

Language as Queries for Referring Video Object Segmentation

Giannan Wu¹, Yi Jiang², Peize Sun¹, Zehuan Yuan², Ping Luo^{1,3}

¹The University of Hong Kong ²ByteDance

³HKU-TCL Joint Research Centre for Artificial Intelligence

Abstract

Referring video object segmentation (R-VOS) is an emerging cross-modal task that aims to segment the target object referred by a language expression in all video frames. In this work, we propose a simple and unified framework built upon Transformer, termed ReferFormer. It views the language as queries and directly attends to the most relevant regions in the video frames. Concretely, we introduce a small set of object queries conditioned on the language as the input to the Transformer. In this manner, all the queries are obligated to find the referred objects only. They are eventually transformed into dynamic kernels which capture the crucial object-level information, and play the role of convolution filters to generate the segmentation masks from feature maps. The object tracking is achieved naturally by linking the corresponding queries across frames. This mechanism greatly simplifies the pipeline and the end-to-end framework is significantly different from the previous methods. Extensive experiments on Ref-Youtube-VOS, Ref-DAVIS17, A2D-Sentences and JHMDB-Sentences show the effectiveness of ReferFormer. On Ref-Youtube-VOS, ReferFormer achieves 55.6 $\mathcal{J}\&\mathcal{F}$ with a ResNet-50 backbone without bells and whistles, which exceeds the previous state-of-the-art performance by 8.4 points. In addition, with the strong Video-Swin-Base backbone, ReferFormer achieves the best $\mathcal{J}\&\mathcal{F}$ of 64.9 among all existing methods. Moreover, we show the impressive results of 55.0 mAP and 43.7 mAP on A2D-Sentences and JHMDB-Sentences respectively, which significantly outperforms the previous methods by a large margin. Code is publicly available at <https://github.com/wjn922/ReferFormer>.

1. Introduction

Referring video object segmentation (R-VOS) aims to segment the target object in a video given a natural language description. This emerging topic has raised great attention in the research community and is expected to benefit many applications in a friendly and interactive way, *e.g.*, video

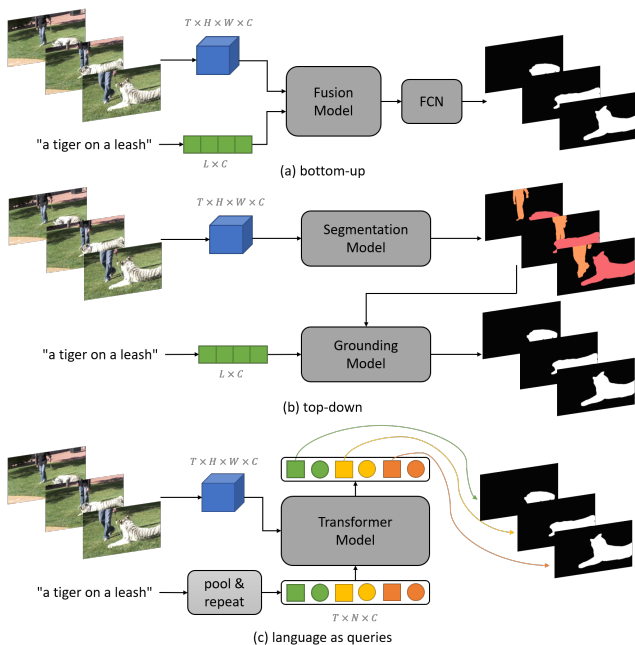


Figure 1. Comparison of current referring video object segmentation (R-VOS) pipelines. (a) Bottom-up. (b) Top-down. (c) Ours.

editing and video surveillance. R-VOS is more challenging than the traditional semi-supervised video object segmentation [33, 46], because it does not only lack the ground-truth mask annotation in the first frame, but also require the comprehensive understanding of the cross-modal sources, *i.e.*, vision and language. Therefore, the model should have a strong ability to infer which object is referred and to perform accurate segmentation.

To accomplish this task, the existing methods can be mainly categorized into two groups: (1) *Bottom-up* methods. These methods incorporate the vision and language features in a early-fusion manner, and then adopt a FCN [28] as decoder to generate object masks, as shown in Figure 1(a). (2) *Top-down* methods. These methods tackle the problem in a top-down perspective and follow a two-stage pipeline. As illustrated in Figure 1(b), they first employ an instance segmentation model to find all the objects in each

frame, and then associate them in the entire video to form the tracklet candidates. Afterwards, they use the expression as the grounding criterion to select the best-matched one.

Although these two streams of methods have demonstrated their effectiveness with promising results, they still have some intrinsic limitations. First, for the *bottom-up* methods, they fail to capture the crucial instance-level information and do not consider the object association across multiple frames. Therefore, these methods can not provide explicit knowledge for cross-modal reasoning and would encounter the discrepancy of predicted object due to scene changes. Second, although *top-down* methods have greatly boost the performance over the *bottom-up* methods, they suffer from heavy workload because of the complex, multi-stage pipeline. Furthermore, the separate optimization on several sub-problems would lead to sub-optimal solution.

These limitations of current methods motivate us to design a simple and unified framework that solves the R-VOS task elegantly. The recent success of Transformer [40] in object detection [4, 55] and video instance segmentation [15, 43, 44] demonstrates a promising solution. However, it is non-trivial to apply such models to the R-VOS task. These models [4, 55] use a fixed number (*e.g.*, 100) of learnable queries to detect all the objects in an image. Under this circumstance, it would be confused for the model to distinguish which object is referred due to the randomness of the expression. Here raises a natural question: “*Is it possible for a unified model to know where to look using queries?*”

This work answers the question by proposing the notion of *language as queries*, as shown in Figure 1(c). We put the linguistic restriction on all object queries and use these *conditional queries* as input for the model. In this manner, the expression will make the queries focus on the referred object only, and thus greatly reducing the query number (*e.g.*, 5 in our experiments). The next challenge lies in how to decode the object mask from query representations. As the queries contain rich instance characteristics, we view them as instance-aware dynamic kernels to filter out the segmentation masks from feature maps.

