

Temporally Grounding Instructional Diagrams in Unconstrained Videos

Zhang, Jiahao; Zhang, Frederic; Rodriguez, Cristian; Ben-Shabat, Itzik; Cherian, Anoop; Gould,
Stephen

TR2025-002 January 08, 2025

Abstract

We study the challenging problem of simultaneously localizing a sequence of instructional diagram queries in a video. This requires understanding not only the individual diagram queries but also their interrelationships. However, most existing methods focus on grounding one query at a time, ignoring the inherent structures among queries such as the general mutual exclusiveness and the temporal order. Consequently, the predicted timespans of different step diagrams may overlap considerably or violate the temporal order, thus harming the accuracy. In this paper, we tackle this issue by simultaneously grounding a sequence of step diagrams. Specifically, we propose composite queries, constructed by exhaustively pairing up the visual content features of the step diagrams and a fixed number of learnable positional embeddings. Our insight is that self-attention among composite queries carrying different content features suppress each other to reduce timespan overlaps in predictions, while the cross-attention corrects the temporal misalignment via content and position joint guidance. We demonstrate the effectiveness of our approach on the IAW dataset for grounding step diagrams and the YouCook2 benchmark for grounding natural language queries, significantly outperforming existing methods while simultaneously grounding multiple queries.

IEEE Winter Conference on Applications of Computer Vision (WACV) 2024

Temporally Grounding Instructional Diagrams in Unconstrained Videos

Jiahao Zhang¹ Frederic Z. Zhang² Cristian Rodriguez²
Yizhak Ben-Shabat^{1,3} Anoop Cherian⁴ Stephen Gould¹

¹The Australian National University ²The Australian Institute for Machine Learning

³Technion Israel Institute of Technology ⁴Mitsubishi Electric Research Labs

¹{first.last}@anu.edu.au ²{first.last}@adelaide.edu.au

³sitzikbs@gmail.com ⁴cherian@merl.com

Abstract

We study the challenging problem of simultaneously localizing a sequence of instructional diagram queries in a video. This requires understanding not only the individual diagram queries but also their interrelationships. However, most existing methods focus on grounding one query at a time, ignoring the inherent structures among queries such as the general mutual exclusiveness and the temporal order. Consequently, the predicted timespans of different step diagrams may overlap considerably or violate the temporal order, thus harming the accuracy. In this paper, we tackle this issue by simultaneously grounding a sequence of step diagrams. Specifically, we propose composite queries, constructed by exhaustively pairing up the visual content features of the step diagrams and a fixed number of learnable positional embeddings. Our insight is that self-attention among composite queries carrying different content features suppress each other to reduce timespan overlaps in predictions, while the cross-attention corrects the temporal misalignment via content and position joint guidance. We demonstrate the effectiveness of our approach on the IAW dataset for grounding step diagrams and the YouCook2 benchmark for grounding natural language queries, significantly outperforming existing methods while simultaneously grounding multiple queries.

1. Introduction

Instructional content has become increasingly prevalent, spanning a wide array of applications from DIY projects to educational tutorials. One area where this is particularly evident is in instructional videos, which have surged in popularity on platforms like YouTube. These videos often provide step-by-step guides on a variety of tasks, including assembly [3, 52], cooking [7, 27, 33, 54], applying makeup [40], how-to [23] topics and general tasks [37].

These guides often boil down to a simple breakdown of complex procedures into manageable and often sequential steps, enabling viewers to easily follow and replicate the processes. Nevertheless, it could sometimes be difficult to navigate these lengthy videos—some lasting longer than an hour!—if one intends to find a specific instruction or step.

Among the above-mentioned datasets, the recently proposed Ikea Assembly in the Wild (IAW) dataset [52] offers a challenging and real-world setting for research. The IAW dataset distinct itself from others in two main aspects. First, the IAW provides temporal correspondences between segments and step diagrams from instruction manual book, rather than textual descriptions. Second, there are inherent structures among these step-by-step diagrams within a manual book, where the corresponding timespans are naturally expected to be mutually exclusive and sequentially ordered. However, the videos are obtained from YouTube and annotated via crowd-sourcing such that there is no guarantee that the video creators would strictly follow the instruction manual during assembly like those under lab environment. Additionally, corresponding video segments can overlap with each other considerably (e.g., multiple assemblers, two steps concurrently), and some steps might have multiple segments (e.g., repeating two steps alternately) or none (e.g., optional steps). Therefore, the structural prior between queries can only be served as weak general guidance rather than strict constraints. Consequently, we formulate the task as predicting timestamps for each step diagram from a publicly accessible manual, given an unconstrained video of assembling a piece of furniture, as shown in Fig. 1.

In this work, we propose a novel approach to simultaneously ground a sequence of instructional step diagrams in a video. Drawing insights from DETR-based temporal grounding methods [11, 18], we argue that given the interdependencies among the steps in the instruction diagrams of a manual, it may be beneficial to jointly localize all the steps to the respective segments in the associated video rather than localizing one step at a time. To ground




Figure 1. An illustration of the temporal instructional diagram grounding task between a YouTube video (top) [xPNkHAii3fU](#) and an Ikea furniture manual (bottom) [00352894](#). This task aims to predict the start and end timestamps for all step diagrams simultaneously.

all step diagrams at the same time, we first directly concatenate feature representation of one step diagram (content query) with one learnable positional query (position query) to create one composite query. Then, the final composite queries, which fed into the decoder, are constructed by exhaustively pairing up all possible combinations of content and position queries. Such that, these composite queries are able to interact with each other via self-attention, allowing the modeling of the inherent structures of a step diagram sequence. Additionally, since each composite query already carries the respective step diagram, the correspondence is known, eliminating the need for applying classification [18] or matching [41] at the final output layer. It is also worth noting that the position queries of composite queries act as input-agnostic learned temporal bias, similar to templates [4, 22] in object detection transformers, which guide the cross-attention temporally in conjunction with the semantics of the video and diagram.

