

# Tree-Augmented Cross-Modal Encoding for Complex-Query Video Retrieval

Xun Yang  
xunyang@nus.edu.sg  
National University of Singapore

Jianfeng Dong\*  
dongjf24@gmail.com  
Zhejiang Gongshang University

Yixin Cao  
caoyixin2011@gmail.com  
National University of Singapore

Xun Wang  
wx@zjgsu.edu.cn  
Zhejiang Gongshang University

Meng Wang  
wangmeng@hfut.edu.cn  
Hefei University of Technology

Tat-Seng Chua  
chuats@comp.nus.edu.sg  
National University of Singapore

## ABSTRACT

The rapid growth of user-generated videos on the Internet has intensified the need for text-based video retrieval systems. Traditional methods mainly favor the concept-based paradigm on retrieval with simple queries, which are usually ineffective for complex queries that carry far more complex semantics. Recently, embedding-based paradigm has emerged as a popular approach. It aims to map the queries and videos into a shared embedding space where semantically-similar texts and videos are much closer to each other. Despite its simplicity, it forgoes the exploitation of the syntactic structure of text queries, making it suboptimal to model the complex queries.

To facilitate video retrieval with complex queries, we propose a Tree-augmented Cross-modal Encoding method by jointly learning the linguistic structure of queries and the temporal representation of videos. Specifically, given a complex user query, we first recursively compose a latent semantic tree to structurally describe the text query. We then design a tree-augmented query encoder to derive structure-aware query representation and a temporal attentive video encoder to model the temporal characteristics of videos. Finally, both the query and videos are mapped into a joint embedding space for matching and ranking. In this approach, we have a better understanding and modeling of the complex queries, thereby achieving a better video retrieval performance. Extensive experiments on large scale video retrieval benchmark datasets demonstrate the effectiveness of our approach.

## CCS CONCEPTS

• **Information systems** → **Multimedia and multimodal retrieval**; **Video search**.

\*Corresponding Author.

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](mailto:permissions@acm.org).

SIGIR '20, July 25–30, 2020, Virtual Event, China

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8016-4/20/07...\$15.00

<https://doi.org/10.1145/3397271.3401151>

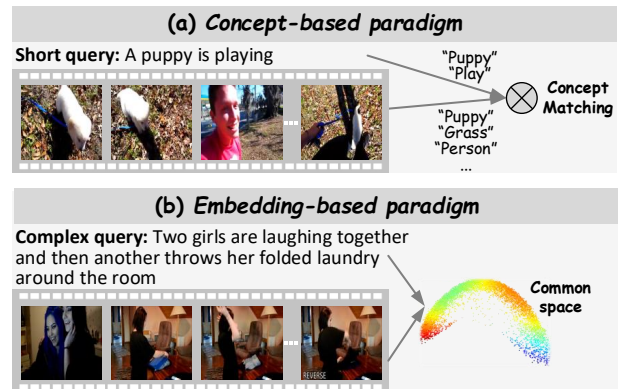


Figure 1: Concept-based paradigm vs. Embedding-based paradigm for text-based video retrieval.

## KEYWORDS

Multimedia retrieval, Video Search, Natural Language Understanding, Latent Tree Structure

### ACM Reference Format:

Xun Yang, Jianfeng Dong, Yixin Cao, Xun Wang, Meng Wang, and Tat-Seng Chua. 2020. Tree-Augmented Cross-Modal Encoding for Complex-Query Video Retrieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3397271.3401151>

## 1 INTRODUCTION

With the exponential growth of user-generated videos on the Internet, searching the videos of interest has been an indispensable activity in people’s daily lives. Meanwhile, text-based video retrieval has attracted world-wide research interests and achieved promising progress for retrieval with keyword-based simple queries [37]. However, the expression of text query has been transformed from the keyword-based mechanism to complex queries in recent years. A complex query is usually defined as a natural language query, e.g., “Two girls are laughing together and then another throws her folded laundry around the room”, which carries far more complex semantics than short queries. How to correctly understand the complex queries has become one of the key challenges in the multimedia information retrieval community.

Existing efforts on video retrieval with complex queries can be roughly categorized into two groups: 1) **Concept-based** paradigm [18, 24, 25, 31, 41, 52, 53], as shown in Figure 1 (a). It usually uses a

large set of visual concepts to describe the video content, then transforms the text query into a set of primitive concepts, and finally performs video retrieval by aggregating the matching results from different concepts [53]. Despite its efficiency, it is usually ineffective for complex long queries, since they carry complex linguistic context and cannot be simply treated as an aggregation of extracted concepts. Besides, it is also quite challenging to effectively train concept classifiers and select the relevant concepts. 2) **Embedding-based** paradigm [1, 7, 21, 27, 28, 43, 49], as shown in Figure 1 (b). Recent efforts proposed to learn a joint text-video embedding space [7, 27, 28, 30] to support video retrieval by leveraging the strong representation ability of deep neural networks [17, 34]. The natural language queries are usually transformed into dense vector representations by Recurrent Neural Networks (RNNs) [34] (e.g., Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU)) that are powerful for modeling sequence data. The videos are usually modeled as a temporal aggregation of frame/clip-level features, extracted from pre-trained Convolutional Neural Networks (CNNs) [17]. Both the queries and videos are mapped into a shared embedding space where semantically-similar videos and text queries are mapped to close points. Although the embedding-based methods have shown much better performance, simply treating queries holistically as one dense vector representations may obfuscate the keywords or phrases that have rich temporal and semantic cues.

Some prior works proposed to transform the complex queries into structured forms, e.g., semantic graph [22], to describe the spatial or semantic relations between concepts. However, such solutions usually require the text query to be well annotated with syntactic labels (e.g., part of speech (POS) tag) and rely on complex predefined rules to construct the structure of text queries, which make it hard to be applied in a new scenario with different linguistic expression patterns. Although so much efforts have been devoted to complex-query video retrieval, it still remains to be a very challenging task.

Towards this research goal, this paper aims to model complex queries in a more flexible structure to facilitate the joint learning of the representations of the queries and videos in a unified framework. Specifically, we develop a Tree-augmented Cross-modal Encoding (TCE) framework for video retrieval with complex queries. As shown in Figure 2, for the modeling of the complex query, we first recursively compose a Latent Semantic Tree (LST) to describe the query (e.g., *A baby plays with a fatty cat*) without any syntactic annotations, where each node (e.g., *a baby plays*) denotes a constituent in the complex query. We also propose a memory-augmented node scoring and selection method to inject linguistic context into the construction of LST. We then design a tree-augmented query encoder that identifies the informative constituent nodes in LST and aggregates the constituent embeddings into the structure-aware query representation. For the modeling of the videos, we introduce a temporal attentive video encoder that first models the temporal dependence and interaction between frames and then attentively aggregates the frame embeddings into the temporal-attentive video representation. Finally, both the user queries and videos are mapped to a text-video joint embedding space where semantically-similar videos and text queries are mapped to close points. All the modules are jointly optimized in an end-to-end fashion using only the paired query-to-video supervision. We evaluate the proposed approach

on two large-scale text-to-video retrieval datasets, which clearly demonstrates the effectiveness of each component in our approach. The contributions of this paper are roughly summarized as follows:

- We develop a novel complex-query video retrieval framework that can automatically compose a flexible tree structure to model the complex query and derive the query and video representations in a joint text-video embedding space.
- We design a memory-augmented node scoring and selection method to explore linguistic context for the tree construction. We also introduce the attention mechanism into the encodings of complex queries and videos, which can identify the informative constituent nodes and frames.
- We conduct extensive experiments on large-scale datasets to demonstrate that our approach can achieve state-of-the-art retrieval performance.

## 2 RELATED WORK

In this section, we briefly introduce two representative research directions in text-based video retrieval. One is the concept based methods and the other one is the embedding based methods.

Concept based methods [18, 24, 25, 31, 41] mainly rely on establishing cross-modal associations via concepts [12]. Markatopoulou *et al.* [24, 25] first utilized relatively complex linguistic rules to extract relevant concepts from a given query and used pre-trained CNNs to detect the objects and scenes in video frames. Then the similarity between a given query and a specific video is measured by concept matching. Ueki *et al.* [41] depended on a much larger concept vocabulary. In addition to pre-trained CNNs, they additionally trained SVM-based classifiers to automatically annotate the videos. Snoek *et al.* [38] trained a more elegant model, called VideoStory, from freely available web videos to annotate videos, while they still represented the textual query by selecting concepts based on part-of-speech tagging heuristically. Despite the promising performance, the concept based methods still face many challenges, e.g., how to specify a set of concepts and how to extract relevant concepts for both textual queries and videos. Moreover, the extraction of concepts from videos and textual queries are usually treated independently, which makes it suboptimal to explore the relations between two modalities. In contrast, our method is concept free and jointly learns the representation of textual queries and videos.