The unified framework can not only produce the segmentation masks for referred objects, but also the classification results and detection boxes. Moreover, the *conditional queries* are linked via instance matching strategy across frames so that the object tracking is achieved naturally without post-process. We hope this framework could serve as a strong baseline for R-VOS task.

The main contributions of this work are as follows.

- We propose a simple and unified framework for referring video object segmentation, termed ReferFormer. Given a video clip and the corresponding language expression, our framework directly detects, segments and tracks the referred object in all frames in an end-to-end manner.
- We present the notion of *language as queries*. We in-

troduce a small set of object queries which conditioned on the text expression to attend the referred object only. These *conditional queries* are shared across different frames in the initial state and they are transformed into dynamic kernels to filter out the segmentation masks from feature maps. This mechanism provides a new perspective for the R-VOS task.

- We design the cross-modal feature pyramid network (CM-FPN) for multi-scale vision-language fusion, which improves the discriminativeness of mask features for accurate segmentation.
- Extensive experiments on Ref-Youtube-VOS, Ref-DAVIS17, A2D-Sentences and JHMDB-Sentences show that ReferFormer outperforms the previous methods on these four benchmarks by a large margin. *E.g.*, on Ref-Youtube-VOS, ReferFormer with a ResNet-50 backbone achieves 55.6 $\mathcal{J}\&\mathcal{F}$ without bells and whistles, showing the significant 8.4 points gain over the previous state-of-the-art methods. And using the strong Video-Swin-Base visual backbone, ReferFormer achieves the impressive results of 64.9 $\mathcal{J}\&\mathcal{F}$.

2. Related Work

Video Object Segmentation. The traditional video object segmentation (VOS) aims to propagate the ground-truth object masks given in the first frame to the entire video. Most recent works [7, 32, 41, 49] lie in the group of matching-based methods, which perform feature matching to track the target objects.

Referring Video Object Segmentation. Referring video object segmentation (R-VOS) provides the language description instead of mask annotation as the object reference, thus it would be a more challenging task. The current methods for R-VOS mainly follow the two pipelines:

(1) *Bottom-up* methods. An intuitive thinking is directly applying the image-level referring methods [8, 14, 29, 53] on the video frames independently, *e.g.*, RefVOS [2]. The obvious drawback of such methods is that they fail to utilize the valuable temporal information across frames, resulting in inconsistent object prediction due to the scene or appearance variations. To address this issue, URVOS [36] casts the task as a joint problem of referring object segmentation in an image and mask propagation in a video. They propose a unified referring VOS framework that employs a memory attention module to leverage the information of mask predictions in previous frames.

(2) *Top-down* methods. The typical top-down method [20] first constructs an exhaustive set of object tracklets by propagating the object masks detected from several key frames to the whole video. Then, a language grounding model is built to select the best object tracklet from the candidate set. Although the method has made breakthrough

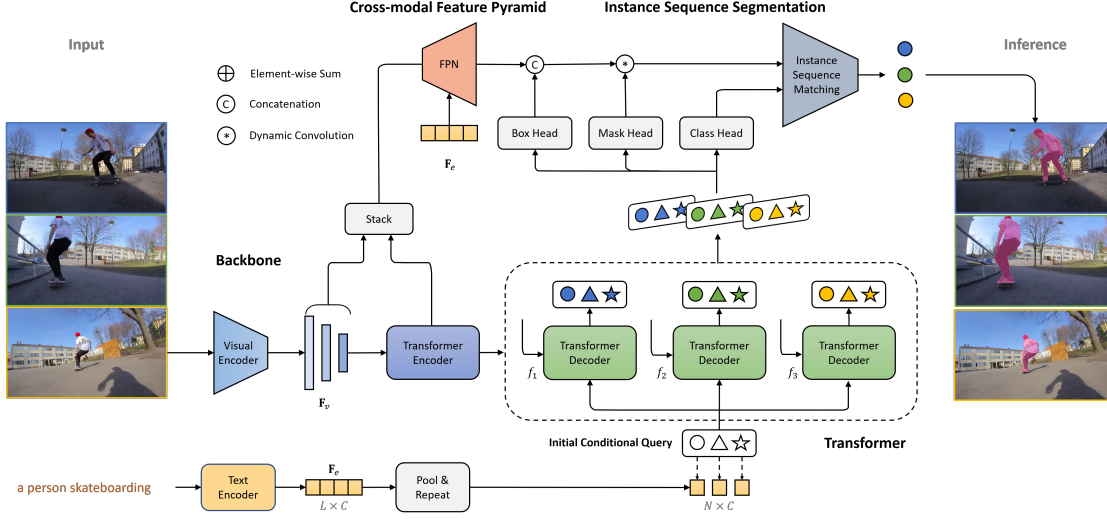


Figure 2. The overall pipeline of ReferFormer. It mainly consists of four parts: Backbone, Transformer, Cross-modal Feature Pyramid and the Segmentation part. The model takes a video clip with the corresponding language expression as input and output the segmentation mask of the referred object in each frame. For the Transformer decoder input, the object queries are conditioned on the language expression to find the referred object. The same colors represent the queries in the same frame and the same shapes represent the queries refer to the same instance. The order of queries inner frame keep consistent for different frames. Best viewed in color.

performance improvement over the previous methods, the complex, multi-stage pipeline is computational-expensive and impractical. In contrast to these two pipelines, we propose a *query-based* method that achieves the strongest performance with a simple and unified framework.

