

VLG-Net: Video-Language Graph Matching Network for Video Grounding

Mattia Soldan*, Mengmeng Xu*, Sisi Qu*, Jesper Tegner, Bernard Ghanem

King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia

{mattia.soldan, mengmeng.xu, sisi.qu, jesper.tegner, bernard.ghanem}@kaust.edu.sa

Abstract

Grounding language queries in videos aims at identifying the time interval (or moment) semantically relevant to a language query. The solution to this challenging task demands understanding videos' and queries' semantic content and the fine-grained reasoning about their multi-modal interactions. Our key idea is to recast this challenge into an algorithmic graph matching problem. Fueled by recent advances in Graph Neural Networks, we propose to leverage Graph Convolutional Networks to model video and textual information as well as their semantic alignment. To enable the mutual exchange of information across the modalities, we design a novel Video-Language Graph Matching Network (VLG-Net) to match video and query graphs. Core ingredients include representation graphs built atop video snippets and query tokens separately and used to model intra-modality relationships. A Graph Matching layer is adopted for cross-modal context modeling and multi-modal fusion. Finally, moment candidates are created using masked moment attention pooling by fusing the moment's enriched snippet features. We demonstrate superior performance over state-of-the-art grounding methods on three widely used datasets for temporal localization of moments in videos with language queries: ActivityNet-Captions, TACoS, and DiDeMo.

1. Introduction

Temporal action understanding is at the forefront of computer vision research. Hendricks *et al.* [2] and Gao *et al.* [16] recently introduced the task of temporally grounding language queries in videos as a generalization of the temporal action localization task, aiming to overcome the constraint of a predefined set of actions. This novel interdisciplinary task has gained momentum within the vision and language communities for its relevance and possible applications in video retrieval [13, 52, 75], video question answering [21, 28], human-computer interaction [82], and video storytelling [19]. Enabling this fine-grained matching of language in videos can be adopted by professional

*equal contribution.

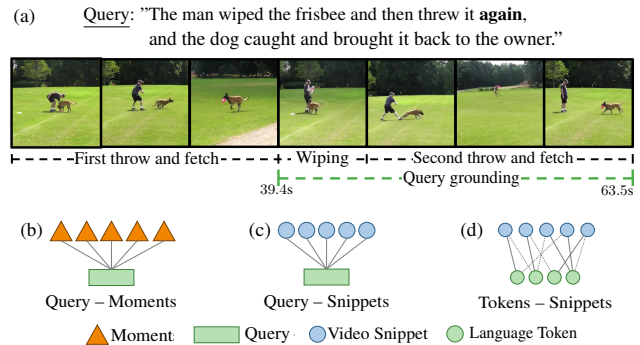


Figure 1. **Temporal video grounding task and multi-modality interaction schemes.** (a) A video grounding example showcasing the importance of fine-grained semantic understanding and proper context modeling. (b,c,d) Approaches for multi-modal interactions. We regard moments and queries as sequences of snippets and tokens respectively. We employ scheme (d), which allows for fine alignment by snippet-token matching.

video content creators during the editing process. For example, video editing often requires searching through many hours of raw, unlabelled video content for specific interesting highlights. Thus, the ability to retrieve such highlights through textual queries could provide a faster experience.

Natural language grounding in videos inherits challenges from temporal action localization such as context modeling and candidate moment generation [1]. Semantic context is a fundamental cue necessary to boost the performance of localization methods [11, 15, 30, 61, 72]. To enrich video representation [2] adopted *global-context*, which is moment independent, leading to sub-optimal performance. Conversely, a moment specific *local-context*, defined as a moment's temporal neighbourhood, was used in [16, 17, 24, 35, 55]. In our view, *non-local context* merits deeper analysis, since it has the potential to identify relevant information not restricted to the temporal neighbourhood within one data modality. For example, in Fig. 1(a), although “*First throw and fetch*” is not in the temporal vicinity of “*Second throw and fetch*”, it is still semantically related with the target moment, showcasing the importance of non-local context modeling for video grounding.

Moreover and as shown by the example, the free-form

nature of the language modality introduces additional challenges. A model must understand the semantic content of both videos and language queries and reason about their multi-modal interactions. Previous studies [16, 18, 55] employ a cross-modal processing unit designed to jointly model text and visual features through simple operations such as element-wise addition, Hadamard product, and direct concatenation of a moment’s representation and the query embedding. A high-level overview of this multi-modal interaction scheme is presented in Fig. 1(b). Instead, recent works such as [44], only employ the Hadamard product to fuse the multi-modal information at the query-snippet level. This scheme, depicted in Fig. 1(c), can determine different correlations between a query and each video snippet, allowing for a finer fusion with respect to Fig. 1(b).

Motivated by the work in [72], we propose to leverage the representational capability of graphs to encode snippet-snippet, token-snippet, and token-token relations as graph edge connections. As such, we design a new architecture referred to as Video-Language Graph Matching Network (VLG-Net) which employs Graph Convolutional Networks (GCN) [57]. First, representation graphs for both video and language are constructed. The video graph models each snippet as a node and takes advantage of two sets of edges to represent both local temporal relations and non-local semantic relations between video snippets. Similarly, we construct a language graph, where each node is a token, and each edge reflects token-to-token relations, *e.g.* syntactic dependencies [40, 41]. These modality-specific graphs are used to model local and non-local intra-modality context through graph convolutions. This sets the stage for addressing modality alignment by recasting inter-modality interactions as an algorithmic graph matching problem. Inspired by [29, 71], we adopt a cross-graph attention-based matching mechanism to enable the mutual exchange of information between modalities, allowing for fine-grained alignment through a specialized set of learnable edges. Unlike some methods that focus on relatively coarse query-moment or query-snippet interactions, and similar to [6, 79], our method performs the matching operation at the level of snippets and tokens, as depicted in Fig. 1(d). With this design, we avoid the need for heuristics of context modeling while learning a successful strategy for multi-modal fusion.

Contributions. (1) We propose VLG-Net, a new deep learning pipeline that consistently adopts graph representation for modeling modality interaction and multi-modal fusion. We address the modality fusion problem by resorting to a graph-matching approach that learns snippet-token connectivity. (2) Through extensive experiments, VLG-Net demonstrates its effectiveness in capturing modality interactions by achieving performance on par or better than state-of-the-art on three standard datasets, showing significant improvements over previously published methods in TACoS [46] and DiDeMo [2] datasets.

2. Related work

2.1. Video Grounding

Moment candidates. Previous works can be categorized into proposal-free and proposal-based methods. Proposal-free approaches [10, 32, 44, 47, 53, 73, 77] aim at directly regressing the temporal boundaries of the queried moment from the multi-modal fused feature. In contrast, proposal-based methods adopt a propose-and-rank pipeline by first generating moment proposals and then ranking them according to their similarity with the textual query [2, 7, 23, 35, 60, 81]. Similar to these approaches, VLG-Net is a proposal based approach.

