

Controllable Video Captioning with an Exemplar Sentence

Yitian Yuan*

yyt18@mails.tsinghua.edu.cn
Tsinghua-Berkeley Shenzhen Institute,
Tsinghua University

Jingwen Wang

jaywongjaywong@gmail.com
Tencent AI Lab

Lin Ma[†]

forest.linma@gmail.com
Meituan

Wenwu Zhu[†]

wwzhu@tsinghua.edu.cn
Department of Computer Science and Technology,
Tsinghua University

ABSTRACT

In this paper, we investigate a novel and challenging task, namely controllable video captioning with an exemplar sentence. Formally, given a video and a syntactically valid exemplar sentence, the task aims to generate one caption which not only describes the semantic contents of the video, but also follows the syntactic form of the given exemplar sentence. In order to tackle such an exemplar-based video captioning task, we propose a novel Syntax Modulated Caption Generator (SMCG) incorporated in an encoder-decoder-reconstructor architecture. The proposed SMCG takes video semantic representation as an input, and conditionally modulates the gates and cells of long short-term memory network with respect to the encoded syntactic information of the given exemplar sentence. Therefore, SMCG is able to control the states for word prediction and achieve the syntax customized caption generation. We conduct experiments by collecting auxiliary exemplar sentences for two public video captioning datasets. Extensive experimental results demonstrate the effectiveness of our approach on generating syntax controllable and semantic preserved video captions. By providing different exemplar sentences, our approach is capable of producing different captions with various syntactic structures, thus indicating a promising way to strengthen the diversity of video captioning. Code for this paper is available at <https://github.com/yytzy/SMCG>.

CCS CONCEPTS

• **Computing methodologies** → **Artificial intelligence**; *Computer vision*; *Natural language processing*.

KEYWORDS

Controllable video captioning; syntax modulated caption generator

*This work was done while Yitian Yuan was a Research Intern at Tencent AI Lab.

[†]Corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

MM '20, October 12–16, 2020, Seattle, WA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-7988-5/20/10...\$15.00

<https://doi.org/10.1145/3394171.3413908>

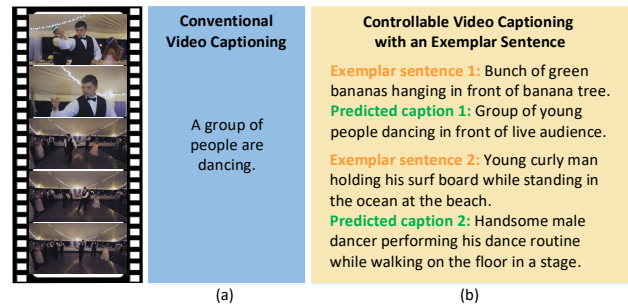


Figure 1: The comparison between the conventional video captioning task (a) and our proposed controllable video captioning with an exemplar sentence task (b).

ACM Reference Format:

Yitian Yuan, Lin Ma[†], Jingwen Wang, and Wenwu Zhu. 2020. Controllable Video Captioning with an Exemplar Sentence. In *Proceedings of the 28th ACM International Conference on Multimedia (MM '20)*, October 12–16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3394171.3413908>

1 INTRODUCTION

Automatically generating sentence descriptions for videos, i.e., video captioning, has emerged as a prominent research problem to bridge computer vision and natural language processing. From early template-based approaches [10, 14, 25, 26, 39] to recent sequence learning approaches [2, 20, 21, 30, 31, 33, 40, 42], remarkable advancements have been achieved on this task. However, most conventional video captioning models mainly focus on understanding video semantics so as to produce captions describing video contents, while the linguistic expressiveness and diversity of the video captions are often ignored. For example, as shown in Figure 1(a), by learning from the caption annotations with simple linguistic characteristics in the training set, the predicted caption is monotonous and plain, which will further influence the user experiences of potential practical applications, such as chat robots [4, 15], automatic video commenting [17], etc.

To enrich the linguistic expressiveness of visual captions, some researchers [9, 19, 41] incorporate specific styles, such as humorous or romantic, into the captioning process. Despite the promising results, such stylized video captioning models still constrain the generated captions in a predefined set of linguistic styles, and may limit

the diversity and flexibility of video captions. Recently, the part-of-speech (POS) tag sequences are further introduced to control the syntactic structure of the predicted captions [6, 32]. However, directly manipulating the POS tag needs specialized knowledge on the sentence syntax and grammar rules, which is hard to be realized practically.

Compared with style labels or POS sequences, one intuitive and straightforward way is to directly leverage one exemplar sentence to control or customize the video caption generation. Therefore, it motivates us to propose and investigate a novel task in this paper, namely controllable video captioning with an exemplar sentence or exemplar-based video captioning. As shown in Figure 1(b), given any syntactically valid exemplar sentence and a video, the task aims to generate one caption, which can not only express the semantic contents of the video, but also follow the syntactic structure of the given exemplar sentence. Since exemplar sentences are easy to acquire and have no specific constraint, this task can generate diverse captions with a variety of syntactic structures, thus strengthening the expressiveness and diversity of the video captions.

The proposed exemplar-based video captioning is a challenging problem. Firstly, how to extract syntactic information from exemplar sentences, and make generated captions absorb the extracted sentence syntaxes is quite crucial. Secondly, since introducing extra exemplar sentences will also bring some disturbances for semantic contents of the input video, how to preserve video semantics in the generated captions is worthy of further considerations.

In order to tackle the aforementioned challenges, we propose a novel Syntax Modulated Caption Generator (SMCG) incorporated in one encoder-decoder-reconstructor architecture. As shown in Figure 2, the constituency parse tree of the exemplar sentence is firstly extracted by a sentence parser. Afterwards, two encoders then process the exemplar sentence parse tree and video feature sequence to obtain the syntactic and semantic representations, respectively. The proposed SMCG, which summarizes the semantic representation as input and utilizes the syntactic representation to modulate the gates and cells in the sentence decoding long short-term memory (LSTM) [11], is thereby able to generate one caption expressing the semantic meaning of the video and meanwhile following the syntactic structure of the exemplar sentence. Moreover, two reconstructors built on the hidden states of the caption generator are established to reproduce the original video feature and exemplar sentence parse tree, respectively. Such reconstructions can further help preserve video semantics in predicted captions and ensure the syntax customization.

