short video clips. However, it totally ignores the order of frames in original video and discards the temporal information which is crucial for video captioning. S2VT (Venugopalan et al., 2015a) firstly introduces sequence-to-sequence model to video captioning task and processes frames and generates words sequentially while preserving the temporal structure. In order to make the model concentrate on relevant features for specific word generation, temporal attention (Yao et al., 2015) is proposed to assign higher weights to relevant features. Apart from the frame-level attention, MAM-RNN (Li et al., 2017) proposes that different

regions in the video frame contribute differently to the word prediction, and designs a two-layers structure with the first layer focusing on the most salient regions in each frame and the second one attending to the most correlated frames.

Recently, inspired by the superior performance in learning long-range relations through the attention mechanism compared with RNNs (*e.g.*, LSTM and GRU (Chung et al., 2014)), Transformer-based models have emerged in this field. Zhou et al. (2018b) proposes an end-to-end trained transformer model, where the encoder is designed to extract semantic representations from the video, and the proposal decoder receives the encoding output with different anchors to form video event proposals. Sun et al. (2019) designs the VideoBERT model to learn bidirectional joint distributions over sequences of visual and linguistic tokens. Lei et al. (2020) develops the memory-augmented recurrent transformer, which uses a highly summarized memory state from the video clips and the sentence history to facilitate better prediction of the next sentence. More recently, by employing the tubelet embedding scheme and factorized encoder architecture from ViViT (Arnab et al., 2021) and elaborately designed bi-directional objective, MV-GPT (Seo et al., 2022) can be applied to raw pixels instead pre-extracted visual features and generates captions directly. SwinBERT (Lin et al., 2021) is also an end-to-end transformer-based video captioning model with video frame patches as input, which uses VidSwin (Liu et al., 2021b) as visual encoder to extract video tokens and promotes cross-modality representation by applying masked text token prediction. They provide new insights for the future development of this field.

Multi-modal cross-attention mechanism. The interactions between different modalities are critical for the video captioning task. Recent transformer based methods (Iashin and Rahtu, 2020; Zhu and Yang, 2020; Tang et al., 2021) use the cross-attention module to learn correlations across different modalities. For example, Iashin and Rahtu (2020) concatenates the learned embeddings from multiple modalities, *e.g.*, video, audio and speech, for event description. Tang et al. (2021) uses frame-level dense captions as an auxiliary text input for better video and language associations, where the constrained attention loss is used to force the model to automatically focus on the best matched caption from a pool of misalignment caption candidates. Recently, Li et al. (2022) proposed a novel approach to strengthen image and text embeddings through the incorporation of action information through cross-attention manner, and achieves state-of-art performance on image-text retrieval task.

Despite sharing some similarities with these approaches in using self-attention to align visual-linguistic interactions, our approach is significantly different. In specific, we design the cross-granularity attention module integrated in the collaborative three-stream transformers to align three types of visual-linguistic interactions of different granularities, as we believe that they could support and complement each other, for example, the model could more easily understand the current motion scene after injecting the information from detection feature of "basketball" and action feature of "play" to video feature, leading to more discriminative semantic cues for better caption generation.

Multi-branch architectures. The idea of multi-branch has been widely applied in various tasks (Liu et al., 2021a; Zhu and Yang, 2020; Kim et al., 2020). SibNet (Liu et al., 2021a) proposes a two-branch architecture to encode the content and semantic of videos separately and then the both branches are combined and fed to decoder for video caption generation. ActBert (Zhu and Yang, 2020) proposes the tangled transformer block to encode three sources of information, *i.e.*, global actions, local regional objects, and linguistic descriptions to learn better video-text representation. MTTSNet (Kim et al., 2020) defines three region features corresponding to part-of-speech (POS) relation and designs a triple-stream network to encode these features separately and merges the processed embeddings to generate caption and predicts POS of each word for image caption task.

Our COST also adopts the multi-branch architecture, but it is different than the above methods from two aspects. First, the design motivations are different. We use three-branch architecture to encode the visual-linguistic interaction of different granularities to disentangle the complex video information by self-attention, and use the cross-granularity attention module to complement and enhance the embeddings extracted from different branches, so as to improve the video analysis ability of the model to generate accurate description. Second, the contributions of branch to task are different. For example, ActBert and MTTSNet only set optimization objective for the whole model instead of each branch, and SibNet designs specific loss for two branches but only to improve the ability of the model to extract feature from the video. Differing from them, we set training objective for each branch to align the semantics of pre-extracted visual features to correspond to POS, and the embeddings in each branch are semantically complementary and directly participate in caption generation after information fusion.

Multi-modal pre-training models. Large-scale pre-training is another effective way to improve the accuracy of captioning models. Specifically, the jointly trained video and language models (Sun et al., 2019; Zhu and Yang, 2020; Huang et al., 2020; Luo et al., 2020; Ging et al., 2020; Seo et al., 2022) on the large-scale datasets, such as YouTube-8M (Abu-El-Haija et al., 2016) and HowTo100M (Miech et al., 2019) with automatic speech recognition¹ transcripts, provide discriminative features for downstream tasks, such as video captioning, action localization and *et al.* VideoBert (Sun et al., 2019) and ActBert (Zhu and Yang, 2020) collect large paired video sequences and text descriptions with the help of off-the-shelf automatic speech recognition (ASR) model, and construct the BERT-style objective to train the video and text encoders simultaneously. Huang et al. (2020) constructs a dense video captioning dataset, *i.e.*, Video Timeline Tags (ViTT), and explores several multi-modal sequence-to-sequence pre-training strategies using transformers (Vaswani et al., 2017). Luo et al. (2020) also applies transformers with two single-modal encoders to encode the video and text separately, a cross encoder to model the interactions between video and text representations, and a decoder to reconstruct or generate text. Ging et al. (2020) develops the Coop-

erative hierarchical Transformer (COOT) to model the interactions between different levels of granularities and modalities, which achieves superior results on video-text retrieval task and provides learned representations to improve the performance of video captioning model (Lei et al., 2020) significantly. Recently, Seo et al. (2022) proposes that recent visual-language pre-training frameworks lack the ability to generate sentences and presents MV-GPT with novel bidirectional objective for generation task.

3. Our Approach

As discussed above, we design the collaborative three-stream transformers to model the interactions of objects, and actions/relations of objects between different modalities, which is formed by three branches of transformers, *i.e.*, Video-Text, Detection-Text, and Action-Text transformers. Specifically, the video and text inputs are firstly encoded to extract the multi-modal feature embeddings. After that, the embeddings are fed into the three-stream transformers to exploit the interactions between linguistic embedding and visual embeddings of different granularities in spatial-temporal domain, *i.e.*, global videos, object regions, and actions. Meanwhile, the cross-granularity attention module is designed to align the interactions modeled by the three branches of transformers. The overall architecture of the proposed method is shown in Fig. 2.

3.1. Multi-Modality Tokens

Three kinds of tokens, *i.e.*, visual tokens, linguistic tokens, and special tokens, are used to express the video and text inputs, which are described as follows.

Visual tokens. For the visual tokens, we use three kinds of tokens for different granularities in spatial-temporal domain, that is the video tokens, the detection tokens, and the action tokens.

- *Video tokens* provide the global semantic information in the video sequence. In contrast to Lei et al. (2020), we only use the appearance features extracted by Temporal Segment Networks (TSN) (Wang et al., 2016) (denoted as TSN-APP in Fig. 2) as the video tokens, *i.e.*, $\{f_1^v, f_2^v, \dots, f_{N_v}^v\}$, where f_i^v is the extracted feature of the i -th video clip, and N_v is the number of video clips. Notably, we can also leverage more powerful multi-modal feature extraction method COOT (Ging et al., 2020) to improve the performance, which is pre-trained on the large-scale HowTo100M dataset (Miech et al., 2019).
- *Detection tokens* are used to enforce the model to focus on the Subjects or Objects in caption sentences. Similar to Park et al. (2019); Lu et al. (2019); Zhu and Yang (2020), we use the Faster R-CNN method with 101-layer residual net (ResNet-101) (He et al., 2016) as backbone to detect the objects in each frame, which is pre-trained on the Visual Genome dataset (Krishna et al., 2017b). After that, the detection features in Faster R-CNN corresponding to

¹<https://developers.google.com/youtube/v3/docs/captions>




Fig. 2. The network architecture of the proposed COST method, which is formed by three branches of transformers, *i.e.*, the Action-Text, Video-Text and Detection-Text transformers. The cross-granularity attention module is designed to align the interactions modeled by the three different branches of transformers, and Y , H and X represent the interactions of these branches separately.

the objects with the highest confidence scores in K categories² are used to generate the detection tokens for each frame. We use $\{f_1^d, f_2^d, \dots, f_{N_d}^d\}$ to denote the set of detection tokens, where f_i^d is the i -th detection feature, and N_d is the total number of detections in the video sequence.

