

End-to-End Referring Video Object Segmentation with Multimodal Transformers

Adam Botach, Evgenii Zheltonozhskii, Chaim Baskin
Technion – Israel Institute of Technology

{botach, evgeniizh}@campus.technion.ac.il chaimbaskin@cs.technion.ac.il

Abstract

The referring video object segmentation task (RVOS) involves segmentation of a text-referred object instance in the frames of a given video. Due to the complex nature of this multimodal task, which combines text reasoning, video understanding, instance segmentation and tracking, existing approaches typically rely on sophisticated pipelines in order to tackle it. In this paper, we propose a simple Transformer-based approach to RVOS. Our framework, termed Multimodal Tracking Transformer (MTTR), models the RVOS task as a sequence prediction problem. Following recent advancements in computer vision and natural language processing, MTTR is based on the realization that video and text can be processed together effectively and elegantly by a single multimodal Transformer model. MTTR is end-to-end trainable, free of text-related inductive bias components and requires no additional mask-refinement post-processing steps. As such, it simplifies the RVOS pipeline considerably compared to existing methods. Evaluation on standard benchmarks reveals that MTTR significantly outperforms previous art across multiple metrics. In particular, MTTR shows impressive +5.7 and +5.0 mAP gains on the A2D-Sentences and JHMDB-Sentences datasets respectively, while processing 76 frames per second. In addition, we report strong results on the public validation set of Refer-YouTube-VOS, a more challenging RVOS dataset that has yet to receive the attention of researchers. The code to reproduce our experiments is available at <https://github.com/mttr2021/MTTR>.

1. Introduction

Attention-based [41] deep neural networks exhibit impressive performance on various tasks across different fields, from computer vision [10, 27] to natural language processing [3, 8]. These advancements make networks of this sort, such as the Transformer [41], particularly interesting candidates for solving multimodal problems. By relying on the self-

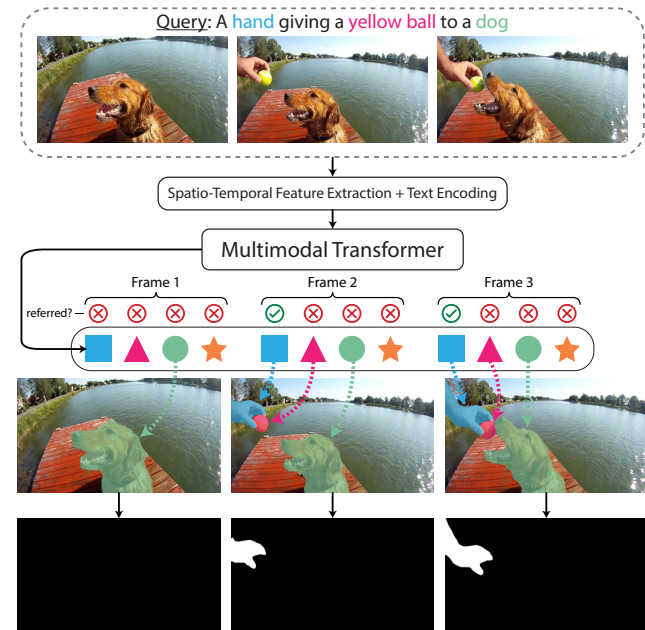


Figure 1. Given a text query and a sequence of video frames, the proposed model outputs prediction sequences for all object instances in the video prior to determining the referred instance. Here predictions with the same color and shape belong to the same sequence and attend to the same object instance in different frames. Note that the order of instance predictions for different frames remains the same. Best viewed in color.

attention mechanism, which allows each token in a sequence to globally aggregate information from every other token, Transformers excel at modeling global dependencies and have become the cornerstone in most NLP tasks [3, 8, 35, 50]. Transformers have also started showing promise in solving computer vision tasks, from recognition [10] to object detection [4] and even outperforming the long-used CNNs as general-purpose vision backbones [27].

The referring video object segmentation task (RVOS) involves the segmentation of a text-referred object instance in the frames of a given video. Compared with the referring image segmentation task (RIS) [29, 52], in which objects are mainly referred to by their appearance, in RVOS objects can

also be referred to by the actions they are performing or in which they are involved. This renders RVOS significantly harder than RIS, as text expressions that refer to actions often cannot be properly deduced from a single static frame. Furthermore, unlike their image-based counterparts, RVOS methods may be required to establish data association of the referred object across multiple frames (tracking) in order to deal with disturbances such as occlusions or motion blur.

To solve these challenges and effectively align video with text, existing RVOS approaches [14, 25, 32] typically rely on complicated pipelines. In contrast, here we propose a simple, end-to-end Transformer-based approach to RVOS. Using recent advancements in Transformers for textual feature extraction [26, 41], visual feature extraction [10, 27, 28] and object detection [4, 45], we develop a framework that significantly outperforms existing approaches. To accomplish this, we employ a single multimodal Transformer and model the task as a sequence prediction problem. Given a video and a text query, our model generates prediction sequences for *all* objects in the video before determining the one the text refers to. Additionally, our method is free of text-related inductive bias modules and utilizes a simple cross-entropy loss to align the video and the text. As such, it is much less complicated than previous approaches to the task.