Our primary contribution is a novel paradigm for a detection transformer model capable of temporally grounding a sequence of step diagrams simultaneously. Building on Moment-DETR [18], we demonstrate that by incorporating the inherent structure of a sequence of step diagrams via constructing composite queries, the model can reduce the overlapping of predicted timespans and enhance the temporal correlation with respect to the ground truth. Visualizing the cross-attention maps reveals that the content and position joint cross-attention enables the model to do temporal corrections. To tailor the transformer decoder module to our problem, we design various attention masks in the self-attention module and experiment with different choices in the transformer decoder. Our method is complementary to those focusing on improving the performances for singleton queries. By integrating our method into a stronger single-query model, EaTR [11], we observed a consistent performance boost. Our method significantly outperforms exist-

ing ones on the IAW [52] dataset for step diagram sequence grounding and on YouCook2 [14] benchmark for sentence sequence grounding, demonstrating its effectiveness in simultaneously grounding multiple queries.

2. Related Work

Temporal sentence grounding in videos (TSGV) aims to locate a segment in a video given a textual query. There has been significant research interests [51] since its introduction by Hendricks *et al.* [1] and Bao *et al.* [9]. Research in this area primarily divides into two approaches: proposal-based and proposal-free methods. Proposal-based methods adapts a *propose-then-match* strategy, initially using a sliding window technique [1, 9] to create candidate proposals. To improve efficiency and reduce redundancy, anchor-based methods are developed [6, 44, 46]. A notable contribution in this category is the 2D-TAN model [53], which utilizes a 2D map to enumerate all possible anchor proposals and consider their temporal relationships. Later, proposal-free methods aim to identify the target moment directly without generating proposals. This includes regression-based methods [10, 45], that directly regress the start and end time of the target moment, and span-based methods [29–32, 49, 50] that calculate the probability of each video frame or clip being the start or end of the target.

DETR-based TSGV methods draw significant inspiration from the success of DETR [4], adapting its concepts to the regression-based, proposal-free domain of TSGV. By conceptualizing the grounding task as a 1D variant of object detection, the pioneering work Moment-DETR [18] first adapted DETR and incorporated saliency for each sentence as guidance. Subsequent works have built on this foundation with various enhancements: QD-DETR [25] improved predictions by detecting not only positive but also negative queries; MS-DETR [13] merged the concepts of 2DTAN

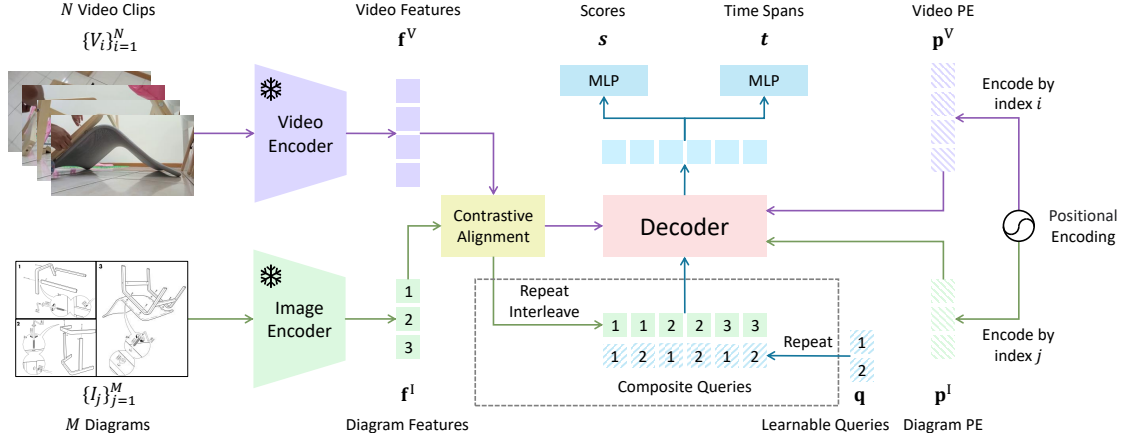


Figure 2. Schematic of our overall pipeline. First, we extract features from the inputs using pre-trained video and image encoders and employ two MLPs to project these features into a unified D -dimensional space, resulting in $\mathbf{f}^V \in \mathbb{R}^{N \times D}$ and $\mathbf{f}^I \in \mathbb{R}^{M \times D}$, respectively. Following Zhang *et al.* [52], first apply contrastive alignment then project these two features into a single joint embedding space. Second, positional encoding (PE) is used to obtain a video \mathbf{p}^V and a diagram \mathbf{p}^I PE vectors. Third, we build composite queries by repeating diagram features and learnable queries. Last, all features, positional encoding, and composite queries are fed into a variant of a transformer decoder, which is specifically designed for this task (Sec. 3.2). The decoder’s output is then processed by another two MLPs, tasked with generating scores \mathbf{s} and timespans \mathbf{t} , respectively.

and DETR, introducing sampling moment-moment interactions, EaTR [11] learned pseudo events from video as priors injected into queries, UniVTG [19], TaskWeave [43], and UVCOM [42] studied the joint learning of moment retrieval and highlight detection. CG-DETR [24] guided the cross-attention via word-level textual conditions to reduce wrong correspondences. RGTR [35] aimed to make the predicted region of each query more distinct and concentrated. However, all the above methods are designed for the single-query scenario. The exception, LVTR [41], specifically addresses the challenge of multi-sentence grounding by concatenating a fixed four sentences with video features before encoding, however performs poorly when the number of queries varies across the dataset like in the IAW. Where we tackle this issue by constructing composite queries.

Dense events grounding in videos was recently introduced by Bao *et al.* [2]. This task extends TSGV tasks by aiming to ground all sentences simultaneously. Bao *et al.* also introduced a model titled DepNet that adopts the *propose-then-match* approach. Another study by Tan *et al.* [36] explores the hierarchical structure of words, sentences, and paragraphs with HSCNet. Jiang *et al.* [12] proposed SVPTTR, advancing the task into the semi-supervised learning domain. The work most similar to ours is PRVG [34], but differs in two key ways: 1) PRVG eliminates the use of learnable queries, restricting the model to make only one prediction per text query; 2) it fuses video and text features before the decoding process using two encoders. All mentioned methods are evaluated on the ActivityNet Captions [16] and TACoS [27] datasets. In these two datasets, the sentences from ground truth and video segments correspond one-to-

one, which means that there is no instance where the same sentence corresponds to multiple timespans, or cases where the one sentence does not have a corresponding timespan. This is in contrast to the IAW [52] dataset we primarily use.

