

Exploiting Visual Semantic Reasoning for Video-Text Retrieval

Zerun Feng¹, Zhimin Zeng^{1,2}, Caili Guo^{1,2*} and Zheng Li¹

¹Beijing Key Laboratory of Network System Architecture and Convergence,
School of Information and Communication Engineering,
Beijing University of Posts and Telecommunications, Beijing, China
²Beijing Laboratory of Advanced Information Networks, Beijing, China
{fengzerun,zengzm, guocaili, lizhengzachary}@bupt.edu.cn

Abstract

Video retrieval is a challenging research topic bridging the vision and language areas and has attracted broad attention in recent years. Previous works have been devoted to representing videos by directly encoding from frame-level features. In fact, videos consist of various and abundant semantic relations to which existing methods pay less attention. To address this issue, we propose a Visual Semantic Enhanced Reasoning Network (ViSERN) to exploit reasoning between frame regions. Specifically, we consider frame regions as vertices and construct a fully-connected semantic correlation graph. Then, we perform reasoning by novel random walk rule-based graph convolutional networks to generate region features involved with semantic relations. With the benefit of reasoning, semantic interactions between regions are considered, while the impact of redundancy is suppressed. Finally, the region features are aggregated to form frame-level features for further encoding to measure video-text similarity. Extensive experiments on two public benchmark datasets validate the effectiveness of our method by achieving state-of-the-art performance due to the powerful semantic reasoning.

1 Introduction

Vision and language are two fundamental bridges for humans to access the real world. In the research community, video retrieval is one of the active topics to make connections between these two counterparts, which aims to measure the semantic similarity between a video and its corresponding caption in an embedding space. Semantic-based video retrieval can reduce the effort of manual annotations to improve productiveness.

Lately, with the dramatic development of visual understanding tasks, various video retrieval methods have been proposed and achieved promising performance. Most of them [Mithun *et al.*, 2018; Dong *et al.*, 2018] utilize pre-trained Convolutional Neural Network (CNN) to extract mid-level features as the representations of frames for further ma-

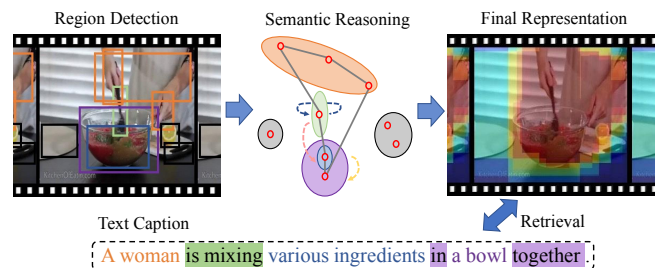


Figure 1: Illustration of the proposed Visual Semantic Enhanced Reasoning Network (ViSERN).

nipulating to calculate similarity scores between videos and texts. More recently, [Dong *et al.*, 2019] encode frame-level features in a multi-level way for video retrieval. Because of the existing huge video-text representation discrepancy, it is still an enormous challenge to predict the accurate video-text retrieval results. To lift this limitation in a principled way, we need to take a deeper insight into videos. The most basic unit of videos is a series of regions (such as persons, objects and landscapes). We human beings understand the videos by observing and reasoning semantic relation evolution between these regions during the timeline. Recall that if we describe video contents using a sentence, we will not only focus on objects, but also care about motions, interactions, relative positions and other high-level semantic concepts (e.g., “mixing”, “in” and “together”, the highlighted words shown in Figure 1), which lack apparent correspondence between visual representations and word descriptions. Moreover, videos have spatial-temporal redundancy information (e.g., “lemon” and “chopping board” denoted by black boxes shown in Figure 1), which is not always mentioned in captions. The intuitionistic phenomena indicate that visual semantic reasoning is fundamental for humans to understand the real world [Katsuki and Constantinidis, 2014]. However, the current existing video retrieval methods lack such kind of reasoning mechanism. Although these representations can obtain global visual information from the frames, they fail to capture high-level semantic interactions and treat discriminative regions equally. Thus, the semantic gap between videos and texts is not well addressed, and it leads to limited performance when matching videos and texts.