The proposed pipeline is schematically depicted in Fig. 1. First, we extract linguistic features from the text query using a standard Transformer-based text encoder, and visual features from the video frames using a spatio-temporal encoder. The features are then passed into a multimodal Transformer, which outputs several sequences of object predictions [45]. Next, to determine which of the predicted sequences best corresponds to the referred object, we compute a text-reference score for each sequence. For this we propose a *temporal segment voting scheme* that allows our model to focus on more relevant parts of the video when making the decision.

Our main contributions are as follows:

- We present a Transformer-based RVOS framework, dubbed **Multimodal Tracking Transformer (MTTR)**, which models the task as a parallel sequence prediction problem and outputs predictions for *all* objects in the video prior to selecting the one referred to by the text.
- Our sequence selection strategy is based on a *temporal segment voting scheme*, a novel reasoning scheme that allows our model to focus on more relevant parts of the video with regards to the text.
- The proposed method is end-to-end trainable, free of text-related inductive bias modules, and requires no additional mask refinement. As such, it greatly simplifies the RVOS pipeline compared to existing approaches.
- We thoroughly evaluate our method. On the A2D-Sentences and JHMDB-Sentences [12], MTTR signifi-

cantly outperforms all existing methods across all metrics. We also show strong results on the public validation set of Refer-YouTube-VOS [39], a challenging dataset that has yet to receive attention in the literature.

2. Related Work

Referring video object segmentation. The RVOS task was introduced by Gavriluk et al. [12], whose goal was to attain pixel-level segmentation of actors and their actions in video content. To effectively aggregate and align visual, temporal and lingual information from video and text, state-of-the-art RVOS approaches typically rely on complicated pipelines [25, 30, 32, 42, 43]. Gavriluk et al. [12] proposed an I3D-based [5] encoder-decoder architecture that generated dynamic filters from text features and convolved them with visual features to obtain the masks. Following them, Wang et al. [42] added spatial context to the kernels with deformable convolutions [7]. For a more effective representation, VT-Capsule [30] encoded each modality in capsules [37], while ACGA [43] utilized a co-attention mechanism to enhance the multimodal features. To improve positional relation representations in the text, PRPE [32] explored a positional encoding mechanism based on polar coordinates. URVOS [39] improved tracking capabilities by performing language-based object segmentation on a key frame and then propagating its mask throughout the video. AAMN [49] utilized a top-down approach where an off-the-shelf object detector is used to localize objects in the video prior to parsing relations between visual and textual features. CMPC-V [25] achieved state-of-the-art results by constructing a temporal graph from video and text features, and applying graph convolution [18] to detect the referred entity.

Transformers. The Transformer [41] was introduced as an attention-based building block for sequence-to-sequence machine translation, and since then has become the cornerstone for most NLP tasks [3, 8, 35, 50]. Unlike previous architectures, the Transformer relies entirely on the attention mechanism to draw dependencies between input and output.

Recently, the introduction of Transformers to computer vision tasks has demonstrated spectacular performance. DETR [4], which utilizes a non-auto-regressive Transformer, simplifies the traditional object detection pipeline while achieving performance comparable to that of CNN-based detectors [36]. Given a fixed set of learned object queries, DETR reasons about the global context of an image and the relations between its objects and then outputs a final set of detection predictions in parallel. VisTR [45] extends the idea behind DETR to video instance segmentation. It views the task as a direct end-to-end parallel sequence prediction problem. By supervising video instances at the sequence level as a whole, VisTR is able to output an ordered sequence of masks for each instance in a video directly (i.e., natural tracking).