Moments in context. For moment context modeling, some methods [7, 20] attempt to use the memory property of LSTM cells [51] to contextualize the video features. Alternatively, attention-based mechanisms [59] adopted in [31, 34, 60] can improve the aggregation of long-range semantic dependencies. Similar to [60], we argue that visual context modeling should be dynamic and query-dependent. [81] claims that neighbouring proposals hold valuable context and thus apply 2D convolutions (with large kernel size) to moment representations to gather context information in the latter stages of their pipeline. Compared to [81], we delegate context gathering to earlier stages of our pipeline and only use a Multi-Layer Perceptron (MLP) network for moment score computation, reducing the overall computation.

Multi-modal fusion. Moving beyond the simple scheme adopted in [16], the work of [67] devises a new cross-modality interaction scheme based on circular matrices. In [7, 35, 60], frame features are concatenated with frame-guided attention-pooled features from the query. Lu *et al.* [39] take advantage of the QANet [74] architecture, which is based on cross-attention and convolution operations, for multi-modal fusion. Dynamic filters generated from language features are used in [47, 80] in order to modulate (through convolutions) the visual information based on the query content. Recently, the Hadamard product has become a popular way to fuse/gate multi-modal information [44, 77, 81]. In contrast to these methods, our graph matching layer specifically models local, non-local, and query-guided context, thereby exploiting the semantic neighbourhood of snippets and tokens to fuse the modalities through graph convolutions. Concurrently to our method, [33] adopted attention based cross-modal graphs for fusing the video and language modality. However, opposed to our formulation, [33] lacks a formal design for the graph edges.

2.2. Graphs and Graph Neural Networks

Graphs in Videos. In various video understanding tasks, such as action recognition [9, 36, 65] and action localization [72, 76], graphs can offer extensive representational power to data sequences. For example, a video can be represented as a space-time region graph [65] or as a 3D point

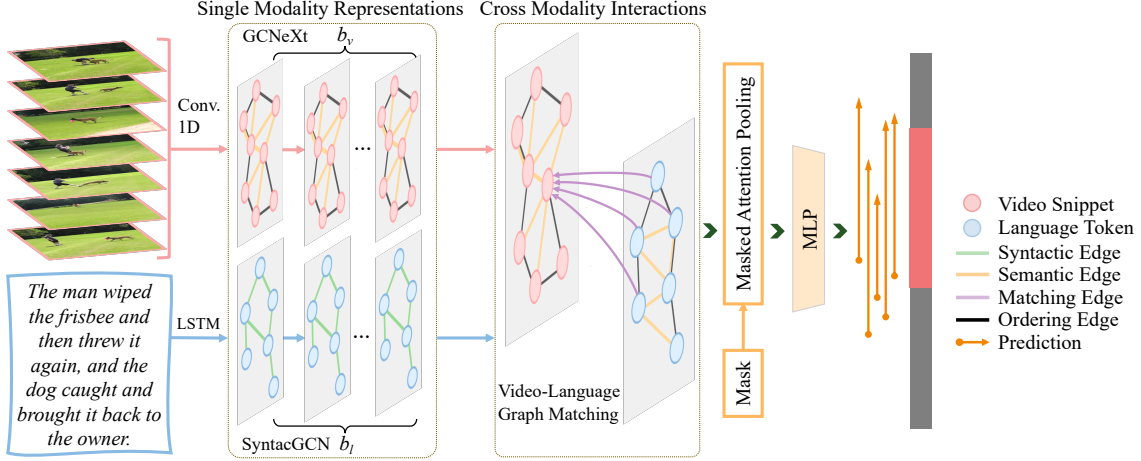


Figure 2. **VLG-Net Architecture.** Inputs are snippets features and tokens embeddings. The video stream comprises: 1D convolutions and b_v GCNeXt operations. Correspondingly, tokens are fed to a stack of LSTM layers and b_l SyntacGCN layers. A graph matching layer is adopted for cross-modal context modeling and multi-modal fusion. Masked attention pooling lists all possible moment candidates, and a Multi-Layer Perceptron (MLP) computes each moment’s score to rank them as final predictions.

cloud in the spatial-temporal space [36]. Moreover, Zeng *et al.* [76] define temporal action proposals as nodes to form a graph, while Xu *et al.* [72] consider video snippets as the graph nodes. Inspired by [72], in VLG-Net, video snippets are represented as nodes in a graph and different specifically designed edges model their relationships.

Graphs in Language. In natural language processing (NLP), both sequential and non-local relations are crucial. The former is usually captured by recurrent neural networks [49], while the latter can be represented using graph neural networks [3, 4, 42, 55]. Moreover, syntactic information has proven useful for language modeling when combined with GCNs [22, 31, 41]. Driven by these findings, we use LSTMs and Syntactic Graph Convolution Networks (SyntacGCN) together to model and enrich the language features in the query.

Graph Neural Networks in Graph Matching. Graph matching is one of the core problems in graph pattern recognition, aiming to find node correspondences between different graphs [5]. Given the ability of Graph Neural Networks (GNNs) to encode graph structure information, approaches leveraging GNNs have recently surfaced to address the graph matching problem [62, 71]. For example, Li *et al.* [29] propose a GNN-based graph matching network to represent graphs as feature vectors, which simplifies measuring their similarity. Following [29], a neighborhood matching network is introduced by [69] to match graph nodes by estimating similarities of their neighborhoods. Due to their superiority in finding consistent correspondences between sets of features, graph matching methods have been widely applied in various tasks [25, 37, 63, 69, 71]. Motivated by these works, we apply graph matching to the video grounding task by specifically employing a cross-graph attention matching mechanism.

3. Methodology

3.1. Problem Formulation

Given an untrimmed video and a language query, the video grounding task aims to localize a temporal moment in the video that matches the query. Each video-query pair has one associated ground-truth moment, defined as a temporal interval with boundary (τ_s, τ_e) . Our method scores m candidate moments, where the k -th moment consists of start time $t_{s,k}$, end time $t_{e,k}$, and confidence score p_k . The video stream is represented as a sequence of n_v **snippets** $V = \{v_i\}_{i=1}^{n_v}$, where each snippet has ϵ consecutive frames. Similarly, a language query is represented by n_l **tokens** $L = \{l_i\}_{i=1}^{n_l}$. The inputs to VLG-Net are n_v snippet features $X_v \in \mathbb{R}^{c_v \times n_v}$ and n_l token features $X_l \in \mathbb{R}^{c_l \times n_l}$ extracted using pre-trained models, where c_v and c_l are the snippet and token feature dimensions. We describe the details of feature extraction in Sec. 4.2.

3.2. VLG-Net Architecture

Our video grounding architecture is illustrated in Fig. 2. First, we feed both the video features X_v and the query embeddings X_l into a stack of computation blocks. On the video path, we use 1D convolutions and GCNeXt [72] blocks to enrich the visual representation with local and non-local intra-modality context. On the language path, we apply LSTM and SyntacGCN [22] to aggregate temporal and syntactic context, which models the grammatical structure of the language query. The two paths converge in the graph matching layer for cross-modal context modeling and multi-modal fusion. After the graph matching layer, we apply masked moment attention pooling to produce the representations of possible moment candidates. Finally, we use an MLP to score the query-moment pair based

on their representation and post-process the score through non-maximum suppression (NMS). We report top- κ ranked moments as the final predictions.