- *Action tokens* are designed to enforce the model to concentrate on the Predicates in caption sentences. Following Lei et al. (2020), the optical flow features of video sequences are extracted by TSN (Wang et al., 2016) (denoted as TSN-MOT is Fig. 2) to generate the action tokens, which are used to describe the actions/relations of objects. The action tokens are denoted as $\{f_1^a, f_2^a, \dots, f_{N_a}^a\}$, where f_i^a is the motion feature of the i -th video clip, and N_a is the total number of action tokens. It is noted that the N_v is always equal to N_a in our method based on the pre-trained model

we use for extracting appearance features and motion features.

Linguistic tokens. We break down the captions of video sequences into individual words and compute the corresponding linguistic tokens using the GloVe model (Pennington et al., 2014). The linguistic tokens are denoted as $\{f_1^l, f_2^l, \dots, f_{N_l}^l\}$, where f_i^l is the extracted features of the i -th word using the GloVe model, and N_l is the total number of words. In light of the emergence of more powerful language representation models, such as BERT (Devlin et al., 2019) and CLIP (Radford et al., 2021), GloVe appears outdated as a model for vector representations for text. Nonetheless, we utilize GloVe features to represent captions for the sake of comparison fairness, as the main methods under comparison, such as MART Lei et al. (2020) and PDVC Wang et al. (2021), rely on GloVe as their text encoding model. Additionally, setting feature vectors to special characters by first random initialization with the same shape as features provided by GloVe, such as "[EOS]", and the semantics of these characters can be learned during the training process, rendering GloVe a suitable choice for the current

²If the category number of the detected objects is less than K in a frame, we select the K detected objects with the highest confidence scores regardless the object categories to generate the detection tokens.

task.

Special tokens. Besides the aforementioned tokens, we also introduce three kinds of special tokens in transformer, similar to BERT (Devlin et al., 2019). The first one is the granularity token [CLS], which is added at the beginning of visual features to denote which granularity the following tokens belong to. The second one is the three kinds of separation token, *i.e.*, [SEP], [BOS], and [EOS]. [SEP] is used at the end of the visual tokens to separate them from the linguistic tokens, [BOS] is used to denote the beginning of linguistic tokens, and [EOS] is used to denote the ending of the linguistic tokens, respectively. And the last one is the padding token [PAD] with two purposes: supplementing the visual token sequence or text token sequence to specified length for training or inference in parallel; substituting for all linguistic tokens in inference phase so no caption information will be leaked to the model. In addition, we use a fully-connected layer to encode the aforementioned tokens to the same dimension. Thus, the inputs for the three-stream transformer are computed as

$$\begin{aligned} & \{[\text{CLS}^{(v)}], f_1^{(v)}, f_2^{(v)}, \dots, f_{N_v}^{(v)}, [\text{SEP}], [\text{PAD}]_{\times N_1}, \\ & [\text{BOS}], f_1', \dots, f_{N_t}', [\text{EOS}], [\text{PAD}]_{\times N_2} \} \end{aligned} \quad (1)$$

where $(\cdot) \in \{v, d, a\}$ indicates the video, detection, and action tokens, respectively. And the subscript of [PAD] means the number of it to supplement. Please note that the purpose of concatenating visual and linguistic tokens is not to directly fuse their information. Rather, the interaction between the features of two modalities will be encoded through a subsequent masked multi-head attention module. Thus, we choose concatenation as it helps to preserve the information of both modalities. We also use the positional encoding strategy (Vaswani et al., 2017) in the Video-Text, Detection-Text, and Action-Text transformers to describe the order information of caption sentences.

3.2. Three-Stream Transformers

As shown in Fig. 2, we feed the aforementioned tokens into the three-stream transformers. The Video-Text, Detection-Text, and Action-Text branches are formed by S basic blocks, and each block mainly consists of a self-attention module and a cross-granularity attention module. Both the self-attention and cross-granularity modules are followed by a feed forward layer. **Self-attention module.** The self-attention module is designed to model the visual-linguistic alignments in each branch of transformers, *i.e.*, Video-Text, Detection-Text, and Action-Text. Following Vaswani et al. (2017), we compute the attention function between different tokens as follows.

$$\mathcal{A}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (2)$$

where Q , K and V are created by inputting tokens into three different fully-connected layers in each branch. And their dimensions are $\mathbb{R}^{N \times d}$, where N and d are the number of tokens and the dimension of embeddings, respectively. We advocate h paralleled heads of scaled dot-product attentions to increase the diversity. It is noted that we also add mask to the scaled dot-product before applying Softmax to prevent the model from seeing future words (Vaswani et al., 2017). We denote the whole module as *Masked Multi-Head Self-Attention* in Fig. 2




Fig. 3. Two cross-granularity attention architectures to merge embeddings from different branches. M-MHA denotes the Masked Multi-Head Attention. X , Y and H indicate the hidden states from Detection-Text branch, Action-Text branch and Video-Text branch separately and both sub-figures demonstrate the process of information fusion from the other two branches to Action-Text branch.

and the output is merged with input using a residual connection and layer norm.

Cross-granularity attention module. Besides the self-attention module, we use the cross-granularity attention module to align the interactions modeled by the three branches of transformers. Specifically, we complement and enhance the embeddings in each branch by injecting the information from other branches based on cross-attention between them. The branch that is incorporated into information and the other two branches are dubbed as injected branch and supply branches for better understanding the following description. We propose two architectures as shown in Fig. 3 and compare their performance by experiments. It is important to note that we only illustrate the implementation of the cross-granularity attention module within the Action-Text branch in Fig 3. Specifically, embeddings from this branch are utilized as the query to fuse information from the other two branches, while the embeddings from the other two branches similarly serve as queries within their respective cross-granularity attention modules.

Concatenated Architecture A simple and straightforward way to inject the related information from supply branches is to concatenate the embeddings from them and use the obtained result as key and value while regarding the embeddings from injected branch as query in the multi-head attention layer. The whole process in Action-Text branch is shown in Fig. 3(a) and the other two branches perform the same computation with their own embeddings as query. The computation process is

$$M_Y = \text{softmax}(Y \odot ([H; X]^T)) \quad (3)$$

$$Y' = \text{FFN}(M_Y \odot ([H; X])) \quad (4)$$

where \odot represents the dot product, $[\cdot; \cdot]$ denotes the concatenation operation, FFN is the feed forward layer, X , Y and H denote the hidden states from Detection-Text branch, Action-Text branch and Video-Text branch separately. We estimate the modality-wise normalization score of the interactions using a softmax layer. It should be noted that we add the mask here to prevent possible linguistic information leakage to the model. The affinity matrix $M_Y \in \mathbb{R}^{(N_a+N_t) \times (N_v+N_d+2N_t)}$ and $M_Y(i, j)$ denotes the normalized interaction score between the i -th entity in the Action-Text embeddings and the j -th entity in the con-

catenated Video-Text and Detection-Text embeddings. Based on the matrix, the feature embeddings of the Action-Text transformer can extract information from other branches. Feed forward layer is used to further encode the extracted information which will be merged to Action-Text transformer.

Parallel Architecture As shown in 3 (b), the other architecture design of cross-granularity attention is to feed embedding from one of supply branches and that from the injected branch to multi-head attention layer parallel, and then fuse the result from the two attention layers. The computation process is

$$\begin{aligned} M_{YH} &= \text{softmax}(Y \odot H^T) & M_{YX} &= \text{softmax}(Y \odot X^T) & (5) \\ Y' &= \text{FFN}(M_{YH} \odot H + M_{YX} \odot X) & & & (6) \end{aligned}$$

Then Y' will be merged with Y using the residual connection and layer norm. Notably, we evaluate our model equipped with different cross-granularity architecture and analyse the result in *Ablation Studies*. The results show that the our model with the parallel cross-granularity structure performs better, and we keep this structure in all comparisons with other methods. We apply the cross-granularity attention in all blocks for each branch of transformers. In this way, the visual entities of different granularities can enhance each other with more discriminative semantics for video captioning. It is noteworthy that we can leverage the video-text features in history to obtain the long-term sentence-level recurrence to generate the next sentences according to Lei et al. (2020), which can further enhance the performance of our model.