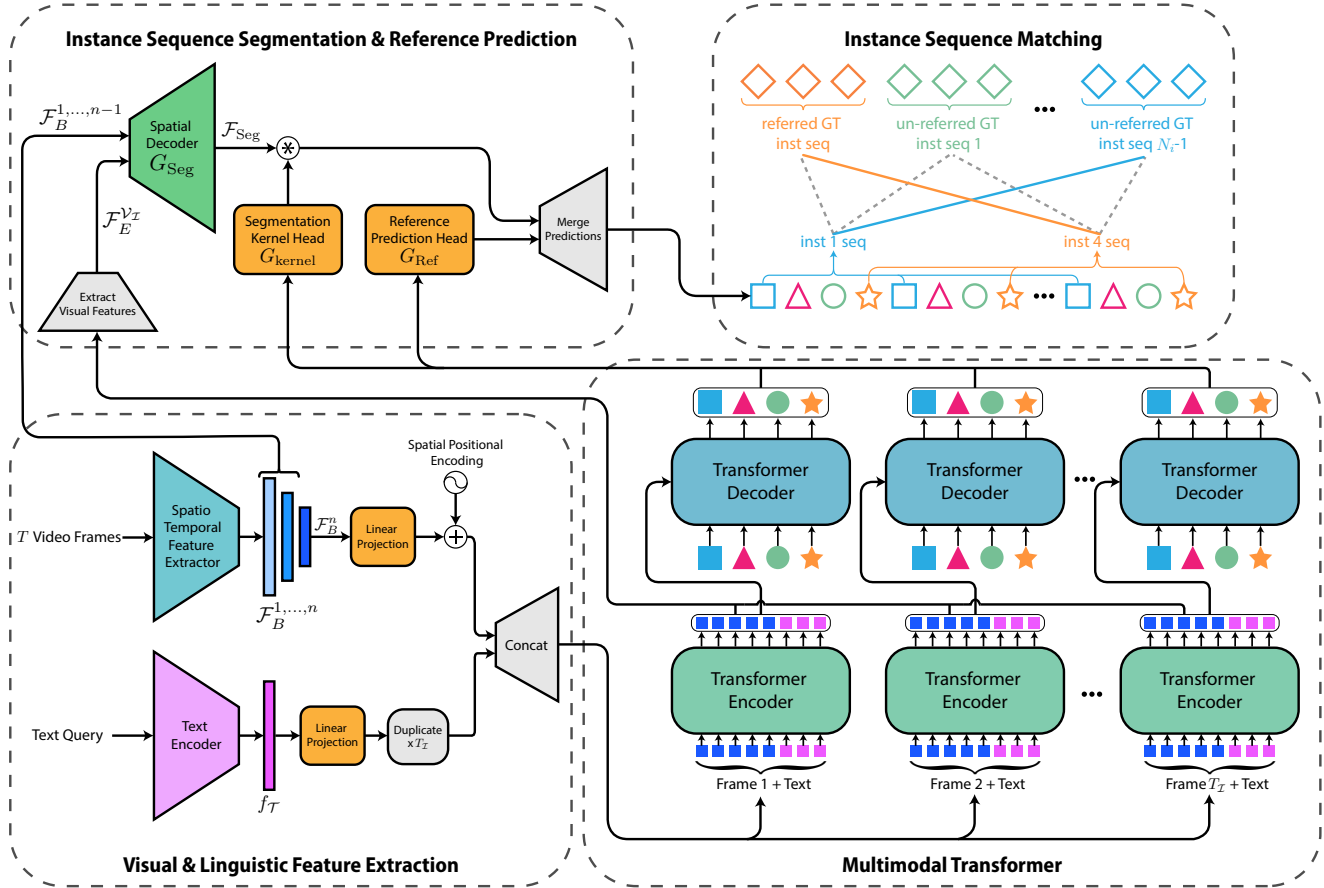


Figure 2. A detailed overview of MTTR. First, the input text and video frames are passed through feature encoders and then concatenated into multimodal sequences (one per frame). A multimodal Transformer then encodes the feature relations and decodes instance-level features into a set of prediction sequences. Next, corresponding mask and reference prediction sequences are generated. Finally, the predicted sequences are matched with the ground truth sequences for supervision (in training) or used to generate the final prediction (during inference).

ViT [10] introduced the Transformer to image recognition by using linearly projected patches as tokens for a Transformer encoder. Swin Transformer [27] proposed a general-purpose backbone for computer vision based on a hierarchical Transformer whose representations are computed inside shifted windows. This architecture was also extended to the video domain [28], which we adapt as our temporal encoder.

Another recent relevant work is MDETR [16], a DETR-based end-to-end multimodal detector that detects objects in an image conditioned on a text query. Different from our method, their approach is designed to work on static images, and its performance largely depends on well-annotated datasets that contain *aligned* text and box annotations, the types of which are not available in the RVOS task.

3. Method

3.1. Method Overview

Task definition. The input of RVOS consists of a frame sequence $\mathcal{V} = \{v_i\}_{i=1}^T$, where $v_i \in \mathbb{R}^{C \times H_0 \times W_0}$, and a text

query $\mathcal{T} = \{t_i\}_{i=1}^L$, where t_i is the i^{th} word in the text. Then, for a subset of frames of interest $\mathcal{V}_{\mathcal{T}} \subseteq \mathcal{V}$ of size $T_{\mathcal{T}}$, the goal is to segment the object referred by \mathcal{T} in each frame in $\mathcal{V}_{\mathcal{T}}$. We note that since producing mask annotations requires significant efforts, $\mathcal{V}_{\mathcal{T}}$ rarely contains all of the frames in \mathcal{V} .

Feature extraction. We begin by extracting features from each frame in the sequence \mathcal{V} using a deep spatio-temporal encoder. Simultaneously, linguistic features are extracted from the text query \mathcal{T} using a Transformer-based [41] text encoder. Then, the spatio-temporal and linguistic features are linearly projected to a shared dimension D .