To summarize, the contributions of this work lie in three folds:

- We propose a novel task – controllable video captioning with an exemplar sentence. Provided with widely accessible and various exemplar sentences, it can strengthen the diversity and expressiveness of video captions in an intuitive and natural way.
- A novel Syntax Modulated Caption Generator (SMCG) is proposed, which relies on the syntactic information of the exemplar sentence to modulate the sequential word decoding procedure, and thus controls the syntactic structures of the predicted sentence while preserving the semantic meanings of the original video contents.

- We conduct the exemplar-based video captioning experiments by collecting auxiliary exemplar sentences for two public video datasets. Extensive experimental results demonstrate that our model can generate captions which not only precisely describe video contents, but also possess similar syntactic structures to the exemplar sentences.

2 RELATED WORK

2.1 Video Captioning

Early works [10, 14, 25, 26, 34, 39, 45] on video captioning adopted template-based methods, which first define some specific rules for language grammar, and then generate captions by associating words detected from visual contents with predefined sentence templates. Currently, with the development of deep neural networks, sequence learning methods with an encoder-decoder architecture [31] have been widely adopted for video captioning. Based on such an architecture, numerous improvements were introduced, such as sequential encoding of video features [30], soft attention [40], visual-semantic embedding [20], dual learning [33], hierarchical recurrent neural network [1, 42], multimodal feature fusion [38], multi-task learning [22], reinforcement learning [23, 36], etc.

Although promising results have been achieved, template-based methods constrain the video captions with fixed syntactic structures, while sequence learning methods generate sentence descriptions by fitting the monotonous linguistic patterns of captions from the training set. All of those methods ignore that one single video can be described in a variety of sentences with diverse expressions. In this paper, we propose a problem of controllable video captioning with an exemplar sentence, which is expected to control the syntactic structures of the predicted captions by imitating different exemplar sentences, and thus strengthen the expressiveness and diversity of video captions.

2.2 Controllable Captioning with Auxiliary Information Guidance

Recently, some research works start to utilize auxiliary information to control the captioning procedure. A part of works focused on generating captions satisfying certain stylistic requirements such as being humorous or exhibiting a particular sentiment [9, 19, 41]. Hence, style indicators or labels were introduced to guide the caption generator. However, those works still typically assume a known, finite set of values that the style attribute can take on, which may still limit the diversity and flexibility of caption generation. Besides, to make the generated captions more diverse and accurate, Deshpande *et al.* leveraged the quantized Part-of-Speech (POS) tag sequence sampled from a given benchmark to condition word prediction at the decoding recurrent model [6]. Wang *et al.* tried to predict the POS sequence tag by tag from the input video, and then embed them as a global POS representation to gate the inputs of the sentence decoder for syntax control [32]. With manually altering the predicted POS tag sequence, Wang *et al.* showed that they can obtain captions with different syntaxes. However, controlling caption generation by manually manipulating POS tag needs specialized linguistic knowledge and is not intuitive, which can hardly be applied in practical scenarios.




Figure 2: An overview of our proposed model, which is realized in an encoder-decoder-reconstructor architecture. The video semantic and sentence syntactic encoders extract visual semantic and sentence syntactic information from the input video and the given exemplar sentence, respectively. The syntax modulated caption generator attentively summarizes the video semantic information and relies on the syntactic information to modulate the semantic word decoding procedure for the exemplar-based caption generation. To further preserve video semantic and exemplar sentence syntactic information, the video and syntax reconstructors reproduce the original video features and exemplar syntax tokens from the generated caption, respectively. Here \textcircled{N} denotes our proposed modulation network, and \textcircled{A} denotes the soft attention mechanism.

Our proposed controllable video captioning with an exemplar sentence task is different from the above works. The syntactic guidance for video captioning is derived from the given exemplar sentences, which are easy to access and have various syntactic structures. Thus, it can naturally extend the diversity of video captioning. Moreover, controlling caption generation with an exemplar sentence is more intuitive for people to realize and evaluate.

3 THE PROPOSED MODEL

In this paper, we propose a novel model to solve the controllable video captioning with an exemplar sentence task. As shown in Figure 2, our proposed model, which is realized in an encoder-decoder-reconstructor architecture, consists of a pair of semantic and syntactic encoders, a syntax modulated caption generator, and a pair of video and syntax reconstructors. Please note that the whole architecture is fully coupled together and can therefore be trained in an end-to-end manner.

3.1 Semantic and Syntactic Encoders

Video Semantic Encoder. Given an input video, we first use one pretrained CNN to encode each video frame into a fixed-length representation. In this way, the given video is encoded as a sequential representation $\mathbf{V} = \{v_1, v_2, \dots, v_m\}$, where m denotes the total number of the video frames. To incorporate the context information, an LSTM is used to aggregate the sequential representation \mathbf{V} as follows:

$$\{h_i^v, c_i^v\} = \text{LSTM}_v^E(v_i, h_{i-1}^v, c_{i-1}^v). \quad (1)$$

Hence, we obtain the context incorporated video sequence representation $\mathbf{H}^v = \{h_1^v, h_2^v, \dots, h_m^v\}$, with \mathbf{H}^v encoding the semantic information of the video sequence.

Sentence Syntactic Encoder. To syntactically control the target caption generation, we first use a sentence parser to extract the constituency parse tree [13, 18, 28] of the given exemplar sentence. For the convenience of further processing, the leaf nodes (i.e., word tokens) are removed from the extracted parse tree. For example, as shown in Figure 2, the obtained parse tree for the exemplar sentence “The girl on the scooter riding in the summer sun in the park” is “(ROOT (FRAG (NP (DT) (NN)) (PP (IN) (NP (NP (DT) (NN)) (VP (VBG) (PP (IN) (NP (NP (DT) (NN) (NN)) (PP (IN) (NP (DT) (NN))))))))))”. We regard the parse tree of the exemplar sentence as a syntactic sequence $S = \{s_1, s_2, \dots, s_n\}$, where each element, such as “ROOT”, “NP”, and the bracket “(” or “)” is taken as an independent syntax token in S . Another LSTM is then introduced to encode the exemplar syntactic sequence:

$$\{h_j^s, c_j^s\} = \text{LSTM}_s^E(s_j, h_{j-1}^s, c_{j-1}^s), \quad (2)$$

where s_j is the embedding of the j -th syntax token. As such, the syntactic representation $\mathbf{H}^s = \{h_1^s, h_2^s, \dots, h_n^s\}$ is obtained, which will be taken as the syntactic control signal for the further exemplar-based caption generation.