In this paper, we propose a Visual Semantic Enhanced Rea-

*Corresponding author

soning Network (ViSERN) to generate visual representations by exploiting semantic relations. After sampling frames, our model detects frame regions by bottom-up attention [Anderson *et al.*, 2018] and extract region features. In this way, each frame can be represented by several regions. Specifically, the bottom-up attention module is implemented with Faster R-CNN [Ren *et al.*, 2015] pre-trained on Visual Genome [Krishna *et al.*, 2017], an image region relation annotated dataset. Then, we build up a fully-connected graph with these regions as vertices and their semantic relevance as edges. We perform reasoning by novel random walk [Perozzi *et al.*, 2014] rule-based Graph Convolutional Networks (GCN) [Kipf and Welling, 2017] to produce region features with semantic relations.

We suppose that the regions with intensely semantic relations would be beneficial to infer video-text similarity. Therefore, we combine random walks with GCN and perform reasoning to enhance the effect of the most valuable ones. The random walks are a rule for vertices to access their neighbors. The transition probability is determined by the weights of edges. The neighbor vertices with higher weight are more likely to be visited. By this means, we can choose the discriminative regions and drop the insignificant ones by performing reasoning according to the random walk rule. Finally, the region features are passed by mean pooling layer to generate frame-level features for further manipulating to obtain retrieval results.

To illustrate intuitively, we also design a visualization map to analyze what has been learned after reasoning, in addition to the quantitative assessment of our method on standard benchmarks. Correlations between the frame-level feature and each reasoning region feature belonging to the same frames are calculated and visualized in an attention form. As shown in Figure 1, we find that pivotal semantic interactions can be well captured, while redundant frame regions cause low responses.

The main contributions of our work are summarized as follows:

- We propose a novel Visual Semantic Enhanced Reasoning Network (ViSERN) based on GCN to explore semantic correlations between regions within frames for measuring video-text similarity. As far as we know, ViSERN is the first video-text retrieval model which focuses on visual semantic reasoning.
- To further enhance the reasoning capacity, we observe that it is powerful to introduce the random walk rule into GCN. The regions with high semantic relevance are more possibly to access each other for better reasoning by preserving the structure graph topology.
- Experimental results conducted on two publicly available datasets demonstrate the effectiveness of the proposed model by achieving a new state-of-the-art due to the powerful semantic reasoning.

2 Related Work

Video Retrieval. Over the past several years, with big advances of deep learning in computer vision and natural

language processing, we observe that embedding based methods have been proposed for video retrieval. [Dong *et al.*, 2018] leverage pre-trained CNN models to extract features from frames and aggregate the frame-level features into a video-level feature by mean pooling. [Dong *et al.*, 2019] use mean pooling, bidirectional GRU, and 1-D CNN to encode global, temporal and local pattern information respectively for both video and text ends. However, to the best of our knowledge, no study has attempted to incorporate visual semantic reasoning when learning embedding for video retrieval.

Visual Semantic Reasoning. As an emerging research area, graph-based methods have been popular in recent years and shown to be an efficient approach to semantic relation reasoning. Graph Convolution Network (GCN) [Kipf and Welling, 2017] is first proposed for semi-supervised classification. Graph convolution extends the applicable scope of standard convolution from regular grids to more general pre-defined graphs by manipulating in the spectral domain, which is also proved to be powerful in video understanding. [Wang and Gupta, 2018] represent videos as similarity graphs and bi-directional spatial-temporal graphs to perform reasoning via GCN for video classification. [Liu *et al.*, 2019] design a graph model to capture social relations from videos and perform temporal reasoning with multi-scale receptive fields on the graph by GCN. Inspired by the above studies, we propose to represent the frames as fully-connected semantic correlation graphs, on which reasoning is performed by novel random walk rule-based graph convolutional networks for video retrieval.

Random Walks. Random walks on a graph is a rule to visit a sequence of vertices together with a sequence of edges and widely applied in network representation learning. [Perozzi *et al.*, 2014] treat nodes as words and sentences by truncated random walks to bridge the gap between network embeddings and word embeddings. [Grover and Leskovec, 2016] define a flexible notion of a node’s network neighborhood and design a biased random walk procedure for learning representations in complex networks. Since networks and graphs have similar topological structures, we argue that it is reasonable to introduce random walk statistics into GCN to enhance its visual reasoning ability.