3.3. Video and Language Representations

Here, we detail the set of operations performed on each modality. The stack of computation blocks of each path is specifically designed to model intra-modality context to enrich the snippet and token features.

Video Representation. We add 1D positional encoding, as formulated in [12], to each input visual feature and apply 1D convolutions to map them to a desired dimension. The video is then cast as a graph, where each node represents a snippet and each edge represents a dependency between a snippet pair. We design two types of edges: (i) *Ordering Edges* and (ii) *Semantic Edges*. Static *Ordering Edges* connect consecutive snippets and model the temporal order. Conversely, *Semantic Edges* are dynamically constructed, using the k-nearest neighbors algorithm. They connect semantically similar snippets based on their current feature representations. Specifically, an ordering or semantic snippet neighborhood is determined, and its aggregated representation is computed through edge convolutions \mathcal{F} , similar to [66]. Each edge convolution employs a split-transform-merge strategy [70] to increase the diversity of transformations. These graph operations (called GCNeXt) were proposed in [72] to enrich video snippet representations for the purpose of temporal action localization. In our architecture, we stack b_v GCNeXt blocks together and refer to the input of each block as $X_v^{(i)}$ such that

$$X_v^{(i+1)} = \text{GCNeXt}(X_v^{(i)}) = \sigma \left(\mathcal{F}(X_v^{(i)}, \mathcal{A}_o^{(i)}, W_o^{(i)}) + \mathcal{F}(X_v^{(i)}, \mathcal{A}_s^{(i)}, W_s^{(i)}) + X_v^{(i)} \right), \quad (1)$$

where $X_v^{(0)}$ is the output of the convolutional layer, $\mathcal{A}_o^{(i)}$ and $\mathcal{A}_s^{(i)}$ are the adjacency matrices of *Ordering Edges* and *Semantic Edges*, respectively, and $W_o^{(i)}$ and $W_s^{(i)}$ are the trainable weights for the i -th GCNeXt block. We use Rectified Linear Unit (ReLU) as the activation function σ . Refer to [72] for additional details about GCNeXt. The output of the last block is referred to as $X_v^{(b_v)}$, which is the input to the graph matching layer.

Language Representation. The query token features X_l are fed through an LSTM of b_s layers to capture semantic information and the sequential context. Moreover, given that language follows a predefined set of grammatical rules, we set out to leverage syntactic information [41, 78] to model grammatical inter-word relations. For this purpose, we adopt SyntacGCN, as shown in Fig. 2. Syntactic graphs are preferred over fully connected graphs, since the former’s sparsity property offers more robustness against noise in language [22]. Our SyntacGCN represents a query as a sparse directed graph, in which each output of the last

LSTM layer, referred to as $X_l^{(0)}$, is viewed as a node, and each syntactic relation as an edge. The adjacency matrix \mathcal{A}_l is directly constructed from the query’s syntactic dependencies [40] and the graph convolution is formulated as:

$$X_{l,j}^{(i+1)} = \sigma \left(X_{l,j}^{(i)} + \sum_{k \in \mathcal{N}(j)} \alpha_{jk}^{(i)} \mathcal{A}_{l,jk} W_l^{(i)} X_{l,k}^{(i)} \right), \quad (2)$$

where $X_{l,j}^{(i)}$ is the j -th token feature of previous layer’s output, $\mathcal{N}(j)$ is the syntactic neighbourhood of node j , $W_l^{(i)}$ is the learnable weight in the i -th layer, and σ is ReLU. Moreover, $\alpha_{jk}^{(i)}$ is the edge weight learned from the feature of paired nodes $X_{l,j}^{(i)}$ and $X_{l,k}^{(i)}$, defined as:

$$\alpha_{jk}^{(i)} = \text{SoftMax}(\mathbf{w}_\alpha^{(i)\top} \sigma(W_\alpha^{(i)}(X_{l,j}^{(i)} || X_{l,k}^{(i)}))), \quad (3)$$

where \mathbf{w}_α and W_α are learnable parameters and $||$ denotes vector concatenation. We refer to the last output of the SyntacGCN as $X_l^{(b_l)}$, which will be used to match the video representation $X_v^{(b_v)}$.

3.4. Video-Language Graph Matching

The enriched video representation $X_v^{(b_v)}$ and query representation $X_l^{(b_l)}$ meet and interact in the graph matching layer. This layer models the cross-modal context and allows for multi-modal fusion. It does so by evaluating the intra-modality correlation between each video snippet and between each query token and the inter-modality correlation between each snippet-token pair. The process is shown in Fig. 3. First, we create a video-language matching graph, where each node can be either a video snippet or a query token. We include three types of edges: (i) *Ordering Edge* (\mathcal{O}), (ii) *Semantic Edge* (\mathcal{S}), and (iii) *Matching Edge* (\mathcal{M}).

As depicted in Fig. 3, we use *Ordering Edge* and *Semantic Edge* in the video-language matching graph (as defined in Sec. 3.3), while the *Matching Edge* reflects the inter-modality relation. *Ordering Edge* models the sequential nature of both modalities. For example, if an *Ordering Edge* links two tokens, the words corresponding to the two tokens are consecutive in the input query. *Semantic Edge* is used to connect graph nodes in the same modality according to their feature similarity, providing non-local dependencies over the entire graph. Importantly, *Matching Edge* is employed to explicitly model and learn the cross-modality interaction, to extract meaningful alignment information and learn an aggregation policy. The *Matching Edge* weights are referenced as \mathcal{B} . We use *Matching Edge* to densely connect all possible snippet-token pairs, and set the edge weight proportional to the correlation between the matched node features. Similar to *Semantic Edges*, *Matching Edges* are dynamic and evolve in the training process.

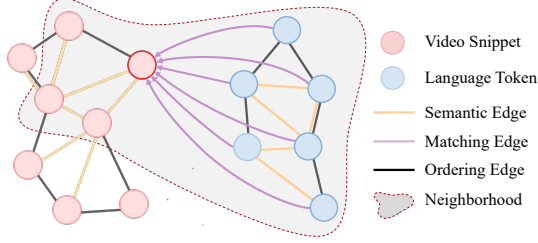


Figure 3. **The video-language matching graph.** The nodes represent video snippets and query tokens. Ordering Edge models the sequential nature of both modalities. Semantic Edge connects graph nodes in the same modality according to their feature similarity. Matching Edge captures the cross-modality relationships. We apply graph convolution on the video-language graph for cross-modal context modeling and multi-modal fusion. The neighborhood is specific for the node highlighted in red.