3.2 Syntax Modulated Caption Generator

Conventional video captioning models rely on a recurrent neural network (for example, an LSTM) to decode the sentence word by word based on the video semantic inputs. However, for the proposed exemplar-based video captioning task, one important thing is to introduce the exemplar syntactic information into the recurrent model so as to generate syntax customized captions, and meanwhile preserve video semantics.

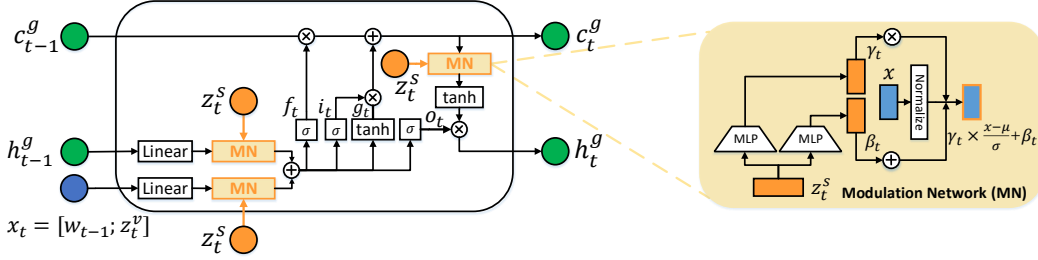


Figure 3: The architecture of our proposed syntax modulated caption generator and the detailed modulation network. Conditioned on the syntactic information, the modulation network manipulates the gate and cell units of the LSTM, and therefore makes the predicted captions imitate the syntactic structure of the exemplar sentence.

Inspired by the existing approaches which modulate the normalization parameters of ConvNet under the conditional input guidance for specific tasks (e.g., VQA [5], temporal sentence grounding [43], and image style transfer [7]), in this work, we propose a novel Syntax Modulated Caption Generator (SMCG) to modulate the gate and cell units of the LSTM with the syntactic cues extracted from one exemplar sentence. The purpose of SMCG is to let the exemplar syntax to manipulate the LSTM gates and cells by scaling them up or down, negating them, or shutting them off [5], so as to control the states of the sentence decoding LSTM as well as the syntactic structures of the decoded sentences. While in this procedure, the video semantic contents are barely modified or mixed with the syntactic information, and are therefore able to be maintained in the predicted captions.

Specifically, suppose that the hidden state of the caption decoding LSTM at timestep $t - 1$ is \mathbf{h}_{t-1}^g , we first attentively summarize the encoded video semantic and sentence syntactic representations [40] as follows:

$$\begin{aligned} z_t^v &= \text{Attention}_v(\mathbf{h}_{t-1}^g, \mathbf{H}^v), \\ z_t^s &= \text{Attention}_s(\mathbf{h}_{t-1}^g, \mathbf{H}^s). \end{aligned} \quad (3)$$

As shown in Figure 3, the attentively summarized semantic feature z_t^v and the embedding w_{t-1} of the previously predicted word w_{t-1} are concatenated as the input $\mathbf{x}_t = [w_{t-1}; z_t^v]$ for updating the states of the LSTM. With the proposed syntactic modulation strategy, the computing flow of the sentence decoding LSTM in SMCG is defined as:

$$\begin{pmatrix} f_t \\ i_t \\ o_t \\ g_t \end{pmatrix} = \text{MN}(\mathbf{W}^h \mathbf{h}_{t-1}^g; \gamma_t^h, \beta_t^h) + \text{MN}(\mathbf{W}^x \mathbf{x}_t; \gamma_t^x, \beta_t^x) + \mathbf{b}, \quad (4)$$

$$\begin{aligned} c_t^g &= \sigma(f_t) \odot c_{t-1}^g + \sigma(i_t) \odot \tanh(g_t), \\ h_t^g &= \sigma(o_t) \odot \tanh(\text{MN}(c_t^g; \gamma_t^c, \beta_t^c)), \end{aligned}$$

where $\text{MN}(\mathbf{x}; \boldsymbol{\gamma}_t, \boldsymbol{\beta}_t)$ denotes the modulation network, which firstly performs lay normalization of the input vector and then scales and shifts the normalized vector with regard to the attentively summarized syntactic feature z_t^s :

$$\begin{aligned} \text{MN}(\mathbf{x}; \boldsymbol{\gamma}_t, \boldsymbol{\beta}_t) &= \boldsymbol{\gamma}_t \cdot \frac{\mathbf{x} - \boldsymbol{\mu}(\mathbf{x})}{\boldsymbol{\sigma}(\mathbf{x})} + \boldsymbol{\beta}_t \\ \boldsymbol{\gamma}_t &= f_\gamma(z_t^s), \quad \boldsymbol{\beta}_t = f_\beta(z_t^s), \end{aligned} \quad (5)$$

where $\boldsymbol{\mu}(\mathbf{x})$ and $\boldsymbol{\sigma}(\mathbf{x})$ denote the mean and standard deviation, respectively. As illustrated in the right part of Figure 3, f_γ and f_β

denote the two independent multi-layer perceptrons (MLPs) with tanh activation, which yield the modulation vectors $\boldsymbol{\gamma}$ and $\boldsymbol{\beta}$ to control the feature scaling and shifting, respectively. Since there are three pairs of modulation vectors $(\boldsymbol{\gamma}_t^h, \boldsymbol{\beta}_t^h)$, $(\boldsymbol{\gamma}_t^x, \boldsymbol{\beta}_t^x)$, and $(\boldsymbol{\gamma}_t^c, \boldsymbol{\beta}_t^c)$ in Eq. (4), three pairs of MLPs with independent parameters are used in our syntax modulated caption generator accordingly.