Deep learning technologies have been popularly explored for video retrieval recently [7, 20, 21, 27, 28, 30, 43, 45, 49]. Most works proposed to embed textual queries and videos into a common space, and their similarity is measured in this space by distance metric, e.g., cosine distance. For textual query embedding, the word2vec models pre-trained on large-scale text corpora are increasingly popular [27, 28, 35, 43]. However, they ignored the sequential order in textual queries. To alleviate this, Mithun *et al.* [30] utilized GRU for modeling the word orders. Further, Dong *et al.* [21] and Li *et al.* [21] jointly employed multiple text embedding strategies including bag-of-words, word2vec, and GRU, to obtain robust query representation. In a follow-up work [7], Dong *et al.* proposed a multi-level text encoding to capture the global, local, and temporal patterns in the textual queries. Despite their effectiveness, these methods simply treating queries holistically as one dense vector representations, which may obfuscate the keywords or phrases that

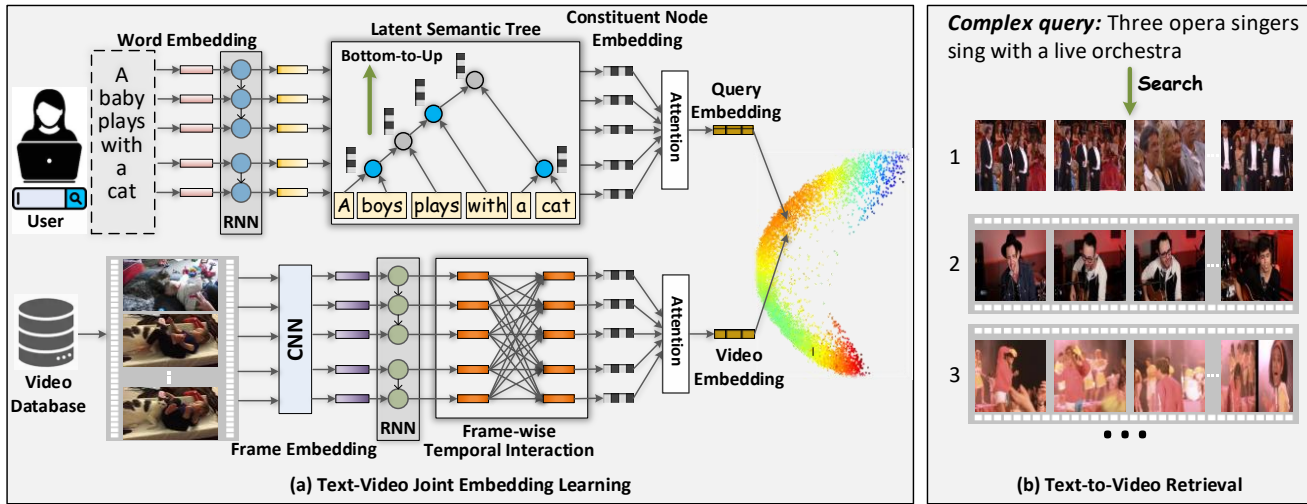


Figure 2: An illustration of our tree-augmented cross-modal encoding method for complex-query video retrieval.

have rich semantic cues and are less interpretable than the concept-based paradigm. In this work, we explicitly explore the syntactic structure of natural language query, thus will help to better understand the search intention. Lin *et al.* [22] and Xu *et al.* [45] have made attempts in this direction. Lin *et al.* first obtained the parse tree of the textual query, and modeled the word dependency based on a series of manually derived rules. In [45], Xu *et al.* constructed the dependency-tree structure based on subject-verb-object triplets extracted from a sentence and modeled the structure by a recursive neural network.

For video embedding, a typical approach is to first extract the frame-level features by pre-trained CNNs and subsequently aggregate them into a video representation. To obtain the video-level feature, mean pooling and max pooling are common choices [6, 28, 30, 45]. Yu *et al.* [51] used LSTM to model the temporal information, where frame-level features are sequentially fed into LSTM, and the hidden vector at the last step is used as the video feature. Dong *et al.* [7] also explicitly exploited the global and local patterns in videos to obtain a multi-level video representation. In this work, we design a temporal attentive video encoder that jointly models the temporal dependence between consecutive frames by RNNs and frame-wise temporal interaction by using the multi-head attention mechanism.

Natural language query-based retrieval techniques have also been successfully applied for domain-specific object retrieval in the field of video surveillance or E-commerce, such as text-based person search [19] and its application in person re-identification [32, 46, 47] and dialog-based fashion retrieval [9]. In [19], Li *et al.* collected a large-scale person description dataset with detailed natural language annotations and person samples from various sources and proposed a novel recurrent neural network with gated neural attention for person search. Niu *et al.* [32] designed a multi-granularity image-text alignments model for better modeling the similarity between text description and person images. In [9], Guo *et al.* introduced the reinforcement learning techniques to the task of dialog-based interactive image search that enables users to provide feedback via natural language.

### 3 THE PROPOSED APPROACH

This paper proposes to tackle the content-based complex-query video retrieval task, in which the query is a natural language sentence that describes a video. We basically follow the embedding-based paradigm that embeds the queries and videos into a joint embedding space where texts and videos can be easily matched and ranked. In this section, we first introduce an approach to recursively compose a latent semantic tree to model the complex query in Section 3.1. Then we introduce how to obtain the vector representations of the query and videos in Section 3.2 and 3.3, followed by the joint optimization of the query and video embeddings in the same space in Section 3.4.

#### 3.1 LST: Latent Semantic Tree

To better understand the complex query, this work proposes to use the Tree-structured LSTM (TreeLSTM) [39] to recursively compose a latent semantic tree (LST) in a bottom-to-up fashion to structurally describe each given query. Following [4], the LST structure is formulated as a binary recursive tree with two kinds of nodes: *leaf nodes* (i.e., words) and *parent nodes* (i.e., constituents). A parent node takes in two adjacent child nodes and describes more complex semantics than its child nodes. In this section, we first briefly describe how to apply TreeLSTM to compute the parent node representation from its two child nodes and then describe how to select the parent node at each layer for recursively building the LST.

**TreeLSTM.** Given the representations of two adjacent child nodes  $(\mathbf{h}_i, \mathbf{c}_i)$  and  $(\mathbf{h}_{i+1}, \mathbf{c}_{i+1})$  as inputs, the parent node representation  $(\mathbf{h}_p, \mathbf{c}_p)$  is computed by

$$\begin{bmatrix} \mathbf{i} \\ \mathbf{f}_l \\ \mathbf{f}_r \\ \mathbf{o} \\ \mathbf{g} \end{bmatrix} = \begin{bmatrix} \sigma \\ \sigma \\ \sigma \\ \sigma \\ \tanh \end{bmatrix} \left( \mathbf{W}^p \begin{bmatrix} \mathbf{h}_i \\ \mathbf{h}_{i+1} \end{bmatrix} + \mathbf{b}^p \right), \quad (1)$$

$$\mathbf{c}_p = \mathbf{f}_l \odot \mathbf{c}_i + \mathbf{f}_r \odot \mathbf{c}_{i+1} + \mathbf{i} \odot \mathbf{g}, \quad (2)$$

$$\mathbf{h}_p = \mathbf{o} \odot \tanh(\mathbf{c}_p), \quad (3)$$

where  $\mathbf{W}^p \in \mathbb{R}^{5d_t \times 2d_t}$  and  $\mathbf{b}^p \in \mathbb{R}^{5d_t}$  are trainable parameters,  $\sigma(\cdot)$  denotes the activation function *sigmoid*, and  $\odot$  denotes the element-wise product. Similar to the standard LSTM, each node is represented by a hidden state  $\mathbf{h} \in \mathbb{R}^{d_t}$  and a cell state  $\mathbf{c} \in \mathbb{R}^{d_t}$ .