3.3. Optimization Objective

As we stated before, we introduce specified loss to each stream to guide the training of our COST method. This approach helps in preserving the semantics of visual features at different granularities and facilitate their complementary nature. The loss is formed by three terms, *i.e.*, $\mathcal{L}_v(\cdot, \cdot)$ for the Video-Text transformer, $\mathcal{L}_d(\cdot, \cdot)$ for the Detection-Text transformer, and $\mathcal{L}_a(\cdot, \cdot)$ for the Action-Text transformer, *i.e.*,

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_v(\ell_v, [\{f_1^t, \dots, f_{N_t}^t, [PAD]_{\times N_t}\}]) \\ &+ \lambda_d \cdot \mathcal{L}_d(\ell_d, [\{f_1^d, f_2^d, \dots, f_{N_d}^d\}]) & (7) \\ &+ \lambda_a \cdot \mathcal{L}_a(\ell_a, [CLS^a]) \end{aligned}$$

where λ_d and λ_a are the pre-set parameters used to balance these three terms. $\mathcal{L}_v(\cdot, \cdot)$ is the cross-entropy loss used to penalize the errors of the predicted captions by the linguistic tokens from Video-Text branch comparing to the ground-truth descriptions ℓ_v , which are the indexes of words in caption sentences. $\mathcal{L}_d(\cdot, \cdot)$ is also the cross-entropy loss used to penalize the errors of the predicted categories of objects by the visual tokens from Detection-Text branch comparing to the pseudo category labels generated by the Faster R-CNN detector³. $\mathcal{L}_a(\cdot, \cdot)$ is designed as the multi-label classification loss because there may

be multiple actions in a video clip. Specifically, we first aggregate all action tokens as the granularity token $[CLS^a]$ and then compute the confidence score using a fully-connected layer, *i.e.*, $\text{FC}([CLS^a])$. Following Zhang et al. (2021) and Sun et al. (2020), the loss function is computed as

$$\mathcal{L}_a(\ell_a, [CLS^a]) = \log(1 + \sum_{i \in \Omega_{\text{pos}}(\ell_a)} e^{-s_i}) + \log(1 + \sum_{j \in \Omega_{\text{neg}}(\ell_a)} e^{s_j}) \quad (8)$$

which expects the confidence scores s_i of the existing actions $\Omega_{\text{pos}}(\ell_a)$ in video sequences are greater than the predefined threshold 0 while the confidence scores s_j of non-existing actions $\Omega_{\text{neg}}(\ell_a)$ are less than 0. The ground-truth action labels are the most common verbs in caption sentences, which are retrieved by the off-the-shelf part-of-speech tagger method (Sun et al., 2019).

4. Experiments

4.1. Datasets And Evaluation Metrics

Datasets. We conduct several experiments on two challenging datasets, *i.e.*, YouCookII (Zhou et al., 2018a) and ActivityNet Captions (Krishna et al., 2017a).

1) *YouCookII* includes 2,000 long untrimmed videos describing 89 cooking recipes, where each video contains one reference paragraph and is further split into several event segments with annotated sentences. 1,333 and 457 video sequences are used for training and validation, respectively.

2) *ActivityNet Captions* is a large-scale dataset formed by 10,009 videos for training and 4,917 videos for validation and testing. Notably, since the testing set is not publicly available, following Zhou et al. (2019), the original validation set is split into the *ae-val* subset with 2,460 videos for validation and the *ae-test* subset with 2,457 videos for testing.

3) *MSVD* consists of 1970 video clips collected from YouTube and each clip is spanning over 10 to 25 seconds, with 40 sentences as annotations. Following Chen and Jiang (2021); Ye et al. (2022), we split the dataset into 1,200, 100, and 670 video clips for training, validation and testing, respectively.

Data Preprocessing. For fair comparisons, we try to keep consistent with the previous works. For YouCookII and ActivityNet Captions datasets, we use the features extracted provided by Zhou et al. (2018b) as video features and motion features to represent videos. Specifically, firstly the video is down-sampled to extract frame every 0.5s, then the video feature is extracted from the 'Flatten-673' layer in ResNet-200 (He et al., 2016) with the dimension of 2048 and motion feature is extracted from the 'global pool' layer in BN-Inception (Ioffe and Szegedy, 2015) with the dimension of 1024. Both networks are pre-trained on ActivityNet dataset (Heilbron et al., 2015) for the action recognition task. Following previous works (Ye et al., 2022; Zhang et al., 2020), we utilize InceptionResNetV2 (Szegedy et al., 2017) and C3D (Tran et al., 2015) to extract 1536-dim video features and 2048-dim motion features, respectively. The two models are pretrained on ImageNet dataset (Russakovsky et al., 2015) and Sports 1M dataset Karpathy

³Notably, we do not use the annotated objects for model training, but use the pseudo category label ℓ_d generated by the Faster R-CNN detector as the ground-truth label to enforce the network to maintain the original encoded semantic information of detector. The reason is that the ground-truth of video captioning does not explicitly provide the categories of items in the scene and there is a one-to-one correspondence between pseudo categories and extracted features.

et al. (2014) separately. And we use 16-frame clips with 8-frame overlap as input to C3D. For all the three datasets, we uniformly use Faster R-CNN (Ren et al., 2017) with ResNet-101 as the backbone to extract the object representation as detection features with the dimension of 2048. Considering the trade-off between accuracy and complexity, we give preference to 5 detection features with highest confidence and different classes per frame. As we predefine that at most 100 frames are sampled from one video clips, *i.e.*, $N_v = N_a = 102, N_d = 502$ with the two special tokens [CLS] and [SEP]. Meanwhile, we exploit the first 20 words in the caption sentences and compute the 300-dim GloVe features, *i.e.*, $N_t = 22$ with the two special tokens [BOS] and [EOS]. It is worth noting that, we selected 20 as the truncation number firstly to be consistent with the previous method. Concurrently, we also conducted a statistical analysis on the distribution of caption lengths across the three datasets. The results indicate that over 90% of captions fall within this range, and even more than 97% of YouCookII satisfies this criterion. Therefore, this choice is a reasonable decision. For the COOT features, we concatenate the local clip-level (384-dim) and the global video-level (768-dim) features to describe the videos. After the fully-connected layer, all the tokens are converted into the 768-dim features.

Evaluation metrics. Similar to former works (Park et al., 2019; Lei et al., 2020; Zhu and Yang, 2020), we use several standard metrics to evaluate our method, including BLEU@ n (B@ n) (Papineni et al., 2002) for n -gram precision, METEOR (M) (Denkowski and Lavie, 2014) for n -gram with synonym matching, CIDEr-D (C) (Vedantam et al., 2015) for consensus measurement, Rouge(R) (Lin, 2004) for longest subsequence similarities and Repetition@4 (R@4) (Xiong et al., 2018; Park et al., 2019) for n repetition in the description. Notably, two evaluation modes are considered, *i.e.*, micro-level and paragraph-level. The micro-level evaluation reports the average score on all video sequences separately; while the paragraph-level evaluation first concatenates the caption sentences of all video sequences and then computes the scores averaged across all videos based on the ground-truth paragraph caption sentences, which aims to preserve the story flow with coherence and conciseness. Following existing works (Xiong et al., 2018; Lei et al., 2020), we apply the paragraph-level evaluation to YouCookII and ActivityNet Captions datasets as they annotate multiple clips in each videos. In both modes, CIDEr is used as the primary metric for ranking.

4.2. Implementation Details

Our COST algorithm is implemented using PyTorch. All the experiments are conducted on a machine with 2 NVIDIA RTX-3090 GPUs. We train the model using the strategies similar to BERT (Devlin et al., 2019). Specifically, we use Adam (Kingma and Ba, 2015) with an initial learning rate of $1e - 4$, $\beta_1=0.9$, $\beta_2=0.999$, L^2 weight decay of 0.01, and the learning rate warmup over the first 2 epochs. We train the model at most 20 epochs with early-stop strategy based on CIDEr-D and the batch size is set to 64. For each branch of transformers, we set the dimension of the feature embeddings $d = 768$, the number of transformer blocks $S = 2$, and the number of attention heads

$h = 12$. The loss weights λ_q and λ_d in equation 7 are set to 2.0 and 0.02 which are verified by multiple experiments. It is noted that text tokens are all initialized as [PAD] at the time of evaluation and the description is generated word by word: one word token is generated at a time and substitutes for [PAD] in the corresponding position, and then sent to the model with previously generated tokens to predict the next one until [EOS] token is generated or the length of sentence reaches the maximum.