Instance prediction. In the next step, the features of each frame of interest are flattened and separately concatenated with the text embeddings, producing a set of $T_{\mathcal{T}}$ multimodal sequences. These sequences are fed in parallel into a Transformer [4, 41]. In the Transformer’s encoder layers, the textual embeddings and the visual features of each frame exchange information. Then, the decoder layers, which are

fed with N_q object queries per input frame, query the multimodal sequences for entity-related information and store it in the object queries. Corresponding queries of different frames share the same trainable weights and are trained to attend to the same instance in the video (each one in its designated frame). We refer to these queries (represented by the same unique color and shape in Figs. 1 and 2) as queries belonging to the same *instance sequence*. This design allows for natural tracking of each object instance in the video [45].

Output generation. For each output instance sequence, we generate a corresponding mask sequence using an FPN-like [22] spatial decoder and dynamically generated conditional convolution kernels [40, 44]. Finally, we use a novel text-reference score function that, based on text associations, determines which of the object query sequences has the strongest association with the object described in \mathcal{T} , and returns its segmentation sequence as the model’s prediction.

3.2. Temporal Encoder

A suitable temporal encoder for the RVOS task should be able to extract *both* visual characteristics (e.g., shape, size, location) and action semantics for each instance in the video. Several previous works [12, 25, 32] utilized the Kinetics-400 [17] pre-trained I3D network [5] as their temporal encoder. However, since I3D was originally designed for action classification, using its outputs as-is for tasks that require fine details (e.g., instance segmentation) is not ideal as the features it outputs tend to suffer from spatial misalignment caused by temporal downsampling. To compensate for this side effect, past state-of-the-art approaches came up with different solutions, from auxiliary mask refinement algorithms [19, 25] to utilizing additional backbones that operate alongside the temporal encoder [14]. In contrast, our end-to-end approach does not require any additional mask refinement steps and utilizes a single backbone.

Recently, the Video Swin Transformer [28] was proposed as a generalization of the Swin Transformer [27] to the video domain. While the original Swin was designed with dense predictions (such as segmentation) in mind, Video Swin was tested mainly on action recognition benchmarks. To the best of our knowledge, we are the first to utilize it (with a slight modification) for video segmentation. As opposed to I3D, Video Swin contains just a single temporal downsampling layer and can be easily modified to output *per-frame* feature maps (we refer to the supplement for more details). As such, it is a much better choice for processing a full sequence of consecutive video frames for segmentation purposes.

3.3. Multimodal Transformer

For each frame of interest, the temporal encoder generates a feature map $f_t^{V\mathcal{I}} \in \mathbb{R}^{H \times W \times C_V}$ and the text encoder outputs a linguistic embedding vector $f_{\mathcal{T}} \in \mathbb{R}^{L \times D_{\mathcal{T}}}$ for the

text. These visual and linguistic features are linearly projected to a shared dimension D . The features of each frame are then flattened and *separately* concatenated with the text embeddings, resulting in a set of $T_{\mathcal{T}}$ multimodal sequences, each of shape $(H \times W + L) \times D$. The multimodal sequences along with a set of N_q instance sequences are then fed in parallel into a Transformer as described earlier. Our Transformer architecture is similar to the one used in DETR [4]. Accordingly, the problem now comes down to finding the instance sequence that attends to the text-referred object.

3.4. The Instance Segmentation Process

Our segmentation process, as shown in Fig. 2, consists of several steps. First, given \mathcal{F}_E , the updated multimodal sequences output by the last Transformer encoder layer, we extract and reshape the video-related part of each sequence (i.e., the first $H \times W$ tokens) into the set $\mathcal{F}_E^{V\mathcal{I}}$. Then, we take $\mathcal{F}_B^{1, \dots, n-1}$, the outputs of the first $n - 1$ blocks of our temporal encoder, and hierarchically fuse them with $\mathcal{F}_E^{V\mathcal{I}}$ using an FPN-like [22] spatial decoder G_{Seg} . This process results in semantically-rich, high resolution feature maps of the video frames, denoted as \mathcal{F}_{Seg} .

$$\mathcal{F}_{\text{Seg}} = \{f_{\text{Seg}}^t\}_{t=1}^{T_{\mathcal{I}}}, f_{\text{Seg}}^t \in \mathbb{R}^{D_s \times \frac{H_0}{4} \times \frac{W_0}{4}} \quad (1)$$

Next, for each instance sequence $\mathcal{Q} = \{q_t\}_{t=1}^{T_{\mathcal{I}}}, q_t \in \mathbb{R}^D$ output by the Transformer decoder, we use a two-layer perceptron G_{kernel} to generate a corresponding sequence of conditional segmentation kernels [40, 44].

$$G_{\text{kernel}}(\mathcal{Q}) = \{k_t\}_{t=1}^{T_{\mathcal{I}}}, k_t \in \mathbb{R}^{D_s} \quad (2)$$

Finally, a sequence of segmentation masks \mathcal{M} is generated for \mathcal{Q} by convolving each segmentation kernel with its corresponding frame features, followed by a bilinear upsampling operation to resize the masks into ground-truth resolution,

$$\mathcal{M} = \{m_t\}_{t=1}^{T_{\mathcal{I}}}, m_t = \text{Upsample}(k_t * f_{\text{Seg}}^t) \in \mathbb{R}^{H_0 \times W_0}. \quad (3)$$

3.5. Instance Sequence Matching

During the training process we need to determine which of the predicted instance sequences best fits the referred object. However, if the video sequence \mathcal{V} contains additional annotated instances, we found that supervising their detection (as negative examples) alongside that of the referred instance helps stabilize the training process.

