

SWINBERT: End-to-End Transformers with Sparse Attention for Video Captioning

Kevin Lin*, Linjie Li*, Chung-Ching Lin*, Faisal Ahmed, Zhe Gan, Zicheng Liu, Yumao Lu, Lijuan Wang
Microsoft

{keli, lindsey.li, chungching.lin, fiaahmed, zhe.gan, zliu, yumaolu, lijuanw}@microsoft.com

Abstract

The canonical approach to video captioning dictates a caption generation model to learn from offline-extracted dense video features. These feature extractors usually operate on video frames sampled at a fixed frame rate and are often trained on image/video understanding tasks, without adaption to video captioning data. In this work, we present SWINBERT, an end-to-end transformer-based model for video captioning, which takes video frame patches directly as inputs, and outputs a natural language description. Instead of leveraging multiple 2D/3D feature extractors, our method adopts a video transformer to encode spatial-temporal representations that can adapt to variable lengths of video input without dedicated design for different frame rates. Based on this model architecture, we show that video captioning can benefit significantly from more densely sampled video frames as opposed to previous successes with sparsely sampled video frames for video-and-language understanding tasks (e.g., video question answering). Moreover, to avoid the inherent redundancy in consecutive video frames, we propose adaptively learning a sparse attention mask and optimizing it for task-specific performance improvement through better long-range video sequence modeling. Through extensive experiments on 5 video captioning datasets, we show that SWINBERT achieves across-the-board performance improvements over previous methods, often by a large margin. The learned sparse attention masks in addition push the limit to new state of the arts, and can be transferred between different video lengths and between different datasets. Code is available at <https://github.com/microsoft/SwinBERT>.

1. Introduction

Video captioning [1, 10, 25, 28, 35, 39, 44, 45, 56] is the task of describing the visual content of a given video in natural language. As such, it requires an algorithm to understand

* Equal contribution.




Figure 1. Comparison between previous works and SWINBERT. Different from prior works that use offline-extracted 2D/3D features, we propose to adopt the video transformer as our video encoder, and present an end-to-end fully Transformer-based model for video captioning. We further propose to adaptively learn a sparse attention mask to improve long-range video sequence modeling.

and model the spatial-temporal dynamics in video, as well as the relationships between visual and textual elements, and to generate a sequence of output words. This has usually been tackled with transformer-based models that learn from offline extracted video representations [21, 25, 31, 45] (Figure 1 (a)). Specifically, multiple feature extractors, usually trained on image/video understanding tasks (e.g., image classification or action recognition), are employed to extract 2D appearance features and 3D motion features from densely sampled video frames. Although achieving promising results, there exists a discrepancy in both data domain and task formulation between these off-the-shelf feature extractors and downstream video captioning. However, end-to-end training with multiple feature extractors on such dense video frames is computationally intensive, or even infeasible.

Method	MSVD \uparrow	YouCook2 \uparrow	MSRVTT \uparrow	TVC \uparrow	VATEX \uparrow
SOTA	95.2 [57]	53.6 [25]	52.9 [56]	51.0 [25]	58.1 [25]
SWINBERT	160.0	109.0	55.9	56.9	73.0

Table 1. Comparison with state-of-the-art methods across all video captioning datasets considered on CIDEr [47] metric.

More recently, CLIPBERT [21] points out the repetitive information presented in consecutive video frames is not necessary for downstream video-and-language tasks, and proposes a sparse sampling strategy that enables affordable end-to-end training to the raw pixel inputs. Although it has shown great success in video-and-language understanding tasks, such as video question answering [22] and text-to-video retrieval [23, 53], it remains unclear whether these sparsely sampled video frames are sufficient to generate rich and descriptive captions. Moreover, CLIPBERT leverages a 2D Convolutional Neural Network together with mean pooling that operates directly on the raw video frames to learn video representations, which may lose temporal information that is essential to describe visual events in chronological order.

In this work, we aim to find an end-to-end solution to the video captioning task. Inspired by the recent successes of Transformer-based models in computer vision [2, 5, 14, 29], especially for video understanding tasks [6], we propose SWINBERT (Figure 1 (b)), a pure Transformer-based model that directly takes raw video frames as inputs for end-to-end caption generation. Unlike previous methods leveraging off-the-shelf 2D/3D feature extractors at a fixed frame rate, we employ a video Transformer capable of learning from variable lengths of video frame sequence without dedicated design for different frame rates. Based on this specific model design, we investigate *how many video frames are sufficient for the video captioning task?*. Our experiments show that the captioning performance (*i.e.*, CIDEr score) can be greatly lifted by more densely sampled frames (*e.g.*, Ours: 64 frames, vs. CLIPBERT: 16 frames), in contrast to previous success with sparsely sampled frames for video-and-language understanding tasks. Lastly, to avoid the redundancy that comes naturally in consecutive video frames, we further introduce a learnable Sparse Attention Mask as a regularizer that allows the model to focus more on video frame patches that contain more spatial-temporal movements (*e.g.*, the main moving objects) than those staying unchanged for the entire video duration (*e.g.*, the background). In contrast to prior models [21, 25, 31] with predefined attention structures, our model can learn adaptive attention maps to optimize for task-specific performance improvements through better video sequence modeling.

Our extensive experimental results on 5 video captioning datasets demonstrate that our proposed model is effective in learning sparse attention patterns to improve long-range video sequence modeling, and consequently outperforms

previous state-of-the-art approaches by a large margin. To the best of our knowledge, SWINBERT is the first end-to-end pure Transformer-based architecture for video captioning. Additionally, the proposed Sparse Attention Mask effectively regularizes model training and brings further performance improvements across all 5 datasets, which opens a new direction in removing redundancy in video inputs for video-and-language modeling.

In summary, our contributions are three-fold.