Transformer Transformer [40] was first introduced for sequence-to-sequence translation in natural language processing community and has achieved marvelous success in most computer vision tasks [10, 17, 26] such as object detection [4, 55], tracking [5, 30, 37, 47] and segmentation [6, 15, 54]. DETR [4] introduces the new *query-based* paradigm [38, 55] for object detection, which employs a set of object queries as candidates and inputs them to the Transformer decoder. Beyond image field, VisTR [43] extends the framework for video instance segmentation (VIS) [48] task and solves the problem in a direct end-to-end parallel sequence decoding manner. SeqFormer [44] decouples the content query and box query to aggregates temporal information from each frame and achieves the state-of-the-art performance on VIS task. Inspired by these works, our work also relies on the *query* mechanism of Transformer but considers an additional modality, *i.e.*, language, as the object reference. Thus, we propose the notion of *language as queries* and build the simple and unified framework that segments and tracks the referred object simultaneously.

3. Approach

Given a video clip $\mathcal{I} = \{I_t\}_{t=1}^T$ with T frames and a referring expression $\mathcal{E} = \{e_l\}_{l=1}^L$ with L words, we aim to produce T -frame binary segmentation masks of referred

object $\mathcal{S} = \{s_t\}_{t=1}^T, s_t \in \mathbb{R}^{H \times W}$ in an end-to-end manner. To this end, we propose a simple and unified framework named ReferFormer, as shown in Figure 2. It mainly consists of four key components: Backbone, Transformer, Cross-modal Feature Pyramid network (CM-FPN) and the Instance Sequence Segmentation process. A small set of object queries conditioned on the language is introduced to find the referred object. During inference, we directly output the mask predictions by selecting the queries with the highest average score as the final results.

3.1. Backbone

Visual Encoder. We start by adopting a visual backbone to extract the multi-scale feature maps for each frame in the video clip independently, resulting in the visual feature sequence $\mathcal{F}_v = \{f_t\}_{t=1}^T$. It is noteworthy that both the 2D and 3D visual encoder could be adopted.

Linguistic Encoder. Given the language description with L words, we use off-the-shelf linguistic embedding model, RoBERTa [25], to extract the text feature $\mathcal{F}_e = \{f_i\}_{i=1}^L$. And we also obtain the sentence-level feature $f_e^s \in \mathbb{R}^C$ by pooling the features of each word.

3.2. Language as Queries

The key design comes from that we use a set of object queries conditioned on the language expression, termed *conditional queries*, as the Transformer decoder input. These queries are obligated to focus on the referred object only and produce the instance-aware dynamic kernels. The final segmentation masks are obtained by per-

forming dynamic convolution between the dynamic kernels and their corresponding feature maps. Here, we adopt the Deformable-DETR [55] as our Transformer model due to its effectiveness and efficiency for global attention.

Transformer Encoder. First, a 1×1 convolution is applied on the multi-scale visual features \mathcal{F}_v to reduce the channel dimension of all feature maps to $C = 256$. To enrich the information of visual features, we then incorporate projected visual features with the text feature \mathcal{F}_e in a multiplication way and form the new multi-scale feature maps $\mathcal{F}'_v = \{f'_t\}_{t=1}^T$. Afterwards, the fixed 2D positional encoding is added to feature maps of each frame and the summed features are fed into the Transformer encoder. To utilize the Transformer process the video frames independently, we flatten the spatial dimensions and move the temporal dimension to batch dimension for efficiency. Finally, the output of the Transformer encoder, *i.e.*, *encoded memory*, is then input to the decoder.

Transformer Decoder. We introduce N object queries to represent the instances for each frame similar to [43], the difference lies in that the query weights are shared across video frames to handle the length-variable videos. Meanwhile, we repeat the sentence feature f_e^s for N times to fit the query number. Both the object queries and repeated sentence features are fed into the decoder as input. In this manner, all the queries will use the language expression as guidance and try to find the referred objects only. These *conditional queries* are duplicated to serve as the decoder input for all the frames and they are turned into instance embeddings by the decoder eventually, resulting in the set of $N_q = T \times N$ predictions. It should be noted the queries keep the same order across different frames and we refer to the queries in the same relative position (represented as the same shape in Figure 2) as *instance sequence*. Therefore, the temporal coherence of referred object could be achieved easily by linking the corresponding queries.

Prediction Heads. Three lightweight heads are built on top of the decoder to further transform the N_q instance embeddings. The class head outputs the binary probability which indicates whether the instance is referred by the text sentence and this instance is visible in the current frame. It could also be modified to predict the referred object category by simply changing the output class number. The mask head is implemented by three consecutive linear layers. It produces the parameters of N_q dynamic kernels $\Omega = \{\omega_i\}_{i=1}^{N_q}$, which is similar to the conditional convolutional filters in [39]. These parameters will be reshaped to form the three 1×1 convolution layers with the channel number as 8. The box head is a 3-layer feed forward network (FFN) with ReLU activation except for the last layer. It will predict the box location of the referred object and



Figure 3. We visualize the predicted boxes from all the queries. It can be seen that these boxes will locate near the referred object only even if there are other objects in the video.

thus the position of dynamic kernels could be determined by the center of corresponding boxes.