To combine all three types of edges, we employ relation graph convolution [50] on the constructed video-language matching graph. Eq. 4 shows the high level representation of the convolutions in this layer. Refer to the supplementary material for the full formulation.

$$\mathbf{X}^{(GM)} = \mathcal{A}_O \mathbf{X} \mathbf{W}_O + \mathcal{A}_S \mathbf{B} \mathbf{X} \mathbf{W}_S + \mathcal{A}_M \mathbf{\Gamma} \mathbf{X} \mathbf{W}_M + \mathbf{X} \quad (4)$$

Here, $\mathbf{X} = \{X_{v,1}^{(b_v)}, \dots, X_{v,n_v}^{(b_v)}, X_{l,1}^{(b_l)}, \dots, X_{l,n_l}^{(b_l)}\}$ is the feature representation of all the nodes in the video-language matching graph. \mathcal{A}_r and \mathbf{W}_r for $r \in \{O, S, M\}$ represent the binary adjacency matrix and learnable weights for each set of edges. \mathbf{B} and $\mathbf{\Gamma}$ scale the adjacency matrices \mathcal{A}_S and \mathcal{A}_M , respectively, such that $\beta_{i,j} \in \mathbf{B}$ and $\gamma_{i,j} \in \mathbf{\Gamma}$ are proportional to $\mathbf{x}_i^\top \mathbf{x}_j$. We stack together all video and language node features to form $X_v^{(GM)} \in \mathbb{R}^{c \times n_v}$ and $X_l^{(GM)} \in \mathbb{R}^{c \times n_l}$, and we pass them to the masked moment pooling layer.

3.5. Masked Attention Pooling

The graph matching layer returns a new video graph and a new language graph fused with information from the other modality. Then, a masked attention pooling operation is applied to the new video graph to list the relevant sub-graph representations as candidate moments. The output of this module is denoted as $Y = [\mathbf{y}_k]_{k=1}^m, \mathbf{y}_k \in \mathbb{R}^c$, where m is the number of candidate moments, and c is the feature dimension of each moment. For efficiency purposes, the operation is implemented as a masked attention, allowing us to process each snippet feature only once, while computing each moment’s representation.

We implement three different schemes, namely: (i) learnable self-attention, (ii) cross-attention, and (iii) learnable cross-attention. In (i), we obtain the unnormalized attention weights by applying a 1D convolutional layer that maps each snippet feature to a single score. In (ii) and (iii), we compute the query representation $X_l^{(att)}$ by

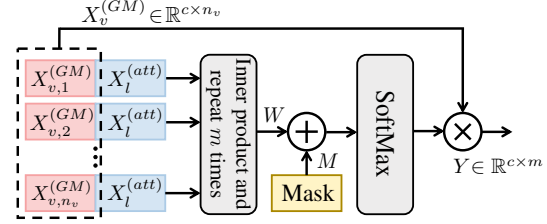


Figure 4. **Masked attention pooling.** Sequence of operations for the learnable cross-attention configuration. Inputs are video nodes $X_v^{(GM)}$ from the graph matching layer and the query embedding $X_l^{(att)}$ computed through self-attention pooling atop the graph matching output. The output Y represents all moment candidates.

applying self-attention pooling atop the graph matching output $X_l^{(GM)}$. Cross-attention obtains the unnormalized weights by computing the inner product between the snippet and query feature, while learnable cross-attention concatenates each snippet feature with the query feature and uses a 1D convolutional layer to obtain the weights. Configuration (iii) is depicted in Fig. 4. In all cases, the unnormalized weight vector has shape $w \in \mathbb{R}^{n_v \times 1}$ for each video. The vector is repeated m times to obtain the matrix $W \in \mathbb{R}^{n_v \times m}$, and a fixed mask $M \in \mathbb{R}^{n_v \times m}$ is applied to it. Similar to Songyang *et al.* [81], we generate m moment candidates and apply a sparse sampling strategy to discard redundant moments. The value m is dataset dependent. The mask is generated according to the sampled moments, highlighting for each of them, which are the snippets that must be taken into account when computing the moment’s pooled feature. The attention scores are then obtained by a softmax operation. Thanks to the masking operation, snippets not related to the n -th moment will not be considered. Finally, the moments’ features are obtained simply as a matrix multiplication $Y = X^{(GM)} \text{SoftMax}(W + M)$.

3.6. Moment Localization

Each candidate moment representation, output of previous module, is endowed with an additive 2D positional embedding encoding the start and end timestamps ($t_{s,k}, t_{e,k}$ with $k \in [1, m]$). Then each moment feature is fed to an MLP to compute the alignment confidence score p_k . This score predicts the Intersection-over-Union (IoU) of each moment with the ground truth of the corresponding video-query pair. For training, we supervise this process using a cross-entropy loss, shown in Eq. 5. Similar to [81], we assign the label $t_k = 1$ if the $\text{IoU} \geq \theta_2$, $t_k = 0$ if the $\text{IoU} \leq \theta_1$, and otherwise $t_k = (\text{IoU} - \theta_1) / (\theta_2 - \theta_1)$.

$$\mathcal{L} = \frac{1}{m} \sum_{k=1}^m t_k \log p_k + (1 - t_k) \log(1 - p_k), \quad (5)$$

During inference, moment candidates are ranked based on their predicted scores and NMS is used to discard highly overlapping moments.

4. Experiments

4.1. Datasets

ActivityNet-Captions [27] is a popular dataset initially collected for the task of dense captioning, and recently adopted for the task of moment localization with natural language queries [7, 31]. The dataset is subdivided into four splits: train, val_1, val_2, and test. The test set is withheld for competition purposes, while leaving the rest publicly available. Refer to Table 1 for details about the publicly available splits. Following the setting in [31], we use val_1 as the validation set and val_2 as the testing set.

TACoS [46] consists of videos selected from the MPII Cooking Composite Activities video corpus [48]. It comprises 18818 video-query pairs of different cooking activities. Each video contains an average of 148 queries, some of which are annotations of short video segments.

DiDeMo [2] contains unedited video footage from Flickr with sentences aligned to unique moments in its 10642 videos. It contains more than 40k video-query pairs with coarse temporal annotations. Moment start and end points are aligned to five-second intervals and the maximum annotated moment length is 30 seconds.

Concerns regarding the Charades-STA [16] dataset discouraged us from evaluating our method on it. Refer to the supplementary material for a detailed discussion.

4.2. Implementation

Evaluation Metrics. We follow the commonly used setting in [14], where the Rank@ κ for IoU= θ serves as our evaluation metric. For example, given a video-query pair, the result is positive if any of the top- κ predictions has IoU with the ground-truth larger or equal to θ ; otherwise the result is negative. We average the results across all testing samples. Following standard practice, we set $\kappa \in \{1, 5\}$ with $\theta \in \{0.3, 0.5, 0.7\}$ for ActivityNet Captions, $\kappa \in \{1, 5\}$ with $\theta \in \{0.1, 0.3, 0.5\}$ for TACoS, and $\kappa \in \{1, 5\}$ with $\theta \in \{0.5, 0.7, 1.0\}$ for DiDeMo.