- We present SWINBERT, the first end-to-end fully Transformer-based model for video captioning.
- We introduce the Sparse Attention Mask as a regularizer for improving long-range video sequence modeling, and quantitatively validate the effectiveness of the learnable sparse attention mask in caption generation.
- Our method outperforms previous state-of-the-art methods by a large margin on 5 popular video captioning benchmarks. As shown in Table 1, SWINBERT achieves an absolute CIDEr improvement of +64.8 on MSVD, +55.4 on YouCook2, +3.0 on MSRVTT, +5.9 on TVC and +14.9 on VATEX.

2. Related Work

Video Captioning. Recent researches [1, 31, 35, 39, 44] mainly focus on modeling the relationship between fixed video representations and the output textual descriptions via an encoder-decoder framework for video captioning. Specifically, these methods [10, 25, 28, 31, 56] employ an encoder to refine video representations from a set of fixed video frame features, and a language decoder operates on top of these refined video representations to learn visual-textual alignment for caption generation. Researchers [1, 25, 35] have focused on exploring different 2D/3D video representations, including IncepResNetV2 [46], ResNet [17], CLIP-ViT [14, 40], SlowFast [15], C3D [16] and S3D [33, 52], for improving video captioning. In addition, object-level representations [20, 55, 57] have been explored to enrich captions with fine-grained objects and actions. Prior works [11] also studied frame selection schemes to capture informative visual inputs. Unlike previous studies that learn from multiple offline-extracted 2D/3D features with a fixed sampling rate, we introduce Video Swin Transformer [29] as the video encoder in our framework to encode spatial-temporal representations from raw video frames. Benefiting from the flexibility of the transformer architecture, our model can learn with variable number of video tokens and can be trained end-to-end.

Video transformers. Dosovitskiy *et al.* [14] demonstrate that a pure-transformer based architecture can outperform its convolutional counterparts in ImageNet classification task [42]. Since then, there has been a growing interest

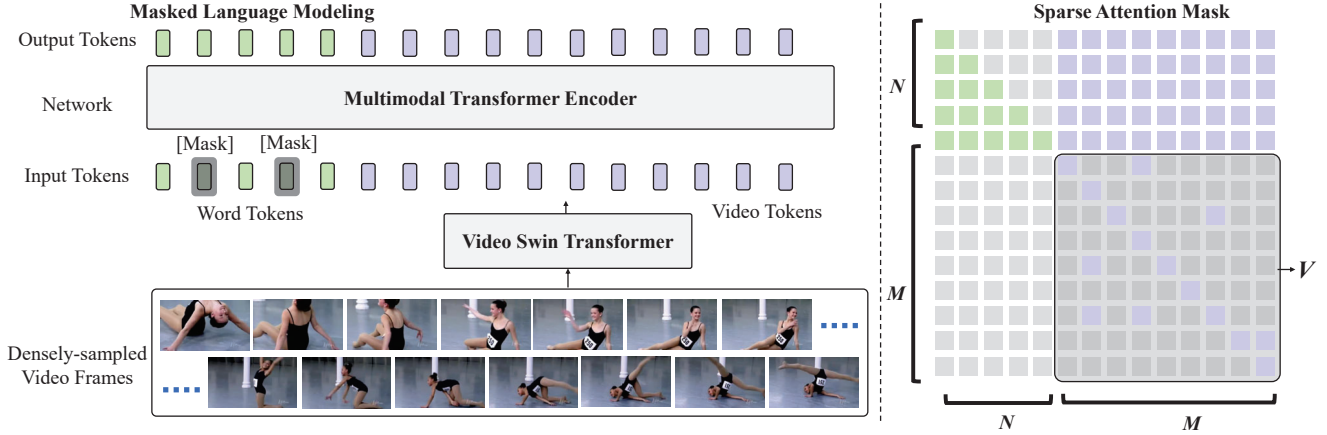


Figure 2. Overview of the proposed framework. Our model takes a sequence of video frames as inputs, and extracts a set of video tokens using a Video Swin Transformer (VidSwin). Given the word tokens and video tokens, we perform self-attention through multiple layers of a multimodal transformer encoder, and predict the word tokens via masked language modeling. As shown on the right, we propose a learnable Sparse Attention Mask as a regularizer for multimodal transformer encoder to reduce redundancy among the video tokens. During inference, the model takes a testing video sample (single-modality) to generate natural language descriptions.

in applying vision transformer (ViT) to the video domain. For example, ViViT [2] and TimeSformer [5] propose a new transformer architecture that can leverage spatial-temporal attention for improving representation learning. Video Swin Transformer (VidSwin) [29] further introduces locality inductive bias into the transformer self-attention, and achieves state-of-the-art performance on action recognition benchmark [6]. While recent studies [2, 5, 29] mainly focus on developing video transformer architecture for action recognition, video captioning has not been explored along this research direction, which is the focus of this work.

Video and language. Recent studies [21, 25, 32–34, 54] have shown great success on multimodal representation learning for video-and-language understanding. Popular downstream tasks include video question answering [22], text-video retrieval [23, 53] and video captioning [50]. Among the literature, Frozen-in-time [3] is a relevant study that explores pure transformer-based model design, but they focus on text-video retrieval. Specifically, they employ two independent transformer encoders for visual and textual inputs, respectively. Retrieval is conducted by estimating the similarity between the outputs of their visual and textual encoders. With a similar spirit, CLIP4Clip [32] studied using the pre-trained CLIP [40] as a feature extractor for video retrieval. While existing architectures [3, 32] are effective for video retrieval, it cannot be directly applied to video captioning, which is the focus of this work.

3. Method

In this section, we present SWINBERT, a new video-based pure-Transformer architecture for caption generation. We first detail the model architecture in Section 3.1, then introduce Sparse Attention Mask in Section 3.2.