Dynamic Convolution. Suppose now we have obtained the semantically-rich feature maps $\mathcal{F}_{seg} = \{f_{seg}^t\}_{t=1}^T$ (will be discussed in Sec. 3.3) for each frame, the question is how we perform the instance sequence segmentation and obtain the masks of referred object from them. Since the dynamic kernels have captured the object-level information, we use them as convolution filters on the feature maps for instance decoding. Considering that the location prior of dynamic kernels Ω provides a strong and robust reference for the referred object, we concatenate the feature maps \mathcal{F}_{seg} with relative coordinates for each dynamic kernel. Finally, the binary segmentation masks are generated by performing dynamic convolution between the conditional convolutional weights and their corresponding feature maps:

$$\hat{s}_i = \left\{ \hat{f}_i \otimes \omega_i \right\}_{i=1}^{N_q} \quad (1)$$

where ω_i and \hat{f}_i are the i -th dynamic kernel weights and its exclusive feature map, respectively. We reshape the output masks in frame-order sequence, resulting in a set as $\hat{S} \in \mathbb{R}^{T \times N \times \frac{H}{4} \times \frac{W}{4}}$.

Illustration of conditional queries. It is well known that the decoder embedding and position embedding in Transformer decoder encode the content and spatial information respectively. In our framework, these two parts are fed with the text sentence feature and learnable queries parameters, so that all the queries are restricted by the language expression. As shown in figure 3, these queries will focus on the referred object only even if other objects exist in the video. And there will be one query with much higher score while the scores of other queries will be suppressed.

3.3. Cross-modal Feature Pyramid Network

Feature pyramid network (FPN) [22] is adopted to produce multi-scale feature maps for video frames. We construct a 4-level pyramid with the spatial stride from $4 \times$ to $32 \times$. Specifically, the Transformer memory (spatial strides $\{8, 16, 32\}$) and the $4 \times$ feature from visual backbone are stacked to form the hierarchical features. And we design a cross-modal feature pyramid network (CM-FPN) to perform multi-scale cross-modal fusion for finer interaction, see the architecture figure in supplementary materials.

In each level, the interaction process is achieved by the vision-language fusion module. And we take the l -th level feature of FPN as an example to clarify the process. Here, we use $f_v^l \in \mathbb{R}^{T \times H_l \times W_l \times C}$ to represent the l -th level visual feature for simplicity. Before the MHSA module, the spatial size of vision feature F_v^l is downsampled by a factor of σ while the temporal dimension is kept unchanged. Thus, the complexity of self-attention [40] operation would be greatly reduced, making the fusion module can be inserted into each level of FPN. Then, the spatial size of vision feature is recovered to $H_l \times W_l$ for maintaining fine-grained information. We set the downsample factors as [8, 4, 2, 1] for the 4-level features maps. Next, f_v^l interact with word-level feature f_e in a *cross-attention* way, where the *query*, *key* are vision and language feature, respectively:

$$\text{Interact}(f_v^l, f_e) = \text{Softmax}\left(\frac{f_v^l W^Q \cdot (f_e W^K)^T}{\sqrt{d_{\text{head}}}}\right) f_e W^V \quad (2)$$

where $W^Q, W^K, W^V \in \mathbb{R}^{C \times d_{\text{head}}}$ are learnable parameters. Finally, we apply an additional 3×3 convolutional layer on the feature maps with spatial stride 4 to get the final feature maps $\mathcal{F}_{\text{seg}} = \{f_{\text{seg}}^t\}_{t=1}^T$, where $f_{\text{seg}}^t \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C_d}$.

3.4. Instance Sequence Matching and Loss

Using N *conditional queries*, we generate the set of $N_q = T \times N$ predictions, which can be regarded as the trajectories of N instances on T frames. As described previous, the predictions across frames maintain the same relative positions. Therefore, we can supervise the instance sequence as a whole using instance matching strategy [43]. Let us denote the prediction set as $\hat{y} = \{\hat{y}_i\}_{i=1}^N$, and the predictions for the i -th instance is represented by:

$$\hat{y}_i = \{\hat{p}_i^t, \hat{b}_i^t, \hat{s}_i^t\}_{t=1}^T \quad (3)$$

For the t -th frame, $\hat{p}_i^t \in \mathbb{R}^1$ is a probability scalar indicating whether the instance corresponds to the referred object and this object is visible in the current frame. $\hat{b}_i^t \in \mathbb{R}^4$ is the normalized vector defines the center coordinates as well as the height and width of predicted box. $\hat{s}_i^t \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4}}$ is the predicted binary segmentation mask.

Since there is only one referred object in the video, the ground-truth instance sequence is represented as $y = \{c^t, b^t, s^t\}_{t=1}^T$. c^t is an one-hot value and it equals 1 when the ground-truth instance is visible in the frame I_t otherwise 0. To train the network, we first find the best prediction as the positive sample via minimizing the matching cost:

$$\hat{y}_{\text{pos}} = \arg \min_{\hat{y}_i \in \hat{y}} \mathcal{L}_{\text{match}}(y, \hat{y}_i) \quad (4)$$

$$\mathcal{L}_{\text{match}}(y, \hat{y}_i) = \lambda_{\text{cls}} \mathcal{L}_{\text{cls}}(y, \hat{y}_i) + \lambda_{\text{box}} \mathcal{L}_{\text{box}}(y, \hat{y}_i) + \lambda_{\text{mask}} \mathcal{L}_{\text{mask}}(y, \hat{y}_i) \quad (5)$$

The matching cost is computed from each frame and normalized by the frame number. Here, $\mathcal{L}_{\text{cls}}(y, \hat{y}_i)$ is the focal loss [23] that supervises the predicted instance sequence reference results. The box-related loss sums up the L1 loss and GIoU loss [35]. And the mask-related loss is the combination of DICE loss [31] and binary mask focal loss. Both the two mask losses are spatio-temporally calculated over the entire video clip. The network is optimized by minimizing the total loss $\mathcal{L}_{\text{match}}$ for positive samples while letting the negative samples predict the \emptyset class.