**Layer-wise Node Transformation.** At the bottom layer, given a query with  $N$  words as inputs, we first represent it as a sequence of word embeddings and then transform the word embeddings to the representations of leaf nodes at the corresponding locations. Assume the  $t$ -th layer of the LST consists of  $N_t$  nodes  $\{\mathbf{r}_i^t = (\mathbf{h}_i^t, \mathbf{c}_i^t)\}_{i=1}^{N_t}$ . If two adjacent nodes  $\mathbf{r}_i^t$  and  $\mathbf{r}_{i+1}^t$  are selected to be merged, then we can utilize the above-mentioned TreeLSTM to compute the representation of the parent node  $\mathbf{r}_i^{t+1} = \text{TreeLSTM}(\mathbf{r}_i^t, \mathbf{r}_{i+1}^t)$  at the  $t+1$  layer. The representations of the unselected nodes are directly copied to the corresponding positions at the  $t+1$  layer.

**Memory-augmented Node Scoring and Selection.** The key step to the building of LST is how to accurately select the parent node at each layer. Previous work [4] proposed to enumerate all adjacent two nodes (e.g.,  $\mathbf{r}_i^t$  and  $\mathbf{r}_{i+1}^t$ ) to compose the parent node candidates  $\{\mathbf{r}_i^{t+1}\}_{i=1}^{N_{t+1}}$  and compute their representations by feeding the two consecutive child nodes into the TreeLSTM, and then select the best parent node candidate based on a node scoring module. Choi *et al.* [4] implemented the scoring module by first introducing a global *query vector* and then computing the inner-product between the query vector and the hidden states of parent node candidates, followed by a softmax operation. Despite its simplicity, it is still difficult to effectively decide the best candidate, due to the ambiguity of language and the limited capacity of hidden state [5, 48] to remember the input history, especially when the given query is very long. To address this issue, we design a memory-augmented scoring module  $f_{score}(\cdot; \Theta_{score})$  to select the parent node:

$$s_i^{t+1} = f_{score}(\mathbf{r}_i^{t+1}, \mathbf{u}_i^{t+1}; \Theta_{score}), \quad (4)$$

where  $s_i^{t+1}$  denotes the probability of the  $i$ -th parent node candidate being selected and  $\mathbf{u}_i^{t+1}$  is a node-specific context vector derived from a **global memory**  $\mathbf{M}$  which stores the semantic context. The global memory is defined as the set of leaf node hidden state representations  $\mathbf{M} = [\mathbf{h}_1^1, \mathbf{h}_2^1, \dots, \mathbf{h}_N^1] \in \mathbb{R}^{N \times d_t}$  at the bottom layer which preserves the original semantic context in the given sentence. To obtain the context vector  $\mathbf{u}_i^{t+1}$ , we use the hidden state  $\mathbf{h}_i^{t+1}$  of each parent node candidate to query the global memory and then attentively aggregate the global memory  $\mathbf{M}$ :

$$a_{ij}^{t+1} = \text{Softmax}\left(\frac{(\mathbf{h}_i^{t+1})^T \sigma(\mathbf{W}_m \mathbf{h}_j^1 + \mathbf{b}_m)}{\sqrt{d_t}}\right), \mathbf{u}_i^{t+1} = (\mathbf{a}_i^{t+1})^T \mathbf{M}, \quad (5)$$

where  $\mathbf{a}_i^{t+1} = [a_{i1}^{t+1}, a_{i2}^{t+1}, \dots, a_{iN}^{t+1}]$  is the normalized attention vector over the memory and  $\mathbf{W}_m \in \mathbb{R}^{d_t \times d_t}$  and  $\mathbf{b}_m \in \mathbb{R}^{d_t}$  are trainable parameters. We then implement our scoring module as:

$$s_i^{t+1} = \text{Softmax}\left(\mathbf{w}_s^T \sigma\left(\mathbf{W}_s \begin{bmatrix} \mathbf{h}_i^{t+1} \\ \mathbf{u}_i^{t+1} \end{bmatrix} + \mathbf{b}_s\right) / \sqrt{2d_t}\right), \quad (6)$$

where  $\sigma(\cdot)$  is the nonlinear activation function ReLU and  $\mathbf{w}_s \in \mathbb{R}^{2d_t}$ ,  $\mathbf{b}_s \in \mathbb{R}^{2d_t}$ , and  $\mathbf{W}_s \in \mathbb{R}^{2d_t \times 2d_t}$  are trainable parameters. The main intuition of this node scoring module is to inject the semantic context into each decision for a better parent node selection. In

such a recursive process, we select the candidate with the maximum validity score using Eq. (6) based on the Straight-Through (ST) Gumbel-Softmax estimator [8]. In the forward pass, the ST Gumbel-Softmax estimator discretizes the continuous signal, while in the backward pass, the continuous signals are used for stable training. Note that only the representation of the selected node is updated using the outputs of Eq. (1), (2), and (3). The other nodes that are not selected are not updated.

The above procedure is recursively repeated until only a single node is left. By this procedure, we can automatically compose a  $N$ -layers binary latent semantic tree with semantically-meaningful constituents to better understand the complex query without any syntactic annotations.

### 3.2 Tree-augmented Query Encoder

**One-hot Representation.** Given a query  $Q = \{q_1, q_2, \dots, q_N\}$ , we first represent it as a sequence of one-hot vectors  $\{\mathbf{q}'_1, \mathbf{q}'_2, \dots, \mathbf{q}'_N\}$ , where  $\mathbf{q}'_t$  indicates the vector of the  $t$ -th word. We further convert the word vectors to word dense representations  $\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_N\}$  based on pretrained word embedding matrix [29].

**Leaf Node LSTM.** We use RNNs as the basic sequence modeling block. For keeping consistency with the TreeLSTM in our LST module, we use LSTM to transform the word embeddings to the leaf node representations at the bottom layer. More formally, the LSTM unit, at the  $i$ -th time step, takes the features of the current word  $\mathbf{q}_i$ , previous hidden state  $\mathbf{h}_{i-1}^1$ , and cell state  $\mathbf{c}_{i-1}^1$  as inputs, and yields the current hidden state  $\mathbf{h}_i^1$  and cell state  $\mathbf{c}_i^1$ :

$$(\mathbf{h}_i^1, \mathbf{c}_i^1) = \text{LSTM}(\mathbf{q}_i, \mathbf{h}_{i-1}^1, \mathbf{c}_{i-1}^1). \quad (7)$$

Eq. (7) functions as the leaf node transformation module.

**Tree Construction.** The outputs of Eq. (7) are directly fed into the TreeLSTM module for the transformation of parent node candidates, as detailed in Eq. (1), (2), and (3). After  $N$  steps of the transformation, scoring, and selection, as described in the Section 3.1, we recursively compose a  $N$ -layers latent semantic tree, consisting of  $N-1$  constituent nodes (i.e., parent nodes), formulated by

$$\{\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_{N-1}\} = \text{LSTree}(\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_N\}), \quad (8)$$

where LSTree indicates the overall tree construction procedure and  $\mathbf{e}_i \in \mathbb{R}^{d_t}$  denotes the representation of the  $i$ -th constituent node. The tree can clearly describe the syntactic structure of complex queries, which is helpful to better understand the user query. A similar procedure of tree construction can be found in [4].

**Structure-aware Query Representation.** The next step is to derive the query representation based on the recursively extracted constituent nodes in the LST. In previous work [4, 36], only the last constituent node is used for task-specific inference. However, as mentioned previously, the complex query usually consists of multiple visual concepts and their reference descriptions, in which some concepts or reference descriptions may not have clear visual evidence or just have very short temporal durations in the videos. The last constituent node may not effectively cover the full linguistic context of the complex queries. In this work, we introduce an attention network to explore the importance of each constituent and then derive the structure-aware query representation by attending

to the informative constituent nodes:

$$\beta_i = \text{Softmax}\left(\mathbf{u}_{ta}^T \sigma(\mathbf{W}_{ta} \mathbf{e}_i + \mathbf{b}_{ta}) / \sqrt{d_{ta}}\right), \quad \bar{\mathbf{q}} = \sum_{i=1}^N \beta_i \mathbf{e}_i, \quad (9)$$

where  $\sigma(\cdot)$  is the non-linear activation function ReLU and  $\mathbf{W}_{ta} \in \mathbb{R}^{d_{ta} \times d_t}$ ,  $\mathbf{b}_{ta} \in \mathbb{R}^{d_{ta}}$ , and  $\mathbf{u}_{ta} \in \mathbb{R}^{d_{ta}}$  are trainable parameters,  $\beta_i$  denotes the normalized importance score of the node  $\mathbf{e}_i$ , and  $\bar{\mathbf{q}} \in \mathbb{R}^{d_t}$  denotes the query representation that aggregates the representations of all constituent nodes.