With the updated syntactically modulated hidden state \mathbf{h}_t^g , the t -th word in the video caption is predicted by:

$$P(w_t | w_{<t}, V, S; \theta) = \text{softmax}(\mathbf{W}^g \mathbf{h}_t^g + \mathbf{b}^g). \quad (6)$$

3.3 Video and Syntax Reconstructors

The generated syntax customized video captions in our task should not only express the video semantic meanings but also follow the exemplar sentence syntaxes. Therefore, we propose a pair of video and syntax reconstructors stacking on the caption generator, aiming to reconstruct the video features and the syntax tokens from the hidden states of the caption generator, respectively. The proposed reconstructors are expected to bridge the semantic gaps between the generated captions and the video contents, and meanwhile to close the syntactic distances between the captions and the exemplar sentences.

In this paper, the two reconstructors are both realized in an LSTM architecture, and we only take the video reconstructor [33] as an example for demonstration. Supposing that the hidden state of the video reconstruction LSTM at the $(i-1)$ -th timestep is $\mathbf{h}_{i-1}^{\bar{v}}$, the key hidden states from the caption generator are firstly attentively summarized as follows:

$$z_i^{\bar{v}} = \text{Attention}_{\bar{v}}(\mathbf{h}_{i-1}^{\bar{v}}, \mathbf{H}^g), \quad (7)$$

where $\mathbf{H}^g = \{\mathbf{h}_1^g, \mathbf{h}_2^g, \dots\}$ denotes the collected hidden state sequence of the caption generator. The attentively summarized hidden state $z_i^{\bar{v}}$ is then taken as input to reproduce the original video feature \mathbf{v}_i at the i -th timestep:

$$\mathbf{h}_i^{\bar{v}}, \mathbf{c}_i^{\bar{v}} = \text{LSTM}_{\bar{v}}^R(z_i^{\bar{v}}, \mathbf{h}_{i-1}^{\bar{v}}, \mathbf{c}_{i-1}^{\bar{v}}). \quad (8)$$

Here we regard the yielded hidden state $\mathbf{h}_i^{\bar{v}}$ as the reconstruction of the video feature \mathbf{v}_i .

For the syntax reconstructor, the corresponding hidden state $\mathbf{h}_i^{\bar{s}}$ is used to predict the syntax token s_i in the exemplar syntactic sequence S . With the introduced two reconstructors, the aforementioned encoders are encouraged to embed more useful semantic and syntactic information, and the generator is encouraged to absorb

more valid semantic and syntactic information for the exemplar-based video captioning task.

3.4 Training and Inference

Training. Formally, the proposed architecture is trained by minimizing three loss terms defined in Eq. (9), which involves the caption generation loss, the video reconstruction loss, and the syntax reconstruction loss:

$$\mathcal{L}(\theta, \theta_{rec}) = \sum_{(V, S) \in \Gamma_{train}} \underbrace{(-\alpha \log P(W|V, S; \theta))}_{\text{caption generation}} + \underbrace{\lambda \mathcal{L}_{rec}^v(V, H^{\vec{v}}; \theta, \theta_{rec}^v)}_{\text{video reconstruction}} - \underbrace{\eta \log P(S|V, S; \theta, \theta_{rec}^s)}_{\text{syntax reconstruction}}. \quad (9)$$

α , λ , and η are the hyper-parameters to balance the contributions of different loss terms. The video feature reconstruction loss \mathcal{L}_{rec}^v is calculated by averaging the Euclidean distances between the original and reconstructed video features:

$$\mathcal{L}_{rec}^v(V, H^{\vec{v}}) = \frac{1}{m} \sum_{i=1}^m \text{Euclidean}(v_i, h_i^{\vec{v}}). \quad (10)$$

The caption generation and syntax reconstruction losses are realized by the typical negative log-likelihood loss adopted by most video captioning methods.

Please note that we do not have the corresponding video caption which shares the same sentence syntax as the given exemplar sentence. Therefore, when training our model, the ground-truth caption accompanied with the video is also used as the exemplar sentence. The proposed SMCG, the syntax reconstructor, and the video reconstructor are then trained together to generate the ground-truth caption, reproduce its syntactic information, and reconstruct the video features, respectively.

Inference. During inference, we take the video and one randomly sampled exemplar sentence as the inputs of our model, then forward them through the encoders and the caption generator, and finally generate the exemplar-based video caption.

4 EXPERIMENTS

4.1 Datasets and Exemplar Sentence Collection

Since there is no available benchmark datasets for exemplar-based video captioning, we augment two public video captioning datasets MSRVT [37] and ActivityNet Captions [16] by collecting auxiliary exemplar sentences for each video in these datasets. We will first give a brief introduction of the two datasets and then present our exemplar sentence collection procedure.

MSRVT [37]. The MSRVT is a large-scale dataset for video captioning. In this paper, we use the initial version of MSRVT, which contains 10k videos from 20 categories. Each video is annotated with 20 ground-truth captions, resulting in a total of 200K video-caption pairs. We use the public splits for training and testing, i.e., 6,513 videos for training, 497 for validation, and 2,990 for testing.

ActivityNet Captions [16]. The ActivityNet Captions is originally exploited as the benchmark for dense video captioning. This dataset contains 20K videos in total, and for each video, the temporal segments and their paired caption sentences are annotated. Since we do not consider dense video captioning in this paper, we split the caption-annotated segments from the original videos, and

perform video captioning on them. In this way, 54,926 segment-caption pairs are collected, in which 37,421 segments split from the public training set are used for training, and 17,505 segments from the validation set are used for testing.

Exemplar Sentence Collection. To perform our exemplar-based video captioning, the exemplar sentences should have various syntactic structures. We find that previous work [8] crawled sentence descriptions from online stock photography website Shutterstock for unsupervised image captioning. The image descriptions they crawled are written by the image composer and free of syntactic constraints, which are suitable for the exemplar sentences we need. We download their collected 2,322,628 image descriptions, and filter the descriptions that are too short (< 8 words) or too long (> 30 words), yielding a total of 761,582 exemplar sentences. For each video/segment in MSRVT and ActivityNet Captions, we randomly sample 20 collected descriptions as its exemplar sentences for captioning.