3. Temporal Diagram Grounding in Videos

3.1. Problem Definition

Temporal instructional diagram grounding in untrimmed videos is an extension of TSGV task. Instead of grounding a single descriptive sentence, we aim to simultaneously ground a sequence of diagrams extracted from an instruction manual. Formally, given an untrimmed video that has been evenly divided into N clips for feature extraction, we denote these clips by $\{V_i\}_{i=1}^N$, and the set of M instructional diagrams corresponding to the video by $\{I_j\}_{j=1}^M$, our objective is to develop a model capable of accurately predicting the timespan of each diagram $\mathbf{t} = (t_s, t_e)$, where t_s and t_e are the normalized start and end time of a segment.

3.2. Proposed Method

We propose a method that adaptable to DETR-based methods, as depicted in Fig. 2, which includes a description of the overall architecture. In the following, we focus specifically on introducing the proposed decoder.

Decoder-Based Architecture. As depicted in Fig. 3, inspired by the DETR [4] architecture for object detection, models for TSGV, such as Moment-DETR [18] (Fig. 3a), typically feature with one encoder for fusing multimodal inputs and one decoder dedicated to processing learnable queries and predicting the timespans. However, for the




Figure 3. Illustration of existing DETR-based temporal sequence grounding models for comparing with our method. We use three learnable queries and two diagrams as an example ($K = 3, M = 2$). Note the repeat strategies for learnable queries and diagram features in (d) are different, so that they can be intersecting fused.

task of temporal action localization, Liu *et al.* [20] showed through ablation studies that the presence of the encoder may harm the performance. Kim *et al.* [15] delved deeper into this issue, attributing the negative effect to what they termed *temporal collapse*, where the self-attention only focus on a few key elements. Motivated by these findings, we propose to use the diagram features as content priors for the queries and values, video features as keys to the decoder, and remove the encoder, as depicted in Fig. 3c.

Composite Query via Duplication. As shown in Fig. 3b, LVTR [41], albeit being able to ground multiple sentences simultaneously, employs a fixed number of learnable queries, each of which predicts the timespan for one sentence. To accommodate a varying sequence length, it has to use a large number (≥ 80 in IAW) of learnable queries, which results in most of such queries not being matched to any ground truth instances during training, thus not receiving enough gradients to be trained effectively. We address this issue by using the feature representations of diagrams as content priors, and only employ a relative small number (≤ 10 in IAW) of learnable queries. Each diagram will then be exhaustively paired up with the learnable queries, which can be easily implemented via duplication as shown in the dashed area in Fig. 2. Formally, given M diagram features $\mathbf{f}^l = \{\mathbf{f}^l_i\}_{i=1}^M$ and K learnable positional embeddings $\mathbf{q} = \{\mathbf{q}_j\}_{j=1}^K$, the *composite query* is defined as

$$\{(\mathbf{f}^l_i \oplus \mathbf{q}_j) \mid (\mathbf{f}^l_i, \mathbf{q}_j) \in \mathbf{f}^l \times \mathbf{q}\}, \quad (1)$$

where \oplus denotes the concatenation operation and \times is the Cartesian product. Drawing intuitions from previous work [4, 22], queries in the decoder typically encapsulate two key aspects: semantic and positional information. As Meng *et al.* [22] pointed out, the learnable queries, which are implemented as learnable positional embeddings, serve as positional priors, while the subsequent cross-attention layers inject the relevant semantics into these queries. The composite queries have a distinct advantage, in that the semantic information is explicitly encoded into their representations as content priors, before the cross-attention layers. Consequently, the model is relieved from the challenging task of aggregating semantics, particularly when the num-




Figure 4. Illustration of the decoder architecture. We use \odot to represent first concatenate then project back operation. The concatenation order is always first features then positional encodings.

ber of step diagrams is dynamic, allowing it to concentrate more on the grounding task. During training, we perform Hungarian matching [4] only between the set of composite queries that share the same content prior, *i.e.*, diagram features and the ground truth for that step diagram, significantly simplifying the matching process.

Content and Position Jointly Guided Cross-Attention. Recall that the attention mechanism in transformers [38] uses dot-product between the keys and queries (QK^T) as a similarity measurement to direct information aggregation. We delve into the intuitions behind this mechanism and tailor the design to our specific problem. Following previous practice [22], we denote the keys and queries by K_c and Q_c respectively, as they primarily contain content features, and their positional embeddings by K_p and Q_p . In our model, the content aspect is captured by the features extracted from the frozen encoders (Fig. 2), while positional information is embedded using the sinusoidal functions. Drawing inspiration from [22, 48], the content features are concatenated to the positional embeddings, instead of being added, as it removes two cross-terms ($Q_c K_p^T, Q_p K_c^T$) that may be noisy. Thus, the attention weights before softmax and scaling becomes

$$(Q_c \oplus Q_p)(K_c \oplus K_p) = Q_c K_c^T + Q_p K_p^T, \quad (2)$$

where $Q_c K_c^T$ calculates the semantic similarity between




Figure 5. Illustration of different types of self-attention mask applied in multi-head attention module in Fig. 4a. We use three queries, two diagrams ($K = 3, M = 2$) in this example. We use white to indicate the attention weights that are masked and ignored.

video and diagram features while $Q_p K_p^T$ measures the similarity between video positional encodings and the learnable positional embeddings. As such, the cross-attention mechanism is guided by both the semantic and positional cues.