Let us denote by y the set of ground-truth sequences that are available for \mathcal{V} , and by $\hat{y} = \{\hat{y}_i\}_{i=1}^{N_q}$ the set of the predicted instance sequences. We assume that the number of predicted sequences (N_q) is chosen to be strictly greater than the number of annotated instances (denoted N_i) and that the ground-truth sequences set is padded with \emptyset (no object) to fill any missing slots. Then, we want to find a matching

between the two sets [4, 45]. Accordingly, we search for a permutation $\hat{\sigma} \in S_{N_q}$ with the lowest total cost:

$$\hat{\sigma} = \arg \min_{\sigma \in S_{N_q}} \sum_{i=1}^{N_q} \mathcal{C}_{\text{Match}}(\hat{y}_{\sigma(i)}, y_i), \quad (4)$$

where $\mathcal{C}_{\text{Match}}$ is a pair-wise matching cost. The optimal permutation $\hat{\sigma}$ can be computed efficiently using the Hungarian algorithm [20]. Each ground-truth sequence is of the form

$$y_i = (m_i, r_i) = \left(\{m_i^t\}_{t=1}^{T_x}, \{r_i^t\}_{t=1}^{T_x} \right), \quad (5)$$

where m_i^t is a ground-truth mask, and $r_i^t \in \{0, 1\}^2$ is a one-hot referring vector, i.e., the positive class means that y_i corresponds to the text-referred object *and* that this object is visible in the corresponding video frame v_t . Note that if y_i is a padding sequence then $m_i = \emptyset$.

To allow our model to produce reference predictions in the form of Eq. (5), we use a reference prediction head, denoted G_{Ref} , which consists of a single linear layer of shape $D \times 2$ followed by a softmax layer. Given a predicted object query $q \in \mathbb{R}^D$, this head takes q as input and outputs a reference prediction $\hat{r} \equiv G_{\text{Ref}}(q)$.

Thus, each prediction of our model is a pair of sequences:

$$\hat{y}_j = (\hat{m}_j, \hat{r}_j) = \left(\{\hat{m}_j^t\}_{t=1}^{T_x}, \{\hat{r}_j^t\}_{t=1}^{T_x} \right). \quad (6)$$

We define the pair-wise matching cost function as the sum

$$\mathcal{C}_{\text{Match}}(\hat{y}_j, y_i) = \mathbb{1}_{\{m_i \neq \emptyset\}} \left[\lambda_d \mathcal{C}_{\text{Dice}}(\hat{m}_j, m_i) + \lambda_r \mathcal{C}_{\text{Ref}}(\hat{r}_j, r_i) \right], \quad (7)$$

where $\lambda_d, \lambda_r \in \mathbb{R}$ are hyperparameters. $\mathcal{C}_{\text{Dice}}$ supervises the predicted mask sequence using the ground-truth mask sequence by averaging the negation of the Dice coefficients [31] of each pair of corresponding masks at every time step. We refer to the supplement for the full definition of this cost function. \mathcal{C}_{Ref} supervises the reference predictions using the corresponding ground-truth sequence as follows

$$\mathcal{C}_{\text{Ref}}(\hat{r}_j, r_i) = -\frac{1}{T_x} \sum_{t=1}^{T_x} \hat{r}_j^t \cdot r_i^t. \quad (8)$$

3.6. Loss Functions

Let us denote (with a slight abuse of notation) by \hat{y} the set of predicted instance sequences permuted according to the optimal permutation $\hat{\sigma} \in S_{N_q}$. Then, we can define our loss function as follows:

$$\mathcal{L}(\hat{y}, y) = \sum_{i=1}^{N_q} \mathbb{1}_{\{m_i \neq \emptyset\}} \mathcal{L}_{\text{Mask}}(\hat{m}_i, m_i) + \mathcal{L}_{\text{Ref}}(\hat{r}_i, r_i). \quad (9)$$

Following VisTR [45], the first term, dubbed $\mathcal{L}_{\text{Mask}}$, ensures mask alignment between the predicted and ground-truth sequences. As such, this term is defined as a combination of the Dice [31] and the per-pixel Focal [23] loss functions:

$$\mathcal{L}_{\text{Mask}}(\hat{m}_i, m_i) = \lambda_d \mathcal{L}_{\text{Dice}}(\hat{m}_i, m_i) + \lambda_f \mathcal{L}_{\text{Focal}}(\hat{m}_i, m_i), \quad (10)$$

where $\lambda_d, \lambda_f \in \mathbb{R}$ are hyperparameters. Both $\mathcal{L}_{\text{Dice}}$ and $\mathcal{L}_{\text{Focal}}$ are applied on corresponding masks at every time step, and are normalized by the number of instances inside the training batch. We refer to the supplement for the full definitions of these functions.