Language and Video Features. After lower-case conversion and tokenization, we use the pretrained GloVe model [45] to obtain the initial query embedding for every token and extract the syntactic dependencies using the Stanford CoreNLP 4.0.0 parser [40]. The b_s layers of LSTM with 512 hidden units are used as the query encoder. Then, the syntactic GCN encodes syntactic information of the queries and returns a new embedding with 512 dimensions. For visual features, we use pretrained C3D [58] for ActivityNet Captions and TACoS, and VGG16 [54] for DiDeMo, while holding their parameters fixed during training, as they are readily available and commonly used by state-of-the-art methods. We use 1D convolutions to project the input visual features to a fixed dimension (512), and the GCNeXt blocks' hyper-parameters are set as in [72].

Dataset	Num. Videos	Video-Sentence pairs			Vocab. Size
		train	val	test	
ActivityNet Captions [27]	14926	37421	17505	17031	15406
TACoS [46]	127	10146	4589	4083	2255
DiDeMo [2]	10642	33005	4180	4021	7523
Charades-STA [16]	6670	12404	—	3720	1289

Table 1. **Datasets statistics.** We report relevant information for each datasets available for the grounding task.

Implementation details. We use Adam [26] with a StepLR scheduler [38]. We adopt learning rates ranging from 10^{-3} to 10^{-4} for different datasets. The number of sampled snippets n_v is set to 64 for ActivityNet Captions, and 256 for TACoS, and 48 for DiDeMo. The values of (b_v, b_s, b_l) are equal to (1, 3, 4), (4, 5, 2), (2, 3, 4), respectively for the three datasets. In post-processing, we apply NMS with values respectively to the m predictions to filter out highly overlapping moments. Values of m for each dataset are: 1104, 3101, 505 while the NMS thresholds are: 0.5, 0.3, 0.5. We adapt BCE with logits loss to make the training process more numerically stable. The IoU thresholds (θ_1, θ_l) for the three datasets are: (0.7, 0.71), (0.5, 0.7), (0.69, 1.0).

4.3. Comparison with State-of-the-Art

Comparisons are carried out only against methods using the same input features as VLG-Net. In the Tables, we highlight the top-1 and top-2 performance values by bold and underline, respectively.

ActivityNet Captions (Table 2). VLG-Net offers the highest performance for the tight IoU=0.7. However, it falls short against [33] on the lousier IoU=0.5. In practical terms, the two methods are to be considered on par. Nonetheless, notice how tighter IoU translates to better retrieval quality in a real use-case scenario. In these terms, VLG-Net is to be preferred over the competitive [33]. Finally, notice how VLG-Net achieves a significant boost against the recently released 2D-TAN [81] and DRN [77].

TACoS (Table 3). Our model outperforms state-of-the-art methods and achieves the highest scores for all IoU thresholds with significant improvements. In particular, VLG-Net exceeds the previous art [33, 77, 81] by a large margin, ranging from 7.10% to 11.52%, across all evaluation settings, showcasing the excellent design of the architecture.

DiDeMo (Table 4). Our proposed technique outperforms the top-ranked methods [33, 34, 35] with respect to R@1 and R@5 for IoU0.5 and 0.7 with evident increases. It also reaches the highest performance in regards to R@1 IoU1.0. For R@5 IoU1.0, VLG-Net ranks second, falling short with respect to TMN [7] by 1.32%. For completeness, we report the performances of TGN [32] and TMN [7]; with the caveat that their performance could not be verified as the code is not made publicly available. Hence the different colour for the corresponding rows in Table 4.

	R@1		R@5	
	IoU0.5	IoU0.7	IoU0.5	IoU0.7
MCN [2]	21.36	6.43	53.23	29.70
CTRL [16]	29.01	10.34	59.17	37.54
TGN [7]	27.93	—	44.20	—
ACRN [32]	31.67	11.25	60.34	38.57
CMIN [31]	44.62	24.48	69.66	52.96
ABLR [73]	36.79	—	—	—
TripNet [43]	32.19	13.93	—	—
PMI [53]	38.28	17.83	—	—
2D-TAN (P) [81]	44.51	26.54	77.13	61.96
2D-TAN (C) [81]	44.05	27.38	76.65	<u>62.26</u>
DRN [77]	45.45	24.36	77.97	50.30
CSMGAN [33]	49.11	<u>29.15</u>	<u>77.43</u>	59.63
VLG-Net	<u>46.32</u>	29.82	77.15	63.33

Table 2. **State-of-the-art comparison on ActivityNet Captions.** We report the results at different Recall@ κ and different IoU thresholds. VLG-Net reaches the highest scores for IoU0.7 for both R@1 and R@5.

4.4. Ablation Study

To motivate our design choices, we present two ablations that focus on relevant aspects of our method. The first ablation showcases the importance of context modeling. The second investigates VLG-Net’s performance when other commonly adopted multi-modal fusion operations replace the graph matching module. In Table 5, we report our best result (first column) and summarize all ablated variants. For simplicity, we evaluate on the TACoS dataset and specifically focus on its most challenging setups: R@1 IoU0.5 and R@5 IoU0.5. To promote fair comparison, we report each model’s capacity in millions (M) of parameters. More ablations are reported in the supplementary material.

Context Ablation. First, we investigate the impact of different context modeling strategies and compare six variants with our VLG-Net. VLG-Net_{NC} represents a “No Context” architecture, in which we replace the GCNeXt and SyntacGCN modules with fully connected layers that do not model any intra-modality context. Moreover, we switch off the *Ordering* and *Semantic Edges* in the graph matching module. Although the model capacity for VLG-Net_{NC} only drops 0.21M (1.6%), its performance degrades up to 8.82% with respect to VLG-Net. Following [2], we devise VLG-Net_{GM} (“Global on Moments”) and VLG-Net_{GI} (“Global on Input”) experiments. The first one extends VLG-Net_{NC} by concatenating each moment feature with a global video feature after the matching operation. Instead, in VLG-Net_{GI}, we concatenate each snippet and token feature with an average pooled version of their respective raw input features. Differently, following [16], VLG-Net_{LM} (“Local on Moments”) models local context by extending the moment’s boundaries when computing the moment’s features in the masked attention pooling module. The simple context modeling adopted in these architectures

	R@1			R@5		
	IoU0.1	IoU0.3	IoU0.5	IoU0.1	IoU0.3	IoU0.5
MCN [2]	14.42	—	5.58	37.35	—	10.33
CTRL [16]	24.32	18.32	13.30	48.73	36.69	25.42
MCF [68]	25.84	18.64	12.53	52.96	37.13	24.73
TGN [7]	41.87	21.77	18.90	53.40	39.06	31.02
ACRN [34]	24.22	19.52	14.62	47.42	34.97	24.88
ROLE [35]	20.37	15.38	9.94	45.45	31.17	20.13
VAL [56]	25.74	19.76	14.74	51.87	38.55	26.52
ACL-K [18]	31.64	24.17	20.01	57.85	42.15	30.66
CMIN [31]	36.68	27.33	19.57	64.93	43.35	28.53
SM-RL [64]	26.51	20.25	15.95	50.01	38.47	27.84
SLTA [24]	23.13	17.07	11.92	46.52	32.90	20.86
SAP [8]	31.15	—	18.24	53.51	—	28.11
TripNet [43]	—	23.95	19.17	—	—	—
2D-TAN (P) [81]	<u>47.59</u>	<u>37.29</u>	25.32	70.31	<u>57.81</u>	<u>45.04</u>
2D-TAN (C) [81]	46.44	35.22	25.19	<u>74.43</u>	56.94	44.21
DRN [77]	—	—	23.17	—	—	33.36
CSMGAN [33]	42.74	33.90	<u>27.09</u>	68.97	53.98	41.22
VLG-Net	57.21	45.46	34.19	81.80	70.38	56.56

Table 3. **State-of-the-art comparison on TACoS.** Our model outperforms all previous methods achieving significantly higher performance with great margins on all metrics.