### 3.3 Temporal-Attentive Video Encoder

Given a video clip  $\mathcal{V}$ , we first sample uniformly a sequence of video frames  $\{v_1, v_2, \dots, v_M\}$  from  $\mathcal{V}$  with a pre-specified interval. We extract the frame features using pre-trained CNNs and represent the video clip as  $\mathbf{V} = \{\mathbf{v}_t\}_{t=1}^M$  where  $\mathbf{v}_t \in \mathbb{R}^{d_v}$  denotes the frame vector of the  $t$ -th frame. In this paper, we deal with two types of video characteristics: 1) temporal dependence between consecutive frames along the sequence, and 2) frame-wise temporal interaction over the whole video space.

**Temporal Dependence Modeling.** We leverage the GRU to model the temporal dependence between consecutive frames. At each time step, GRU takes the feature vector of the current frame and the hidden state of the previous frame as inputs and yields the hidden state of the current frame:

$$\mathbf{h}'_t = \text{GRU}(\mathbf{v}_t, \mathbf{h}'_{t-1}), \quad (10)$$

where  $\mathbf{h}'_t \in \mathbb{R}^{d_v}$  denotes the hidden state of the  $t$ -th frame. By the operation in Eq. (10), we can effectively capture the dependence between adjacent frames. For representing a video (clip), previous works either leverage the last hidden state or aggregate all the hidden states of frames using average-pooling, forgoing modeling the frame interaction over the whole video space.

**Frame-wise Temporal Interaction Modeling.** To further enhance of the video sequence representation, we propose to leverage the frame-wise correlation based on the multi-head self-attention mechanism [42]. Given a video sequence  $\mathbf{V} = \{\mathbf{h}'_t\}_{t=1}^M$  produced by the GRU operation in Eq. (10), the basic idea is to first project the video sequence representation  $\mathbf{V}$  into multiple embedding spaces and perform scaled dot-product attention between *query* frame and *key* frame, followed by a softmax operation to obtain the normalized weights on the *value* frames. We finally concatenate the outputs from multiple attention spaces as the final *value*:

$$\hat{\mathbf{V}}^i = \text{Softmax}\left(\frac{1}{\sqrt{d_i}} (\mathbf{W}_Q^i \mathbf{V})^T \mathbf{W}_K^i \mathbf{V}\right) \mathbf{W}_V^i \mathbf{V}, \quad (11)$$

$$\hat{\mathbf{V}} = \text{Norm}\left(\mathbf{V} + \mathbf{W}^P \left(\text{Concat}\left(\hat{\mathbf{V}}^1, \hat{\mathbf{V}}^2, \dots, \hat{\mathbf{V}}^Z\right)\right)\right), \quad (12)$$

where  $\mathbf{W}_Q^i \in \mathbb{R}^{d_i \times d_v}$ ,  $\mathbf{W}_K^i \in \mathbb{R}^{d_i \times d_v}$ , and  $\mathbf{W}_V^i \in \mathbb{R}^{d_i \times d_v}$  are three trainable parameters that transform the original input  $\mathbf{V}$  to the *query*, *key*, and *value* matrices in the  $i$ -th attention space with dimension  $d_i$ .  $\hat{\mathbf{V}}^i \in \mathbb{R}^{d_i \times M}$  denotes the attended *value* in the  $i$ -th attention space.  $\text{Concat}(\cdot)$  denotes the concatenation operation.  $\mathbf{W}^P \in \mathbb{R}^{d_v \times d_v}$  is a trainable parameter that projects the concatenated features into original space.  $\text{Norm}(\cdot)$  denotes the LayerNorm operation.  $\hat{\mathbf{V}} = \{\hat{\mathbf{v}}_t \in \mathbb{R}^{d_v}\}_{t=1}^M$  is the final video sequence representation. The above multi-head attention mechanism allows the

model to jointly attend to information from different representation spaces at different positions, which effectively captures the feature interaction among frames.

**Temporal-attentive Video Representation.** To make informative frames (e.g., foreground frames) contribute more to the final video representation, we design a temporal attention neural network with three trainable parameters  $\mathbf{u}_{va} \in \mathbb{R}^{d_{va}}$ ,  $\mathbf{b}_{va} \in \mathbb{R}^{d_{va}}$ , and  $\mathbf{W}_{va} \in \mathbb{R}^{d_{va} \times d_v}$ :

$$\eta_t = \text{Softmax}\left(\mathbf{u}_{va}^T \sigma(\mathbf{W}_{va} \hat{\mathbf{v}}_t + \mathbf{b}_{va}) / \sqrt{d_{va}}\right), \quad \bar{\mathbf{v}} = \sum_{t=1}^M \eta_t \hat{\mathbf{v}}_t, \quad (13)$$

where  $\eta_t$  denotes the normalized importance score of the  $t$ -th frame, and  $\bar{\mathbf{v}} \in \mathbb{R}^{d_v}$  denotes the final video representation. Eq. (13) has a similar formulation as Eq. (9), both of which are easy to implement and effective to exploit the informative frame/word features for representation.

### 3.4 Text-Video Joint Embedding

Formally, given a natural language query  $Q = \{q_1, q_2, \dots, q_N\}$  and a video sequence  $\mathcal{V} = \{v_1, v_2, \dots, v_M\}$ , we transform  $Q$  and  $\mathcal{V}$  to low-dimensional vector representations  $\bar{\mathbf{q}} \in \mathbb{R}^{d_t}$  and  $\bar{\mathbf{v}} \in \mathbb{R}^{d_v}$  using the query encoder described in Section 3.2 and the video encoder in Section 3.3, respectively. Then we map the text query and video into a joint embedding space by two linear projection matrices:  $f^t: \mathbb{R}^{d_t} \rightarrow \mathbb{R}^{d^*}$  and  $f^v: \mathbb{R}^{d_v} \rightarrow \mathbb{R}^{d^*}$ , where we define the cross modal matching score as the cosine similarity:

$$s(Q, \mathcal{V}) = \frac{f^t(\bar{\mathbf{q}})^T f^v(\bar{\mathbf{v}})}{\|f^t(\bar{\mathbf{q}})\|_2 \|f^v(\bar{\mathbf{v}})\|_2}, \quad (14)$$

where  $f^t(\bar{\mathbf{q}})$  and  $f^v(\bar{\mathbf{v}})$  are implemented by

$$f^t(\bar{\mathbf{q}}) = \mathbf{W}_t^* \bar{\mathbf{q}} + \mathbf{b}_t^*, \quad f^v(\bar{\mathbf{v}}) = \mathbf{W}_v^* \bar{\mathbf{v}} + \mathbf{b}_v^*, \quad (15)$$

where  $\mathbf{W}_t^* \in \mathbb{R}^{d^* \times d_t}$ ,  $\mathbf{b}_t^* \in \mathbb{R}^{d^*}$ ,  $\mathbf{W}_v^* \in \mathbb{R}^{d^* \times d_v}$ , and  $\mathbf{b}_v^* \in \mathbb{R}^{d^*}$  are trainable parameters. We expect Eq. (14) to yield a higher score when the video  $\mathcal{V}$  is matched with the complex query  $Q$  or a lower score if not match. We also apply a batch normalization [13] followed by a non-linear activation  $\text{Tanh}(\cdot)$  on  $f^t(\bar{\mathbf{q}})$  and  $f^v(\bar{\mathbf{v}})$ , respectively, for stable training. Note that both  $f^t(\cdot)$  and  $f^v(\cdot)$  are not indispensable if we enforce the output of the query encoder to have the same dimension as the output of the video encoder, i.e.,  $d_t = d_v$ . We introduce  $f^t(\cdot)$  and  $f^v(\cdot)$  in Eq. (14) just for more formal expression and also make the section 3.4 self-contained. Besides, with  $f^t(\cdot)$  and  $f^v(\cdot)$ , we can derive much lower dimensional embeddings for fast retrieval without modifying the parameters in the two encoders.