The second loss term, denoted \mathcal{L}_{Ref} , is a cross-entropy term that supervises the sequence reference predictions:

$$\mathcal{L}_{\text{Ref}}(\hat{r}_i, r_i) = -\lambda_r \frac{1}{T_x} \sum_{t=1}^{T_x} r_i^t \cdot \log(\hat{r}_i^t), \quad (11)$$

where $\lambda_r \in \mathbb{R}$ is a hyperparameter. In practice we further downweight the terms of the negative (“unreferred”) class by a factor of 10 to account for class imbalance [4]. Also, note that the same λ_r and λ_d are used as weights in the matching cost (7) and loss functions. Intriguingly, despite \mathcal{L}_{Ref} ’s simplicity and lack of explicit text-related inductive bias, it was able to deliver equivalent or even better performance compared with more complex loss functions [16] that we tested. Hence, and for the sake of simplicity, no additional loss functions are used for text supervision in our method.

3.7. Inference

For a given sample of video and text, let us denote by $\mathcal{R} = \{\hat{r}_i\}_{i=1}^{N_q}$ the set of reference prediction sequences output by our model. Additionally, we denote by $p_{\text{ref}}(\hat{r}_i^t)$ the probability of the positive (“referred”) class for a given reference prediction \hat{r}_i^t . During inference we return the segmentation mask sequence $\mathcal{M}_{\text{pred}}$ that corresponds to \hat{r}_{pred} , the predicted reference sequence with the highest positive score:

$$\hat{r}_{\text{pred}} = \arg \max_{\hat{r}_i \in \mathcal{R}} \sum_{t=1}^{T_x} p_{\text{ref}}(\hat{r}_i^t). \quad (12)$$

This sequence selection scheme, which we term the “*temporal segment voting scheme*” (TSVS), grades each prediction sequence based on the total association of its terms with the text referred object. Thus, it allows our model to focus on more relevant parts of the video (in which the referred object is visible), and disregard less relevant parts (which may depict irrelevant objects or in which the referred object is occluded) when making the decision. We refer to the supplement for further analysis of the effect of TSVS.

4. Experiments

To evaluate our approach, we conduct experiments on three referring video object segmentation datasets. The first two, A2D-Sentences and JHMDB-Sentences [12], were created by adding textual annotations to the original A2D [47] and JHMDB [15] datasets. Each video in A2D has 3–5 frames annotated with pixel-level segmentation masks, while in JHMDB, 2D articulated human puppet masks are available for all frames. More details are in the supplement. We

Method	Precision					IoU		mAP
	50%	60%	70%	80%	90%	Overall	Mean	
Hu et al. [13]	34.8	23.6	13.3	3.3	0.1	47.4	35.0	13.2
Gavrilyuk et al. [12] (RGB)	47.5	34.7	21.1	8.0	0.2	53.6	42.1	19.8
RefVOS [1]	57.8	–	–	–	9.3	67.2	49.7	–
AAMN [49]	68.1	62.9	52.3	29.6	2.9	61.7	55.2	39.6
CMSA+CFSA [51]	48.7	43.1	35.8	23.1	5.2	61.8	43.2	–
CSTM [14]	65.4	58.9	49.7	33.3	9.1	66.2	56.1	39.9
CMPC-V (I3D) [25]	65.5	59.2	50.6	34.2	9.8	65.3	57.3	40.4
MTTR ($w = 8$, ours)	72.1	68.4	60.7	45.6	16.4	70.2	61.8	44.7
MTTR ($w = 10$, ours)	75.4	71.2	63.8	48.5	16.9	72.0	64.0	46.1

Table 1. Comparison with state-of-the-art methods on A2D-Sentences [12].

Method	Precision					IoU		mAP
	50%	60%	70%	80%	90%	Overall	Mean	
Hu et al. [13]	63.3	35.0	8.5	0.2	0.0	54.6	52.8	17.8
Gavrilyuk et al. [12] (RGB)	69.9	46.0	17.3	1.4	0.0	54.1	54.2	23.3
AAMN [49]	77.3	62.7	36.0	4.4	0.0	58.3	57.6	32.1
CMSA+CFSA [51]	76.4	62.5	38.9	9.0	0.1	62.8	58.1	–
CSTM [14]	78.3	63.9	37.8	7.6	0.0	59.8	60.4	33.5
CMPC-V (I3D) [25]	81.3	65.7	37.1	7.0	0.0	61.6	61.7	34.2
MTTR ($w = 8$, ours)	91.0	81.5	57.0	14.4	0.1	67.4	67.9	36.6
MTTR ($w = 10$, ours)	93.9	85.2	61.6	16.6	0.1	70.1	69.8	39.2

Table 2. Comparison with state-of-the-art methods on JHMDB-Sentences [12].

adopt Overall IoU, Mean IoU, and precision@K to evaluate our method on these datasets. *Overall IoU* computes the ratio between the total intersection and the total union area over all the test samples. *Mean IoU* is the averaged IoU over all the test samples. *Precision@K* considers the percentage of test samples whose IoU scores are above a threshold K , where $K \in [0.5, 0.6, 0.7, 0.8, 0.9]$. We also compute mean average precision (mAP) over 0.50:0.05:0.95 [24].