4.3. Evaluation Results

We compare the proposed COST method with the state-of-the-art methods on the three challenging datasets, *i.e.*, YouCookII, ActivityNet Captions and MSVD. As shown in Table 1, our method obtains the competitive results on both the YouCookII *val* subset and the ActivityNet Captions *ae-test* subset. Without using the COOT features, our method approaches the state-of-art work, *i.e.*, VLTinT (Yamazaki et al., 2022) on the YouCookII *val* subset. Beside using 3D-CNN network and a human detector to extract global and local visual features, VLTinT utilizes Language-Image Pre-training (CLIP) (Radford et al., 2021) to obtain additional linguistic features. And compare our work with other works which don't introduce additional linguistic features, it can be observed that our method improves near 10% CIDEr score compared to the second best method, *i.e.*, MART (Lei et al., 2020). This is attributed to the fact that we use the multi-branch structure to process visual information at different granularities and make them complementary to each other through the proposed cross-granularity attention module, which make it easier for the model to understand the video content. Besides, our method exploits the local appearance information from the Detection-Text transformer for more accuracy caption generation. We also give experiment results in ablation part to verify their effects respectively. Using the COOT features, the overall video captioning results are significantly improved, which demonstrates that the video features extracted by different pre-trained models have a great impact on the performance of the captioning model, and our COST method also performs favorably against other algorithms by improving over 3% CIDEr score. We observe that the similar trend appears in the ActivityNet Captions *ae-test* subset. It is worth noting that PDVC (Wang et al., 2021) performs better on ActivityNet but relatively poorly on YouCookII, probably due to the different characteristics of the two datasets: each video of ActivityNet lasts 120s with 3.65 temporally-localized sentences on average while the duration of each video in YouCookII is 320s with 7.7 annotated segments and associated sentences. However, COST has a balanced and satisfactory performance on both datasets, which further proves the effectiveness of our method. As presented in Table 2, we also compare the proposed method to some LSTM based methods with input detection features on the ActivityNet Captions *ae-val* subset. Note that Table 1 and 2 don't list same number of compared methods is because these LSTM-based methods had only evaluated on the *ae-val* subset of Activity Captions and provided the corresponding results. Compared to AdvInf (Park et al., 2019), which also inputs multi-modal features (concatenation of image recognition, action recognition and object detection features) and designs

Table 1. Experimental Results on the YouCookII *Val* Subset and ActivityNet Captions *Ae-Test* Subset in the Paragraph-Level Evaluation mode. *COOT* Indicates That the Evaluated Methods Use the Feature Extracted by *COOT* (Ging et al., 2020) Pre-Trained on HowTo100M (Miech et al., 2019). * indicates that VLTinT utilizes extra linguistic features obtained through CLIP (Radford et al., 2021). We Report BLEU@4 (B@4), METEOR (M), CIDEr-D (C) and Repetition@4 (R@4).

Method	<i>COOT</i>	YouCookII (<i>val</i>)				ActivityNet Captions (<i>ae-test</i>)			
		B@4	M	C	R@4 ↓	B@4	M	C	R@4 ↓
Vanilla Transformer (Zhou et al., 2018b)	✗	7.62	15.65	32.26	7.83	9.31	15.54	21.33	7.45
Transformer-XL (Dai et al., 2019)	✗	6.56	14.76	26.35	6.30	10.25	14.91	21.71	8.79
Transformer-XLRG (Lei et al., 2020)	✗	6.63	14.74	25.93	6.03	10.07	14.58	20.34	9.37
MART (Lei et al., 2020)	✗	8.00	15.90	35.74	4.39	9.78	15.57	22.16	5.44
PDVC (Wang et al., 2021)	✗	7.11	15.05	26.03	7.65	11.36	15.73	25.03	10.92
VLTinT* (Yamazaki et al., 2022)	✓	9.40	17.94	48.70	4.29	14.50	17.97	31.13	4.75
COST	✗	9.47	17.67	45.54	4.04	11.14	15.91	24.77	5.86
Vanilla Transformer (Zhou et al., 2018b)	✓	11.05	19.79	55.57	5.69	10.47	15.76	25.90	19.14
Transformer-XL (Dai et al., 2019)	✓	-	-	-	-	10.57	14.76	22.04	15.85
MART (Lei et al., 2020)	✓	11.30	19.85	57.24	6.69	10.85	15.99	28.19	6.64
SART (Man et al., 2022)	✓	11.43	19.91	57.66	8.58	11.35	16.21	28.35	7.18
COST	✓	11.56	19.67	60.78	6.63	11.88	15.70	29.64	6.11

Table 2. Comparison With the State-of-the-Art Methods on ActivityNet Captions *Ae-Val* Subset in the Paragraph-Level Evaluation Mode. *Det.* And *Re.* Indicate Whether the Model Uses Detection Features And the Sentence-Level Recurrence. * indicates that VLTinT utilizes extra linguistic features obtained through CLIP (Radford et al., 2021).

	<i>Det.</i>	<i>Re.</i>	B@4	M	C	R@4 ↓
LSTM based methods:						
MFT (Xiong et al., 2018)	✗	✓	10.29	14.73	19.12	17.71
HSE (Zhang et al., 2018)	✗	✓	9.84	13.78	18.78	13.22
LSTM based methods with detection feature:						
GVD (Zhou et al., 2019)	✓	✗	11.04	15.71	21.95	8.76
GVDsup (Zhou et al., 2019)	✓	✗	11.30	16.41	22.94	7.04
AdvInf (Park et al., 2019)	✓	✓	10.04	16.60	20.97	5.76
Transformer based methods:						
Vanilla Transformer (Zhou et al., 2018b)	✗	✗	9.75	15.64	22.16	7.79
Transformer-XL (Dai et al., 2019)	✗	✓	10.39	15.09	21.67	8.54
Transformer-XLRG (Lei et al., 2020)	✗	✓	10.17	14.77	20.40	8.85
MART (Lei et al., 2020)	✗	✓	10.33	15.68	23.42	5.18
PDVC (Wang et al., 2021)	✗	✗	11.80	15.93	27.27	10.68
SART (Man et al., 2022)	✗	✓	11.35	16.21	28.35	7.18
VLTinT* (Yamazaki et al., 2022)	✗	✗	14.93	18.16	33.07	4.87
COST	✓	✓	11.22	16.58	25.70	7.09

three discriminators to enhance the fluency and relevance in the generated captions, our method produces higher scores for both B@4 and CIDEr, demonstrating the superiority of the collaborative transformers to learn multi-modal representations over LSTM. Meanwhile, MART (Lei et al., 2020) without input detections performs inferior than our method in terms of B@4, M and C metrics, even though it utilizes complex memory module to store video-text features in history to generate more coherent and accurate captions. It indicates that detection features are important entity for video captioning task which is neglected by recent transform-based approaches, and aligning the interactions between these visual features of different granularities can contribute greatly to generating accurate captions. PDVC (Wang et al., 2021) adds multiple temporal convolutional layers to obtain feature sequences across multiple resolutions as input and performs better than our method in terms of B@4 and C metrics, but the sentences generated by it have significant redundancy (with high R@4 metric) due to complex input. It is worth noting that both SART (Man et al., 2022) and VLTinT

Table 3. Comparison With State-of-the-Art Methods on YouCookII *Val* Subset in the Micro-Level Evaluation Mode. * indicates the evaluated method had been pre-trained on large-scale datasets and † indicates the evaluate method using multi-modal features extracted by *COOT* (Ging et al., 2020).

Method	B@3	B@4	M	R	C
Masked Trans. (Zhou et al., 2018b)	7.53	3.84	11.55	27.44	0.38
S3D (Xie et al., 2017)	6.12	3.24	9.52	26.09	0.31
VideoBERT* (Sun et al., 2019)	6.80	4.04	11.01	27.50	0.49
VideoBERT+S3D* (Sun et al., 2019)	7.59	4.33	11.94	28.80	0.50
ActBERT* (Zhu and Yang, 2020)	8.66	5.41	14.30	30.56	0.65
SwinBERT* (Lin et al., 2021)	13.80	9.00	15.60	37.30	1.09
COST	10.69	6.63	12.61	31.09	0.71
COST†	13.50	8.65	15.62	36.53	1.05

(Yamazaki et al., 2022) both performs well on this subset while based on different solutions: the former one proposes a scenario understanding module to enhance the corresponding ability of their model while the latter captures coherent semantics by introducing more powerful linguistic features, which provides insights for future exploration in this dataset.

According to Table 3, our method outperforms several BERT based methods which pre-trained on large-scale datasets in terms of the micro-level evaluation. In particular, ActBERT (Zhu and Yang, 2020) pre-trains on HowTo100M (Miech et al., 2019) firstly and then relies on better visual modalities than our method (for example, its extraction network for action feature pre-trained on Kinetics dataset (Kay et al., 2017)). However, it directly focuses on the learning the alignment between text and other visual modalities, which is difficult to exploit the discriminative semantic information. In contrast, our cross-granularity attention module learns from three visual-linguistic interactions of different granularities using the three-stream transformers, producing better CIDEr score. SwinBERT (Lin et al., 2021) is proposed as a pure Transformer-based end-to-end architecture for video captioning, which leverages VidSwin (Liu et al., 2021b) pre-trained on Kinetics-600 to extract spatial-temporal video representations and then proposes Multi-modal Trans-