	R@1			R@5		
	IoU0.5	IoU0.7	IoU1.0	IoU0.5	IoU0.7	IoU1.0
MCN [2]	—	—	13.10	—	—	44.82
TMN [32]	—	—	18.71	—	—	72.97
TGN [7]	—	—	<u>24.28</u>	—	—	71.43
ACRN [34]	27.44	16.65	—	69.43	29.45	—
ROLE [35]	29.40	15.68	—	70.72	33.08	—
CSMGAN [33]	<u>29.44</u>	<u>19.16</u>	—	<u>70.77</u>	<u>41.61</u>	—
VLG-Net	33.35	25.57	25.57	88.86	71.72	<u>71.65</u>

Table 4. **State-of-the-art comparison on DiDeMo.** Our proposed model outperforms the top ranked method ROLE and ACRN with respect to IoU0.5 and 0.7 for R@1 and R@5 with clear margins. It also reaches the highest performance in regards to R@1 IoU1.0.

allows them to improve their performance with respect to VLG-Net_{NC} up to 2.32%. Nonetheless, they fall short of VLG-Net by 6.5 – 7.4%. Note that VLG-Net_{GM} and VLG-Net_{GI} have a larger model capacity, 0.07M and 0.41M, respectively, compared to VLG-Net. VLG-Net_T (“Temporal Context Only”) is as VLG-Net, where we remove each *Semantic Edge* and replace the SyntacGCN layers with GCNeXt ones. In this variant, only temporal dependencies are modeled. VLG-Net_S (“Semantic Context Only”) does not model temporal dependencies but just the semantic ones. Both models’ performance surpasses 30% for R@1 IoU0.5, which indicates that the introduction of GCNeXt and SyntacGCN layers can boost the performance. Our final architecture takes advantage of both modules, achieving the best results. These ablations demonstrate that temporal and semantic context are complementary and showcase the benefits of the proposed context modeling strategy.

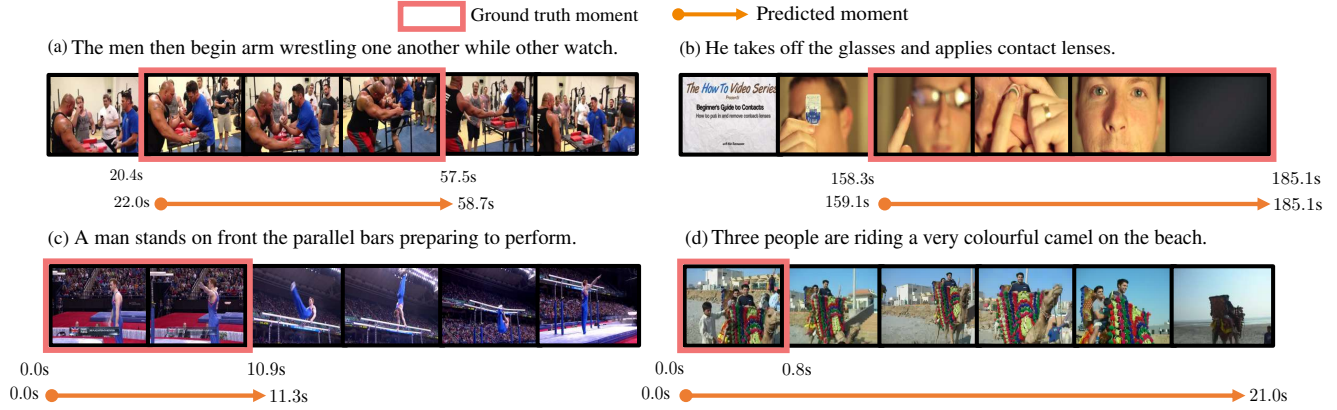


Figure 5. **Qualitative Results.** Examples of grounding results, we compare ground truth annotations (box) and predicted temporal endpoints (arrow). See Section 4.5 for more details.

	VLG-Net	Context Ablation						Multi-modal Fusion Ablation		
		VLG-Net _{NC}	VLG-Net _{GM}	VLG-Net _{LM}	VLG-Net _{GI}	VLG-Net _T	VLG-Net _S	VLG-Net _{HM}	VLG-Net _{HS}	VLG-Net _{CS}
R@1 IoU0.5	34.19	25.37	27.12	26.82	27.69	30.34	30.94	27.27	27.54	28.39
R@5 IoU0.5	56.56	48.24	48.56	47.11	49.39	51.84	53.14	49.99	48.54	51.31
# of Parameters (M)	13.29	13.08	13.35	13.08	13.70	13.22	13.22	13.28	13.28	13.28

Table 5. **Ablation on context modelling and multi-modal fusion approaches.** This ablation shows how the pipeline takes advantage of context modelling and our graph matching module to achieving state-of-the-art performance.

Multi-modal Fusion Ablation. To evaluate our graph matching module’s capabilities for multi-modality fusion, we replace it with other commonly used operations in the literature. Following [81], we use a Hadamard product between the video moment’s features and an average pooled feature for the language query. We adopt several fully connected layers before and after the fusion to keep the model size close to ours. As compared to [81], we adopt the learnable masked attention pooling for generating the moment’s features, which allows for interactions between query tokens and video snippets before the fusion operation. We refer to this model as VLG-Net_{HM} (“Hadamard on Moments”). We also apply the Hadamard product at the snippet level [44, 47, 77] and train VLG-Net_{HS} (“Hadamard on Snippets”). Finally, following [7, 10, 60], we concatenate each snippet feature with the query feature and use linear layers for the projection in VLG-Net_{CS} (“Concatenate on Snippets”). We can conclude that fusing the modalities at the snippet level tends to perform better. The Hadamard operation has gained quite some traction for its good performance and absence of trainable parameters, making it efficient to compute. However, we argue that the complexity of multi-modal alignment calls for more elaborate strategies for multi-modal fusion. Our graph matching module offers a perspective in this research direction.