4.2 Evaluation Metrics

Different from the conventional video captioning, our exemplar-based video captions should describe the semantic contents of videos while imitate the syntactic structures of the exemplar sentences. Therefore, the evaluation of the exemplar-based captions should also be established in two aspects, i.e., semantic aspect and syntactic aspect.

For semantic evaluation, we firstly remove the stop words from each sentence and compute the averaged GloVe [24] embeddings of the remaining words as its sentence embeddings. In the embedding space, the cosine similarity (COS) between the predicted exemplar-based caption and the original ground-truth video captions within the dataset is used to evaluate their semantic similarity. In addition, we also report the typical video captioning metrics BLEU@4, METEOR, ROUGE-L, and CIDEr as a reference.

For syntactic evaluation, we first parse the exemplar sentence and its corresponding predicted caption using Stanford CoreNLP [18], and then compute the syntactic Tree Edit Distance (TED) [44] between their constituency parse trees after removing word tokens. The smaller TED value means higher syntactic similarity. Although there are 20 exemplar sentences for each video, not each exemplar syntactic structure is suitable for describing the video. Therefore, we take the averaged TED between the 20 (exemplar, prediction) pairs as the result for each video, and diminish the impact of some unreliable exemplar sentences.

4.3 Implementation Details

We use the public Stanford parser [18] toolkit to process each exemplar sentence, and get the off-the-shelf constituency parse tree. Also note that the stanford parser will not be tuned in our overall architecture, but other available sentence parsing networks [13, 28] can also be applied in our work, and be jointly trained in an end-to-end fashion.

To extract video features, we feed the static video frames to the Inception-Resnet-v2 network [29] pretrained on the ILSVRC-2012-CLS image classification dataset [27], thus yielding a 1,536 dimensional feature vector for each frame. Considering both the video length distribution and model memory footprint, we take evenly

Table 1: Performance comparisons on the MSRVT dataset (*).

Method	TED↓	COS↑	B@4↑	R↑	M↑	C↑
Exemplar-based	0.00	52.38	0.39	16.65	6.02	0.26
Caption-based	15.91	69.49	37.20	58.64	26.39	40.24
Concat-based	3.76	66.18	4.54	29.69	15.24	8.92
SMCG	3.46	67.02	4.89	29.62	15.64	10.67
SMCG+VideoRec	3.25	67.26	5.07	29.92	15.95	12.07
SMCG+SyntaxRec	3.19	67.15	4.65	29.45	15.49	10.50
SMCG+AllRec	3.12	67.97	5.01	30.41	16.25	12.22
SMCG(BN)+AllRec	23.09	61.22	5.00	28.31	13.92	1.92
SMCG(POS)+AllRec	6.33	67.14	5.08	29.92	15.83	12.03

*: TED is reported with absolute value, while other metrics are reported in percentage % values. ↓ means smaller values are better, and ↑ is on the contrary.

spaced 30 and 100 features to represent videos in the MSRVT and ActivityNet Captions datasets, respectively. Shorter videos of less than 30 or 100 features are padded with zero vectors.

In our model, we set the word embedding size and all the LSTMs’ hidden size as 512, and the embedding size of the syntax token is set as 256. The trade-off parameters α , λ , and η of our model to balance different loss terms are set as 1.0, 1.0, and 4.0, respectively. For training, the model is optimized by the Adam [12] optimizer.

4.4 Performance Comparison

Since there is no existing work for exemplar-based video captioning, we compare our proposed method with three baseline methods as follows. “**Exemplar-based**” method directly outputs the exemplar sentence as the predicted caption, while does not consider the video contents. “**Caption-based**” method applies the general sequence-to-sequence video captioning model [40], while does not consider the exemplar syntax information. “**Concat-based**” method takes the main architecture of “Caption-based” method, but also leverages another encoder to sequentially encode syntax information as we did. Then, the video context features and syntactic context features are directly concatenated as the input to the caption decoder. For comparison, our proposed full model is denoted by SMCG+AllRec.

As shown in the upper part of Table 1 and Table 2, Exemplar-based method gets a TED score 0.0, which accords with that it simply takes the given exemplar sentence as the predicted caption. However, since it ignores video contents, the semantic metrics such as COS, BLEU@4, ROUGE-L, METEOR, and CIDEr are much lower than other methods. In contrast, Caption-based method just considers describing video contents while neglects syntactic requirements of the predicted caption, making higher semantic metric scores but much larger TED values.

For exemplar-based video captioning, the small TED scores of the Concat-based method verify that its predicted captions have indeed imitated the syntactic structures of the given exemplar sentences. Since general caption evaluation metrics like ROUGE take n-gram similarity between sentences into consideration, they will be inevitably influenced by the exemplar syntactic constraints. As such, compared with the Caption-based method, the general captioning metric values of the Concat-based method significantly decrease. Sentence cosine similarity only considers word semantics in sentences, so its values are more stable. However, directly concatenating the syntactic information with video features in

Table 2: Performance comparisons on the ActivityNet Captions dataset.

Method	TED↓	COS↑	B@4↑	R↑	M↑	C↑
Exemplar-based	0.00	57.84	0.07	6.74	2.57	0.98
Caption-based	20.12	73.17	3.36	20.44	8.93	23.92
Concat-based	3.71	64.27	0.49	10.90	4.86	7.83
SMCG	3.41	64.91	0.54	10.91	4.86	8.95
SMCG+VideoRec	3.64	66.46	0.54	10.91	4.98	9.14
SMCG+SyntaxRec	3.30	67.50	0.53	10.88	4.92	8.06
SMCG+AllRec	3.05	68.23	0.53	11.11	5.11	9.31
SMCG(BN)+AllRec	8.48	61.11	0.39	9.36	3.87	3.32
SMCG(POS)+AllRec	6.02	68.17	0.54	11.10	5.09	9.29

the captioning procedure will cause disturbances in describing semantic video contents, thus yielding smaller COS values of the Concat-based method than the Caption-based method.

Compared with the Concat-based method, our proposed model SMCG+AllRec steps further in generating syntax customized and semantic preserved video captions. As shown in Table 1 and Table 2, SMCG+AllRec achieves smaller TED values and higher semantic metric scores on both two datasets. On one hand, our method leverages the syntactic information to meticulously modulate the LSTM updating procedure while does not directly alter the semantic inputs of the decoder. On the other hand, the reconstruction of video features and syntax token sequences can further help preserve the video semantic contents and exemplar syntactic structures in the predicted captions, respectively. By incorporating both of the above aspects, our proposed SMCG+AllRec consistently achieves the best results.