We want to note that we found inconsistencies in the mAP metric calculation in previous studies. For example, examination of published code revealed incorrect calculation of the metric as the average of the precision@K metric over several K values. To avoid further confusion and ensure a fair comparison, we suggest adopting the COCO API¹ for mAP calculation. For reference, a full implementation of the evaluation that utilizes the API is released with our code.

We further evaluate MTTR on the more challenging Refer-YouTube-VOS dataset, introduced by Seo et al. [39], who provided textual annotations for the original YouTube-VOS dataset [48]. Each video has pixel-level instance segmentation annotations for every fifth frame. The original release of Refer-YouTube-VOS contains two subsets. One subset contains *first-frame expressions* that describe only the first frame. The other contains *full-video expressions* that are

based on the whole video and are, therefore, more challenging. Following the introduction of the RVOS competition², only the more challenging subset of the dataset is publicly available now. Since ground-truth annotations are available only for the training samples and the test server is currently inaccessible, we report results on the validation samples by uploading our predictions to the competition’s server³. We refer to the supplement for more details. The primary evaluation metrics for this dataset are the average of the region similarity (\mathcal{J}) and the contour accuracy (\mathcal{F}) [34].

4.1. Implementation Details

As our temporal encoder we use the smallest (“tiny”) Video Swin Transformer [28] pretrained on Kinetics-400 [17]. The original Video Swin consists of four blocks with decreasing spatial resolution. We found the output of the fourth block to be too small for small object detection and hence we only utilize the first three blocks. We use the output of the third block as the input of the multimodal Transformer, while the outputs of the earlier blocks are fed into the spatial decoder. We also modify the encoder’s single temporal downsampling layer to output *per-frame* feature

²<https://youtube-vos.org/dataset/rvos/>

³<https://competitions.codalab.org/competitions/29139>

¹<https://github.com/cocodataset/cocoapi>

Method	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
URVOS [39]	47.23	45.27	49.19
CMPC-V (I3D) [25]	47.48	45.64	49.32
Ding et al. [9] ⁺	54.8	53.7	56.0
MTTR (ours)	55.32	54.00	56.64

Table 3. Results on Refer-YouTube-VOS. The upper half is evaluated on the original validation set, while the bottom half is evaluated on the public validation set. ⁺ – ensemble.

maps as required by our model. As our text encoder we use the Hugging Face [46] implementation of RoBERTa-base [26]. For A2D-Sentences [12] we feed the model windows of $w = 8$ frames with the annotated target frame in the middle. Each frame is resized such that the shorter side is at least 320 pixels and the longer side is at most 576 pixels. For Refer-YouTube-VOS [39], we use windows of $w = 12$ consecutive annotated frames during training, and full-length videos (up to 36 annotated frames) during evaluation. Each frame is resized such that the shorter side is at least 360 pixels and the longer side is at most 640 pixels. We do not use any segmentation-related pretraining, e.g., on COCO [24], which is known to boost segmentation performance [45]. We refer the reader to the supplement for more implementation details.

4.2. Comparison with State-of-the-Art Methods

We compare our method with existing approaches on the A2D-Sentences dataset. For fair comparison with existing works [14, 25], our model is trained and evaluated for this purpose with windows of size 8. As shown in Tab. 1, our method significantly outperforms existing approaches across all metrics. For example, our model shows a 4.3 mAP gain over current state of the art, and an absolute improvement of 6.6% on the most stringent metric P@0.9, which demonstrates its ability to generate high-quality masks. We also note that our top configuration ($w = 10$) achieves a massive 5.7 mAP gain and 6.7% absolute improvement on both Mean and Overall IoU compared to the current state of the art. Impressively, this configuration is able to do so while processing 76 frames per second on a single RTX 3090 GPU.

Following previous works [12, 25], we evaluate the generalization ability of our model by evaluating it on JHMDB-Sentences without fine-tuning. We uniformly sample three frames from each video and evaluate our best model on these frames. As shown in Tab. 2, our method generalizes well and outperforms all existing approaches. Note that all methods (including ours) produce low results on P@0.9. This can be attributed to JHMDB’s [15] imprecise mask annotations which were generated by a coarse human puppet model.

Finally, we report our results on the public validation set of Refer-YouTube-VOS [39] in Tab. 3. As mentioned earlier,

this subset contains only the more challenging *full-video expressions* from the original release of Refer-YouTube-VOS. Compared with existing methods [25, 39] which trained and evaluated on the full version of the dataset, our model demonstrates superior performance across all metrics despite being trained on less data and evaluated exclusively on a more challenging subset. Additionally, our method shows competitive performance compared with the methods that led in the 2021 RVOS competition [9, 21]. We note, however, that these methods use ensembles and are trained on additional segmentation and referring datasets [24, 29, 48, 52].