3 Proposed Method

In order to capture visual semantic correlations for video feature representations, we propose a Visual Semantic Enhanced Reasoning Network (ViSERN). The architecture of ViSERN is shown in Figure 2. Concretely, our goal is to measure the similarity between video-text pairs by common space learning. We utilize the bottom-up attention model to generate frame regions and extract features from frames (Sec. 3.1). For the regions, we construct a graph to model semantic correlations (Sec. 3.2). Subsequently, we do semantic reasoning between these regions by leveraging random walk rule-based Graph Convolutional Networks (GCN) to generate re-

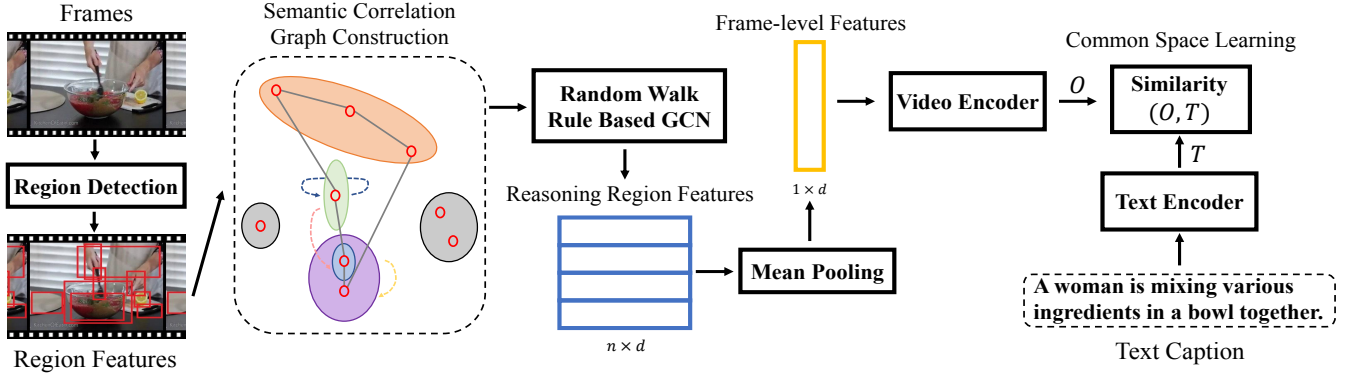


Figure 2: Architecture of the proposed Visual Semantic Enhanced Reasoning Network.

gion features with relation information (Sec. 3.3). Finally, video and text features are generated, and the whole model is trained with common space learning (Sec. 3.4).

3.1 Region Features Extraction

After sampling frames from videos, we use the bottom-up attention model [Anderson *et al.*, 2018] to extract a set of salient frame regions (e.g., person, cars and sky). We implement the bottom-up attention model with the same configuration in [Lee *et al.*, 2018]. The model is pre-trained on Visual Genomes [Krishna *et al.*, 2017], a large-scale image dataset annotating the interactions/relationships between objects. Objects and their relationships within each image in the dataset are densely annotated. Specifically, with the benefit of the bottom-up attention model, each frame can be represented by a set of features with semantic implications. The top n regions with the highest detection confidence scores are selected. We extract each region feature with d dimensions after *pool5* layer of ResNet-101 [He *et al.*, 2016] as backbone in the model.

In this way, we represent each frame by using its region features $V = \{v_1, \dots, v_n\}, v_i \in \mathbb{R}^d$.

3.2 Semantic Correlation Graph Construction

In our work, a relation graph is designed to conduct semantic reasoning between regions. Particularly, we construct a fully-connected semantic correlation graph $G = (V, E)$ on a frame with n region features.

In the graph, $V = \{v_1, \dots, v_n\}, v_i \in \mathbb{R}^d$ is the collection of region features as depicted above. As for edge set E , each two vertices exist a link, we suppose $\forall (v_i, v_j) \in E$. To construct an adjacency matrix R representing the edge set E , dot-product attention mechanism [Vaswani *et al.*, 2017] is implemented in an embedding space. We formulate R in a form of symmetric adjacency matrix following the definition of undirected graph to measure semantic relations between regions:

$$R(v_i, v_j) = \phi(v_i)^T \theta(v_j) + \theta(v_i)^T \phi(v_j), \quad (1)$$

where $\phi(\cdot) = W_\phi(\cdot) + b_\phi$ and $\theta(\cdot) = W_\theta(\cdot) + b_\theta$ are two linear embeddings with bias. In this way, the two regions can be connected with a higher weight edge if they have more intensely semantic correlations.

3.3 Reasoning by Random Walk Rule-Based GCN

At the heart of our approach, we propose a module based on Graph Convolutional Networks (GCN) [Kipf and Welling, 2017] to perform semantic reasoning between regions. To process region features in a spectral domain, we consider spectral convolutions on the graph G . For concise descriptions, we start with the set of vertices $x \in \mathbb{R}^n$, i.e., a scalar for all n vertices, then generalize to multi-input channels.