3.1. Model Architecture

Figure 2 shows the overview of the proposed model. SWINBERT takes a sequence of raw video frames as inputs, and then outputs a natural language description describing the input video. SWINBERT consists of two modules: *Video Swin Transformer* (VidSwin), and *Multimodal Transformer Encoder*. First, we leverage VidSwin to extract spatial-temporal video representations from the raw video frames. Then, our Multimodal Transformer Encoder takes as inputs the video representations and outputs a natural language sentence via sequence-to-sequence (seq2seq) generation. We describe each module in detail as below.

Video Swin Transformer. As discussed in [13, 49], video understanding benefits from long-range temporal modeling. A simple way is to stack a large number of frames to capture long-range structures. However, it would greatly increase the computational cost. Recently, VidSwin [29] is designed to leverage the spatial-temporal locality inherent in videos, and achieves a favorable speed-accuracy trade-off. In the first module of our framework, we propose to use VidSwin as our visual encoder to encode the raw video frames as video feature tokens. VidSwin is pre-trained on the Kinetics action recognition task [6].

Given the raw video frames which are of size $T \times H \times W \times 3$, consisting of T frames and each has $H \times W \times 3$ pixels. We feed them to VidSwin, and extract grid features from the last encoder block of VidSwin. The grid features of VidSwin is defined to be of size $\frac{T}{2} \times \frac{H}{32} \times \frac{W}{32} \times 8C$, where C is the channel dimension. We then tokenize the grid features along the channel dimension, resulting in a total of $\frac{T}{2} \times \frac{H}{32} \times \frac{W}{32}$ video tokens. Each token is a $8C$ -dim feature vector. After that, we input the video tokens to the multimodal transformer encoder for caption generation.

With our generic design, it enables end-to-end training for video captioning from the raw video frames. Moreover, benefiting from the flexibility of the transformer architecture, our model is able to process variable lengths of video sequences. As we will show in experiments, the caption performance (*i.e.*, CIDEr scores) can be improved with longer video sequence inputs (*i.e.*, densely-sampled video frames).

Multimodal Transformer Encoder. In our second module, we use a transformer encoder to generate natural language description. To be specific, it has textual and visual modality inputs, including the tokenized caption description and the video tokens computed from VidSwin. We then perform seq2seq generation to form a natural language sentence. In the same spirit as in image captioning literature [19, 26], we use a causal self-attention mask where a caption token can only attend to the existing output tokens. This effectively simulates a uni-directional seq2seq generation process. In addition, all the textual tokens have full attentions to the video tokens.

3.2. Learning with Sparse Attention Mask

In general, longer inputs across multiple video segments contain more information. However, the computational demand of attention are proportional to input length, which limits the number of input frames. On the other hand, considering the essence of the video properties, the dense-sampling scheme with consecutive video frames contains redundant and perhaps irrelevant information, which may compromise performance. Hence, how to effectively model a long sequence of video tokens is a unique challenge in our proposed framework. We address it by introducing a learnable Sparse Attention Mask as a regularizer to our multimodal transformer encoder.

As shown to the right of Figure 2, the input to the Transformer is split into two parts: N word tokens and M video tokens. The entire attention mask can be defined of size $(N+M) \times (N+M)$, where N is 50 and $M = \frac{T}{2} \times \frac{H}{32} \times \frac{W}{32}$ in our experiments. We denote V as the learnable attention mask of size $M \times M$ governing the attentions among the video tokens. For more accurate video captioning, we allow the text tokens with unrestricted attention so they can take advantage of visual details. To address the redundancy among the video tokens, we impose the sparsity constraint overlay on top of V by:

$$\mathcal{L}_{\text{SPARSE}} = \lambda \times \sum_{i=1}^M \sum_{j=1}^M |V_{i,j}|, \quad (1)$$

where λ is the regularization hyperparameter, and $V_{i,j}$ are the activation values of the learnable attention mask V .

During learning, the sparsity constraint will regularize model training to discover the underlying structure of the video sequences. Through sparse attention, the model

learns to strengthen the most important relationships among different tokens by reducing the likelihood of meaningless connections, while focusing more on the active video tokens that contain rich spatial-temporal information. In this way, the model can produce more expressive and descriptive natural language sentences.

In our implementation, we apply the sigmoid activation function on the sparse attention mask. Therefore, the sparse attention mask consists of continuous activation between 0 and 1. As we will show in our experiments, we can realize a binary mask by simply using a threshold of 0.5.

Training. We train SWINBERT in an end-to-end manner by applying Masked Language Modeling (\mathcal{L}_{MLM}) [12] on top of our multimodal transformer encoder. We mask a percentage of word tokens by replacing them with a pre-defined special token [MASK]. We then ask the multimodal transformer to predict the masked ones. In order to predict a masked word token, the model will have to resort to the video tokens and other word tokens. This facilitates cross-modality representation learning to help ground the caption descriptions in the video context. Moreover, we apply the proposed sparsity constraint on the learnable attention mask to enhance the modeling of the video token sequence.

In summary, our loss function includes \mathcal{L}_{MLM} [12] and $\mathcal{L}_{\text{SPARSE}}$, and we train SWINBERT by simply minimizing the sum of them.

Inference. During inference, our model takes a video sequence as input (single visual modality), and outputs a natural language sentence. We generate the output sentence in an auto-regressive manner. In other words, our model generates one word token at a time, consuming the previously generated tokens as the inputs of the multimodal transformer encoder. We perform generation until our model outputs a pre-defined ending token [EOS] or reaches the maximum output length.

4. Experiments

4.1. Experimental Setup

Datasets. We conduct experiments on 5 video captioning datasets, detailed below.