4.5. Visualization

We show several qualitative grounding results from ActivityNet Captions in Fig. 5. Our VLG-Net can generate

precise moment boundaries that match the query well in different scenarios. Worth mentioning, our method can sometimes give predictions that are more meaningful than the ground truth annotation. As shown in Fig. 5(d), although the ground truth only aligns to the beginning of the video, the query “Three people are riding a very colourful camel on the beach.” can semantically match the whole video. In this case, our VLG-Net gives a more reasonable grounding result. Additional visualizations are reported in the supplementary material.

5. Conclusion

This paper addresses the problem of text-to-video temporal grounding, where we cast the problem as an algorithmic graph matching. We propose Video-Language Graph Matching Network (VLG-Net) to match the video and language modalities. We represent each modality as graphs and explore four types of edges, *Syntactic Edge*, *Ordering Edge*, *Semantic Edge*, and *Matching Edge*, to encode local, non-local, and cross-modality relationships to align the video-query pair. Extensive experiments show that our VLG-Net can model inter- and intra-modality context, learn multi-modal fusion and surpass the current state-of-the-art performance on three widely used datasets.

Acknowledgments. This work was supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research through the Visual Computing Center (VCC) funding.

References

- [1] Humam Alwassel, Fabian Caba Heilbron, Victor Escorcia, and Bernard Ghanem. Diagnosing error in temporal action detectors. In *The European Conference on Computer Vision (ECCV)*, September 2018.
- [2] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing Moments in Video With Natural Language. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [3] Joost Bastings, Ivan Titov, Wilker Aziz, Diego Marcheggiani, and Khalil Sima'an. Graph convolutional encoders for syntax-aware neural machine translation. *arXiv preprint arXiv:1704.04675*, 2017.
- [4] Daniel Beck, Gholamreza Haffari, and Trevor Cohn. Graph-to-sequence learning using gated graph neural networks. *arXiv preprint arXiv:1806.09835*, 2018.
- [5] Tibério S Caetano, Julian J McAuley, Li Cheng, Quoc V Le, and Alex J Smola. Learning graph matching. *IEEE transactions on pattern analysis and machine intelligence*, 31(6):1048–1058, 2009.
- [6] Jingyuan Chen, Xinpeng Chen, Lin Ma, Zequn Jie, and Tat-Seng Chua. Temporally grounding natural sentence in video. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 162–171, 2018.
- [7] Jingyuan Chen, Xinpeng Chen, Lin Ma, Zequn Jie, and Tat-Seng Chua. Temporally Grounding Natural Sentence in Video. In *Proceedings of the conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- [8] Shaoxiang Chen and Yu-Gang Jiang. Semantic Proposal for Activity Localization in Videos via Sentence Query. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.
- [9] Yunpeng Chen, Marcus Rohrbach, Zhicheng Yan, Yan Shuicheng, Jiashi Feng, and Yannis Kalantidis. Graph-based global reasoning networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [10] Chen Shaoxiang, Jiang Yu-Gang. Hierarchical Visual-Textual Graph for Temporal Activity Localization via Language. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [11] Xiyang Dai, Bharat Singh, Guyue Zhang, Larry S. Davis, and Yan Qiu Chen. Temporal Context Network for Activity Localization in Videos. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [13] Jianfeng Dong, Xirong Li, Chaoxi Xu, Shouling Ji, Yuan He, Gang Yang, and Xun Wang. Dual encoding for zero-example video retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [14] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. TALL: Temporal activity localization via language query. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [15] Jiyang Gao, Zhenheng Yang, Kan Chen, Chen Sun, and Ram Nevatia. TURN TAP: Temporal Unit Regression Network for Temporal Action Proposals. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [16] Gao Jiyang, Sun Chen, Yang Zhenheng, Nevatia, Ram. TALL: Temporal Activity Localization via Language Query. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [17] R. Ge, J. Gao, K. Chen, and R. Nevatia. MAC: Mining Activity Concepts for Language-Based Temporal Localization. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019.
- [18] Runzhou Ge, Jiyang Gao, Kan Chen, and Ram Nevatia. MAC: Mining Activity Concepts for Language-based Temporal Localization. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2019.
- [19] Spandana Gella, Mike Lewis, and Marcus Rohrbach. A dataset for telling the stories of social media videos. In *Proceedings of the conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.
- [20] Soham Ghosh, Anuva Agarwal, Zarana Parekh, and Alexander Hauptmann. ExCL: Extractive Clip Localization Using Natural Language Descriptions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.
- [21] Deng Huang, Peihao Chen, Runhao Zeng, Qing Du, Mingkui Tan, and Chuang Gan. Location-Aware Graph Convolutional Networks for Video Question Answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [22] Qingbao Huang, Jielong Wei, Yi Cai, Changmeng Zheng, Junying Chen, Ho-fung Leung, and Qing Li. Aligned Dual Channel Graph Convolutional Network for Visual Question Answering. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020.
- [23] Xu Huijuan, He Kun, Sigal Leonid, Sclaroff Stan, and Saenko Kate. Multilevel Language and Vision Integration for Text-to-Clip Retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.
- [24] Bin Jiang, Xin Huang, Chao Yang, and Junsong Yuan. Cross-Modal Video Moment Retrieval with Spatial and Language-Temporal Attention. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval (ICMR)*, 2019.
- [25] Chenchen Jing, Yuwei Wu, Mingtao Pei, Yao Hu, Yunde Jia, and Qi Wu. Visual-semantic graph matching for visual grounding. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4041–4050, 2020.
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [27] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-Captioning Events in Videos. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [28] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L Berg. TVQA: Localized, Compositional Video Question Answering. In *Proceedings of the conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.