3.5. Inference

As mentioned previously, ReferFormer can handle the videos of arbitrary length in a single forward pass since all the frames share the same initial *conditional queries*. Given the video and language expression, ReferFormer will predict N instance sequences. For each instance sequence, we average the predicted reference probabilities over all the frames and obtain the reference score set $\mathcal{P} = \{p_i\}_{i=1}^N$. We select the instance sequence with the highest average score and its index is denoted as σ :

$$\sigma = \arg \max_{i \in \{1, 2, \dots, N\}} p_i \quad (6)$$

The final segmentation masks for each frame $\mathcal{S} = \{s_t\}_{t=1}^T$ is obtained from the mask candidates set $\hat{\mathcal{S}}$ by selecting the corresponding queries indexed with σ . No post-process is needed for associating objects since the linked queries naturally track the same instance.

4. Experiments

4.1. Datasets and Metrics

Datasets. The experiments are conducted on the four popular R-VOS benchmarks. Ref-Youtube-VOS [36] is a large-scale benchmark which covers 3,978 videos with $\sim 15\text{K}$ language descriptions. Ref-DAVIS17 [18] is built upon DAVIS17 [34] by providing the language description for a specific object in each video and contains 90 videos. A2D-Sentences and JHMDB-Sentences are created by providing the additional textual annotations on the original A2D [45] and JHMDB [16] datasets. A2D-Sentences contains 3,782 videos and each video has 3-5 frames annotated with the pixel-level segmentation masks. JHMDB-Sentences has 928 videos with the 928 corresponding sentences in total.

Evaluation Metrics. We use the standard evaluation metrics for Ref-Youtube-VOS and Ref-DAVIS17: region similarity (\mathcal{J}), contour accuracy (\mathcal{F}) and their average value ($\mathcal{J\&F}$). For Ref-Youtube-VOS, as the annotations of validation set are not released publicly, we evaluate our method on the official challenge server¹. Ref-DAVIS17 is evaluated by the official evaluation code².

¹<https://competitions.codalab.org/competitions/29139>

²<https://github.com/davisvideochallenge/davis2017-evaluation>

Method	Backbone	Ref-Youtube-VOS			Ref-DAVIS17		
		$\mathcal{J} \& \mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J} \& \mathcal{F}$	\mathcal{J}	\mathcal{F}
Spatial Visual Backbones							
CMSA [50]	ResNet-50	34.9	33.3	36.5	34.7	32.2	37.2
CMSA + RNN [50]	ResNet-50	36.4	34.8	38.1	40.2	36.9	43.5
URVOS [36]	ResNet-50	47.2	45.3	49.2	51.5	47.3	56.0
ReferFormer	ResNet-50	55.6	54.8	56.5	58.5	55.8	61.3
ReferFormer*	ResNet-50	58.7	57.4	60.1	-	-	-
PMINet [9]	ResNeSt-101	48.2	46.7	49.6	-	-	-
PMINet + CFBI [9]	ResNeSt-101	53.0	51.5	54.5	-	-	-
CITD* [20]	ResNet-101	56.4	54.8	58.1	-	-	-
ReferFormer	ResNet-101	57.3	56.1	58.4	-	-	-
ReferFormer*	ResNet-101	59.3	58.1	60.4	-	-	-
PMINet + CFBI [9]	Ensemble	54.2	53.0	55.5	-	-	-
CITD [20]	Ensemble	61.4	60.0	62.7	-	-	-
ReferFormer	Swin-L	62.4	60.8	64.0	60.5	57.6	63.4
ReferFormer*	Swin-L	64.2	62.3	66.2	-	-	-
Spatio-temporal Visual Backbones							
MTTR [†] ($\omega = 12$) [3]	Video-Swin-T	55.3	54.0	56.6	-	-	-
ReferFormer [†] ($\omega = 5$)		56.0	54.8	57.3	-	-	-
ReferFormer	Video-Swin-T	59.4	58.0	60.9	-	-	-
ReferFormer*		62.6	59.9	63.3	-	-	-
ReferFormer	Video-Swin-S	60.1	58.6	61.6	-	-	-
ReferFormer*		63.3	61.4	65.2	-	-	-
ReferFormer	Video-Swin-B	62.9	61.3	64.6	61.1	58.1	64.1
ReferFormer*		64.9	62.8	67.0	-	-	-

Table 1. Comparison with the state-of-the-art methods on Ref-Youtube-VOS and Ref-DAVIS17. * means joint trainig with Ref-COCO dataset. † indicates the spatio-temporal visual backbone is trained from scratch.

On A2D-Sentences and JHMDB-Sentences, the model is evaluated with the criteria of Precision@K, Ovrall IoU, Mean IoU and mAP over 0.50:0.05:0.95. The Precision@K measures the percentage of test samples whole IoU scores are higher than the threshold K. Following standard protocol, the thresholds are set as 0.5:0.1:0.9.

4.2. Implementation Details.

Model Settings. We test our models under different visual backbones including: ResNet [12], Swin Transformer [26] and Video Swin Transformer [27]. The text encoder is selected as RoBERTa [25] and its parameters are frozen during the entire training stage. Following [55], we use the last three stage features from the visual backbone as the input to Transformer, their corresponding spatial strides are {8, 16, 32}. In the Transformer model, we adopt 4 encoder layers and 4 decoder layers and the hidden dimension is $C = 256$. The number of *conditional query* is set as 5.

Training Details. During training, we use sliding-windows to obtain the clips from a video and each clip consist of 5 randomly sampled frames. Following [43], the data aug-

mentation includes random horizontal flip, random resize, random crop and photometric distortion. All frames are downsampled so that the short side has the size of 360 and the maximum size for the long side is 640 to fit GPU memory. The coefficients for losses are set as $\lambda_{cls} = 2$, $\lambda_{L1} = 5$, $\lambda_{giou} = 2$, $\lambda_{dice} = 5$, $\lambda_{focal} = 2$.