**Loss Function:** To train the model, we use the margin ranking loss to optimize the network with a batch-hard negative sampling strategy. More formally, during training, we sample a batch of query-video pair  $\mathcal{X} = \{(Q_i, \mathcal{V}_i)\}_{i=1}^B$ . We wish to enforce that, for any given  $(Q_i, \mathcal{V}_i)$ , the similarity score  $s(Q_i, \mathcal{V}_i)$  between a query  $Q_i$  and its ground truth video  $\mathcal{V}_i$  is larger than the score of any negative pairs  $s(Q_i, \mathcal{V}_j)$  by a large margin, when video  $\mathcal{V}_j$  does not match with query  $Q_i$ . The loss on the batch is defined as

$$L(\mathcal{X}) = \frac{1}{|\mathcal{N}^h|} \sum_{i=1}^B \sum_{j \in \mathcal{N}^h} \max(0, \delta + s(Q_i, \mathcal{V}_j) - s(Q_i, \mathcal{V}_i)), \quad (16)$$

where  $\delta$  is the margin ( $\delta \in (0, 1)$ ).  $|\mathcal{N}^h|$  denotes the number of hard negative videos in the set  $\mathcal{N}^h$ . We found that the hardest negative sample may result in unstable training especially in a large batch, but averaging the costs on all negative samples in a batch will result in slow training. Therefore, in this work, we use a trade-off strategy: we just take into consideration the top  $|\mathcal{N}^h|$  negative samples (e.g., 5) and average the costs for stable and efficient training.

## 4 EXPERIMENT

A key contribution of this work is to develop a new complex-query video retrieval approach with a tree-augmented cross-modal encoding method. We aim to answer the following research questions via extensive experiments: (1) **R1**: How does the proposed method perform compared with state-of-the-art methods? (2) **R2**: What are the impacts of different components on the overall performance of our approach? (3) **R3**: How does the proposed method perform on different types of complex queries (e.g., different lengths and different categories)? Can the latent semantic tree help to better understand the complex query and drive stronger query representation?

### 4.1 Experimental Settings

**4.1.1 Datasets.** We use two public datasets: MSR-VTT video caption dataset [44] and LSMDC movie description dataset [33].

**MSR-VTT [44]:** It is an increasingly popular dataset for text-to-video retrieval, consisting of 10K YouTube video clips. Each of them is annotated with 20 crowd-sourced English sentences, which results in a total of 200K unique video-caption pairs. We notice that there are three different dataset partitions for this dataset. The first one is the official partition from [44] with 6,513 clips for training, 497 clips for validation, and the remaining 2,990 clips for testing. The second one is from [27] with 6,656 clips for training and 1000 test clips for testing. The last one is from [49], 7,010 and 1K video clips are used for training and testing respectively. Note for the last two data partitions, only one sentence associated with each video clip is used as the testing query. For a comprehensive evaluation, we evaluate our proposed model on all data partitions.

**LSMDC [33]:** It is another popular dataset that contains 118,081 short video clips extracted from 202 movies. Each video clip has only one caption, either extracted from the movie script or from the transcribed audio description. It is originally used for evaluation in the Large Scale Movie Description Challenge (LSMDC). In this work, we only consider the text-to-video task in LSMDC: given a natural language query, the system retrieves the video of interest from the 1,000 test video set.

**4.1.2 Implementation Details.** On MSR-VTT, for the word features, we initialize the word embedding matrix using a 500-D word2vec model provided by [6] which optimized word2vec on English tags of 30 million Flickr images. The textual sequence is fed into a unidirectional LSTM with the hidden size of  $d_t=512$  for leaf node transformation. The hidden sizes of the TreeLSTM and query attention modules are set to  $d_t=512$  and  $d_{ta}=256$ , respectively. The final query representation has the dimension of  $d_t=512$ . For the video features, we use the frame-level visual features provided by [7], where the 2048-D features are extracted with ResNet-152 [10] pre-trained on ImageNet. The video frame sequences are fed in a

**Table 1: State-of-the-art performance comparison (%) on MSR-VTT with different dataset splits. Note that TCE uses bidirectional GRU and LSTM for better performance in this experiment based on 1024-D query and video embeddings.**

Method	R@1	R@5	R@10	MedR
<b>Data split from [44]</b>				
Dong <i>et al.</i> [6]	1.8	7.0	10.9	193
Mithun <i>et al.</i> [30]	5.8	17.6	25.2	61
DualEncoding [7]	7.7	22.0	31.8	32
TCE	<b>7.7</b>	<b>22.5</b>	<b>32.1</b>	<b>30</b>
<b>Data split from [27]</b>				
Random	0.3	0.7	1.1	502
CCA [43]	7.0	14.4	18.7	100
MEE [27]	12.9	36.4	51.8	10.0
MMEN (Caption) [43]	13.8	36.7	50.7	10.3
JPoSE [43]	14.3	38.1	53.0	<b>9</b>
TCE	<b>17.1</b>	<b>39.9</b>	<b>53.7</b>	<b>9</b>
<b>Data split from [49]</b>				
Random	0.1	0.5	1.0	500
C+LSTM+SA+FC7 [40]	4.2	12.9	19.9	55
VSE-LSTM [15]	3.8	12.7	17.1	66
SNUVL [50]	3.5	15.9	23.8	44
Kaufman <i>et al.</i> [14]	4.7	16.6	24.1	41
CT-SAN [51]	4.4	16.6	22.3	35
JSFusion [49]	10.2	31.2	43.2	13
Miech <i>et al.</i> [28]	12.1	35.0	48.0	12
TCE	<b>16.1</b>	<b>38.0</b>	<b>51.5</b>	<b>10</b>

unidirectional GRU with the hidden size of  $d_v=512$ . The output of GRU is further fed into an 8-head attention module. The dimension of each head subspace is 64. The temporal attention module with the hidden size of  $d_{va}=256$  aggregates the outputs of multi-head attention module and produces a video representation with the dimension of  $d_v=512$ . As mentioned previously, since video encoder and query encoder have the same dimension, we omit the two projection matrices in Eq. (15) for compressing the size of parameters. The number of hard negative samples used in Eq. (16) is 5.

On LSMDC, following [27], we use 300-D *GoogleNews* pre-trained word2vec word embeddings as the input of a unidirectional LSTM with the hidden size of 512, followed by our tree construction and the query attention network using the same setting with MSR-VTT. Since [27] did not release the frame-level features, we directly use the provided multi-modal video-level features (appearance, motion, audios and face) to evaluate the effectiveness of our complex-query modeling module. Note that we do not use the gated embedding module and weighted-fusion of similarity scores in [27]. We first transform the multiple-modal features to the default embedding spaces in [27] with multiple projection matrices and concatenate the multiple features into a long vector, followed by a feature transformation into the 512-D joint embedding space.

**4.1.3 Evaluation Metrics.** Following the setting of [7, 27, 49], we report the rank-based performance metrics, namely R@K ( $K = 1, 5, 10$ ) and Median rank (MedR). R@K is the percentage of test queries for which at least one relevant item is found among the top-K retrieved results. MedR is the median rank of the first relevant



**Table 2: State-of-the-art performance comparison (%) on LSMDC [33]. Our TCE performs the best with a much lower-dimensional embedding (512-D). The *Mot.* and *Aud.* refer to the motion feature and audio feature, respectively.**

Method	R@1	R@5	R@10	MedR
C+LSTM+SA+FC7 [40]	4.3	12.6	18.9	98
VSE-LSTM [15]	3.1	10.4	16.5	79
SNUVL [50]	3.6	14.7	23.9	50
Kaufman <i>et al.</i> [14]	4.7	15.9	23.4	64
CT-SAN [51]	5.1	16.3	25.2	46
Miech <i>et al.</i> [26]	7.3	19.2	27.1	52
CCA (FV HGLMM) [16]	7.5	21.7	31.0	33
JSFusion [49]	9.1	21.2	34.1	36
Miech <i>et al.</i> [28]	7.2	18.3	25.0	44
MEE [27]	10.2	25.0	33.1	<b>29</b>
TCE (Visual)	7.9	20.8	27.8	46
TCE (Visual+Mot.)	9.7	23.3	34.8	32
TCE (Visual+Mot.+Aud.)	<b>10.6</b>	<b>25.8</b>	<b>35.1</b>	<b>29</b>

item in the search results. Higher R@K and lower MedR indicate better performance.

**4.1.4 Training Details.** Our work is implemented using the PyTorch framework. We train our model using the ADAM optimizer and use an initialized learning rate of 0.0005 with a batch size of 128. Each epoch training is just performed using a single GPU and takes no more than 10 minutes.