Masked Self-Attention. The self-attention module among queries in each layer of the decoder, as shown in Fig. 4a, has been shown to exhibit suppressive behaviors [4, 47], which eliminate the needs of non maximum suppression (NMS) on predictions. A composite query by design could suppress other competing composite queries to secure the highest score at the end. But with the duplication mechanism mentioned above, the vanilla self-attention (type A in Fig. 5a) may introduce some undesired pairs of suppression. Specifically, for a given diagram associated with K queries, the intended behavior is for one query to effectively suppress the others, mirroring the block diagonal pattern seen in type B (Fig. 5b). Additionally, when the same query is used across M different diagrams, it is necessary for one diagram to suppress the others to reduce overlaps, as depicted by the sub-diagonals in type C (Fig. 5c). However, attention scores between different queries and diagrams might inadvertently decrease the model’s adaptability by suppressing all potential timespan prediction pairs. To mitigate this, we implement the self-attention mask type D (Fig. 5d) to filter out these undesirable pairings.

Diagram Order Prior. We incorporate the positional encoding of the diagram indices \mathbf{p}^l , into the output from the self-attention module, as depicted in Fig. 4b. This is to ensure coherence with the positional encoding of the video \mathbf{p}^v , effectively aligning both modalities in terms of their temporal ordering information.

Choice of the Attention Value Feature. The core learning objective of the decoder is to derive a mapping function that translates the input Q, K, V into some outputs that implicitly represent scores \mathbf{s} and timespans \mathbf{t} . The outputs are essentially weighted aggregates of the values V , gathering from both self-attention V^s and cross-attention V^c . For the timespan predictions, the mapping can be learned either explicitly from the positional encoding space ($V \subseteq \{\mathbf{p}^v, \mathbf{p}^l, \mathbf{q}\}$) or implicitly from the feature space ($V \subseteq$

$\{\mathbf{f}^v, \mathbf{f}^l\}$). To maximize adaptability and leverage information from both these spaces, we opt to concatenate these two then project them back to their original dimensionality as shown in Fig. 4, thereby affording the model greater flexibility in integrating diverse cues.

3.3. Losses

To train our model, we adopt the moment retrieval loss from Moment-DETR [18]. Formally, denote the predicted timespan for a composite query by $\hat{\mathbf{t}}_{i,j} \in [0, 1]^2$, where $i \in \{1, \dots, M\}$ indexes a step diagram and $j \in \{1, \dots, K\}$ indexes a learnable positional embedding. Denote the normalized prediction scores by $\hat{\mathbf{s}}_{i,j} \in \mathbb{R}$. Following previous practice, we pad the ground truth with \emptyset to size MK and use the Hungarian algorithm [17] for bipartite matching. For brevity of exposition, we use $\mathbf{t}_{i,j}^*$ and $\mathbf{s}_{i,j}^*$ to denote the ground truth for composite query (i, j) , where $\mathbf{s}_{i,j}^* = 0$ if the matched ground truth is an empty set and otherwise $\mathbf{s}_{i,j}^* = 1$ and there are C_i out of K matched for i -th diagram. We then employ a combination of three discrepancy measurements: a L_1 loss and a generalized IoU loss [28] between predicted and ground truth timespans,

$$\mathcal{L}_{L_1} = \frac{1}{M} \sum_{i=1}^M \frac{1}{C_i} \sum_{j=1}^K \mathbb{1}_{\{\mathbf{s}_{i,j}^* = 1\}} \|\hat{\mathbf{t}}_{i,j} - \mathbf{t}_{i,j}^*\|_1, \quad (3)$$

$$\mathcal{L}_{\text{gIoU}} = \frac{1}{M} \sum_{i=1}^M \frac{1}{C_i} \sum_{j=1}^K \mathbb{1}_{\{\mathbf{s}_{i,j}^* = 1\}} \text{gIoU}(\hat{\mathbf{t}}_{i,j}, \mathbf{t}_{i,j}^*); \quad (4)$$

and a cross entropy loss to classify each predicted timespan being foreground or background, depending on whether it is matched by Hungarian algorithm,

$$\mathcal{L}_{\text{score}} = -\frac{1}{MK} \sum_{i=1}^M \sum_{j=1}^K \mathbf{s}_{i,j}^* \log(\hat{\mathbf{s}}_{i,j}). \quad (5)$$

Thus, the overall target is

$$\mathcal{L} = \lambda_{L_1} \mathcal{L}_{L_1} + \lambda_{\text{gIoU}} \mathcal{L}_{\text{gIoU}} + \lambda_{\text{score}} \mathcal{L}_{\text{score}}, \quad (6)$$

where $\lambda_{L_1}, \lambda_{\text{gIoU}}, \lambda_{\text{score}}$ are hyperparameters used for balancing each loss term.

4. Experiments and Results

4.1. Experiment Setup

Datasets. The primary dataset used for designing and evaluating our model is the Ikea Assembly in the Wild (IAW) [48], which features 1,007 untrimmed assembly videos obtained from YouTube across 420 different furniture, each with corresponding instructional step diagrams. The IAW dataset is divided into training, validation, and testing splits with 9,859, 2,210, and 3,512 timespans, respectively. Additionally, we also evaluate our model on YouCook2 [54], which has 2000 untrimmed videos from 89 cooking recipes. The procedure steps for each video are annotated with temporal boundaries and described by imperative English sentences. YouCook2 has 10,337 and 3,492 ground truth timespans for train and validate, respectively.

Preprocessing. For the IAW dataset, we leverage the recent VideoMAEv2 [39] and DINOv2 [26] to extract features from videos and diagrams, respectively. Then, following Zhang *et al.* [52], contrastive learning is applied at the feature level to enhance the representational alignment (refer to supplementary for details). For the YouCook2 dataset, I3D [5] video and BERT [8] text features (averaged word features as a sentence feature) are extracted and established benchmarks [14, 29, 30] are used to ensure a fair and consistent comparison.

Sliding Window Sampler. A sliding window sampler is implemented to temporally sub-samples the videos into chunks with given window size w and step stride s . Furthermore, by aggregating sampled segments from various sliding window configurations (varying both sizes and strides) we enhance the diversity of our training data.