We recall that a fundamental operator in spectral graph analysis is the graph Laplacian. The combinatorial definition of graph Laplacian matrix is $L = D - R \in \mathbb{R}^{n \times n}$ where $D \in \mathbb{R}^{n \times n}$ is diagonal degree matrix with $D_{ii} = \sum_j R_{ij}$. Moreover, the symmetric normalized graph Laplacian matrix is $L^{sym} = I_n - D^{-1/2} R D^{-1/2}$, where $I_n \in \mathbb{R}^{n \times n}$ is the identity matrix. The convolution operator on vertices $\star x$ are filtered by $g_\theta = \text{diag}(\theta)$ [Bruna *et al.*, 2014], parameterized by learnable $\theta \in \mathbb{R}^n$ in the spectral domain:

$$g_\theta \star x = U g_\theta U^T x, \quad (2)$$

where $L^{sym} = U \Lambda U^T$ is diagonalized by the spectral basis $U \in \mathbb{R}^{n \times n}$, with a diagonal matrix of its eigenvalues $\Lambda \in \mathbb{R}^{n \times n}$. We can understand this convolution operation by three steps: First, x is transformed from the spatial domain into the spectral domain by $U^T x$. Next, the computation in the spectral domain is evaluated through $g_\theta U^T x$. Finally, $U g_\theta U^T x$ is utilized to convert x back into the spatial domain.

It was introduced in [Defferrard *et al.*, 2016] that g_θ is a function of Λ and can be approximated as a K -order polynomial filter:

$$g_{\theta'}(\Lambda) \approx \sum_{k=0}^K \theta'_k \Lambda^k, \quad (3)$$

where $\theta'_k \in \mathbb{R}^K$ is a vector of polynomial coefficients. However, computing Eq. 3 might be prohibitively expensive for enormous number of graphs. To solve this problem, the filter is reformulated in terms of Chebyshev polynomials up to K^{th} order [Hammond *et al.*, 2011]. Specifically, the Chebyshev polynomials can be calculated by a recurrence function:

$T_k(a) = 2aT_{k-1}(a) - T_{k-2}(a)$ with $T_0 = 1$ and $T_1 = a$. We now have the truncated expansion of the approximated filter:

$$g_{\theta''}(\Lambda) \approx \sum_{k=0}^K \theta''_k T_k(\tilde{\Lambda}), \quad (4)$$

where $\theta''_k \in \mathbb{R}^K$ is a vector of Chebyshev coefficients following a scaled diagonal matrix $\tilde{\Lambda} = 2\Lambda/\lambda_{max} - I_n$. λ_{max} denotes the spectral radius of L^{sym} . It's easy to check that $(U\Lambda U^T)^k = U\Lambda^k U^T$. Specifically, the convolution in spectral domain of vertices x with a filter $g_{\theta''}$ can be an equation of \tilde{L} :

$$g_{\theta''} \star x \approx \sum_{k=0}^K \theta''_k T_k(\tilde{L})x, \quad (5)$$

with $\tilde{L} = 2L^{sym}/\lambda_{max} - I_n$, a similar scaled strategy of $\tilde{\Lambda}$. To circumvent the problem of overfitting and build a deeper network model, [Kipf and Welling, 2017] approximate $\lambda_{max} \approx 2$ and limit $K = 1$. By this means, Eq. 5 further simplifies to:

$$g_{\theta''} \star x \approx \theta''_0 x + \theta''_1 (L^{sym} - I_n)x, \quad (6)$$

where θ''_0 and θ''_1 serve as two free parameters. Inspired by recent work [Perozzi *et al.*, 2014] that random walks can be effective for learning explicit representations of vertices by preserving the structure of graph topology. The transition probability of a random walker on the graph is defined as $P = D^{-1}R \in \mathbb{R}^{n \times n}$. We normalize P to obtain random walk normalization $L^{rw} \in \mathbb{R}^{n \times n}$ of graph Laplacian matrix for preventing gradient explosion:

$$L^{rw} = I_n - P = I_n - D^{-1}R = D^{-1}L, \quad (7)$$

so that $0 \leq \Lambda_{ii}^{rw} \leq 2$, where $\Lambda^{rw} \in \mathbb{R}^{n \times n}$ is a diagonal matrix which comprises eigenvalues of L^{rw} . Notice that:

$$L^{rw} = D^{-1/2}L^{sym}D^{1/2}, \quad (8)$$

where $D^{1/2}$ is an invertible square matrix. It's obvious that L^{rw} is similar to L^{sym} . For this reason, L^{rw} has real eigenvalues Λ_{ii}^{rw} , even if L^{rw} is in general not symmetric. Particularly, eigenvalues Λ_{ii}^{rw} are equivalent to eigenvalues Λ_{ii} of L^{sym} and thus the approximation of Eq. 3 is suitable for random walk normalization. Substituting L^{rw} for L^{sym} in Eq. 6, we can obtain the expression of spectral graph convolutions involved with the random walk rule:

$$\begin{aligned} g_{\theta''} \star x &\approx \theta''_0 x + \theta''_1 (L^{rw} - I_n)x \\ &= \theta''_0 x - \theta''_1 D^{-1}Rx. \end{aligned} \quad (9)$$

Instead of minimizing the number of parameters by simply hypothesizing $\theta''_0 = -\theta''_1$ in [Kipf and Welling, 2017], we consider Eq. 9 as a potential residual architecture comparable with ResNet [He *et al.*, 2016]. Concretely, x is added as a shortcut connection directly without multiplying θ''_0 . Neglecting the negative sign in $-\theta''_1$, we rewrite Eq. 9 as:

$$g_{\theta''} \star x \approx \theta''_1 D^{-1}Rx + x. \quad (10)$$

Finally, we extend the spectral graph convolutions to multi-input channels with $V \in \mathbb{R}^{n \times d}$ as follows:

$$Z = D^{-1}RV\Theta + V, \quad (11)$$

where $\Theta \in \mathbb{R}^{d \times d}$ is the filter parameter matrix. Θ can be reorganized as few stacked layers and the final output of spectral graph convolution is defined as:

$$Z = (D^{-1}RVW_g)W_r + V, \quad (12)$$

where $W_g \in \mathbb{R}^{d \times d}$ is the GCN weight layer and $W_r \in \mathbb{R}^{d \times d}$ is the residual weight layer. The output $Z = \{z_1, \dots, z_n\}$, $z_i \in \mathbb{R}^d$ is the relation enhanced representations for vertices. Z is further processed by mean pooling to form frame-level feature $\tilde{Z} \in \mathbb{R}^d$.

3.4 Common Space Learning

For both video and text ultimate embedding representations, we use the multi-level encoder [Dong *et al.*, 2019] to map them in the same D -dimensional common space. We denote the affined video and text features as O and T , respectively.

We adopt a hinge-based triplet ranking loss [Faghri *et al.*, 2018] which is proved to be powerful for common space learning with concentration on hard negatives. The loss is expressed as:

$$\begin{aligned} L(O, T) &= [\alpha - S(O, T) + S(\hat{O}, T)]_+ \\ &\quad + [\alpha - S(O, T) + S(O, \hat{T})]_+, \end{aligned} \quad (13)$$

where α is a margin constant, while $S(\cdot)$ computes cosine similarity between two features and $[x]_+ \equiv \max(x, 0)$. Given a positive pair (O, T) within a mini-batch, $\hat{O} = \arg \max_{O \neq i} S(i, T)$ and $\hat{T} = \arg \max_{T \neq j} S(O, j)$ are the hardest negative samples.

4 Experiment

To demonstrate the effectiveness of the proposed ViSERN model, we conduct extensive experiments on the video-text retrieval task (i.e., video query and text query as input respectively to search the counterparts). Two publicly available datasets are performed in our experiments. We analyze the results comparing with recent state-of-the-art methods. Moreover, we estimate the performance of each component in our model by ablation studies.

4.1 Datasets and Evaluation Metric

We evaluate our model on two benchmark datasets: the MSR-VTT dataset [Xu *et al.*, 2016] and the MSVD dataset [Chen and Dolan, 2011]. The MSR-VTT dataset is composed of 10k video clips with 20 sentence captions per clip, one of the largest datasets in terms of video-text pairs and word vocabulary. The MSVD dataset consists of 1970 video clips with around 40 multilingual sentence captions per clip. We only consider the English captions and stochastically choose five captions from it. Both two datasets are collected from Youtube in the range of multiple categories. We follow the identical partition strategy in [Dong *et al.*, 2019] for training, testing and validation in our experiments.