Most of our experiments follow the pretrain-then-finetune process. And some models are trained from scratch for fair comparison. Additionally, on Ref-Youtube-VOS, we also reports the results by training the mixed data from Ref-Youtube-VOS and Ref-COCO [52]. The joint training technique has proven the effectiveness in many VIS tasks [1, 21, 44]. Please see more in the supplementary materials.

Inference Details. During inference, the video frames are downscaled to 360p. The model receives the whole video frames and directly outputs the predicted segmentation masks without post-process.

4.3. Main Results

Ref-Youtube-VOS & Ref-DAVIS17 We compare our method with other state-of-the-art methods in Table 1.

Method	Backbone	Precision					IoU		mAP
		P@0.5	P@0.6	P@0.7	P@0.8	P@0.9	Overall	Mean	
Hu <i>et al.</i> [13]	VGG-16	34.8	23.6	13.3	3.3	0.1	47.4	35.0	13.2
Gavrilyuk <i>et al.</i> [11]	I3D	47.5	34.7	21.1	8.0	0.2	53.6	42.1	19.8
CMSA + CFSA [51]	ResNet-101	48.7	43.1	35.8	23.1	5.2	61.8	43.2	-
ACAN [42]	I3D	55.7	45.9	31.9	16.0	2.0	60.1	49.0	27.4
CMPC-V [24]	I3D	65.5	59.2	50.6	34.2	9.8	65.3	57.3	40.4
ClawCraneNet [19]	ResNet-50/101	70.4	67.7	61.7	48.9	17.1	63.1	59.9	-
MTTR ($\omega = 8$) [3]	Video-Swin-T	72.1	68.4	60.7	45.6	16.4	70.2	61.8	44.7
MTTR ($\omega = 10$) [3]	Video-Swin-T	75.4	71.2	63.8	48.5	16.9	72.0	64.0	46.1
ReferFormer [†] ($\omega = 6$)	Video-Swin-T	76.0	72.2	65.4	49.8	17.9	72.3	64.1	48.6
ReferFormer ($\omega = 5$)	Video-Swin-T	82.8	79.2	72.3	55.3	19.3	77.6	69.6	52.8
ReferFormer ($\omega = 5$)	Video-Swin-S	82.6	79.4	73.1	57.4	21.1	77.7	69.8	53.9
ReferFormer ($\omega = 5$)	Video-Swin-B	83.1	80.4	74.1	57.9	21.2	78.6	70.3	55.0

Table 2. Comparison with the state-of-the-art methods on A2D-Sentences. [†] means our model is trained from scratch.

CITD [20] and PMINet [9] are the top-2 solutions in 2021 Ref-Youtube-VOS Challenge. It can be observed that ReferFormer outperforms previous methods on the two datasets under all metrics and with a large margin. On Ref-Youtube-VOS, ReferFormer with a ResNet-50 backbone achieves the overall \mathcal{J} & \mathcal{F} of 55.6, which is 8.4 points higher than the previous state-of-the-art work URVOS [36], and even beats PMINet [9] using the ensemble models and adopting post-process (55.6 vs 54.2). Using the strong Swin-Large [26] backbone, ReferFormer reaches the surprising 64.2 \mathcal{J} & \mathcal{F} , which obviously exceeds the ensemble results of the complicated, multi-stage method CITD [20] and creates a fairly high new record.

Additionally, we also test the Video Swin Transformer [27] as the backbones. It is well known that the spatio-temporal visual encoder has strong ability to capture both the spatial characteristics and the temporal cues. For a fair comparison with MTTR [3], we train our model with the Video-Swin-Tiny backbone from scratch. It can be seen that our method outperforms MTTR under all the metrics with the smaller window size (5 vs 12). Comparing the results of ReferFormer under Video-Swin-Tiny backbone, it proves that the model benefits from the pretraining stage and joint training process to address the overfitting issue.

On Ref-DAVIS17, our method also achieves the best results under the same ResNet-50 setting (58.5 \mathcal{J} & \mathcal{F}). And the performance consistently improves by using stronger backbones, which proves the generality of our method.

A2D-Sentences & JHMDB-Sentences We further evaluate our method on the A2D-Sentences dataset and compare the performance with other state-of-the-art methods in Table 2. It is obvious that our method achieves the impressive improvement over the previous methods. Compared with the recent MTTR [3], our method exhibits the clear performance advantage (+2.5 mAP) with smaller window size (6 vs. 10). Incorporating the pretraining stage, ReferFormer

Components	\mathcal{J}	\mathcal{F}
Baseline	47.2 (-7.6)	50.1 (-7.2)
w/o Visual-language Fusion	53.0 (-1.8)	56.2 (-1.1)
w/o Relative Coordinates	53.7 (-1.1)	55.9 (-1.4)
Full Model	54.8	57.3

Table 3. Ablation study on the components of ReferFormer. The visual backbone is Video-Swin-Tiny.

with Video-Swin-Base visual backbone achieves 55.0 mAP which shows a significant gain of 8.9 mAP over previous best result. And ReferFormer also demonstrates its strong ability to produce high-quality masks via the stringent metrics (*e.g.*, 57.9 for P@0.8 and 21.2 for P@0.9).

We also evaluate the models on JHMDB-Sentences without finetuning to further prove the generality of our method, and the results table is placed in the supplementary materials. ReferFormer significantly outperforms all the existing methods and achieves superior 43.7 mAP using Video-Swin-Base backbone.