Training and Inference. In line with the approach used by 2DTAN [53], we normalize the length of video features to 256 through interpolation. For the sampling process, we configure the window sizes $w \in \{128, 256, 512, 1024, \infty\}$ and strides $s \in \{32, 64, 128, 256, \infty\}$, with ∞ indicating the use of the entire video. The number of queries is set to 3 based on empirical findings. We use AdamW [21] as the optimizer with a multi-step learning rate scheduler starting at 10^{-4} and a weight decay of 10^{-4} . Training is conducted on a single A100 GPU for 60 epochs, using a batch size of 16, which roughly takes 20 hours. The model demonstrating the best performance on the validation split is selected for reporting results on the test split of the IAW dataset. For YouCook2, since they only provide ground truth for validation split, we use that for both model selection and result report. The code will be made available upon acceptance.

Evaluation Metrics. To evaluate our model, we adapt standard recall at k given a temporal intersection over union threshold m ($R@k, IoU = m$) and the mean intersection over union (mIoU) metrics. Conventionally, we choose $k = 1$ and $m \in \{0.3, 0.5, 0.7\}$.

Table 1. Evaluation results of different methods on IAW [52] dataset test split. Mode *One* denotes one diagram at a time, and *All* means all diagrams at the same time by a single run. All the models listed are trained with the same aligned features unless specified otherwise. The values in bold represent the best results, while underlined indicate the second.

Method	Mode	R@1, IoU=			mIoU
		0.3	0.5	0.7	
Random	-	1.809	0.254	0.057	4.801
LVTR [41]	All	11.26	4.591	1.112	7.515
2D-TAN [53] conv	One	31.24	18.94	8.030	20.51
2D-TAN [53] pool	One	32.94	20.02	8.170	21.21
Moment DETR [18]	One	34.00	18.34	7.290	16.60
Ours w/ Moment DETR	All	37.79	22.74	9.140	23.86
EaTR [11]	One	38.48	22.77	9.540	24.75
Ours w/ EaTR	All	42.02	26.45	11.54	27.27

Table 2. Evaluation results on the val split of YouCook2 [54].

Method	Text	R@1, IoU=			mIoU
		0.3	0.5	0.7	
DORi [30]	-	43.36	30.47	18.24	30.46
LocFormer [31]	BERT [8]	<u>46.76</u>	<u>31.33</u>	15.81	<u>30.92</u>
ExCL [10]	BERT * [14]	26.63	16.15	8.51	18.87
TMLGA [29]		34.77	23.05	12.49	24.42
DORi [30]		42.27	29.90	18.38	29.92
Ours w/ EaTR		52.95	36.28	18.50	35.32

4.2. Main Results

Results on IAW. We report the performance of our model on the IAW [52] dataset in Tab. 1. When grounding all diagrams at the same time (mode *All*), LVTR’s performance drops significantly. It is perhaps due to the challenges posed by its Explore-and-Match (Fig. 3b) module, which seeks to align a mismatched number of learnable queries and diagrams. The queries in LVTR need to be able to aggregate both the semantic and positional information such that they can be matched to the inputs and also be decoded to timespans. Where ours already carry semantics. This is even harder in IAW dataset, since for different videos, the manual books are different. Thus, the performance of LVTR is even lower than 2D-TAN and Moment-DETR. Built upon Moment-DETR, our multi-query version denoted as “Ours w/ Moment DETR” further improves performance by a significant margin comparing with its single-query version, showing the importance to incorporate the inherent structure within a sequence of queries, which benefits from the design of composite queries and content position joint guided cross-attention. Besides that, we tried to integrate ours techniques to a stronger single-query method EaTR [11] as “Ours w/ EaTR¹”, the consistent observation of the performance improvement shows our design could be also beneficial to other DETR-based single-query methods.

Results on YouCook2. We report the performance of our model based on EaTR on the YouCook2 [54] dataset in Tab. 2. Notably, unlike our approach, DORi [30] utilizes

¹We refer the details of modification to the supplementary.

Table 3. Ablation results for different types of self-attention mask. The type name of the mask corresponds to that in Fig. 5.

Mask Type	R@1, IoU=			mIoU
	0.3	0.5	0.7	
A	36.25	21.48	8.655	23.07
B	35.19	20.14	8.398	22.37
C	<u>36.56</u>	<u>21.77</u>	<u>8.712</u>	<u>23.43</u>
D	37.79	22.74	9.140	23.86

Table 4. Ablation results for different value choices for self-attention V^s and cross-attention V^c . We use \oplus to represent first concatenate then project back operation. \dagger : The self-attention output becomes Q_c , and Q_p is then the learnable queries q .

#	V^s	V^c	R@1, IoU=			mIoU
			0.3	0.5	0.7	
1	f^{\dagger}	f^{\dagger}	24.82	12.43	4.456	16.34
2	q	p^V	35.68	20.85	8.255	23.29
3	q	$p^V + f^V$	<u>36.65</u>	20.99	8.026	23.18
4	$q + f^{\dagger}$	$p^V + f^V$	35.56	20.51	<u>8.769</u>	22.81
5	q	$p^V \oplus f^V$	36.33	<u>21.37</u>	8.283	<u>23.37</u>
6	$q \oplus f^{\dagger}$	$p^V \oplus f^V$	37.79	22.74	9.140	23.86

additional object relationship features. Additionally, our sentence features are obtained by averaging pre-extracted word features, whereas LocFormer [31] and DORi [30] use text encoders trained alongside their entire models. Despite these disadvantages due to the lack of text-related enhancements, our model still demonstrates state-of-the-art performance using the same video and text features (BERT \ast). This highlights the effectiveness of our design.

4.3. Ablation Study

Unless specified otherwise, we conducted the ablation study using our moment-DETR based model and evaluated on the test split of the IAW dataset.

Choice of Self-Attention Masks. As show in Tab. 3, the fact that type D actually outperforms type A suggests that those attentions crossing between different diagrams and different learnable queries tend to act as noise for the model. By masking them out, type D achieves the best performance. Notably, the implementation of type B, which essentially mirrors the single-query methods, has inferior performance than others, underscores the effeteness of constructing composite query via duplication.