Methods	Text-to-Video Retrieval					Video-to-Text Retrieval					Sum of Recalls
	R@1	R@5	R@10	Med R	Mean R	R@1	R@5	R@10	Med R	Mean R	
VSE [Kiros <i>et al.</i> , 2014]	5.0	16.4	24.6	47	215.1	7.7	20.3	31.2	28	185.8	105.2
VSE++ [Faghri <i>et al.</i> , 2018]	5.7	17.1	24.8	65	300.8	10.2	25.4	35.1	25	228.1	118.3
JEMC [Mithun <i>et al.</i> , 2018]	5.8	17.6	25.2	61	296.6	10.5	26.7	35.9	25	266.6	121.7
W2VV [Dong <i>et al.</i> , 2018]	6.1	18.7	27.5	45	-	11.8	28.9	39.1	21	-	132.1
Dual encoding [Dong <i>et al.</i> , 2019]	7.7	22.0	31.8	32	-	13.0	30.8	43.3	15	-	148.6
ViSERN	7.9	23.0	32.6	30	178.7	13.1	30.1	43.5	15	119.1	151.1

Table 1: Video-to-Text and Text-to-Video retrieval results on the MSR-VTT dataset.

Methods	Text-to-Video Retrieval					Video-to-Text Retrieval					Sum of Recalls
	R@1	R@5	R@10	Med R	Mean R	R@1	R@5	R@10	Med R	Mean R	
VSE [Kiros <i>et al.</i> , 2014]	12.3	30.1	42.3	14	57.7	15.8	30.2	41.4	12	84.8	171.8
VSE++ [Faghri <i>et al.</i> , 2018]	15.4	39.6	53.0	9	43.8	21.2	43.4	52.2	9	79.2	224.8
JEMC [Mithun <i>et al.</i> , 2018]	16.1	41.1	53.5	9	42.7	23.4	45.4	53.0	8	75.9	232.5
Dual encoding [Dong <i>et al.</i> , 2019]	17.6	47.1	59.5	7	34.8	20.6	42.8	58.8	8	38.9	244.0
ViSERN	18.1	48.4	61.3	6	28.6	24.3	46.2	59.5	7	34.6	257.8

Table 2: Video-to-Text and Text-to-Video retrieval results on the MSVD dataset.

The evaluation criteria in our experiments are the rank-based performance generally listed as Recall at K (R@K, K = 1, 5, 10), Median Rank (Med R) and Mean Rank (Mean R). Given the retrieval results, R@K is the percentage of queries whether at least one corresponding item shows in the top-K. Med R calculates the median rank of the first correct item among the search results. Similarly, Mean R computes the mean rank of whole correct results. The sum of R@1, R@5 and R@10, noted as Sum of Recalls is also reported.

4.2 Implementation Details

We uniformly sample 16 frames from videos with the same time interval between every two frames. The number n of regions within a frame is 36, identical to [Anderson *et al.*, 2018]. The dimension d of region features extracted from ResNet-101 is 2048. The detailed configuration of the bottom-up attention model is the same as [Anderson *et al.*, 2018]. For text feature initial representations, each word is firstly represented by a one-hot vector and then multiplied by a word embedding matrix. We set the word embedding size to 500 and the dimension of the common space D to 2048, similar to [Dong *et al.*, 2019]. The margin parameter α is empirically chosen to be 0.2. The size of a mini-batch is 64. The optimizer in the training procedure is Adam with 50 epochs at most. We start training with an initial learning rate 0.0001, and the adjustment schedule is that once the validation loss does not decrease in three consecutive epochs, the learning rate is divided by 2.

4.3 Comparisons with the State-of-the-art

Quantitative results on the MSR-VTT and the MSVD datasets are shown in Table 1 and Table 2, respectively. We can observe that the proposed ViSERN outperforms recent state-of-the-art methods in terms of most indicators on both datasets,

especially for text-to-video retrieval subtask, which has more actual applications in daily life. For example, our ViSERN improves 2.6%, 4.5% and 2.5% at R@1 and R@5 and R@10 respectively on the MSR-VTT dataset comparing to the current state-of-the-art. In addition, we notice that ViSERN obtains much gain on the MSVD dataset relatively. The reasons are probably considered as two folds: First, the size of the MSVD dataset is smaller than the MSR-VTT dataset. Indeed, it potentially results from the existence of fewer distractors for a given query in the MSVD dataset. Second, the MSVD dataset only provides the visual signal, while the MSR-VTT dataset offers both visual and auditory information. It is non-trivial to exploit comprehensive elements in the MSR-VTT dataset for better representations. The results demonstrate that our ViSERN model can usefully implement semantic visual reasoning and measure the video-text similarity more accurately.