## 4.2 Experimental Results and Analysis

**4.2.1 Comparison with State-of-the-Arts.** To answer the research question **R1**, we compare our proposed Tree-augmented Cross-modal Encoding (TCE) with recently proposed state-of-the-art methods: (1) RNN-based methods: DualEncoding [7], Kaufman *et al.* [14], CT-SAN [51], SNUVL [50], C+LSTM+SA+FC7 [40], and VSE-LSTM [15], (2) Multimodal Fusion methods: Mithun *et al.* [30], MEE [27], MMEN [43], and JPoSE [43], and (3) other methods: JSFusion [49], CCA (FV HGLMM) [16], and Miech *et al.* [26]. The experimental results on MSR-VTT and LSMDC are summarized, respectively, in Table 1 and Table 2. Note that there are different dataset splitting strategies of the MSR-VTT dataset. To fairly compare with the reported results of state-of-the-art methods, we first report our results based on the standard split from the official paper [44] and then evaluate our method on the other two splits from [27] and [49], respectively. Unless otherwise stated, we use unidirectional RNNs (512-D) in our experiments by default.

**MSR-VTT:** Table 1 clearly shows that our proposed TCE outperforms all other available methods in all three dataset splits. Specifically, on the first split [44], we surpass the results of DualEncoding [7] w.r.t. R@5, R@10, and MedR. DualEncoding is the best reported state-of-the-art method on the first split that fuses multi-levels textual and video features for joint embedding learning with the embedding size of 2048. While, our TCE just uses the temporal/sequential features (i.e., the 2nd level features in DualEncoding) with the final embedding size of 1024 for retrieval. Hence, TCE is able to report higher retrieval accuracy while using a much smaller embedding

size. TCE also outperforms the multimodal fusion (object, activity, and audio) method in Mithun *et al.* [30] by a large margin, which indicates the effectiveness of our proposed tree-augmented query modeling and temporal-attentive video sequence modeling methods. Note that our proposed TCE can achieve consistent performance improvement if we integrate some other modalities, like motion features or audio features into the video embedding. We evaluate our method in the multi-modality setting on LSMDC (See Table 2). For the second split [27], we observe a large improvement over the state-of-the-art JPoSE which disentangles the text query into multiple semantic spaces (Verb, Noun) for score-level fusion. Compared with JPoSE, TCE directly composes a latent semantic tree to describe the user complex query in an end-to-end manner and also includes an attention mechanism to capture the most informative constituent nodes in the tree. A similar improvement can also be observed in the third split [49], which further validates the effectiveness of TCE.

**LSMDC:** Table 2 compares the performance of TCE with nearly all reported results on the LSMDC video clip retrieval task. The results again show that our proposed TCE performs the best on this challenging benchmark dataset. Specifically, we outperform the MEE method by a relative improvement of 6% w.r.t. R@10. We use the same multi-modal video features with MEE, but with a simpler feature fusion strategy, i.e., concatenation. We can observe a more significant improvement over the JSFusion method, which is the winner of the LSMDC 2017 Text-to-Video and Video-to-Text retrieval challenge. Besides, in Table 2, we also investigate the effect of multi-modal fusion in our proposed TCE. Specifically, when we just use the 2048-D globally-pooled appearance features to describe the video, our model still outperforms most of the listed methods in Table 2. By augmenting the video representation with the motion feature, we can obtain a relative improvement of 25% w.r.t. R@10. The audio features can further stably improve the performance. That is to say, TCE has the potential of improving its performance on MSR-VTT by leveraging more informative features. Since the multi-modal fusion is not our focus in this paper, we leave the fusion experiment on MSR-VTT for future study.

**4.2.2 Ablation Studies.** To effectively answer the research question **R2**, we conduct extensive ablation studies on MSR-VTT based on the standard split. Specifically, we mainly organize the ablation studies into two groups: one for query encoder and the other for video encoder. The batch normalization is used to normalize the query/video representation in the following counterparts.

**On Query Encoder:** We use the following baselines and variants to transform the natural language queries to vector representations.

- **WordEmb+AvgP** and **WordEmb+MaxP:** Add a fully connected (FC) layer after the word embedding layer and aggregate its output with average-pooling (AvgP) operation or max-pooling operation (MaxP).
- **LSTM** and **LSTM+AvgP:** Instead of composing the latent semantic tree, we directly use the last hidden state (LSTM) or apply an average-pooling over the output of LSTM (LSTM+AP).
- **TCE (w/o-Cxt):** Remove the memory-augmented context vector  $u_t$  in the score module (Eq. (6)) and directly normalize the scaled dot-product between a global query vector and the hidden state of nodes. It is the standard implementation in [4].

- **TCE (w/o-LSTM)**: Remove the leaf node LSTM module. Instead, we use an FC layer to transform the word embedding to the default input of TreeLSTM.
- **TCE (w/o-TAtt)+AvgP**: Remove the text attention module in Eq. (9), instead, we use the average-pooling operation.

**On Video Encoder:** We use the following baselines and variants to transform videos to vector representations.

- **Frame+AvgP** and **Frame+MaxP**: Use a FC layer to first transform the frame features, followed by average-pooling (AvgP) or max-pooling (MaxP).
- **GRU** and **GRU+AvgP**: Directly use the last hidden state of GRU or apply average-pooling over the output of GRU (GRU+AvgP).
- **TCE (w/o-Mha)**: Remove the multihead attention module.
- **TCE (w/o-GRU)**: Replace the GRU module with a FC layer to transform the frame embedding.
- **TCE (w/o-VAtt)+AvgP**: Remove the temporal video attention module in Eq. (13) with an average-pooling operation instead.

Note that in the ablation studies, we only change one (e.g., query) part of our proposed TCE to the above baselines or variants, while keeping the rest of TCE (e.g., video) unchanged.


Table 3 shows the performance comparison of our proposed full TCE model with different ablations on the MSR-VTT dataset.

- Overall, we observe that our full model performs the best except in terms of MedR. Removing each component from TCE, such as *Cxt*, *Mha*, *LSTM/GRU*, and *TAtt/VAtt*, would result in relative performance degeneration, but not dramatically. It not only reflects the effectiveness of each component of our TCE, but also shows the robustness of our method. Each module can effectively complement each other, but is not very sensitive to each other.
- There are also some interesting findings: the RNNs do not play a much more important role than we wish in this task. Compared with the baselines WordEmb+AvgP and Frame+AvgP, the LSTM and GRU help to improve the accuracy by a small margin, due to the modeling of the dependence between words/frames. However, if we remove LSTM or GRU from our query encoder or video encoder, the model exhibits a minor performance degenerates. TCE (w/o-LSTM) and TCE (w/o-GRU) still report high accuracy. This indicates the effectiveness of our latent semantic tree in capturing the structure information of the complex queries. It also reveals the necessity of the temporal interaction module that models the frame-wise feature interaction beyond the dependence between consecutive frames.
- We also observe that the widely used average-pooling strategy does not performs well for the complex-query video retrieval task. Our introduced attention mechanisms in Eq. (9) and (13) performs well by attending to the informative constituent nodes and frames.

**4.2.3 Analysis on Different Types of Queries.(R3).** To investigate how our proposed TCE perform on different groups of complex queries, we group 59,800 test queries of the MSR-VTT dataset (Data split from [44]) according to their query lengths and categories. We compare TCE on different groups with the baseline model DualGRU which utilizes the bidirectional GRU with average pooling for both text encoding and video encoding. The performance comparison

**Table 3: Ablation studies on the MSR-VTT dataset using the standard dataset split [44] to investigate the effects of the tree-based query encoder and the temporal-attentive video encoder. The proposed method performs the best.**