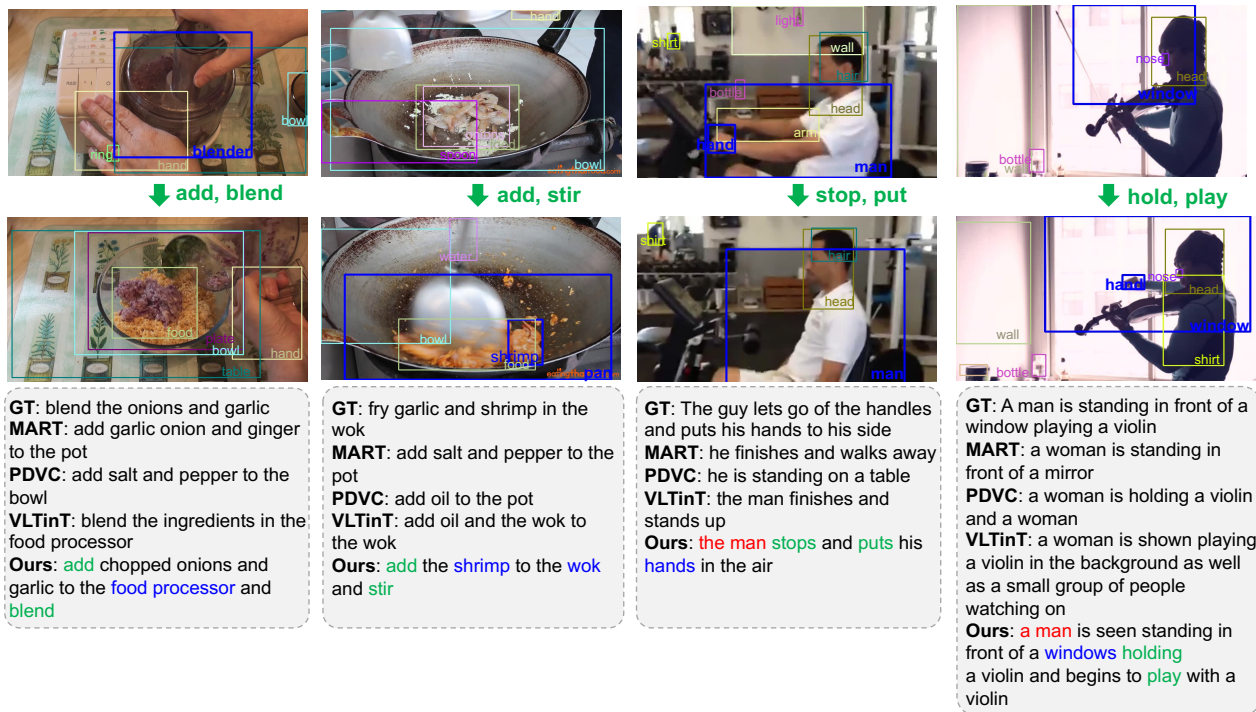


Fig. 4. Qualitative comparison with MART (Lei et al., 2020), PDVC (Wang et al., 2021) and VLTinT (Yamazaki et al., 2022) on YouCookII *val* and ActivityNet Captions *ae-val* split. The Subject, Predicate, and Object in a sentence are highlighted in the red, green and blue fonts, respectively. Best viewed in color.

Table 4. Comparison With State-of-the-Art Methods on MSVD Test Subset in the Micro-Level Evaluation Mode.

Method	B	M	R	C
OA-BTG (Zhang and Peng, 2019)	56.9	36.2	-	90.6
MSGa (Chen and Jiang, 2019)	52.5	34.1	71.3	88.7
ORG-TRL (Zhang et al., 2020)	54.3	36.4	73.9	95.2
JCRR (Hou et al., 2020)	57.0	36.8	-	96.8
SGN (Ryu et al., 2021)	52.8	35.5	72.9	94.3
MGRMP (Chen and Jiang, 2021)	55.8	<u>36.9</u>	74.5	<u>98.5</u>
BFSD (Zhong et al., 2022)	51.2	35.7	72.9	96.7
COST	56.8	37.2	<u>74.3</u>	99.2

former Encoder takes them as input to generate a natural sentence, and achieves significant improvements compared to the previous methods. In order to make a relatively fair comparison as we had not pre-trained our model on large-scale datasets, such as HowTo100M or Kinetics, we use the visual features provided by COOT (Ging et al., 2020) as input, whose video embedding network pre-trained on HowTo100M, and obtain results comparable to SwinBERT, which demonstrates the potential of our method.

We also compare our methods with the state-of-art methods on MSVD dataset, and the results are presented in Table 4. As shown in Table 4, our COST could approach or surpass the state-of-art method. Considering the significant differences between the characteristics of this dataset and the former two datasets, for example, the duration of video and the number of annotations corresponding to each video, which further proves the robustness of our method.

Furthermore, the qualitative results of MART (Lei et al., 2020), PDVC (Wang et al., 2021), VLTinT (Yamazaki et al., 2022) and our COST are shown in Fig. 4. From Fig. 4, it is evident that our COST approach outperforms the compared methods in terms of generating accurate and concise captions, regardless of the complexity of the scenarios. This is attributed to two reasons. First, using the Detection-Text transformer, the Objects in caption sentences can be learned explicitly (e.g., *shrimp*, *hand*, and *windows*) or implicitly (e.g., from *bowl* to *food processor*, and from *pan* to *wok*). Second, the Action-Text transformer in our method can perceive the key actions in caption sentences such as *add*, *blend*, *put* and *hold*. In comparison, the existing methods often fail to recognize these crucial subjects, objects, and verbs. The results indicate that under the constraint of our objective, these two branches of transformers can enforce the network to learn the key elements and supply visual-linguistic interactions with discriminative semantics for captions generation.

Lastly, as there exists many videos shot by moving cameras in our evaluated datasets, for example, YouCookII contains 2000 long untrimmed videos describing 89 cooking recipes, and many videos are shot with varying perspectives for providing good viewing experience, and we wonder whether the performance of our model decreases significantly under such moving camera situation. We have randomly chosen several video clips shot by apparent moving cameras without cherry pick and evaluate our COST on it, and the results are shown on 5. It can be observed that there is no significant decrease in evaluation metrics, which further confirms the robustness of our method.

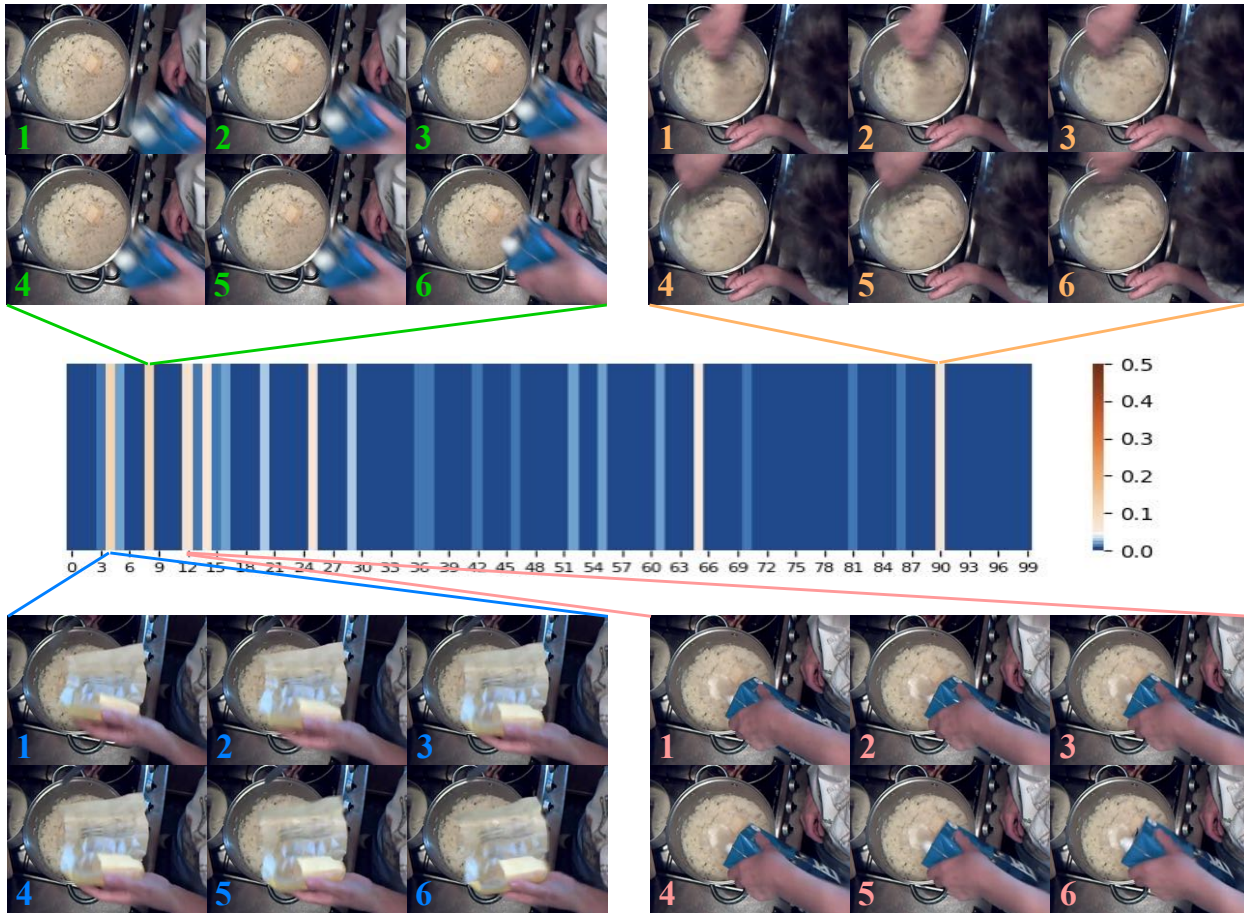


Fig. 5. Visualization of attention score map of all action tokens to $[CLS^a]$ token in Action-Text transformer for given video clips with ground-truth caption "add butter and milk to the mashed potatoes and mix". More crimson color indicates high attention score. We also show the video segments corresponding to the top four attention values. The number in the lower left corner of each picture indicates the time order. Best viewed in color.