Choice of Attention Values. By altering the input value feature for both self-attention V^s and cross-attention V^c , we attempt to gain insight into the model’s learning process for timespan predictions. As discussed in Sec. 3.2, setting V^s to f^{\dagger} and V^c to f^V , as shown in the first row of Tab. 4, suggests that the decoder operates within the feature space, producing an output that is a weighted combination of f^{\dagger} and f^V . Similarly, for the second line where $V^s = p^{\dagger}$ and $V^c = p^V$, the decoder functions in the positional encoding space, leading to an output that integrates only p^{\dagger} and p^V . Because the timespan predictions are semantically more re-

Table 5. Ablation results for component designs.

Method	R@1, IoU=			mIoU
	0.3	0.5	0.7	
Ours	37.79	22.74	9.140	23.86
w/o Duplication	34.65	20.00	6.655	21.95
w/o Diagram PE	33.99	19.42	8.226	21.63
w/o Contrastive Alignment	33.82	18.48	7.112	22.03
w/o Auxiliary Loss	34.93	19.11	7.541	22.00
w/o Average Video Length	36.08	21.19	8.283	22.92

lated to the positional encoding space, we hypothesize that the model can learn temporal information explicitly from those positional encoding. The fact that the first line outperforms random guessing in Tab. 1 indicates the model’s ability to also indirectly learn from the feature space. The comparison between the first two lines depicts the space where learning is more effective, with positional encoding space being the clear winner. The lines 2 and 4 show that simply adding the two inputs do not improve the performance. A more effective approach—as demonstrated in the sixth row—involves first concatenating then projecting, granting the model greater flexibility in merging these two spaces rather than relying on their sum, for information extraction. Lines 3 and 5 suggest that removing diagram features from the value feature of self-attention do not further improves the performance.

Effect of Component Design. We also ablate on the design of components in our model in Tab. 5. First, we disable the duplication mechanism in constructing composite queries by setting the number of learnable positional embeddings to one, resulting in a significant performance drop. Second, we found that incorporating order information into the inference model improves performance. This is achieved through diagram PE p^{\dagger} , which ideally should highlight the diagonal in the similarity matrix $p^{\dagger}p^V$, indicating the alignment between the two sequences. Third, contrastively aligning features prior to the decoder is vital. This necessity arises because video and image encoders we used are trained separately, leading to their feature representations laying in different spaces. Furthermore, our evaluation confirms the benefit of applying the same loss on the output of each layer (the auxiliary loss) in enhancing the overall performance. Last, due the limitation of 2D-TAN, it averages the length of video clip to 256 by interpolation. We find that doing this also helps improve performance, but our model can also handle varying duration using padding thanks to the nature of the attention mechanism.

4.4. Analysis, Qualitative Results and Limitations

Timespan Analysis. We investigated two factors that influence performance by analyzing the top-1 timespan predictions along with their ground truth, as shown in Fig. 6. The left diagram demonstrates the relationship between the average overlapping time among timespans for a sequence




Figure 6. Two factors that affect the performance on IAW test.

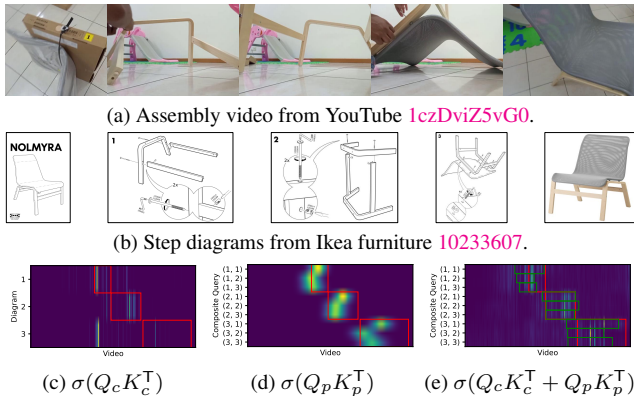


Figure 7. Visualization of content and position jointly guided cross-attention for an example with three diagrams and three learnable queries ($M = 3$, $K = 3$). The top two rows ((a), (b)) demonstrate the video and the corresponding step diagrams. The third row ((c), (d), (e)) visualizes each term and the outcome of Eq. (2), with ground truth timespans marked in red and predictions of each query highlighted in green. Composite queries are denoted as (diagram, learnable query). We use σ to denote the softmax operation applied with scaling.

of queries and the performance (R1@0.5). The right diagram shows the Kendall’s Tau coefficient between the middle points of predicted and ground truth timespans, where a higher score indicates a stronger positive correlation. It can be observed that superior models tend to have shorter overlapping times and higher correlation coefficients. By transitioning from single-query (One) to multi-query (All) using our approach, we consistently observe that as overlapping time decreases and the correlation coefficient increases, performance improves for both Moment-DETR and EaTR.

Visualization of Cross Attention. To elucidate the operation of our joint content-and-position-guided cross-attention mechanism, we qualitatively analyze each term in Eq. (2). As Fig. 7c reveals, videos sourced from YouTube can frequently introduce noise, such as the distinct outlier slightly to the left of the center in the timeline depicted in the third row. Nonetheless, the positional guidance, illustrated in Fig. 7d, effectively self-corrects this misalignment, ensuring that the final predictions are not in the outlier location. Additionally, different queries tend to look at different locations. For example, the second and third queries focus on earlier positions compared to the first query. This tendency enables the model to navigate the temporal dimension

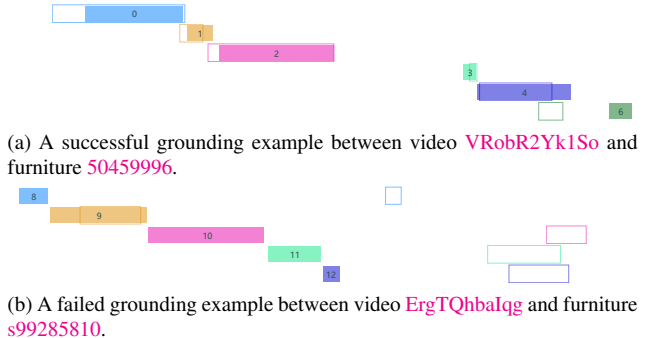


Figure 8. Qualitative results. The horizontal axis represents the timeline of the video. Each row corresponds to a step diagram where the solid rectangle denotes the ground truth and top-1 timespan predictions are represented by bounding boxes with the same color. Notably, the video in (b) starts from step 8.

effectively. Finally, the interplay between these dynamics culminates in jointly guided cross-attention, as illustrated in Fig. 7e, which significantly influences the final timespan predictions, showcasing the integrated effect of semantic and positional cues.