- **MSVD** [7] is a collection of $2K$ open-domain video clips downloaded from YouTube. Each video clip has 35 ground-truth captions written by human. We use the standard split which contains $1.2K$ training videos and 670 test videos.
- **YouCookII** [59] is a cooking domain dataset covering 89 recipes. There are $15.4K$ video clips, and each has 1 ground-truth caption. We use the standard training/validation split in the experiments.
- **MSRVTT** [53] consists of $10K$ open-domain video clips. Each video clip has 20 ground-truth captions. We use the

Method	Features			MSVD				MSRVTT			
	2D Appearance	3D Motion	Object Detection	B4	M	R	C	B4	M	R	C
PickNet [11]	ResNet152	-	-	52.3	33.3	69.6	76.5	41.3	27.7	59.8	44.1
SibNet [28]	GoogleNet	-	-	54.2	34.8	71.7	88.2	40.9	27.5	60.2	47.5
OA-BTG [55]	ResNet200	-	MaskRCNN	56.9	36.2	-	90.6	41.4	28.2	-	46.9
GRU-EVE [1]	IncepResnetV2	C3D	YOLO	47.9	35.0	71.5	78.1	38.3	28.4	60.7	48.1
MGSA [9]	IncepResnetV2	C3D	-	53.4	35.0	-	86.7	42.4	27.6	-	47.5
POS+CG [48]	IncepResnetV2	OpticalFlow	-	52.5	34.1	71.3	88.7	42.0	28.2	61.6	48.7
POS+VCT [18]	IncepResnetV2	C3D	-	52.8	36.1	71.8	87.8	42.3	29.7	62.8	49.1
SAAT [58]	IncepResnetV2	C3D	-	46.5	33.5	69.4	81.0	39.9	27.7	61.2	51.0
STG-KD [35]	ResNet101	I3D	FasterRCNN	52.2	36.9	73.9	93.0	40.5	28.3	60.9	47.1
PMI-CAP [8]	IncepResnetV2	C3D	-	54.6	36.4	-	95.1	42.1	28.7	-	49.4
ORG-TRL [57]	IncepResnetV2	C3D	FasterRCNN	54.3	36.4	73.9	95.2	43.6	28.8	62.1	50.9
OpenBook [56]	IncepResnetV2	C3D	-	-	-	-	-	33.9	23.7	50.2	52.9
SWINBERT	VidSwin	-	-	66.3	42.4	80.9	149.4	45.4	30.6	64.1	55.9

Table 2. Comparison with state-of-the-art methods on MSVD and MSRVTT.

VATEX					
Method	Mod.	B4	R	M	C
VaTeX [50]	V	28.4	47.0	21.7	45.1
ORG-TRL [57]	V	32.1	48.9	22.2	49.7
Support-set [38]	V	32.8	49.1	24.4	51.2
Support-set [38]	V	32.5	48.9	24.1	50.5
OpenBook [56]	V+T	33.9	50.2	23.7	57.5
VALUE [25]	V+T	-	-	-	58.1
SWINBERT	V	38.7	53.2	26.2	73.0

(a) VATEX.

TVC					
Method	Mod.	B4	R	M	C
MMT [23]	V	9.9	30.4	15.2	36.0
MMT [23]	T	6.3	7.7	13.9	33.7
MMT [23]	V+T	10.8	32.8	16.9	45.3
HERO [24]	V+T	12.3	34.1	17.6	49.9
VALUE [25]	V+T	11.6	33.9	17.6	50.5
SWINBERT	V	14.5	36.1	18.5	55.4

(b) TVC.

YouCook2						
Method	Mod.	B3	B4	M	R	C
Masked Trans. [60]	V	7.5	3.8	10.6	-	37.9
DPC [43]	V	-	2.2	17.6	-	-
DPC [43]	V+T	-	2.8	18.1	-	-
VideoBERT [45]	V	7.5	4.3	11.9	-	55.0
ActBERT [61]	V	8.6	5.4	13.3	-	65.0
AT [43]	T	-	8.5	16.9	-	106.0
AT+Video [43]	V+T	-	9.0	17.7	-	112.0
VALUE [25]	V	-	-	-	-	53.6
VALUE [25]	V+T	-	12.4	18.8	40.4	130.3
SWINBERT	V	13.8	9.0	15.6	37.3	109.0

(c) YouCook2.

(a) VATEX.

(b) TVC.

(c) YouCook2.

Table 3. Comparison with state-of-the-art methods on YouCook2, TVC, and VATEX. We gray out models that adopt vision-and-language pre-training on large-scale datasets for a fair comparison.

standard captioning split [31], which has 6.5K training videos and 2.9K testing videos.

- **TVC** [23] is a TV domain dataset. There is a total of 262K caption descriptions paired with 108K video segments. The captions in TVC not only describe the video contents, but it may also describe the subtitles.
- **VATEX** [50] is a relative large open-domain dataset, which contains 41.3K videos. Each video clip has 20 ground-truth captions. We use the official training set for training, and evaluate the results using the public test set.

Implementation Details. We implement our model using Pytorch [37], Huggingface transformer [51], and DeepSpeed library [41]. The VidSwin is initialized with Kinetics-600 pre-trained weights [29], and the multimodal transformer encoder is randomly initialized. In order to ensure that the video tokens have the same embedding size as that of the word tokens, we transform the video tokens using a learnable MLP. Following [25], we employ AdamW optimizer [30] and use a learning rate warm-up during the early 10% training steps followed by linear decay. Additional details can be found in the supplementary material.

4.2. Main Results

We compare SWINBERT with previous state-of-the-art methods on 5 public benchmark datasets. Following the literature [25, 50, 56, 57], we provide detailed comparisons using a diverse set of performance metrics, including BLEU4 [36], METEOR [4], ROUGE-L [27] and CIDEr [47].