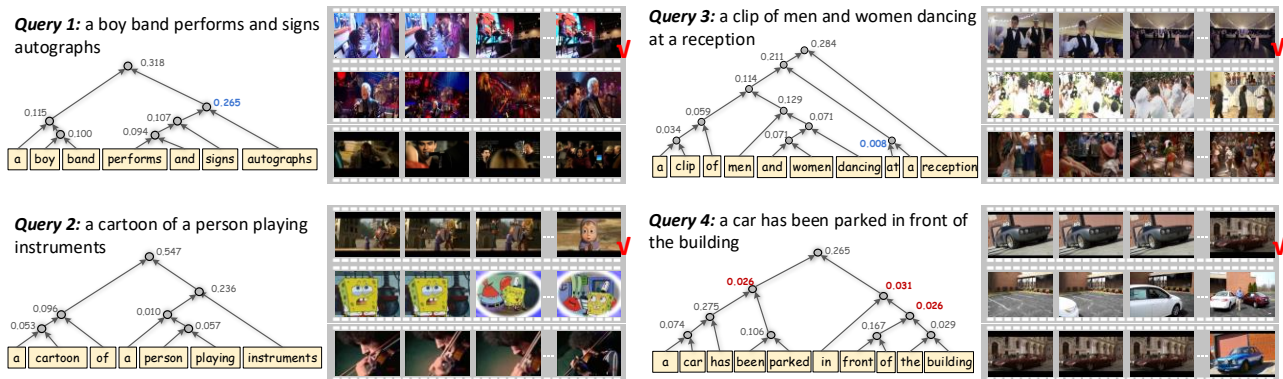
Method	R@1	R@5	R@10	MedR
<b>On Query Encoder</b>				
WordEmb+AvgP	6.79	20.98	30.68	32
WordEmb+MaxP	5.92	18.90	27.82	40
LSTM	6.91	21.31	31.17	31
LSTM+AvgP	6.95	21.28	30.68	35
TCE (w/o-Cxt)	6.98	21.46	31.49	30
TCE (w/o-LSTM)	7.09	21.86	31.67	31
TCE (w/o-TAtt)+AvgP	6.59	20.57	30.48	34
<b>On Video Encoder</b>				
Frame+AvgP	6.67	20.41	29.89	36
Frame+MaxP	6.20	20.24	29.87	35
GRU	6.75	21.03	30.91	31
GRU+AvgP	6.17	19.51	28.71	38
TCE (w/o-Mha)	6.97	21.59	31.19	31
TCE (w/o-GRU)	7.08	21.96	31.86	30
TCE (w/o-VAtt)+AvgP	6.73	21.38	31.74	<b>29</b>
<b>TCE</b>	<b>7.16</b>	<b>21.96</b>	<b>32.04</b>	30



**Figure 3: Performance comparison of DualGRU and our proposed TCE on MSR-VTT. Queries have been grouped in terms of (a) query lengths and (b) query categories.**

in each group by query lengths and categories are shown in Figure 3(a) and 3(b), respectively. In Figure 3(a), our proposed TCE consistently outperforms the DualGRU in all groups with different query lengths, showing its effectiveness in complex query modeling. Especially, we clearly observe that with increasing query lengths (from left to right in Figure 3(a)), the performance gain of TCE over DualGRU becomes much more significant. Generally, the longer queries are more complex than the shorter queries. As demonstrated in Figure 3(a), the query "a crowd appears then a hockey game is





**Figure 4: Four examples obtained by our TCE model on MSR-VTT. The composed latent semantic trees are presented under the corresponding queries. The normalized node weights are also shown. Videos with red marks are the correct ones.**

being played followed by a basketball team playing on a court” is more complex than the query “a music video”. Hence, the results validates that our proposed TCE is better in handling the complex queries, mainly benefiting from the latent semantic tree. In Figure 3(b), the performance in different query categories varies greatly, showing the varying difficulty of queries in different categories. For instance, the performance on the query group of *sports/actions* is higher than 0.25, while only about 0.13 on the query group of *news/politics*. The observation is reasonable since the *sports/actions* scenes are easy to be visually distinguished, while the *news/politics* scenes are much more diverse, and it is hard to learn the relation between news words and the visual scenes with limited training data. Despite the varied difficulty for each group, our proposed TCE model consistently beats the baseline on all groups.

**4.2.4 Qualitative Analysis.** Figure 4 shows four cases of qualitative results about the composed tree structures and the retrieved videos. For each query, the top three videos retrieved from the MSR-VTT are showed. Although only one correct video is annotated for each query, the retrieved three videos in Figure 4 are typically semantically relevant to the given query to some extent, showing the effectiveness of TCE. We observe that our approach is able to construct syntactically reasonable tree structures (e.g., *Query 1* and *Query 2*) and also identify the informative constituent nodes based on the attention mechanism, thus being helpful to better understand the complex query. For example, in *Query 1*, “performs and signs autographs” describes the *action* of the video clip, which is easy to be visually distinguished and usually reflects the main search intention, whose corresponding node in the tree was assigned a relatively large weight of 0.265. In *Query 2*, “a person playing instruments” refers to the key search intention, whose corresponding node in the tree was assigned a relatively large weight of 0.236. For *Query 3*, the composed tree is far from perfect, while the node “at a” contains less semantic information and was reasonably assigned the smallest attention weight of 0.008. For *Query 4*, the composed tree is syntactically reasonable, but some relative important nodes were assigned with small attention weights. Although some noises have been introduced in the latent semantic tree construction, our model still finds relevant videos for *Query 4*, which shows the robustness of our model.

## 5 CONCLUSION

In this work, we proposed a novel framework for complex-query video retrieval, which consists of a tree-based complex query encoder and a temporal attentive video encoder. Specifically, it first automatically composes a latent semantic tree from words to model the user query based on a memory-augmented node scoring and selection strategy and then encodes the tree into a structure-aware query representation based on an attention mechanism. Besides, it jointly models the temporal dependence between frames and frame-wise temporal interaction in the temporal attentive video encoder, followed by an attentive pooling mechanism to vectorize the video. This work provides a novel direction for complex-query video retrieval by automatically transforming the complex query into an easy-to-interpret structure without any syntactic rules and annotations. In the future, we will explore the proposed approach for other language-guided video tasks, such as video moment retrieval with natural language [23]. We also plan to integrate the multimedia indexing technique [11] with our approach for large-scale retrieval. We are also interested in exploring the external knowledge to enhance the text representation learning and the tree construction [2, 3] in the future study.

## 6 ACKNOWLEDGMENTS

This research is supported by The National Key Research and Development Program of China under grant 2018YFB0804205, the National Research Foundation, Singapore under its International Research Centres in Singapore Funding Initiative, the National Natural Science Foundation of China under grant 61902347, and the Zhejiang Provincial Natural Science Foundation under grant LQ19F020002. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

## REFERENCES

- [1] Da Cao, Zhiwang Yu, Hanling Zhang, Jiansheng Fang, Liqiang Nie, and Qi Tian. 2019. Video-Based Cross-Modal Recipe Retrieval. In *MM*. ACM, 1685–1693.
- [2] Yixin Cao, Lei Hou, Juanzi Li, Zhiyuan Liu, Chengjiang Li, Xu Chen, and Tiansi Dong. 2018. Joint Representation Learning of Cross-lingual Words and Entities via Attentive Distant Supervision. In *EMNLP*. 227–237.
- [3] Yixin Cao, Lifu Huang, Heng Ji, Xu Chen, and Juanzi Li. 2017. Bridge text and knowledge by learning multi-prototype entity mention embedding. In *ACL*. 1623–1633.