Visualization of Examples and Limitations. As shown in Fig. 8a our model accurately aligns the predicted timespans with the ground truth except for the last diagram. In contrast, Fig. 8b illustrates a missed grounding result in which only one out of five predictions finds the correct location. This could be attributed to the diagram PE we introduce to the model. Notably, the video starts from step 8, leading to the majority of our predicted timespans occurring in the latter half of the video.

5. Conclusion

In this paper, we present a novel approach for the temporal grounding of instructional diagrams within unconstrained videos. Our method specifically addresses the challenges posed by the inherent sequential structure of instructional content, such as assembly manuals. We introduced composite queries by pairing diagram features with learnable positional queries, allowing our model to leverage both semantic and positional information effectively via content and position joint guided cross-attention. This design not only improves the accuracy of temporal grounding but also provides insights into the attention mechanisms for temporal grounding task.

There are several avenues for future research. For instance, further exploration into the handling of learnable queries and the integration of additional modalities (e.g., audio cues from videos) could enhance the model’s robustness and accuracy. Additionally, extending our approach to other domains where sequential events are prevalent, such as medical procedure videos or educational content, presents an exciting opportunity for broader impact.

References

- [1] 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*, pages 5803–5812, 2017. [2](#)
- [2] Peijun Bao, Qian Zheng, and Yadong Mu. Dense events grounding in video. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 920–928, 2021. [3](#)
- [3] Yizhak Ben-Shabat, Xin Yu, Fatemeh Saleh, Dylan Campbell, Cristian Rodriguez-Opazo, Hongdong Li, and Stephen Gould. The Ikea ASM dataset: Understanding people assembling furniture through actions, objects and pose. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 847–859, 2021. [1](#)
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. [2](#), [3](#), [4](#), [5](#)
- [5] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. [6](#)
- [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. [2](#)
- [7] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision*, pages 1–23, 2022. [1](#)
- [8] 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. [6](#)
- [9] 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*, pages 5267–5275, 2017. [2](#)
- [10] Soham Ghosh, Anuva Agarwal, Zarana Parekh, and Alexander Hauptmann. ExCL: Extractive Clip Localization Using Natural Language Descriptions. In Jill Burstein, Christy Doran, and Tamar Solorio, editors, *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1984–1990, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. [2](#), [6](#)
- [11] Jinhyun Jang, Jungin Park, Jin Kim, Hyeongjun Kwon, and Kwanghoon Sohn. Knowing where to focus: Event-aware transformer for video grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13846–13856, 2023. [1](#), [2](#), [3](#), [6](#)
- [12] Xun Jiang, Xing Xu, Jingran Zhang, Fumin Shen, Zuo Cao, and Heng Tao Shen. Semi-supervised video paragraph grounding with contrastive encoder. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2466–2475, 2022. [3](#)
- [13] Wang Jing, Aixin Sun, Hao Zhang, and Xiaoli Li. MS-DETR: Natural language video localization with sampling moment-moment interaction. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1387–1400, Toronto, Canada, July 2023. Association for Computational Linguistics. [2](#)
- [14] Erica Kido Shimomoto, Edison Marrese-Taylor, Hiroya Takamura, Ichiro Kobayashi, Hideki Nakayama, and Yusuke Miyao. Towards parameter-efficient integration of pre-trained language models in temporal video grounding. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13101–13123, Toronto, Canada, July 2023. Association for Computational Linguistics. [2](#), [6](#)
- [15] Jihwan Kim, Miso Lee, and Jae-Pil Heo. Self-feedback DETR for temporal action detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10286–10296, 2023. [4](#)
- [16] 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)*, Oct 2017. [3](#)
- [17] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. [5](#)
- [18] Jie Lei, Tamara L Berg, and Mohit Bansal. Detecting moments and highlights in videos via natural language queries. *Advances in Neural Information Processing Systems*, 34:11846–11858, 2021. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [19] Kevin Qinghong Lin, Pengchuan Zhang, Joya Chen, Shraman Pramanick, Difei Gao, Alex Jinpeng Wang, Rui Yan, and Mike Zheng Shou. Univtg: Towards unified video-language temporal grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2794–2804, 2023. [3](#)
- [20] Xiaolong Liu, Qimeng Wang, Yao Hu, Xu Tang, Shiwei Zhang, Song Bai, and Xiang Bai. End-to-end temporal action detection with transformer. *IEEE Transactions on Image Processing*, 31:5427–5441, 2022. [4](#)
- [21] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [6](#)
- [22] Depu Meng, Xiaokang Chen, Zejia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang. Conditional DETR for fast training convergence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3651–3660, 2021. [2](#), [4](#)
- [23] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *Proceedings of the*