4.4 Ablation Studies

To validate the contribution of each component in our model, we carry out several ablation experiments. We start with a fundamental option which is simply implemented without the random walk rule-based GCN as a baseline model (denoted as ViSERN_{no-rea}). The region features V are straightforwardly processed by mean pooling to represent the whole frame without performing reasoning. To exam the benefit of reasoning based on graph topology, we use row-wise normalization (denoted as ViSERN_{row-nor}) in Eq. 12, a similar method avoiding gradient explosion implemented in [Li *et al.*, 2019]. Furthermore, we confirm the effectiveness of random walk statistics by comparing with symmetric normalization (denoted as ViSERN_{sym-nor}).

The detailed ablation results on the MSR-VTT dataset are shown in Table 3. From the results of ViSERN_{no-rea}, we ob-

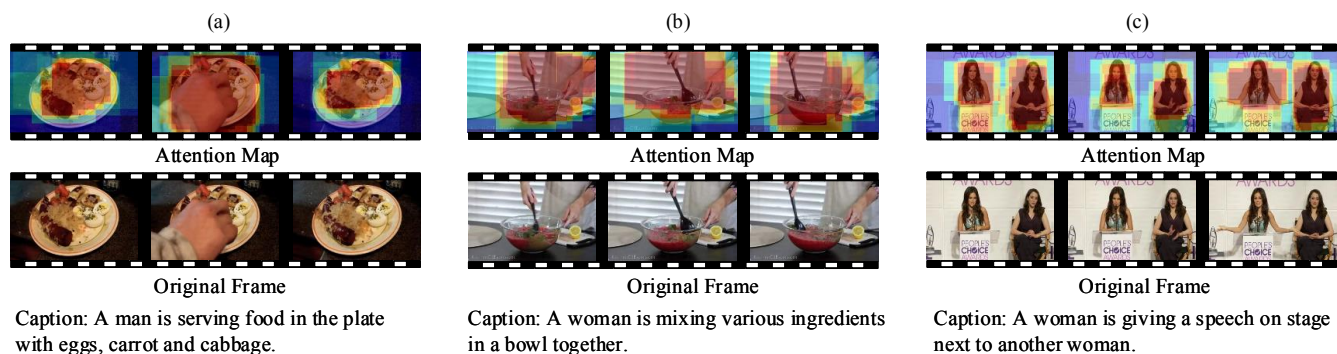


Figure 3: Visualization of attention maps and original frames above corresponding captions.

Methods	Text-to-Video Retrieval					Video-to-Text Retrieval					Sum of Recalls
	R@1	R@5	R@10	Med R	Mean R	R@1	R@5	R@10	Med R	Mean R	
ViSERN _{no-rea}	7.5	22.6	32.4	30	187.2	12.5	30.8	41.8	16	121.0	147.6
ViSERN _{row-nor}	7.1	22.1	31.8	32	193.4	12.1	30.6	42.2	16	128.6	145.9
ViSERN _{sym-nor}	7.6	22.8	32.3	30	189.5	12.1	30.7	42.6	16	119.6	148.0
ViSERN	7.9	23.0	32.6	30	178.7	13.1	30.1	43.5	15	119.1	151.1

Table 3: Ablation studies on the MSR-VTT dataset.

serve that it is reasonable to represent the whole frame by 36 region features. The results of ViSERN_{row-nor} indicate that reasoning is substantially dependent on graph topology-based normalization. We also notice that random walk statistics help to improve the performance by comparing the results of ViSERN_{sym-nor} with ViSERN. One possible reason could be that the random walk rule is effective to maintain the structure of graph topology generated from each individual frame and maximize the contribution of shared GCN layer weights and residual layer weights.

4.5 Visualization and Analysis

The motivation of our visual semantic reasoning is to highlight the regions with strongly semantic information by exploiting relations within a frame. In order to validate our thoughts, we visualize the correlations between the feature of the whole frame and the region features included in this frame by applying an attention map. Specifically, we compute the cosine similarity scores (same as $S(\cdot)$ in Eq. 13) between each region feature z_i and the feature of entire frame \bar{Z} . Then we rank the regions in the descending order according to similarity scores. To generate the attention map, we assign an attention score $s_i \in \mathbb{R}$ to each region based on its rank $r_i \in \{0, 1, \dots, n-1\}$. The function of attention score is defined as $s_i = (n - r_i)^2$. We normalize the set of attention scores to generate the final attention map conveniently.