Table 5. Evaluation Results on Moving Scenes in YouCookII. * indicates the subset picked from YouCookII shot with apparent camera movements.

Dataset	B@4	M	C	R@4 ↓
YouCookII*	9.21	17.43	43.61	5.62
YouCookII	9.47	17.67	45.54	4.04

4.4. Visualization of Self-Attention in Action-Text Transformer

In Fig. 5, the central heatmap shows the visualization of the attention scores in self-attention between $[CLS^a]$ and all action tokens in Action-Text transformer. We obtain the scores by averaging attention scores from all heads of the multi-head self-attention module in our last transformer block. We choose this branch because its granularity token $[CLS^a]$ is used as the input for action classification, and whether it can utilize action tokens effectively partially reflects the ability of our model to understand the context of video. The ground-truth caption of chosen video is "add butter and milk to the mashed potatoes and mix" and the segments corresponding to top four attention values are presented, where each one consists of 6 figures as setting in motion feature extraction network. According to the chronological order of the action, we can see that the four video segments are: finishing adding butter, readying to pour milk,

pouring the milk and stirring the mixture, which proves that our Action-Text Transformer can capture the key actions effectively via self-attention although there exists significant redundancy in the action tokens (vast majority of this video's content is mixing the mixture).

4.5. Visualization of Cross-Granularity Attention in Video-Text Transformer

To better understand the cross-granularity attention module, we use the heatmap to visualize the cross-granularity attention scores of the Video-Text transformer in Fig. 6. For the convenience of display, we concatenate the affinity matrix $\mathcal{M}_{\mathcal{H},\mathcal{Y}}$ and $\mathcal{M}_{\mathcal{H},\mathcal{X}}$ to get $\mathcal{M}_{\mathcal{H}}$. At epoch 1, all values in the heatmap are similar and the maximal value is only 0.001. The corresponding caption results are noisy with the false predictions of nouns and verbs, *e.g.*, *pan* instead of *butter*. It indicates that the cross-granularity attention is randomly initialized and interactions between different modalities are not learned to align well. After training for several epochs, a few entities dominate the heatmap with the maximal value of 0.3 (see the bright vertical lines in the heatmap). It indicates that our cross-granularity attention module can successfully exploit the most relevant entities from other modalities to inject the information from other




Fig. 6. Heatmap used to indicate the affinity matrix M_{VT} in the Video-Text transformer, where the row denotes the video tokens and the column denotes the action and detection tokens. The false predictions of nouns and verbs are denoted in red font. For clarity, we only show a few epochs in the training phase. Best viewed in color and zoom in.

Table 6. Ablation on multi-modality features and three-stream architecture.

COST Variants	YouCookII (<i>val</i>)			
	B@4	M	C	R@4 ↓
COST-1 (v)	6.59	14.29	29.20	6.49
COST-1 (v+a)	7.72	15.45	33.98	4.60
COST-1 (v+a+d)	8.73	16.90	38.63	4.66
COST-2 (v+a)	7.79	15.74	34.99	5.82
COST-2 (v+d)	9.04	17.31	43.09	4.59
COST-3 (v+a+d)	9.47	17.67	45.54	4.04

branches of transformers. This way, the verb *stir* and the noun *pan* can be corrected to *mix* and *butter* for more accurate caption results.

4.6. Ablation Studies

To study the influence of different components in the proposed method, we conduct detailed ablation study on the YouCookII *val* subset. Notably, we use the appearance features extracted by TSN (Wang et al., 2016) in all COST- k variants, where k denotes the number of transformer branches retained in our COST method.

Effectiveness of multi-modality features. To verify the effectiveness of the multi-modality features, we construct three COST-1 variants, *i.e.*, COST-1 (v), COST-1 (v+a) and COST-1 (v+a+d). In particular, we only use the Video-Text transformer, but change the input features as the combinations of the GloVe text features and the concatenated features from video (v), action (a) and detection (d). As shown in Table 6, the scores under all metrics are improved considerably by integrating the action or detection features. Moreover, the CIDEr score is boosted


Table 7. Ablation on training objectives.

COST Variants	YouCookII (<i>val</i>)			
	B@4	M	C	R@4 ↓
COST w/o \mathcal{L}_a	8.95	17.31	42.90	5.27
COST w/o \mathcal{L}_d	9.10	16.58	41.00	5.56
COST w/o \mathcal{L}_d and \mathcal{L}_a	9.01	16.29	40.58	5.91
COST	9.47	17.67	45.54	4.04


from 29.20% to 38.63%, if we include all the three-modality features. It indicates that multi-modality features definitely facilitate to generate more accurate video captions. And it can be observed that the incorporation of detection features into COST has yielded notable performance improvements, highlighting the potential for leveraging more detailed features to advance the field of video captioning.

Effectiveness of three-stream transformers. To demonstrate the effectiveness of the three-stream transformers compared to the simple feature concatenation used in COST-1, we construct three COST- k ($k = 2, 3$) variants, shown in Table 6. It can be seen that the accuracy can be improved by using the three-stream transformers, *i.e.*, the CIDEr score improved from 38.63% to 45.54%. Meanwhile, the accuracy is considerably improved by using more branches of transformers. This is because our cross-granularity attention module is able to capture the most relevant semantic information from different visual-linguistic interactions. Compared to the COST-1 variants, our full model can obtain a more discriminative feature representations for video captioning.

Effectiveness of training objective. To investigate the effectiveness and rationality of proposed training objective, we do



MART: add the dough to the center of the dough
Ours: put the mixture in the bowl and coat them well with the flour



MART: remove the fat from the water and place on a paper towel
Ours: add rice vinegar and rice to the seaweed and roll the seaweed

Fig. 7. Generated Descriptions for Internet videos from the state-of-art MART (Lei et al., 2020) method and our COST method. The blue number in the lower left corner of each picture indicates the time order.

Table 8. Ablation on cross-granularity attention module. *concat* and *para* indicate concatenated and parallel cross-granularity attention module separately. COST w/o cgam. denotes a COST variant without cross-granularity attention module.

COST Variants	YouCookII (<i>val</i>)			
	B@4	M	C	R@4 ↓
COST w/ <i>concat</i>	9.27	17.40	43.43	6.34
COST w/ <i>para</i>	9.47	17.67	45.54	4.04
COST w/o cgam.	7.55	15.11	30.32	8.03

ablation experiments to obtain corresponding results in Table 7. We observe that removing either \mathcal{L}_a or \mathcal{L}_d has a negative impact on the performance, especially when they are all removed, *i.e.*, 40.58 *vs.* 45.54 for CIDEr. This demonstrates that using appropriate objective to make the semantics of visual-linguistic interactions explicit can definitely contribute to accurate captions.

Effectiveness of cross-granularity attention module. To verify the importance of cross-granularity attention module to the success of our method, we compare the performance of COST with the two proposed architectures(concatenated and parallel architectures) separately and the variant with cross-granularity attention module removed. As shown in Table 8, our method with cross-granularity module using parallel structure outperforms that with concatenated structure in all evaluation metrics, it is probably because that there exists feature gap between the embeddings from different branches and the softmax operation may hinder the information of one branch being suppressed during fusion. And after removing the cross-granularity

module, CIDEr drops significantly, *i.e.*, 30.32 *vs.* 45.54, which demonstrate the effectiveness of this module to align the visual-linguistic interactions of different granularities for accurate captions.