Table 2 shows detailed comparisons on MSVD and MSRVTT datasets. SWINBERT outperforms previous state-of-the-art methods in terms of all metrics by a large margin. Specifically, SWINBERT brings significant CIDEr improvements on MSVD (*i.e.*, +54.2 higher than the prior arts). SWINBERT also achieves strong improvements across all metrics on MSRVTT.

In Table 3a, we report detailed comparisons on the VATEX dataset. SWINBERT achieves better performance than the prior works, especially on CIDEr metric. It should be noted that previous state-of-the-art methods (*i.e.*, VALUE [25] and Support-set [38]) perform vision-and-language (VL) pre-training on large-scale datasets for improving multimodal representations, whereas the results

T	MSRVTT	VATEX	T	Learnable Att. Mask	$\mathcal{L}_{\text{SPARSE}}$	MSRVTT	VATEX	Att. Mask	MSRVTT	VATEX
2	36.6	47.4	32	✗	✗	52.3	71.1	Full Attention	52.3	71.1
4	43.7	58.2	32	✓	✗	53.3	70.7	Spatial Window	51.9	71.0
8	47.6	65.2	32	✓	✓	55.1	71.6	Temporal Window	51.0	70.2
16	49.5	68.4						Ours (Learnable, Sparse)	55.1	71.6
32	52.3	71.1								
64	55.3	72.7								

(a) Impact of #Video Frames (T).

(b) Learning of Sparse Attention Mask.

(c) Heuristic vs. Learnable Attention Mask

Table 4. Results on SWINBERT with (a) varying number of video frames (without learnable sparse attention mask), (b) ablation study on learnable attention mask and sparse attention loss, and (c) comparisons between heuristic and learnable attention masks. All experiments are conducted with 32 frames unless specified otherwise. All results are reported on CIDEr metric.

with SWINBERT are not based on VL pre-training. We believe that further integration of VL pre-training will provide additional improvements. This, from another point of view, demonstrates the superior performance of SWINBERT.

We further conduct analysis on the challenging TVC dataset, and the results are shown in Table 3b. Note that captions in TVC are designed to describe not only the visual events but also supplementary information presented in the subtitle sentences. VALUE [25], HERO [24] and MMT [23] are three prior works, that leverage multimodal video inputs, including 2D/3D visual frame features and subtitle sentences from the original TV show scripts. With video frame inputs alone, SWINBERT is able to achieve better performance than all three of them. This superior performance suggests that SWINBERT is effective in exploiting visual representations for video captioning.

Table 3c shows the detailed comparisons on YouCook2. We list the prior works that take visual and/or textual modality signals as inputs. Compared with visual-only approaches, SWINBERT brings significant CIDEr improvements on YouCook2. To be specific, SWINBERT achieves 109.0 CIDEr score, which is +55.4 higher than that of VALUE [25], and +44.0 higher than that of ActBERT [61]. We believe that SWINBERT can be further enhanced with multimodal video inputs by leveraging additional modalities such as subtitle and audio, which is worth exploring in future study.

4.3. Ablation Study

We conduct comprehensive ablation study on multiple datasets to investigate the capability of the proposed model. Following [25], we use CIDEr metric [47] as our primary evaluation metric for video captioning.

Impact of video frames. We first investigate the impact of the sampling rate of the video frames on the task of video captioning. Specifically, we uniformly sample $T = \{2, 4, 8, 16, 32, 64\}$ frames from the given video clip to train and test our SWINBERT. For clarity, we disable the sparse attention mask in this experiment. Table 4a shows the model performance with varying number of video frames on MSRVTT and VATEX. As we increase the number of

T	Attn. Mask	MSVD	YouCook2	MSRVTT	TVC	VATEX
32	Full	127.9	104.2	52.3	53.0	71.1
32	Sparse (soft)	147.6	104.8	55.1	53.8	71.6
32	Sparse (binary)	141.0	101.4	55.3	52.8	71.6
48	Full	144.2	103.1	53.9	53.9	71.7
48	Sparse (soft)	147.8	105.0	54.6	55.2	71.9
48	Sparse (binary)	148.1	103.8	54.9	52.6	72.0
64	Full	144.7	106.1	55.3	54.3	72.7
64	Sparse (soft)	149.4	109.0	55.9	55.4	73.0
64	Sparse (binary)	146.3	106.6	55.0	53.1	71.9

Table 5. Effectiveness of soft/binary sparse attention mask on longer video sequences. T indicates number of video frames. All results are reported on CIDEr metric.

frames, we observe consistent improvements on the CIDEr metric. These results suggest that the performance of video captioning can be greatly lifted by using more densely sampled frames.

Effectiveness of sparse attention mask. One important question is whether adding sparse attention mask to the transformer is helpful. To understand the effect of the sparse attention mask, Table 4b shows the ablation study. First of all, we present a baseline that does not have any learnable attention mask, shown in the first row of Table 4b. In the second row, we show another baseline which uses a learnable attention mask but no sparsity constraints are added. This is equivalent to a random attention mask. Finally, the bottom row shows our proposed method. We observe that the proposed sparsity constraint is helpful in improving video captioning in terms of CIDEr scores (*i.e.*, +2.8 on MSRVTT and +0.5 on VATEX).