4.5 Ablation Studies

To better demonstrate the effectiveness of our model design, we perform several ablation studies as follows:

SMCG: We only keep the semantic and syntactic encoders, and the syntax modulated caption generator in the proposed architecture, with the two reconstructors removed.

SMCG+VideoRec: Based on SMCG, the video reconstructor is included to help the model training.

SMCG+SyntaxRec: Based on SMCG, the syntax reconstructor is included to help the model training.

SMCG+AllRec: Our proposed full model.

SMCG(BN)+AllRec: In this setting, the feature normalization procedure in our proposed syntactic modulation operation is replaced by batch normalization, instead of the originally used layer normalization.

SMCG(POS)+AllRec: In this setting, the sentence syntactic information is represented with the POS sequence of words as used in the previous work [32], instead of the constituency parse tree in our model.

The ablation experimental results are also shown in Table 1 and Table 2. By incorporating video reconstruction in SMCG+VideoRec, the semantic metric values improve on both the datasets, which demonstrates the benefits of the video feature reconstruction for preserving video semantics. In addition, incorporating syntax reconstruction in SMCG+SyntaxRec yields smaller TED scores than SMCG, which indicates the syntax reconstruction can further help the syntax customization. Interestingly, we also find that incorporating video reconstruction makes smaller TED scores on the




Figure 4: Subjective evaluations on the MSRVT datasets.

MSRVTT dataset, and incorporating syntax reconstruction yields higher COS scores on the ActivityNet captions dataset. Such observations show that a better understanding of video contents may help the model organize sentence syntaxes, and incorporating valid syntax constraints may drive the model to present video contents better. Combining both of the two reconstructions together, our full model SMCG+AllRec yields further improvements in both TED and COS scores.

By changing the feature normalization method from channel-wise layer normalization to batch normalization, we find that the performances of SMCG(BN)+AllRec drop drastically compared to SMCG+AllRec. Although batch normalization has made achievements in CNN and RNN architectures [3, 5, 7], such a normalization scheme is not very suitable for our proposed SMCG. Since one batch in our modeling training procedure contains several (video, exemplar syntax) pairs, the normalization across the whole batch will make disturbances between different training instances and degrade the model performance.

Moreover, we also find that replacing sentence syntactic representation from the constituency parse tree to the POS sequence makes the TED of SMCG(POS)+AllRec increase a lot compared to SMCG+AllRec. Such results indicate that constituency parses represent sentence syntaxes better than POS sequences, and are more suitable for our exemplar-based syntax customized video captioning.

4.6 Subjective Evaluation

Although COS and TED metrics can evaluate the quality of the predicted sentences from semantic and syntactic aspects, it is more reliable to conduct a human study to judge whether our method can produce reasonable semantic-preserved and syntax-controlled captions. Therefore, we invite 18 evaluators (9 females and 9 males) from different education backgrounds to perform the subjective evaluation. For both of the two datasets, we randomly choose 300 samples, and each of them contains one video and one exemplar sentence. The evaluators are required to watch the video first, and then evaluate the predicted captions from the semantic, syntactic and fluency aspects by grading it on three scales (0, 0.5, 1). Grade 1 means that the predicted caption presents the video semantics well/follows the exemplar syntax well/has good linguistic fluency, 0.5 means not bad, and 0 means poor. Three groups of captions generated by the Caption-based method, the Exemplar-based method, and our proposed SMCG+AllRec method are graded by these evaluators for comparisons. The ratios of different grades on the MSRVT dataset are illustrated in Figure 4, while the results on the ActivityNet Captions dataset are provided in our supplemental material.

Specifically, we can observe that for the semantic subjective evaluation, our proposed SMCG+AllRec method achieves comparable

Table 3: Captioning Diversity Evaluation.

Method	MSRVTT		ActivityNet Captions	
	LSA	Self-CIDEr	LSA	Self-CIDEr
Caption-based	0.0	0.0	0.0	0.0
Exemplar-based	0.7437	0.9737	0.7431	0.9741
SMCG+AllRec	0.5896	0.8628	0.5862	0.8748

results with the Caption-based baseline, and there are 83.4% and 85% records graded equal or larger than 0.5 for these two methods, respectively. It shows that both of these two methods can generate captions describing the video semantic contents well. However, on the conventional captioning metrics (e.g., B@4) as shown in Table 1 and Table 2, our method and the Caption-based baseline have significant performance gap. The different evaluation results indicate that the conventional captioning metrics, which are influenced by sentence syntaxes/n-gram characteristics, indeed have limitations on evaluating the exemplar-based video captioning task. For the syntactic evaluation, the Caption-based method gets 99.8% records graded at 0.0, which shows that this method cannot make the predicted captions follow the exemplar sentence syntaxes. In contrast, the grade 0.0 only occupies 3% syntactic evaluation records for our SMCG+AllRec method, which verifies that our generated captions can not only present the video semantics well, but also follow the exemplar syntaxes well. As for the Exemplar-based baseline, it gets (99.4%,0.6%,0%) and (0%,0%,100%) records on (0,0.5,1) grades for semantic and syntactic evaluation. It is evident that the given exemplar sentences are totally irrelevant to video semantics.

When evaluating the sentence fluency of the predicted captions, there are 98.3%, 95.3%, and 61.2% records graded at 1.0 for the Caption-based method, the Exemplar-based method, and the proposed SMCG+AllRec method, respectively. The Caption-based method generates captions by fitting the simple and monotonous sentence syntaxes (like “A is doing B”) of the training captions in the dataset. Such simple syntactic structures are easy to learn by the captioning model, therefore making the generated captions have good linguistic fluency. The provided exemplar sentences are human-edited sentences with more various and complex syntactic structures. Even a small fraction (4.7%) of the exemplar sentences are not very easy to follow and are graded at 0.5 in fluency evaluation. By making the generated captions follow such complex exemplar syntaxes, our proposed SMCG+AllRec method can still generate fluent captions to a large extent with only 3.2% records graded at 0.0. It shows that our proposed method can make the captions maintain good linguistic fluency in the meanwhile of imitating the exemplar syntaxes.