5. Generalization

In order to verify the generalization and practicality of our method, we collected some videos outside the trained datasets and evaluate the quality of generated descriptions to them. Because the video included in YouCookII dataset has relatively clear action intention and contains abundant clips recording the interactions between human and surrounding objects, we collected some cooking videos from the Internet and applied model trained on YouCookII to them. The way to extract features from collected videos keeps consistent with that to YouCookII, and we adopt trained models on both short and long video clips to test the performance comprehensively. We list the captions generated by MART (Lei et al., 2020) and our method in Fig. 7. The first clip is one step of making fried milk, which is a dessert of Cantonese cuisine, and lasts about 20 seconds to coat the fried milk with breadcrumbs. It can be seen that our method not only makes a correct judgment on the action(coat), but also correctly recognizes the state of using substances(both breadcrumbs and flour are powdery objects), which is much better than MART. The second clip is one step of making rice ball and lasts about 1 minute to add rice, cooked beef and asparagus to seaweed and wrap them. Because there are many added contents, our model makes wrong recognition (the beef is recognized as vinegar because of the similar

color) and generates redundant words (rice are added twice in the clip). Nevertheless, the generated description is far more accurate in action and much natural than the description generated by MART. Through these two examples, it illustrates that although our model may make mistakes in generating object words due to the interference of similar shape or color, it can make accurate recognition on the actions of the subject and the structure of generated sentence is very natural, which indicates the effectiveness and practical value of our model.

6. Limitation

Although our model has demonstrated satisfactory performance on three benchmark datasets and exhibits strong generalization capabilities, there still exists some limitations. Firstly, in terms of evaluation efficiency and model parameters, we observe certain disadvantages when compared to single-stream methods like MART. Specially, under the same hardware conditions, our model generates captions for approximately 2.47 videos per second, and the duration of each videos is about 5 minutes, whereas MART achieves a higher processing rate of 6.85 videos per second. Furthermore, our model has a larger parameter count of 35.6M compared to MART's 25.5M. Additionally, it is worth noting that our model requires additional data processing steps for extracting region features, which requires additional pre-processing time.

7. Conclusion

In this paper, we propose the collaborative three-stream transformers to exploit the interactions of objects, and the actions/relations of objects between different modalities of different granularities in spatial-temporal domain. Meanwhile, with the help of proposed training objective to specify the semantics of visual-linguistic interactions in each branch, the cross-granularity attention module is designed to align the interactions modeled by the three branches of transformers, which contributes to more accurate captions generation. Several experiments conducted on the YouCookII, ActivityNet Captions and MSVD datasets demonstrate the effectiveness of our proposed method, and we also evaluate its generalization ability through testing on Internet videos. In the future, we intent to further improve the performance of our model with the help of pre-training on large-scale datasets.

Acknowledgments

This work was supported by the Key Research Program of Frontier Sciences, CAS, Grant No. ZDBS-LY-JSC038 for funding this work. Libo Zhang was also supported by Youth Innovation Promotion Association, CAS (2020111). Heng Fan and his employer received no financial support for the research, authorship, and/or publication of this article.

References

- Abu-El-Hajja, S., Kothari, N., Lee, J., Natsev, P., Toderici, G., Varadarajan, B., Vijayanarasimhan, S., 2016. Youtube-8m: A large-scale video classification benchmark. *CoRR abs/1609.08675*.
- Arnab, A., Deghani, M., Heigold, G., Sun, C., Lucic, M., Schmid, C., 2021. Vivit: A video vision transformer, in: *ICCV, IEEE*. pp. 6816–6826.
- Chen, D.L., Dolan, W.B., 2011. Collecting highly parallel data for paraphrase evaluation, in: *ACL, The Association for Computer Linguistics*. pp. 190–200.
- Chen, S., Jiang, Y., 2019. Motion guided spatial attention for video captioning, in: *AAAI*, pp. 8191–8198.
- Chen, S., Jiang, Y., 2021. Motion guided region message passing for video captioning, in: *ICCV, IEEE*. pp. 1523–1532.
- Chen, S., Yao, T., Jiang, Y., 2019. Deep learning for video captioning: A review, in: *IJCAI*, pp. 6283–6290.
- Chung, J., Gülçehre, Ç., Cho, K., Bengio, Y., 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *CoRR abs/1412.3555*.
- Dai, Z., Yang, Z., Yang, Y., Carbonell, J.G., Le, Q.V., Salakhutdinov, R., 2019. Transformer-xl: Attentive language models beyond a fixed-length context, in: *ACL*, pp. 2978–2988.
- Das, P., Xu, C., Doell, R.F., Corso, J.J., 2013. A thousand frames in just a few words: Lingual description of videos through latent topics and sparse object stitching, in: *CVPR*, pp. 2634–2641.
- Denkowski, M.J., Lavie, A., 2014. Meteor universal: Language specific translation evaluation for any target language, in: *WMT@ACL*, pp. 376–380.
- Devlin, J., Chang, M., Lee, K., Toutanova, K., 2019. BERT: pre-training of deep bidirectional transformers for language understanding, in: *NAACL-HLT*, pp. 4171–4186.
- Fan, H., Yang, Y., 2020. Person tube retrieval via language description, in: *AAAI, AAAI Press*. pp. 10754–10761.
- Gao, L., Li, X., Song, J., Shen, H.T., 2020. Hierarchical lstms with adaptive attention for visual captioning. *IEEE Trans. Pattern Anal. Mach. Intell.* 42, 1112–1131.
- Ging, S., Zolfaghari, M., Pirsiavash, H., Brox, T., 2020. COOT: cooperative hierarchical transformer for video-text representation learning, in: *NeurIPS*.
- Guadarrama, S., Krishnamoorthy, N., Malkarnenkar, G., Venugopalan, S., Mooney, R.J., Darrell, T., Saenko, K., 2013. Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition, in: *ICCV*, pp. 2712–2719.
- Hao, W., Zhang, Z., Guan, H., 2018. Integrating both visual and audio cues for enhanced video caption, in: *AAAI*, pp. 6894–6901.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: *CVPR, IEEE Computer Society*. pp. 770–778.
- Heilbron, F.C., Escorcia, V., Ghanem, B., Niebles, J.C., 2015. Activitynet: A large-scale video benchmark for human activity understanding, in: *CVPR, IEEE Computer Society*. pp. 961–970.
- Hori, C., Hori, T., Wichern, G., Wang, J., Lee, T., Cherian, A., Marks, T.K., 2018. Multimodal attention for fusion of audio and spatiotemporal features for video description, in: *CVPRW*, pp. 2528–2531.
- Hou, J., Wu, X., Zhang, X., Qi, Y., Jia, Y., Luo, J., 2020. Joint commonsense and relation reasoning for image and video captioning, in: *AAAI, AAAI Press*. pp. 10973–10980.
- Huang, G., Pang, B., Zhu, Z., Rivera, C., Soricut, R., 2020. Multimodal pre-training for dense video captioning, in: *ACL/IJCNLP*, pp. 470–490.
- Iashin, V., Rahtu, E., 2020. Multi-modal dense video captioning, in: *CVPRW*, pp. 4117–4126.
- Ioffe, S., Szegedy, C., 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift, in: *ICML, JMLR.org*. pp. 448–456.
- Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L., 2014. Large-scale video classification with convolutional neural networks, in: *CVPR, IEEE Computer Society*. pp. 1725–1732.
- Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., Suleyman, M., Zisserman, A., 2017. The kinetics human action video dataset. *CoRR abs/1705.06950*.
- Kim, D., Oh, T.H., Choi, J., Kweon, I.S., 2020. Dense relational image captioning via multi-task triple-stream networks. *CoRR abs/2010.03855*.
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization, in: *ICLR*.
- Krishna, R., Hata, K., Ren, F., Fei-Fei, L., Niebles, J.C., 2017a. Dense-captioning events in videos, in: *ICCV*, pp. 706–715.

- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L., Shamma, D.A., Bernstein, M.S., Fei-Fei, L., 2017b. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV* 123, 32–73.
- Lei, J., Wang, L., Shen, Y., Yu, D., Berg, T.L., Bansal, M., 2020. MART: memory-augmented recurrent transformer for coherent video paragraph captioning, in: *ACL*, pp. 2603–2614.
- Li, J., Niu, L., Zhang, L., 2022. Action-aware embedding enhancement for image-text retrieval, in: *AAAI*, AAAI Press. pp. 1323–1331.
- Li, X., Zhao, B., Lu, X., 2017. MAM-RNN: multi-level attention model based RNN for video captioning, in: *IJCAI*, ijcai.org. pp. 2208–2214.
- Lin, C.Y., 2004. Rouge: A package for automatic evaluation of summaries, in: *Text summarization branches out*, pp. 74–81.
- Lin, K., Li, L., Lin, C., Ahmed, F., Gan, Z., Liu, Z., Lu, Y., Wang, L., 2021. Swinbert: End-to-end transformers with sparse attention for video captioning. *CoRR* abs/2111.13196.
- Liu, S., Ren, Z., Yuan, J., 2021a. Sibnet: Sibling convolutional encoder for video captioning. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 3259–3272.
- Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H., 2021b. Video swin transformer. *CoRR* abs/2106.13230.
- Lu, J., Batra, D., Parikh, D., Lee, S., 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks, in: *NeurIPS*, pp. 13–23.
- Luo, H., Ji, L., Shi, B., Huang, H., Duan, N., Li, T., Chen, X., Zhou, M., 2020. Univilm: A unified video and language pre-training model for multimodal understanding and generation. *CoRR* abs/2002.06353.
- Man, X., Ouyang, D., Li, X., Song, J., Shao, J., 2022. Scenario-aware recurrent transformer for goal-directed video captioning. *ACM Trans. Multim. Comput. Commun. Appl.* 18, 104:1–104:17.
- Miech, A., Zhukov, D., Alayrac, J., Tapaswi, M., Laptev, I., Sivic, J., 2019. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips, in: *ICCV*, pp. 2630–2640.
- Nan, G., Qiao, R., Xiao, Y., Liu, J., Leng, S., Zhang, H., Lu, W., 2021. Interventional video grounding with dual contrastive learning, in: *CVPR*, Computer Vision Foundation / IEEE. pp. 2765–2775.
- Pan, Y., Yao, T., Li, H., Mei, T., 2017. Video captioning with transferred semantic attributes, in: *CVPR*, pp. 984–992.
- Papineni, K., Roukos, S., Ward, T., Zhu, W., 2002. Bleu: a method for automatic evaluation of machine translation, in: *ACL*, pp. 311–318.
- Park, J.S., Rohrbach, M., Darrell, T., Rohrbach, A., 2019. Adversarial inference for multi-sentence video description, in: *CVPR*, pp. 6598–6608.
- Pei, W., Zhang, J., Wang, X., Ke, L., Shen, X., Tai, Y., 2019. Memory-attended recurrent network for video captioning, in: *CVPR*, pp. 8347–8356.
- Pennington, J., Socher, R., Manning, C.D., 2014. Glove: Global vectors for word representation, in: *EMNLP*, pp. 1532–1543.
- 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., 2021. Learning transferable visual models from natural language supervision, in: *ICML*, PMLR. pp. 8748–8763.
- Ren, S., He, K., Girshick, R.B., Sun, J., 2017. Faster R-CNN: towards real-time object detection with region proposal networks. *TPAMI* 39, 1137–1149.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., Fei-Fei, L., 2015. Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* 115, 211–252.
- Ryu, H., Kang, S., Kang, H., Yoo, C.D., 2021. Semantic grouping network for video captioning, in: *AAAI*, AAAI Press. pp. 2514–2522.
- Seo, P.H., Nagrani, A., Arnab, A., Schmid, C., 2022. End-to-end generative pretraining for multimodal video captioning. *CoRR* abs/2201.08264.
- Sun, C., Myers, A., Vondrick, C., Murphy, K., Schmid, C., 2019. Videobert: A joint model for video and language representation learning, in: *ICCV*, pp. 7463–7472.
- Sun, Y., Cheng, C., Zhang, Y., Zhang, C., Zheng, L., Wang, Z., Wei, Y., 2020. Circle loss: A unified perspective of pair similarity optimization, in: *CVPR*, pp. 6397–6406.
- Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to sequence learning with neural networks, in: *NeurIPS*, pp. 3104–3112.
- Szegedy, C., Ioffe, S., Vanhoucke, V., Alemi, A.A., 2017. Inception-v4, inception-resnet and the impact of residual connections on learning, in: *AAAI*, AAAI Press. pp. 4278–4284.
- Tang, Z., Lei, J., Bansal, M., 2021. Decembert: Learning from noisy instructional videos via dense captions and entropy minimization, in: *NAACL-HLT*, pp. 2415–2426.
- Tran, D., Bourdev, L.D., Fergus, R., Torresani, L., Paluri, M., 2015. Learning spatiotemporal features with 3d convolutional networks, in: *ICCV*, IEEE Computer Society. pp. 4489–4497.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need, in: *NeurIPS*, pp. 5998–6008.
- Vedantam, R., Zitnick, C.L., Parikh, D., 2015. Cider: Consensus-based image description evaluation, in: *CVPR*, pp. 4566–4575.
- Venugopalan, S., Rohrbach, M., Donahue, J., Mooney, R.J., Darrell, T., Saenko, K., 2015a. Sequence to sequence - video to text, in: *ICCV*, pp. 4534–4542.
- Venugopalan, S., Xu, H., Donahue, J., Rohrbach, M., Mooney, R.J., Saenko, K., 2015b. Translating videos to natural language using deep recurrent neural networks, in: *HLT-NAACL*, pp. 1494–1504.
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Gool, L.V., 2016. Temporal segment networks: Towards good practices for deep action recognition, in: *ECCV*, pp. 20–36.
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Gool, L.V., 2019. Temporal segment networks for action recognition in videos. *IEEE Trans. Pattern Anal. Mach. Intell.* 41, 2740–2755.
- Wang, T., Zhang, R., Lu, Z., Zheng, F., Cheng, R., Luo, P., 2021. End-to-end dense video captioning with parallel decoding, in: *ICCV*, IEEE. pp. 6827–6837.
- Wang, X., Wang, Y., Wang, W.Y., 2018. Watch, listen, and describe: Globally and locally aligned cross-modal attentions for video captioning, in: *NAACL-HLT*, pp. 795–801.
- Wei, X., Zhang, T., Li, Y., Zhang, Y., Wu, F., 2020. Multi-modality cross attention network for image and sentence matching, in: *CVPR*, pp. 10938–10947.
- Xie, S., Sun, C., Huang, J., Tu, Z., Murphy, K., 2017. Rethinking spatiotemporal feature learning for video understanding. *CoRR* abs/1712.04851.
- Xiong, Y., Dai, B., Lin, D., 2018. Move forward and tell: A progressive generator of video descriptions, in: *ECCV*, pp. 489–505.
- Xu, N., Liu, A., Wong, Y., Zhang, Y., Nie, W., Su, Y., Kankanhalli, M.S., 2019. Dual-stream recurrent neural network for video captioning. *IEEE Trans. Circuits Syst. Video Technol.* 29, 2482–2493.
- Xu, Y., Liu, X., Cao, X., Huang, C., Liu, E., Qian, S., Liu, X., Wu, Y., Dong, F., Qiu, C.W., et al., 2021. Artificial intelligence: A powerful paradigm for scientific research. *The Innovation* 2, 100179.
- Yamazaki, K., Vo, K., Truong, S., Raj, B., Le, N., 2022. Vltint: Visual-linguistic transformer-in-transformer for coherent video paragraph captioning. *CoRR* abs/2211.15103.
- Yao, L., Torabi, A., Cho, K., Ballas, N., Pal, C.J., Larochelle, H., Courville, A.C., 2015. Describing videos by exploiting temporal structure, in: *ICCV*, pp. 4507–4515.
- Ye, H., Li, G., Qi, Y., Wang, S., Huang, Q., Yang, M., 2022. Hierarchical modular network for video captioning, in: *CVPR*, IEEE. pp. 17918–17927.
- Yu, H., Wang, J., Huang, Z., Yang, Y., Xu, W., 2016. Video paragraph captioning using hierarchical recurrent neural networks, in: *CVPR*, pp. 4584–4593.
- Zhang, B., Hu, H., Sha, F., 2018. Cross-modal and hierarchical modeling of video and text, in: *ECCV*, pp. 385–401.
- Zhang, J., Peng, Y., 2019. Object-aware aggregation with bidirectional temporal graph for video captioning, in: *CVPR*, Computer Vision Foundation / IEEE. pp. 8327–8336.
- Zhang, N., Chen, X., Xie, X., Deng, S., Tan, C., Chen, M., Huang, F., Si, L., Chen, H., 2021. Document-level relation extraction as semantic segmentation, in: *IJCAI*, pp. 3999–4006.
- Zhang, Z., Shi, Y., Yuan, C., Li, B., Wang, P., Hu, W., Zha, Z., 2020. Object relational graph with teacher-recommended learning for video captioning, in: *CVPR*, pp. 13275–13285.
- Zheng, Q., Wang, C., Tao, D., 2020. Syntax-aware action targeting for video captioning, in: *CVPR*, pp. 13093–13102.
- Zhong, X., Li, Z., Chen, S., Jiang, K., Chen, C., Ye, M., 2022. Refined semantic enhancement towards frequency diffusion for video captioning. *CoRR* abs/2211.15076.
- Zhou, L., Kalantidis, Y., Chen, X., Corso, J.J., Rohrbach, M., 2019. Grounded video description, in: *CVPR*, pp. 6578–6587.
- Zhou, L., Xu, C., Corso, J.J., 2018a. Towards automatic learning of procedures from web instructional videos, in: *AAAI*, pp. 7590–7598.
- Zhou, L., Zhou, Y., Corso, J.J., Socher, R., Xiong, C., 2018b. End-to-end dense video captioning with masked transformer, in: *CVPR*, pp. 8739–8748.
- Zhu, L., Yang, Y., 2020. Actbert: Learning global-local video-text representa-

tions, in: CVPR, pp. 8743–8752.