Comparison between heuristic and learnable attention masks. We also study the design of attention patterns for constructing our sparse attention mask. To be specific, we explore two heuristic designs including (i) *Spatial Window*: A sliding window attention pattern that attends to its neighbor tokens along the spatial dimension; (ii) *Temporal Window*: A sliding window attention pattern, which attends along the temporal dimension. We use a fixed window size w for both Spatial Window and Temporal Window,

#Video Frames (T)	VATEX	MSVD	YouCook2	MSRVTT	TVC	Dataset	CIDEr	Dataset	CIDEr
32	71.6	147.6	104.8	55.1	53.8	VATEX	71.6	VATEX	71.6
64	73.0	149.4	109.0	55.9	55.4	MSVD	147.6	MSRVTT	55.1
32 \rightarrow 64	73.0	150.0	108.2	55.9	56.9	VATEX \rightarrow MSVD	148.1	VATEX \rightarrow MSRVTT	55.8
(attn. mask only)						(att. mask only)		(att. mask only)	
32 \rightarrow 64	72.4	152.3	106.8	55.6	56.0	VATEX \rightarrow MSVD	160.3	VATEX \rightarrow MSRVTT	54.5
(entire model)						(entire model)		(entire model)	

(a) Transfer between different frame rates within a dataset

(b) Transfer across datasets

Table 6. Transferability of SWINBERT (a) between different frame rates within a dataset and (b) across datasets. We experiment with two settings: (i) transfer the learned sparse attention masks only (attn. mask only) and (ii) transfer the learned model weights along with sparse attention masks (entire model). All results are reported on CIDEr metric.

and we have explored $w = \{10, 20, 50, 100\}$ in our experiments. Table 4c presents results with different sparse attention masks along with a *Full Attention* baseline, that is the original attention mask allowing full attentions among all video tokens. Results show that both Spatial Window and Temporal Window brings performance degradation, compared to the Full Attention baseline. In contrast, our learnable sparse attention mask improves over Full Attention and heuristic sparse attention masks. We conjecture that our sparsity constraint enforces the model to identify more salient video frame patches along both spatial and temporal dimensions for caption generation. Visualizations of the learned sparse attention masks shown later in this section further corroborates our hypothesis.

Longer video sequences. We further examine the capability of the proposed sparse attention mask using longer video sequences. We apply the learnable sparse attention masks to $T = \{32, 48, 64\}$ frames uniformly sampled from the video clips, and the results are shown in Table 5 (*Full vs. Sparse (soft)*). We observe that adding sparse attention mask to SWINBERT consistently improves the CIDEr scores across different video sequence length, and push the limit to new state-of-the-arts on all the 5 benchmarks. These results suggest that the sparse attention mask is effective in regularizing model training for long-range video sequence modeling.

Binary sparse attention mask. Our learnable sparse attention mask can be seen as a *soft* attention mask, which consists of continuous values between 0 and 1. An interesting question we aim to answer is: *can we enforce it into a binary mask?* We test this hypothesis by simply thresholding the learned sparse attention mask with a fixed threshold 0.5. In addition, we fine-tune the model for a few training steps to adapt it to the binarized mask. In Table 5, we observe that converting the mask to a binary one may have a slight performance drop on the CIDEr metric, which is expected as we reduce the capacity of the attention mask. It is worth noting that, with the binarized mask, the caption performance is comparable or better than the Full Attention (Full) baseline. In future, we plan to leverage custom

CUDA implementations to construct this binary sparse attention mask to improve runtime speed.

Generalization capability. Since our sparse attention mask is optimized for task-specific performance improvements, one may wonder its generalizability to different frame rates and different datasets. We study the generalization capability under two configurations: (i) *Across frame rates*: we first train SWINBERT at a slow frame rate, and then move to a faster frame rate for further training. To achieve this, we expand the learned sparse attention mask by linear interpolation along the temporal dimension; (ii) *Across datasets*: we first train SWINBERT on one dataset, and then fine-tune it on another dataset. The experiments are conducted in two settings, transferring the whole model weights or only the sparse attention mask.

Table 6a shows the results of transferring from 32 frames to 64 frames. We observe that it yields a comparable or better CIDEr score compared to using 64 frames directly. It should be noted that, transferring only the sparse attention mask is able to achieve reasonable CIDEr scores on the 5 datasets. The results suggest that linear interpolation along the temporal dimension is effective for transferring between different frame rates.

Table 6b shows the results of transferring across datasets. In this experiment, we first train our model on VATEX dataset, and then fine-tune it on MSRVTT and MSVD datasets, respectively. We observe that such transfer learning scheme improves CIDEr scores for both datasets. As the data domain between VATEX and MSVD is similar, fine-tuning the entire model is more effective for improving CIDEr scores on MSVD. The success in transfer learning suggests that the performance of SWINBERT can be further improved with pre-training on even larger-scale video-text datasets, which we leave as future study.

Visualization of sparse attention mask. We visualize the learned sparse attention pattern in Figure 3. Note that the values are obtained from the *soft* attention mask without thresholding. On the left, we show an example video clip which is randomly sampled from MSVD dataset. Additionally, we denote the patch regions and the correspond-

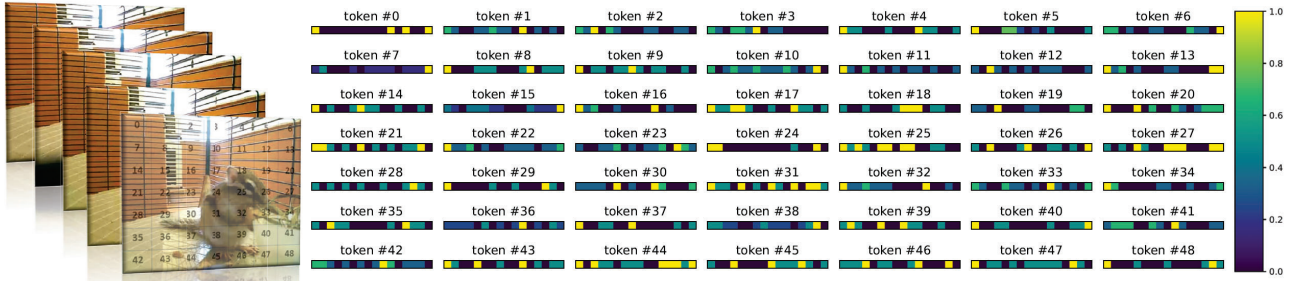


Figure 3. Visualization of sparse attention mask along the temporal dimension. Our sparse attention mask discovers possible principle in the video sequences. We observe that boundary-region tokens can be sparsely sampled along the temporal dimension. This is probably due to similar background in a video clip. On the other hand, as the center-region tokens may contain more pixel variations (such as movements, actions, or scene changes), they thus require denser sampling along the temporal dimension.