Groundtruth	Exemplar sentence	Caption-based	Concat-based	Ours
Man reviewing a new car.	Collagen sheet mask in a blue pack with handmade tulips on the bamboo chair. Curious abyssinian cat sitting on a toilet bowl and looking down.	A man is talking about a car.	Man driving car in a blue car with other cars in a parking car. Black black car driving on a car road and driving down.	Man speaking review of a new car with different features of a car interior. Black haired man sitting in a car seat and driving down.

(a)

Groundtruth	Exemplar sentence	Caption-based	Concat-based	Ours
Then the boy walks and pass on side the toddler.	Photo of washbasin in loft bathroom arranged in warm colors. Fishermen fishing during sea low tide and looking at the horizon when the sun is starting to go down and the fisherman boats are coming back.	A baby is playing on the swing.	Woman in front of monkey swing followed by several times. Theres child in front several swing and walking on the swing when the boy is trying to swing back and the girl hes are going back.	Baby in swing at playground equipment followed by one boy. Children line on monkey high speed and playing in the playground when a girl is trying to climb back and the girl whos are standing again.

(b)

Figure 5: Qualitative results for some exemplar-based video captioning cases. In the first column of each case, we provide the ground-truth captions. In the second column, we give two exemplar sentences for syntactic constraints. The predicted captions by the Caption-based method, the Concat-based method, and our method are presented in the third, fourth, and fifth columns, respectively. Case (a) is from the MSRVTT dataset, and case (b) is from the ActivityNet Captions dataset.

4.7 Captioning Diversity Evaluation

Given different exemplar sentences (we have 20 exemplar sentences for each video in our collected dataset) and one single video, our SMCG+AllRec can generate different captions describing the semantic contents of the video, and thus it can be regarded as diverse video captioning. To evaluate the diversity of the generated captions, we follow the evaluation metrics introduced in the work [35], and report the LSA and Self-CIDEr based diversities of different methods in Table 3. Since the Caption-based method can only generate one single caption for one video no matter what exemplar sentence is provided, the diversity values of its generated captions are always 0.0. The Exemplar-based method directly outputs the exemplar sentences as the caption prediction results, and the provided 20 exemplar sentences in our experiments are different from each other. Therefore, the diversity values of the Exemplar-based method can be seen as the upper bounds for the syntax-controllable video captioning. For our SMCG+AllRec method, it generated captions achieve high Self-CIDEr diversities 0.8628 and 0.8748 on the MSRVTT and ActivityNet Captions dataset, respectively. Meanwhile, the LSA based diversities are also comparable to other diverse video captioning methods as reported in [35]. Therefore, the results verify the contribution of our method on strengthening the diversity of video captioning.

4.8 Qualitative Results

Finally, we show some qualitative results in Figure 5 for the proposed exemplar-based video captioning. It can be observed that by providing different exemplar sentences, our method can generate different video captions with various syntactic structures, which enhances the diversity and expressiveness of the video descriptions compared to the Caption-based method in the third column. Meanwhile, our method can generate more accurate and fluent sentence descriptions than the Concat-based method, which indicates the

effectiveness of our SMCG design in preserving video semantics and controlling sentence syntaxes.

In addition, we also provide a failure case of the exemplar-based video captioning, as shown in the red dotted box in Figure 5. The given exemplar sentence of this case is relatively long and complex, making our predicted video caption unfluent and confusing. It demonstrates that the exemplar sentence selection has a great influence on the final video captioning results. Different videos are suited to different descriptive manners, and we should choose appropriate exemplar sentences for them to match. However, establishing the matching relations between semantic video contents and syntactic sentence structures is a challenging problem, which is not the main focus of this paper, and will be explored in the future. For more qualitative results, please refer our supplemental material.

5 CONCLUSION

In this paper, we proposed a novel exemplar-based video captioning problem, which aims to generate one natural sentence describing the video content and meanwhile following the syntactic structure of the provided exemplar sentence. A novel syntax modulated caption generator (SMCG) was proposed to leverage encoded syntactic information to modulate the LSTM gates and cells for decoding each word, therefore controlling the predicted sentence syntax. Moreover, the video and syntax reconstructors are further introduced to help preserve video semantics in predicted captions and ensure the syntax customization. Extensive experiments demonstrate that our proposed method can yield different video descriptions of various syntax structures with respect to different exemplar sentences, hence enhancing the diversity of video captions.

6 ACKNOWLEDGMENTS

This work was supported by National Natural Science Foundation of China Major Project (No. U1611461) and National Key R&D Program of China under Grand No. 2018AAA0102000.