- [4] Jihun Choi, Kang Min Yoo, and Sang-goo Lee. 2018. Learning to compose task-specific tree structures. In *AAAI*.
- [5] Jasmine Collins, Jascha Sohl-Dickstein, and David Sussillo. 2016. Capacity and trainability in recurrent neural networks. In *ICLR*.
- [6] Jianfeng Dong, Xirong Li, and Cees GM Snoek. 2018. Predicting visual features from text for image and video caption retrieval. *IEEE Transactions on Multimedia* 20, 12 (2018), 3377–3388.
- [7] Jianfeng Dong, Xirong Li, Chaoxi Xu, Shouling Ji, Yuan He, Gang Yang, and Xun Wang. 2019. Dual Encoding for Zero-Example Video Retrieval. In *CVPR*. IEEE, 9346–9355.
- [8] Jiatao Gu, Daniel Jiwoong Im, and Victor OK Li. 2018. Neural machine translation with gumbel-greedy decoding. In *AAAI*.
- [9] Xiaoxiao Guo, Hui Wu, Yu Cheng, Steven Rennie, Gerald Tesauro, and Rogerio Feris. 2018. Dialog-based interactive image retrieval. In *NeurIPS*. 678–688.
- [10] K. He, X. Zhang, S. Ren, and J. Sun. 2016. Deep Residual Learning for Image Recognition. In *CVPR*. IEEE, 770–778.
- [11] Richang Hong, Lei Li, Junjie Cai, Dapeng Tao, Meng Wang, and Qi Tian. 2017. Coherent semantic-visual indexing for large-scale image retrieval in the cloud. *IEEE Transactions on Image Processing* 26, 9 (2017), 4128–4138.
- [12] Richang Hong, Yang Yang, Meng Wang, and Xian-Sheng Hua. 2015. Learning visual semantic relationships for efficient visual retrieval. *IEEE Transactions on Big Data* 1, 4 (2015), 152–161.
- [13] Sergey Ioffe and Christian Szegedy. 2015. Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*. 448–456.
- [14] Dotan Kaufman, Gil Levi, Tal Hassner, and Lior Wolf. 2017. Temporal tessellation: A unified approach for video analysis. In *ICCV*. IEEE, 94–104.
- [15] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. 2014. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539* (2014).
- [16] Benjamin Klein, Guy Lev, Gil Sadeh, and Lior Wolf. 2015. Associating neural word embeddings with deep image representations using fisher vectors. In *CVPR*. IEEE, 4437–4446.
- [17] A. Krizhevsky, I. Sutskever, and G. Hinton. 2012. ImageNet Classification using Deep Convolutional Neural Networks. In *NeurIPS*. 1097–1105.
- [18] D. Le, S. Phan, V.-T. Nguyen, B. Renoust, T. A. Nguyen, V.-N. Hoang, T. D. Ngo, M.-T. Tran, Y. Watanabe, M. Klinkigt, et al. 2016. NII-HITACHI-UIT at TRECVID 2016. In *TRECVID Workshop*.
- [19] Shuang Li, Tong Xiao, Hongsheng Li, Bolei Zhou, Dayu Yue, and Xiaogang Wang. 2017. Person search with natural language description. In *CVPR*. IEEE, 1970–1979.
- [20] Xirong Li. 2019. Deep Learning for Video Retrieval by Natural Language. In *Proceedings of the 1st International Workshop on Fairness, Accountability, and Transparency in MultiMedia*. 2–3.
- [21] Xirong Li, Chaoxi Xu, Gang Yang, Zhineng Chen, and Jianfeng Dong. 2019. W2VV++: fully deep learning for ad-hoc video search. In *MM*. ACM, 1786–1794.
- [22] Dahua Lin, Sanja Fidler, Chen Kong, and Raquel Urtasun. 2014. Visual semantic search: Retrieving videos via complex textual queries. In *CVPR*. IEEE, 2657–2664.
- [23] Meng Liu, Xiang Wang, Liqiang Nie, Qi Tian, Baoquan Chen, and Tat-Seng Chua. 2018. Cross-modal moment localization in videos. In *MM*. ACM, 843–851.
- [24] F. Markatopoulou, D. Galanopoulos, V. Mezaris, and I. Patras. 2017. Query and Keyframe Representations for Ad-hoc Video Search. In *ICMR*. ACM, 407–411.
- [25] F. Markatopoulou, A. Mourtzidou, D. Galanopoulos, T. Mironidis, V. Kaltsa, A. Ioannidou, S. Symeonidis, K. Avgerinakis, S. Andreadis, et al. 2016. ITI-CERTH Participation in TRECVID 2016. In *TRECVID Workshop*.
- [26] Antoine Miech, Jean-Baptiste Alayrac, Piotr Bojanowski, Ivan Laptev, and Josef Sivic. 2017. Learning from video and text via large-scale discriminative clustering. In *ICCV*. IEEE, 5257–5266.
- [27] Antoine Miech, Ivan Laptev, and Josef Sivic. 2018. Learning a text-video embedding from incomplete and heterogeneous data. *arXiv preprint arXiv:1804.02516* (2018).
- [28] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. 2019. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *ICCV*. IEEE, 2630–2640.
- [29] T. Mikolov, K. Chen, G. Corrado, and J. Dean. 2013. Efficient Estimation of Word Representations in Vector Space. In *ICLR*.
- [30] N. C. Mithun, J. Li, F. Metzger, and A. K. Roy-Chowdhury. 2018. Learning Joint Embedding with Multimodal Cues for Cross-Modal Video-Text Retrieval. In *ICMR*. ACM, 19–27.
- [31] P. A. Nguyen, Q. Li, Z.-Q. Cheng, Y.-J. Lu, H. Zhang, X. Wu, and C.-W. Ngo. 2017. VIREO @ TRECVID 2017: Video-to-Text, Ad-hoc Video Search and Video Hyperlinking. In *TRECVID Workshop*.
- [32] K. Niu, Y. Huang, W. Ouyang, and L. Wang. 2020. Improving Description-Based Person Re-Identification by Multi-Granularity Image-Text Alignments. *IEEE Transactions on Image Processing* 29 (2020), 5542–5556.
- [33] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. 2015. A dataset for movie description. In *CVPR*. IEEE, 3202–3212.
- [34] Haşim Sak, Andrew Senior, and Françoise Beaufays. 2014. Long short-term memory recurrent neural network architectures for large scale acoustic modeling. In *Fifteenth annual conference of the international speech communication association*.
- [35] Dian Shao, Yu Xiong, Yue Zhao, Qingqiu Huang, Yu Qiao, and Dahua Lin. 2018. Find and Focus: Retrieve and Localize Video Events with Natural Language Queries. In *ECCV*. 200–216.
- [36] Haoyue Shi, Jiayuan Mao, Kevin Gimpel, and Karen Livescu. 2019. Visually Grounded Neural Syntax Acquisition. In *ACL*.
- [37] Cees GM Snoek, Marcel Worring, et al. 2009. Concept-based video retrieval. *Foundations and Trends® in Information Retrieval* 2, 4 (2009), 215–322.
- [38] C. G. M. Snoek, X. Li, C. Xu, and D. C. Koelma. 2017. University of Amsterdam and Renmin University at TRECVID 2017: Searching Video, Detecting Events and Describing Video. In *Proceedings of TRECVID 2017*.
- [39] Kai Sheng Tai, Richard Socher, and Christopher D Manning. 2015. Improved Semantic Representations From Tree-Structured Long Short-Term Memory Networks. In *ACL*. 1556–1566.
- [40] Atousa Torabi, Niket Tandon, and Leonid Sigal. 2016. Learning language-visual embedding for movie understanding with natural-language. *arXiv preprint arXiv:1609.08124* (2016).
- [41] K. Ueki, K. Hirakawa, K. Kikuchi, T. Ogawa, and T. Kobayashi. 2017. Waseda\_Meisei at TRECVID 2017: Ad-hoc Video Search. In *TRECVID Workshop*.
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*. 5998–6008.
- [43] Michael Wray, Diane Larlus, Gabriela Csurka, and Dima Damen. 2019. Fine-Grained Action Retrieval Through Multiple Parts-of-Speech Embeddings. In *ICCV*. IEEE, 450–459.
- [44] J. Xu, T. Mei, T. Yao, and Y. Rui. 2016. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *CVPR*. IEEE, 5288–5296.
- [45] R. Xu, C. Xiong, W. Chen, and J. J. Corso. 2015. Jointly Modeling Deep Video and Compositional Text to Bridge Vision and Language in a Unified Framework. In *AAAI*.
- [46] Xun Yang, Meng Wang, Richang Hong, Qi Tian, and Yong Rui. 2017. Enhancing person re-identification in a self-trained subspace. *ACM Transactions on Multimedia Computing, Communications, and Applications* 13, 3 (2017), 1–23.
- [47] Xun Yang, Meng Wang, and Dacheng Tao. 2018. Person re-identification with metric learning using privileged information. *IEEE Transactions on Image Processing* 27, 2 (2018), 791–805.
- [48] Dani Yogatama, Yishu Miao, Gabor Melis, Wang Ling, Adhiguna Kuncoro, Chris Dyer, and Phil Blunsom. 2018. Memory architectures in recurrent neural network language models. In *ICLR*.
- [49] Youngjae Yu, Jongseok Kim, and Gunhee Kim. 2018. A joint sequence fusion model for video question answering and retrieval. In *ECCV*. 471–487.
- [50] Youngjae Yu, Hyungjin Ko, Jongwook Choi, and Gunhee Kim. 2016. Video Captioning and Retrieval Models with Semantic Attention. *ArXiv abs/1610.02947* (2016).
- [51] Youngjae Yu, Hyungjin Ko, Jongwook Choi, and Gunhee Kim. 2017. End-to-end concept word detection for video captioning, retrieval, and question answering. In *CVPR*. IEEE, 3165–3173.
- [52] Jin Yuan, Zheng-Jun Zha, Yao-Tao Zheng, Meng Wang, Xiangdong Zhou, and Tat-Seng Chua. 2011. Learning concept bundles for video search with complex queries. In *MM*. ACM, 453–462.
- [53] Jin Yuan, Zheng-Jun Zha, Yan-Tao Zheng, Meng Wang, Xiangdong Zhou, and Tat-Seng Chua. 2011. Utilizing related samples to enhance interactive concept-based video search. *IEEE Transactions on Multimedia* 13, 6 (2011), 1343–1355.