Figure 4. Training behavior of SWINBERT. (a) During training, the proposed sparsity constraint effectively reduces the percentage of non-zero elements in the attention mask. (b) Sparsity constraint does not interfere captioning as CIDEr score keeps increasing.

ing token IDs at the first frame. On the right, for each token, we visualize the weights of the learned sparse attention mask along the temporal dimension using a horizontal bar, where yellow color indicates stronger attention activity. We briefly summarize our findings: (i) Many of the tokens at the boundary are attending to some starting and ending frames. This is possibly because the background does not change much, and therefore for those tokens, the temporal information can be sparsely sampled with respect to the attention mask; (ii) The center-region tokens may contain more movements or scene changes, therefore require denser sampling along the temporal dimension.

Training behavior. In Figure 4, we investigate the learning behavior of our sparse attention mask. In Figure 4a, our proposed sparsity constraint is effective in reducing the number of non-zero elements in the attention mask, and more than 95% of the elements are set to zero in the end. This verifies the sparsity of the attention mask. In addition, as shown in Figure 4b, we find that our sparsity constraint does not interfere the learning of video captioning, as CIDEr scores keep increasing during the learning process.

Qualitative results. Figure 5 shows the qualitative examples of SWINBERT. We find that SWINBERT is capable of recognizing the visual contents (e.g., dog and watermelon), and correctly describes the actions and events (e.g., eating) in the given video. We also note that, while our model gen-

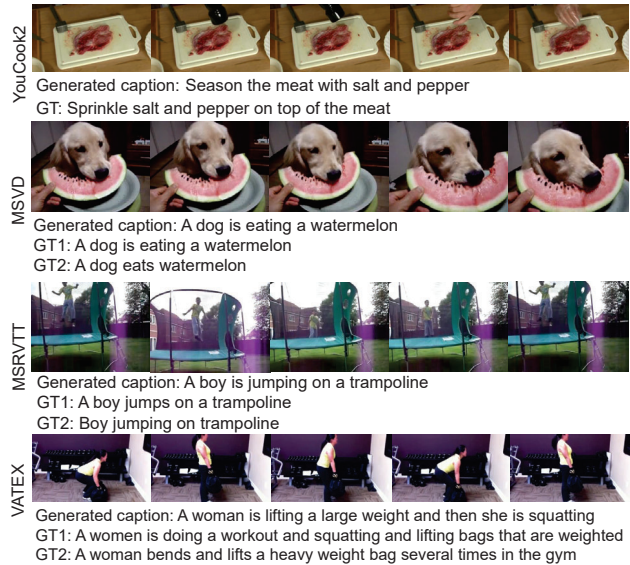


Figure 5. Qualitative examples generated by SWINBERT. The generated captions are semantically reasonable and describe the video contents correctly.

erates semantically reasonable captions, the predicted word sequences may not always equal to the ground truth.

5. Conclusion

We present SWINBERT, a new end-to-end fully Transformer-based architecture for video captioning. We further propose to adaptively learn a sparse attention mask for better video sequence modeling. Extensive experimental results on 5 popular benchmark datasets show that SWINBERT achieves better performance than the previous state-of-the-art methods by a large margin. In future, we plan to investigate large-scale video-language pre-training to further enhance the captioning performance.

Acknowledgements: We thank reviewers for the insightful comments. We thank Jianfeng Wang, Xiaowei Hu, Lin Liang, Zhengyuan Yang, Ehsan Azarnasab, Yue Cao, Lei Ji, Huaishao Luo and Ze Liu for the valuable discussions.