REFERENCES

- [1] Lorenzo Baraldi, Costantino Grana, and Rita Cucchiara. 2017. Hierarchical boundary-aware neural encoder for video captioning. In *CVPR*.
- [2] Yangyu Chen, Shuhui Wang, Weigang Zhang, and Qingming Huang. 2018. Less is more: Picking informative frames for video captioning. In *ECCV*.
- [3] Tim Cooijmans, Nicolas Ballas, César Laurent, Çağlar Gülçehre, and Aaron Courville. 2016. Recurrent batch normalization. *arXiv preprint arXiv:1603.09025* (2016).
- [4] Abhishek Das, Satwik Kottur, José MF Moura, Stefan Lee, and Dhruv Batra. 2017. Learning cooperative visual dialog agents with deep reinforcement learning. In *ICCV*.
- [5] Harm De Vries, Florian Strub, Jérémie Mary, Hugo Larochelle, Olivier Pietquin, and Aaron C Courville. 2017. Modulating early visual processing by language. In *NeurIPS*.
- [6] Aditya Deshpande, Jyoti Aneja, Liwei Wang, Alexander G Schwing, and David Forsyth. 2019. Fast, Diverse and Accurate Image Captioning Guided By Part-of-Speech. In *CVPR*.
- [7] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. 2016. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629* (2016).
- [8] Yang Feng, Lin Ma, Wei Liu, and Jiebo Luo. 2019. Unsupervised Image Captioning. In *CVPR*.
- [9] Chuang Gan, Zhe Gan, Xiaodong He, Jianfeng Gao, and Li Deng. 2017. Stylenet: Generating attractive visual captions with styles. In *CVPR*.
- [10] Sergio Guadarrama, Niveda Krishnamoorthy, Girish Malkarnenkar, Subhashini Venugopalan, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2013. Youtube2text: Recognizing and describing arbitrary activities using semantic hierarchies and zero-shot recognition. In *ICCV*.
- [11] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [12] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [13] Nikita Kitaev and Dan Klein. 2018. Constituency parsing with a self-attentive encoder. In *arXiv preprint arXiv:1805.01052*.
- [14] Atsuhiko Kojima, Takeshi Tamura, and Kunio Fukunaga. 2002. Natural language description of human activities from video images based on concept hierarchy of actions. *International Journal of Computer Vision* 50, 2 (2002), 171–184.
- [15] Satwik Kottur, José MF Moura, Devi Parikh, Dhruv Batra, and Marcus Rohrbach. 2018. Visual coreference resolution in visual dialog using neural module networks. In *ECCV*.
- [16] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In *ICCV*.
- [17] Yehao Li, Ting Yao, Rui Hu, Tao Mei, and Yong Rui. 2016. Video ChatBot: Triggering Live Social Interactions by Automatic Video Commenting. In *ACM MM*.
- [18] Christopher Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven Bethard, and David McClosky. 2014. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*. 55–60.
- [19] Alexander Mathews, Lexing Xie, and Xuming He. 2018. Semstyle: Learning to generate stylised image captions using unaligned text. In *CVPR*.
- [20] Yingwei Pan, Tao Mei, Ting Yao, Houqiang Li, and Yong Rui. 2016. Jointly modeling embedding and translation to bridge video and language. In *CVPR*.
- [21] Yingwei Pan, Ting Yao, Houqiang Li, and Tao Mei. 2017. Video captioning with transferred semantic attributes. In *CVPR*.
- [22] Ramakanth Pasunuru and Mohit Bansal. 2017. Multi-task video captioning with video and entailment generation. *arXiv preprint arXiv:1704.07489* (2017).
- [23] Ramakanth Pasunuru and Mohit Bansal. 2017. Reinforced video captioning with entailment rewards. *arXiv preprint arXiv:1708.02300* (2017).
- [24] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*.
- [25] Anna Rohrbach, Marcus Rohrbach, Wei Qiu, Annemarie Friedrich, Manfred Pinkal, and Bernt Schiele. 2014. Coherent multi-sentence video description with variable level of detail. In *German conference on pattern recognition*. Springer, 184–195.
- [26] Marcus Rohrbach, Wei Qiu, Ivan Titov, Stefan Thater, Manfred Pinkal, and Bernt Schiele. 2013. Translating video content to natural language descriptions. In *ICCV*.
- [27] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, and Michael Bernstein. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [28] Yikang Shen, Zhouhan Lin, Athul Paul Jacob, Alessandro Sordani, Aaron Courville, and Yoshua Bengio. 2018. Straight to the tree: Constituency parsing with neural syntactic distance. *arXiv preprint arXiv:1806.04168* (2018).
- [29] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*.
- [30] Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to sequence-video to text. In *ICCV*.
- [31] Subhashini Venugopalan, Huijuan Xu, Jeff Donahue, Marcus Rohrbach, Raymond Mooney, and Kate Saenko. 2014. Translating videos to natural language using deep recurrent neural networks. *arXiv preprint arXiv:1412.4729* (2014).
- [32] Bairui Wang, Lin Ma, Wei Zhang, Wenhao Jiang, Jingwen Wang, and Wei Liu. 2019. Controllable Video Captioning with POS Sequence Guidance Based on Gated Fusion Network. In *ICCV*.
- [33] Bairui Wang, Lin Ma, Wei Zhang, and Wei Liu. 2018. Reconstruction network for video captioning. In *CVPR*.
- [34] Jingwen Wang, Wenhao Jiang, Lin Ma, Wei Liu, and Yong Xu. 2018. Bidirectional Attentive Fusion with Context Gating for Dense Video Captioning. In *CVPR*.
- [35] Qingzhong Wang and Antoni B Chan. 2019. Describing like humans: on diversity in image captioning. In *CVPR*.
- [36] Xin Wang, Wenhao Chen, Jiawei Wu, Yuan-Fang Wang, and William Yang Wang. 2018. Video captioning via hierarchical reinforcement learning. In *CVPR*.
- [37] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*.
- [38] Jun Xu, Ting Yao, Yongdong Zhang, and Tao Mei. 2017. Learning multimodal attention LSTM networks for video captioning. In *ACM MM*.
- [39] Ran Xu, Caiming Xiong, Wei Chen, and Jason J Corso. 2015. Jointly modeling deep video and compositional text to bridge vision and language in a unified framework. In *AAAI*.
- [40] Li Yao, Atousa Torabi, Kyunghyun Cho, Nicolas Ballas, Christopher Pal, Hugo Larochelle, and Aaron Courville. 2015. Describing videos by exploiting temporal structure. In *ICCV*.
- [41] Quanzeng You, Hailin Jin, and Jiebo Luo. 2018. Image captioning at will: A versatile scheme for effectively injecting sentiments into image descriptions. *arXiv preprint arXiv:1801.10121* (2018).
- [42] Haonan Yu, Jiang Wang, Zhiheng Huang, Yi Yang, and Wei Xu. 2016. Video paragraph captioning using hierarchical recurrent neural networks. In *CVPR*.
- [43] Yitian Yuan, Lin Ma, Jingwen Wang, Wei Liu, and Wenwu Zhu. 2019. Semantic Conditioned Dynamic Modulation for Temporal Sentence Grounding in Videos. In *NeurIPS*.
- [44] Kaizhong Zhang and Dennis Shasha. 1989. Simple fast algorithms for the editing distance between trees and related problems. *SIAM journal on computing* 18, 6 (1989), 1245–1262.
- [45] Wei Zhang, Bairui Wang, Lin Ma, and Wei Liu. 2019. Reconstruct and Represent Video Contents for Captioning via Reinforcement Learning. In *TPAMI*. <https://doi.org/10.1109/TPAMI.2019.2920899>