- IEEE/CVF international conference on computer vision*, pages 2630–2640, 2019. [1](#)
- [24] WonJun Moon, Sangeek Hyun, SuBeen Lee, and Jae-Pil Heo. Correlation-guided query-dependency calibration in video representation learning for temporal grounding. *arXiv preprint arXiv:2311.08835*, 2023. [3](#)
- [25] WonJun Moon, Sangeek Hyun, SangUk Park, Dongchan Park, and Jae-Pil Heo. Query-dependent video representation for moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23023–23033, 2023. [2](#)
- [26] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. [6](#)
- [27] 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*, 1:25–36, 2013. [1](#), [3](#)
- [28] Hamid Rezaatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019. [5](#)
- [29] Cristian Rodriguez, Edison Marrese-Taylor, Fatemeh Sadat Saleh, Hongdong Li, and Stephen Gould. Proposal-free temporal moment localization of a natural-language query in video using guided attention. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2464–2473, 2020. [2](#), [6](#)
- [30] Cristian Rodriguez-Opazo, Edison Marrese-Taylor, Basura Fernando, Hongdong Li, and Stephen Gould. Dori: Discovering object relationships for moment localization of a natural language query in a video. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1079–1088, 2021. [2](#), [6](#), [7](#)
- [31] Cristian Rodriguez-Opazo, Edison Marrese-Taylor, Basura Fernando, Hiroya Takamura, and Qi Wu. Memory-efficient temporal moment localization in long videos. In Andreas Vlachos and Isabelle Augenstein, editors, *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1909–1924, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. [2](#), [6](#), [7](#)
- [32] Cristian Rodriguez-Opazo, Edison Marrese-Taylor, Basura Fernando, Hiroya Takamura, and Qi Wu. Memory-efficient temporal moment localization in long videos. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, May 2023. [2](#)
- [33] Marcus Rohrbach, Michaela Regneri, Mykhaylo Andriluka, Sikandar Amin, Manfred Pinkal, and Bernt Schiele. Script data for attribute-based recognition of composite activities. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7-13, 2012, Proceedings, Part I 12*, pages 144–157. Springer, 2012. [1](#)
- [34] Fengyuan Shi, Limin Wang, and Weilin Huang. End-to-end dense video grounding via parallel regression. *arXiv preprint arXiv:2109.11265*, 2021. [3](#)
- [35] Xiaolong Sun, Liushuai Shi, Le Wang, Sanping Zhou, Kun Xia, Yabing Wang, and Gang Hua. Diversifying query: Region-guided transformer for temporal sentence grounding. *arXiv preprint arXiv:2406.00143*, 2024. [3](#)
- [36] Chaolei Tan, Zihang Lin, Jian-Fang Hu, Wei-Shi Zheng, and Jianhuang Lai. Hierarchical semantic correspondence networks for video paragraph grounding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18973–18982, 2023. [3](#)
- [37] Yansong Tang, Dajun Ding, Yongming Rao, Yu Zheng, Danyang Zhang, Lili Zhao, Jiwen Lu, and Jie Zhou. Coin: A large-scale dataset for comprehensive instructional video analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1207–1216, 2019. [1](#)
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [4](#)
- [39] Limin Wang, Bingkun Huang, Zhiyu Zhao, Zhan Tong, Yinan He, Yi Wang, Yali Wang, and Yu Qiao. Videomae v2: Scaling video masked autoencoders with dual masking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14549–14560, 2023. [6](#)
- [40] Weiyang Wang, Yongcheng Wang, Shizhe Chen, and Qin Jin. Youmakeup: A large-scale domain-specific multimodal dataset for fine-grained semantic comprehension. 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)*, pages 5133–5143, 2019. [1](#)
- [41] Sangmin Woo, Jinyoung Park, Inyong Koo, Sumin Lee, Minki Jeong, and Changick Kim. Explore-and-match: Bridging proposal-based and proposal-free with transformer for sentence grounding in videos. *arXiv preprint arXiv:2201.10168*, 2022. [2](#), [3](#), [4](#), [6](#)
- [42] Yicheng Xiao, Zhuoyan Luo, Yong Liu, Yue Ma, Hengwei Bian, Yatai Ji, Yujiu Yang, and Xiu Li. Bridging the gap: A unified video comprehension framework for moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18709–18719, 2024. [3](#)
- [43] Jin Yang, Ping Wei, Huan Li, and Ziyang Ren. Task-driven exploration: Decoupling and inter-task feedback for joint moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18308–18318, 2024. [3](#)
- [44] Yitian Yuan, Lin Ma, Jingwen Wang, Wei Liu, and Wenwu Zhu. Semantic conditioned dynamic modulation for temporal sentence grounding in videos. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [45] Yitian Yuan, Tao Mei, and Wenwu Zhu. To find where you talk: Temporal sentence localization in video with attention

- based location regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9159–9166, 2019. [2](#)
- [46] Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, and Larry S Davis. 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*, pages 1247–1257, 2019. [2](#)
- [47] Frederic Z Zhang, Dylan Campbell, and Stephen Gould. Efficient two-stage detection of human-object interactions with a novel unary-pairwise transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20104–20112, 2022. [5](#)
- [48] Frederic Z Zhang, Yuhui Yuan, Dylan Campbell, Zhuoyao Zhong, and Stephen Gould. Exploring predicate visual context in detecting of human-object interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10411–10421, 2023. [4](#), [6](#)
- [49] Hao Zhang, Aixin Sun, Wei Jing, Liangli Zhen, Joey Tianyi Zhou, and Siow Mong Rick Goh. Parallel attention network with sequence matching for video grounding. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 776–790, Online, Aug. 2021. Association for Computational Linguistics. [2](#)
- [50] Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. Span-based localizing network for natural language video localization. In Dan Jurafsky, Joyce Chai, Natalie Schluter, and Joel Tetreault, editors, *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6543–6554, Online, July 2020. Association for Computational Linguistics. [2](#)
- [51] Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. Temporal sentence grounding in videos: A survey and future directions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. [2](#)
- [52] Jiahao Zhang, Anoop Cherian, Yanbin Liu, Yizhak Ben-Shabat, Cristian Rodriguez, and Stephen Gould. Aligning step-by-step instructional diagrams to video demonstrations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2483–2492, 2023. [1](#), [2](#), [3](#), [6](#)
- [53] Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. Learning 2d temporal adjacent networks for moment localization with natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12870–12877, 2020. [2](#), [6](#)
- [54] Luwei Zhou, Chenliang Xu, and Jason Corso. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018. [1](#), [6](#)