- [29] Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. Graph matching networks for learning the similarity of graph structured objects. *arXiv preprint arXiv:1904.12787*, 2019.
- [30] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. BSN: Boundary Sensitive Network for Temporal Action Proposal Generation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [31] Zhijie Lin, Zhou Zhao, Zhu Zhang, Zijian Zhang, and Deng Cai. Moment Retrieval via Cross-Modal Interaction Networks With Query Reconstruction. *IEEE Transactions on Image Processing*, 2020.
- [32] Bingbin Liu, Serena Yeung, Edward Chou, De-An Huang, Li Fei-Fei, and Juan Carlos Niebles. Temporal Modular Networks for Retrieving Complex Compositional Activities in Videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [33] Daizong Liu, Xiaoye Qu, Xiao-Yang Liu, Jianfeng Dong, Pan Zhou, and Zichuan Xu. Jointly cross-and self-modal graph attention network for query-based moment localization. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4070–4078, 2020.
- [34] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. Attentive Moment Retrieval in Videos. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, 2018.
- [35] Meng Liu, Xiang Wang, Liqiang Nie, Qi Tian, Baoquan Chen, and Tat-Seng Chua. Cross-Modal Moment Localization in Videos. In *Proceedings of the 26th ACM International Conference on Multimedia*, 2018.
- [36] Xingyu Liu, Joon-Young Lee, and Hailin Jin. Learning Video Representations from Correspondence Proposals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [37] Yongfei Liu, Bo Wan, Xiaodan Zhu, and Xuming He. Learning cross-modal context graph for visual grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11645–11652, 2020.
- [38] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic Gradient Descent with warm Restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [39] Chujie Lu, Long Chen, Chilie Tan, Xiaolin Li, and Jun Xiao. DEBUG: A Dense Bottom-Up Grounding Approach for Natural Language Video Localization. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 2019.
- [40] Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. The Stanford CoreNLP natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, 2014.
- [41] Diego Marcheggiani and Ivan Titov. Encoding Sentences with Graph Convolutional Networks for Semantic Role Labeling. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2017.
- [42] Diego Marcheggiani and Ivan Titov. Encoding sentences with graph convolutional networks for semantic role labeling. *arXiv preprint arXiv:1703.04826*, 2017.
- [43] Meera Hahn, Asim Kadav, James M. Rehg, and Hans Peter Graf. Tripping through time: Efficient Localization of Activities in Videos. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020.
- [44] Jonghwan Mun, Minsu Cho, and Bohyung Han. Local-Global Video-Text Interactions for Temporal Grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [45] Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014.
- [46] Michaela Regneri, Marcus Rohrbach, Dominikus Wetzel, Stefan Thater, Bernt Schiele, and Manfred Pinkal. Grounding Action Descriptions in Videos. *Transactions of the Association for Computational Linguistics (ACL)*, 2013.
- [47] Rodriguez Cristian, Marrese-Taylor Edison, Saleh Fatemeh Sadat, Li Hongdong, Gould Stephen. Proposal-free Temporal Moment Localization of a Natural-Language Query in Video using Guided Attention. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2020.
- [48] Marcus Rohrbach, Michaela Regneri, Mykhaylo Andriluka, Sikandar Amin, Manfred Pinkal, and Bernt Schiele. Script data for attribute-based recognition of composite activities. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012.
- [49] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- [50] Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Ri- anne van den Berg, Ivan Titov, and Max Welling. Modeling Relational Data with Graph Convolutional Networks, 2017.
- [51] Hochreiter Sepp and Schmidhuber Jürgen. Long Short-Term Memory. *Neural Computation*, 1997.
- [52] Dian Shao, Yu Xiong, Yue Zhao, Qingqiu Huang, Yu Qiao, and Dahua Lin. Find and Focus: Retrieve and Localize Video Events with Natural Language Queries. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [53] Shaoxiang Chen, Wenhao Jiang, Wei Liu, Yu-Gang Jiang. Learning Modality Interaction for Temporal Sentence Localization and Event Captioning in Videos, 2020.
- [54] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [55] Xiaomeng Song and Yahong Han. VAL: Visual-Attention Action Localizer. In Richang Hong, Wen-Huang Cheng, Toshihiko Yamasaki, Meng Wang, and Chong-Wah Ngo, editors, *Proceedings of the Advances in Multimedia Information Processing (PCM)*, 2018.
- [56] Xiaomeng Song and Yahong Han. VAL: Visual-Attention Action Localizer. In *Proceedings of the Advances in Multimedia Information Processing (PCM)*, 2018.

- [57] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks, 2017.
- [58] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 2015.
- [59] 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 (NIPS)*, 2017.
- [60] Jingwen Wang, Lin Ma, and Wenhao Jiang. Temporally Grounding Language Queries in Videos by Contextual Boundary-aware Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [61] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal Segment Networks: Towards Good Practices for Deep Action Recognition. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- [62] Runzhong Wang, Junchi Yan, and Xiaokang Yang. Learning combinatorial embedding networks for deep graph matching. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3056–3065, 2019.
- [63] Sijin Wang, Ruiping Wang, Ziwei Yao, Shiguang Shan, and Xilin Chen. Cross-modal scene graph matching for relationship-aware image-text retrieval. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1508–1517, 2020.
- [64] Weining Wang, Yan Huang, and Liang Wang. Language-driven Temporal Activity Localization: A Semantic Matching Reinforcement Learning Model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [65] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [66] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay Sarma, Michael Bronstein, and Justin Solomon. Dynamic Graph CNN for Learning on Point Clouds. *ACM Transactions on Graphics*, 2018.
- [67] Aming Wu and Yahong Han. Multi-modal Circulant Fusion for Video-to-Language and Backward. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- [68] Aming Wu and Yahong Han. Multi-modal Circulant Fusion for Video-to-Language and Backward. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (IJCAI)*, 2018.
- [69] Yuting Wu, Xiao Liu, Yansong Feng, Zheng Wang, and Dongyan Zhao. Neighborhood matching network for entity alignment. *arXiv preprint arXiv:2005.05607*, 2020.
- [70] Saining Xie, Ross Girshick, Piotr Dollar, Zhuowen Tu, and Kaiming He. Aggregated Residual Transformations for Deep Neural Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [71] Kun Xu, Liwei Wang, Mo Yu, Yansong Feng, Yan Song, Zhiguo Wang, and Dong Yu. Cross-lingual knowledge graph alignment via graph matching neural network. *arXiv preprint arXiv:1905.11605*, 2019.
- [72] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-TAD: Sub-Graph Localization for Temporal Action Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [73] Yuan Yitian, Mei Tao, and Zhu Wenwu. To Find Where You Talk: Temporal Sentence Localization in Video with Attention Based Location Regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.
- [74] Adams Wei Yu, David Dohan, Minh-Thang Luong, Rui Zhao, Kai Chen, Mohammad Norouzi, and Quoc V. Le. QANet: Combining Local Convolution with Global Self-Attention for Reading Comprehension. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [75] Youngjae Yu, Jongseok Kim, and Gunhee Kim. A joint sequence fusion model for video question answering and retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [76] Runhao Zeng, Wenbing Huang, Minghui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph Convolutional Networks for Temporal Action Localization. *arXiv preprint arXiv:1909.03252*, 2019.
- [77] Runhao Zeng, Haoming Xu, Wenbing Huang, Peihao Chen, Minghui Tan, and Chuang Gan. Dense Regression Network for Video Grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [78] Yuhao Zhang, Peng Qi, and Christopher D. Manning. Graph Convolution over Pruned Dependency Trees Improves Relation Extraction. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018.
- [79] Zhu Zhang, Zhijie Lin, Zhou Zhao, and Zhenxin Xiao. Cross-modal interaction networks for query-based moment retrieval in videos. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 655–664, 2019.
- [80] Zhang Da, Dai Xiyang, Wang Xin, Wang Yuan-Fang, and Davis Larry S. MAN: Moment Alignment Network for Natural Language Moment Retrieval via Iterative Graph Adjustment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [81] Zhang Songyang, Peng Houwen, Fu Jianlong, Luo, Jiebo. Learning 2D Temporal Adjacent Networks for Moment Localization with Natural Language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [82] Fengda Zhu, Yi Zhu, Xiaojun Chang, and Xiaodan Liang. Vision-Language Navigation With Self-Supervised Auxiliary Reasoning Tasks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.