4.4. Ablation Study

In this section, we perform extensive ablation studies on Ref-Youtube-VOS to study the effect of core components in our model. All models are based on Video-Swin-Tiny visual backbone and we train the models from scratch otherwise specified. The detailed analysis is as follows.

Component Analysis. We build a simple Transformer bottom-up baseline. Specifically, considering a video clip of T frames, we flatten the temporal and spatial dimension into one dimension and then concatenate the visual features with the textual features along length dimension to form the multi-modal feature map $f_m \in \mathbb{R}^{(T \times H \times W + L) \times C}$. The vanilla Transformer encoder builds the global dependencies between the visual and textual features. And a standard

Queries	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
1	53.6	52.7	54.5
3	54.2	53.2	55.2
5	56.0	54.8	57.3
8	55.3	54.1	56.6

(a) The effect of query number.

Frames	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
1	50.0	48.4	51.6
3	54.8	53.6	56.0
5	56.0	54.8	57.3

(b) The effect of frame number.

Class	Box	Mask	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
✓	✓		55.2	54.0	56.4
✓		✓	54.5	53.5	55.5
✓	✓	✓	56.0	54.8	57.3

(c) The effect of label assignment method.

Table 4. Ablation study on different settings of ReferFormer. All the models are using Video-Swin-Tiny as visual backbone.

FPN plays the role of decoder to generate segmentation masks from encoded visual features. The baseline method operates the fixed-length video of 5 frames during the training and inference phases. We report the performance of the baseline method and also study core components in Table 3.

First, from the first row of Table 3, the baseline method only achieves 47.2 \mathcal{J} and 50.1 \mathcal{F} . This inferior behavior attributes to two reasons: (1) The baseline method can not distinguish the similar objects that are close together and tends to segment the most salient region. In contrast, our method performs well with only 1 conditional query (see Table 4(a)), proving that dynamic convolution is essential for segmenting the referred object. (2) Our method uses a set of shared queries to track instances in all frames, and the best query is determined by the *voting* scores of each frame. In this sense, our model can produce a reliable reasoning result and keep the temporal consistency in the entire video. On the contrary, the baseline method could be regarded as a image-level method that independently predicts the results of each frame even though the model is able to aggregate the information from other frames.

Second, comparing the second and last row of Table 3, we can see that the standard FPN has already achieved strong performance and the vision-language fusion process further helps to provide more accurate segmentation. This is because the object mask would be inaccurate due to light variation, whereas the cross-modal fusion uses the text as a complementary to strengthen the object pixel features. Another technique is using the relative coordinate features, this would help the model better determine the location of referred object and lead to performance improvement, as shown in the third row in Table 3.

Number of Conditional Queries. Benefit from the design of *conditional queries*, all the initial object queries tend to find the referred object only. In this situation, we can only use a relatively small number of queries. In Table 4(a), it can be seen that the model achieves considerable results under all these settings, even with $N = 1$. Certainly, more queries enable the model make judgement from a wide range of instance candidates, which could better handle the complicated scenes where the similar objects are clustered together. The performance saturates at $N = 5$ and begins to slightly decrease when the query number gets larger. We

conjecture that it is caused by the imbalance of label assignment as there is only one positive sample in each frame.

Number of Training Clip Frames. We study the effect of training clip frame number in Table 4(b). Note that under $T = 1$, the model can be viewed as an image-level method and the performance of metric $\mathcal{J}\&\mathcal{F}$ is only 50.0. When the frame number increases to 3, the model enjoys a significant $\mathcal{J}\&\mathcal{F}$ gain of 4.8. This is because using more frames to form a clip helps the model better aggregate the temporal action-related information. We choose $T = 5$ by default.

Label Assignment Method. Our framework is able to predict the reference probability, box location and segmentation mask of the referred object. We find the optimal positive sample by minimizing the overall matching cost in Eq.4. There are some variants in the label assignment method and the comparison experiments are shown in Table 4(c). From the first two rows we show that the lack of box or mask cost would both lead to the performance drop. With the segmentation-centric design, the mask cost is the most direct guidance for optimization, and the box provides the location prior for dynamic kernel. Thus, the combination of classification, box and mask cost shows more robustness.

5. Conclusion

In this work, we propose ReferFormer, an extremely simple and unified framework for referring video object segmentation. This framework provides a new perspective for the R-VOS task which views the language as queries. These queries are restricted to attend to the referred object only, and the object tracking is easily achieved by linking the corresponding queries. Given the video and an expression, our framework directly produces the segmentation masks as well as the detected boxes of the referred object in all frames without post-process. We validate our model on Ref-Youtube-VOS, Ref-DAVIS17, A2D-Sentences and JHMDB-Sentences and it shows the state-of-the-art performance on the four benchmarks.

Acknowledgements

Ping Luo is supported by the General Research Fund of HK No.27208720, No.17212120, and the HKU-TCL Joint Research Center for Artificial Intelligence.