In Figure 3, we show the attention maps, original frames and corresponding captions. The regions with strongly semantic meaning are painted red in attention maps. From the attention visualization, we can observe that ViSERN well represents frame-level features by exploiting semantic reasoning and capturing regions with intensely semantic relations. Ad-

ditionally, Figure 3(a) indicates that ViSERN is sensitive to short-term interaction. We also find that long-term motion is also captured from the illustration of Figure 3(b). Furthermore, the regions of both two salient people are highlighted separately in Figure 3(c).

5 Conclusion

In this paper, we propose a Visual Semantic Enhanced Reasoning Network (ViSERN) to learn semantic reasoning by using the novel random walk rule-based graph convolutional networks for video-text retrieval. The enhanced frame representations capture key regions and suppress redundancy in a scene so that it is helpful to align with the corresponding video caption. Extensive experiments on the MSR-VTT and MSVD datasets demonstrate the proposed model consistently outperforms the state-of-the-art methods. In addition, we provide visualization analyses to show how ViSERN can provide more discrimination and interpretability to such vision-language models. Moreover, ViSERN can be further developed to support content-based searches on online video websites.

In the future, we will consider the location information of regions representing in the graph. We would also like to extend the undirected graph to a directed graph for modeling video temporal sequences reasonably.

Acknowledgments

This work was partially supported by the China Telecom Dict Application Capability Center, the Beijing Natural Science Foundation (4202049) and the National Key R&D Program of China (2018YFB1800805).

References

- [Anderson *et al.*, 2018] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, and Stephen Gould Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6077–6086, 2018.
- [Bruna *et al.*, 2014] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann Lecun. Spectral networks and locally connected networks on graphs. In *International Conference on Learning Representations (ICLR2014), CBLIS, April 2014*, 2014.
- [Chen and Dolan, 2011] David L Chen and William B Dolan. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 190–200. Association for Computational Linguistics, 2011.
- [Defferrard *et al.*, 2016] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in neural information processing systems*, pages 3844–3852, 2016.
- [Dong *et al.*, 2018] Jianfeng Dong, Xirong Li, and Cees GM Snoek. Predicting visual features from text for image and video caption retrieval. *TMM*, 20(12):3377–3388, 2018.
- [Dong *et al.*, 2019] Jianfeng Dong, Xirong Li, Chaoxi Xu, Shouling Ji, Yuan He, Gang Yang, and Xun Wang. Dual encoding for zero-example video retrieval. In *CVPR*, pages 9346–9355, 2019.
- [Faghri *et al.*, 2018] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. *BMVC*, 2018.
- [Grover and Leskovec, 2016] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864. ACM, 2016.
- [Hammond *et al.*, 2011] David K Hammond, Pierre Vandergheynst, and Rémi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [Katsuki and Constantinidis, 2014] Fumi Katsuki and Christos Constantinidis. Bottom-up and top-down attention: different processes and overlapping neural systems. *The Neuroscientist*, 20(5):509–521, 2014.
- [Kipf and Welling, 2017] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. 2017.
- [Kiros *et al.*, 2014] Ryan Kiros, Ruslan Salakhutdinov, and Richard S Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014.
- [Krishna *et al.*, 2017] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1):32–73, 2017.
- [Lee *et al.*, 2018] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 201–216, 2018.
- [Li *et al.*, 2019] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4654–4662, 2019.
- [Liu *et al.*, 2019] Xinchun Liu, Wu Liu, Meng Zhang, Jingwen Chen, Lianli Gao, Chenggang Yan, and Tao Mei. Social relation recognition from videos via multi-scale spatial-temporal reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3566–3574, 2019.
- [Mithun *et al.*, 2018] Niluthpol Chowdhury Mithun, Juncheng Li, Florian Metze, and Amit K Roy-Chowdhury. Learning joint embedding with multimodal cues for cross-modal video-text retrieval. In *ICMR*, pages 19–27, 2018.
- [Perozzi *et al.*, 2014] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710. ACM, 2014.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [Wang and Gupta, 2018] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 399–417, 2018.
- [Xu *et al.*, 2016] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296, 2016.