References

- [1] Nayyer Aafaq, Naveed Akhtar, Wei Liu, Syed Zulqarnain Gilani, and Ajmal Mian. Spatio-temporal dynamics and semantic attribute enriched visual encoding for video captioning. In *CVPR*, 2019. 1, 2, 5
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, 2021. 2, 3
- [3] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, 2021. 3
- [4] Satantjeet Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005. 5
- [5] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 2, 3
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 2, 3
- [7] David Chen and William Dolan. Collecting highly parallel data for paraphrase evaluation. In *ACL*, 2011. 4
- [8] Shaoxiang Chen, Wenhao Jiang, Wei Liu, and Yu-Gang Jiang. Learning modality interaction for temporal sentence localization and event captioning in videos. In *ECCV*, 2020. 5
- [9] Shaoxiang Chen and Yu-Gang Jiang. Motion guided spatial attention for video captioning. In *AAAI*, 2019. 5
- [10] Shaoxiang Chen, Ting Yao, and Yu-Gang Jiang. Deep learning for video captioning: A review. In *IJCAI*, 2019. 1, 2
- [11] Yangyu Chen, Shuhui Wang, Weigang Zhang, and Qingming Huang. Less is more: Picking informative frames for video captioning. In *ECCV*, 2018. 2, 5
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 4
- [13] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, 2015. 3
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 2
- [15] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019. 2
- [16] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *CVPR*, 2018. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2
- [18] Jingyi Hou, Xinxiao Wu, Wentian Zhao, Jiebo Luo, and Yunde Jia. Joint syntax representation learning and visual cue translation for video captioning. In *ICCV*, 2019. 5
- [19] Xiaowei Hu, Xi Yin, Kevin Lin, Lijuan Wang, Lei Zhang, Jianfeng Gao, and Zicheng Liu. Vivo: Surpassing human performance in novel object captioning with visual vocabulary pre-training. In *AAAI*, 2021. 4
- [20] Yaosi Hu, Zhenzhong Chen, Zheng-Jun Zha, and Feng Wu. Hierarchical global-local temporal modeling for video captioning. In *ACM MM*, 2019. 2
- [21] Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *CVPR*, 2021. 1, 2, 3
- [22] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. Tvqa: Localized, compositional video question answering. *EMNLP*, 2018. 2, 3
- [23] Jie Lei, Licheng Yu, Tamara L Berg, and Mohit Bansal. Tvr: A large-scale dataset for video-subtitle moment retrieval. In *ECCV*, 2020. 2, 3, 5, 6
- [24] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. Hero: Hierarchical encoder for video+ language omni-representation pre-training. In *EMNLP*, 2020. 5, 6
- [25] Linjie Li, Jie Lei, Zhe Gan, Licheng Yu, Yen-Chun Chen, Rohit Pillai, Yu Cheng, Luowei Zhou, Xin Eric Wang, William Yang Wang, et al. Value: A multi-task benchmark for video-and-language understanding evaluation. In *NeurIPS*, 2021. 1, 2, 3, 5, 6
- [26] Xiujuan Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020. 4
- [27] Chin-Yew Lin and Franz Josef Och. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *ACL*, 2004. 5
- [28] Sheng Liu, Zhou Ren, and Junsong Yuan. Sibnet: Sibling convolutional encoder for video captioning. *IEEE TPAMI*, 2020. 1, 2, 5
- [29] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 2, 3, 5
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [31] Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Jason Li, Taroon Bharti, and Ming Zhou. Univl: A unified video and language pre-training model for multimodal understanding and generation. *arXiv preprint arXiv:2002.06353*, 2020. 1, 2, 5
- [32] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval. *arXiv preprint arXiv:2104.08860*, 2021. 3
- [33] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, 2020. 2, 3

- [34] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *ICCV*, 2019. 3
- [35] Boxiao Pan, Haoye Cai, De-An Huang, Kuan-Hui Lee, Adrien Gaidon, Ehsan Adeli, and Juan Carlos Niebles. Spatio-temporal graph for video captioning with knowledge distillation. In *CVPR*, 2020. 1, 2, 5
- [36] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002. 5
- [37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, 2019. 5
- [38] Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metze, Alexander Hauptmann, Joao Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. In *ICLR*, 2021. 5
- [39] Wenjie Pei, Jiuyan Zhang, Xiangrong Wang, Lei Ke, Xiaoyong Shen, and Yu-Wing Tai. Memory-attended recurrent network for video captioning. In *CVPR*, 2019. 1, 2
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 3
- [41] Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *KDD*, 2020. 5
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 2
- [43] Botian Shi, Lei Ji, Yaobo Liang, Nan Duan, Peng Chen, Zhendong Niu, and Ming Zhou. Dense procedure captioning in narrated instructional videos. In *CoNLL*, 2019. 5
- [44] Botian Shi, Lei Ji, Zhendong Niu, Nan Duan, Ming Zhou, and Xilin Chen. Learning semantic concepts and temporal alignment for narrated video procedural captioning. In *ACM MM*, 2020. 1, 2
- [45] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *ICCV*, 2019. 1, 5
- [46] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, 2017. 2
- [47] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015. 2, 5, 6
- [48] Bairui Wang, Lin Ma, Wei Zhang, Wenhao Jiang, Jingwen Wang, and Wei Liu. Controllable video captioning with pos sequence guidance based on gated fusion network. In *ICCV*, 2019. 5
- [49] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE TPAMI*, 2018. 3
- [50] Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang. VateX: A large-scale, high-quality multilingual dataset for video-and-language research. In *ICCV*, 2019. 3, 5
- [51] Thomas Wolf, Julien Chaumond, Lysandre Debut, Victor Sanh, Clement Delangue, Anthony Moi, Pierric Cistac, Morgan Funtowicz, Joe Davison, Sam Shleifer, et al. Transformers: State-of-the-art natural language processing. In *EMNLP: System Demonstrations*, 2020. 5
- [52] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, 2018. 2
- [53] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, 2016. 2, 3, 4
- [54] Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. Merlot: Multimodal neural script knowledge models. In *NeurIPS*, 2021. 3
- [55] Junchao Zhang and Yuxin Peng. Object-aware aggregation with bidirectional temporal graph for video captioning. In *CVPR*, 2019. 2, 5
- [56] Ziqi Zhang, Zhongang Qi, Chunfeng Yuan, Ying Shan, Bing Li, Ying Deng, and Weiming Hu. Open-book video captioning with retrieve-copy-generate network. In *CVPR*, 2021. 1, 2, 5
- [57] Ziqi Zhang, Yaya Shi, Chunfeng Yuan, Bing Li, Peijin Wang, Weiming Hu, and Zheng-Jun Zha. Object relational graph with teacher-recommended learning for video captioning. In *CVPR*, 2020. 2, 5
- [58] Qi Zheng, Chaoyue Wang, and Dacheng Tao. Syntax-aware action targeting for video captioning. In *CVPR*, 2020. 5
- [59] Luowei Zhou, Chenliang Xu, and Jason J Corso. Towards automatic learning of procedures from web instructional videos. In *AAAI*, 2018. 4
- [60] Luowei Zhou, Yingbo Zhou, Jason J Corso, Richard Socher, and Caiming Xiong. End-to-end dense video captioning with masked transformer. In *CVPR*, 2018. 5
- [61] Linchao Zhu and Yi Yang. Actbert: Learning global-local video-text representations. In *CVPR*, 2020. 5, 6