References

- [1] Ali Athar, Sabarinath Mahadevan, Aljosa Osep, Laura Leal-Taixé, and Bastian Leibe. Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In *European Conference on Computer Vision*, pages 158–177. Springer, 2020. 6
- [2] Miriam Bellver, Carles Ventura, Carina Silberer, Ioannis Kazakos, Jordi Torres, and Xavier Giro-i Nieto. Refvos: A closer look at referring expressions for video object segmentation. *arXiv preprint arXiv:2010.00263*, 2020. 2
- [3] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multi-modal transformers. *arXiv preprint arXiv:2111.14821*, 2021. 6, 7
- [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
- [5] Xin Chen, Bin Yan, Jiawen Zhu, Dong Wang, Xiaoyun Yang, and Huchuan Lu. Transformer tracking. In *CVPR*, 2021. 3
- [6] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [7] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *arXiv preprint arXiv:2106.05210*, 2021. 2
- [8] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. Vision-language transformer and query generation for referring segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16321–16330, 2021. 2
- [9] Zihan Ding, Tianrui Hui, Shaofei Huang, Si Liu, Xuan Luo, Junshi Huang, and Xiaoming Wei. Progressive multimodal interaction network for referring video object segmentation. 2021. 6, 7
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 3
- [11] Kirill Gavrilyuk, Amir Ghodrati, Zhenyang Li, and Cees GM Snoek. Actor and action video segmentation from a sentence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5958–5966, 2018. 7
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [13] Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from natural language expressions. In *European Conference on Computer Vision*, pages 108–124. Springer, 2016. 7
- [14] Shaofei Huang, Tianrui Hui, Si Liu, Guanbin Li, Yunchao Wei, Jizhong Han, Luoqi Liu, and Bo Li. Referring image segmentation via cross-modal progressive comprehension. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10488–10497, 2020. 2
- [15] Sukjun Hwang, Miran Heo, Seoung Wug Oh, and Seon Joo Kim. Video instance segmentation using inter-frame communication transformers. *arXiv preprint arXiv:2106.03299*, 2021. 2, 3
- [16] Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J Black. Towards understanding action recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 3192–3199, 2013. 5
- [17] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021. 3
- [18] Anna Khoreva, Anna Rohrbach, and Bernt Schiele. Video object segmentation with language referring expressions. In *Asian Conference on Computer Vision*, pages 123–141. Springer, 2018. 5
- [19] Chen Liang, Yu Wu, Yawei Luo, and Yi Yang. Clawcranenet: Leveraging object-level relation for text-based video segmentation. *arXiv preprint arXiv:2103.10702*, 2021. 7
- [20] Chen Liang, Yu Wu, Tianfei Zhou, Wenguan Wang, Zongxin Yang, Yunchao Wei, and Yi Yang. Rethinking cross-modal interaction from a top-down perspective for referring video object segmentation. *arXiv preprint arXiv:2106.01061*, 2021. 2, 6, 7
- [21] Huaijia Lin, Ruizheng Wu, Shu Liu, Jiangbo Lu, and Jiayia Jia. Video instance segmentation with a propose-reduce paradigm. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1739–1748, 2021. 6
- [22] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 4
- [23] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 5
- [24] Si Liu, Tianrui Hui, Shaofei Huang, Yunchao Wei, Bo Li, and Guanbin Li. Cross-modal progressive comprehension for referring segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 7
- [25] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 3, 6
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021. 3, 6, 7

- [27] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 6, 7
- [28] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 1
- [29] Gen Luo, Yiyi Zhou, Xiaoshuai Sun, Liujuan Cao, Chenglin Wu, Cheng Deng, and Rongrong Ji. Multi-task collaborative network for joint referring expression comprehension and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10034–10043, 2020. 2
- [30] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. TrackFormer: Multi-object tracking with transformers. *arXiv preprint arXiv:2101.02702*, 2021. 3
- [31] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016. 5
- [32] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9226–9235, 2019. 2
- [33] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016. 1
- [34] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 5
- [35] Hamid Rezatofighi, 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
- [36] Seonguk Seo, Joon-Young Lee, and Bohyung Han. Urvos: Unified referring video object segmentation network with a large-scale benchmark. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XV 16*, pages 208–223. Springer, 2020. 2, 5, 6, 7
- [37] Peize Sun, Yi Jiang, Rufeng Zhang, Enze Xie, Jinkun Cao, Xinting Hu, Tao Kong, Zehuan Yuan, Changhu Wang, and Ping Luo. Transtrack: Multiple-object tracking with transformer. *arXiv preprint arXiv:2012.15460*, 2020. 3
- [38] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14454–14463, 2021. 3
- [39] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 282–298. Springer, 2020. 4
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2, 3, 5
- [41] Paul Voigtlaender, Yuning Chai, Florian Schroff, Hartwig Adam, Bastian Leibe, and Liang-Chieh Chen. Feelvos: Fast end-to-end embedding learning for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9481–9490, 2019. 2
- [42] Hao Wang, Cheng Deng, Junchi Yan, and Dacheng Tao. Asymmetric cross-guided attention network for actor and action video segmentation from natural language query. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3939–3948, 2019. 7
- [43] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8741–8750, 2021. 2, 3, 4, 5, 6
- [44] Junfeng Wu, Yi Jiang, Wenqing Zhang, Xiang Bai, and Song Bai. Seqformer: a frustratingly simple model for video instance segmentation. *arXiv preprint arXiv:2112.08275*, 2021. 2, 3, 6
- [45] Chenliang Xu, Shao-Hang Hsieh, Caiming Xiong, and Jason J Corso. Can humans fly? action understanding with multiple classes of actors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2264–2273, 2015. 5
- [46] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. 1
- [47] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *ICCV*, 2021. 3
- [48] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5188–5197, 2019. 3
- [49] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. In *European Conference on Computer Vision*, pages 332–348. Springer, 2020. 2
- [50] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10502–10511, 2019. 6
- [51] Linwei Ye, Mrigank Rochan, Zhi Liu, Xiaoqin Zhang, and Yang Wang. Referring segmentation in images and videos with cross-modal self-attention network. *arXiv preprint arXiv:2102.04762*, 2021. 7

- [52] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer, 2016. 6
- [53] Yuting Zhang, Luyao Yuan, Yijie Guo, Zhiyuan He, I-An Huang, and Honglak Lee. Discriminative bimodal networks for visual localization and detection with natural language queries. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 557–566, 2017. 2
- [54] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890, 2021. 3
- [55] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 2, 3, 4, 6