4.3. Ablation Studies

We conduct ablation studies on A2D-Sentences to evaluate our model’s design and robustness. Unless stated otherwise, we use window size $w = 6$. An ablation study on the number of object queries can be found in the supplement.

Temporal encoder. To evaluate MTTR’s performance independently of the temporal encoder, we compare it with CMPC-I, the image-targeted version of CMPC-V [25]. Following CMPC-I, we use DeepLab-ResNet101 [6] pretrained on PASCAL-VOC [11] as a visual feature extractor. We train our model using only the target frames (i.e., without additional frames for temporal context). As shown in Tab. 4a, our method significantly surpasses CMPC-I across all metrics, with a 6.1 gain in mAP and 8.7% absolute improvement in Mean IoU. In fact, this configuration of our model surpasses all existing methods *regardless* of the temporal context.

Temporal context. In Tab. 4b we study the effect of the temporal context size on MTTR’s performance. A larger temporal context enables better extraction of action-related information. For this purpose, we train and evaluate our model using different window sizes. As expected, widening the temporal context leads to large performance gains, with an mAP gain of 4.3 and an absolute Mean IoU improvement of 3.7% when gradually changing the window size from 1 to 10. Intriguingly, however, peak performance on A2D-Sentences is obtained using $w = 10$, as widening the window even further (e.g., $w = 12$) results in a performance drop.

Text encoder. To study the effect of the selected word embeddings on our model’s performance, we train our model using two additional widely-used Transformer-based text encoders, namely BERT-base [8] and Distill-RoBERTa-base [38], a distilled version of RoBERTa [26]. Additionally, we experiment with GloVe [33] and fastText [2], two simpler word embedding methods. As shown in Tab. 4c, our model achieves comparable performance when relying on the different Transformer-based encoders, which demonstrates its robustness to this change. Unsurprisingly, however, performance is slightly worse when relying on the simpler methods.



Figure 3. Visual examples of MTTR’s performance on the Refer-YouTube-VOS [39] validation set. Best viewed in color.

Method	IoU			mAP
	Overall	Mean	mAP	
CMPC-I [25]	64.9	51.5	35.1	
MTTR (DeepLab-ResNet101)	67.5	60.2	41.2	
MTTR (Video Swin-T, $w = 1$)	68.9	60.3	41.8	

(a) Non-temporal backbones.

w	IoU			mAP
	Overall	Mean	mAP	
$w = 1$	68.9	60.3	41.8	
$w = 4$	69.7	61.5	43.8	
$w = 6$	69.5	61.8	44.0	
$w = 8$	70.2	61.8	44.7	
$w = 10$	72.0	64.0	46.1	
$w = 12$	69.3	62.0	44.0	

(b) Input window size effect.

Method	IoU			mAP
	Overall	Mean	mAP	
RoBERTa (base)	69.5	61.8	44.0	
BERT (base)	69.7	62.1	44.3	
Distill-RoBERTa (base)	70.5	62.4	43.8	
GloVe	68.3	60.9	43.4	
fastText	67.3	59.5	43.1	

(c) Word embeddings effect.

Table 4. Ablation studies on A2D-Sentences [12] dataset.

This may be explained by the fact that while Transformer-based encoders are able to dynamically encode sentence context within their output embeddings, simpler methods disregard this context and merely rely on fixed pretrained embeddings.

Supervision of un-referred instances. To study the effect of supervising the detections of un-referred instances alongside that of the referred instance in each sample, we train different configurations of our model without supervision of un-referred instances. Intriguingly, in all such experiments our model immediately converges to a local minimum of the text loss (\mathcal{L}_{Ref}), where the same object query is repeatedly matched with all ground-truth instances, thus leaving the rest of the object queries untrained. In some experiments our model manages to escape this local minimum after a few epochs and then achieves comparable performance with our original configuration. Nevertheless, in other experiments this phenomenon significantly hinders its final mAP score.

4.4. Qualitative Analysis

As illustrated in Fig. 3, MTTR can successfully track and segment the referred objects even in challenging situations where they are surrounded by similar instances, occluded, or completely outside of the frame in large parts of the video.

5. Conclusion

We introduced MTTR, a simple Transformer-based approach to RVOS that models the task as a sequence prediction problem. Our end-to-end method considerably simplifies existing RVOS pipelines by simultaneously processing both text and video frames in a single multimodal Transformer. Extensive evaluation of our approach on standard benchmarks reveals that our method outperforms existing state-of-the-art methods by a large margin (e.g., a 5.7 mAP improvement on A2D-Sentences). We hope our work will inspire others to see the potential of Transformers for solving complex multimodal tasks.

References

1. Miriam Bellver, Carles Ventura, Carina Silberer, Jordi Torres Ioannidis Kazakos, and Xavier Giró-i-Nieto. RefVOS: a closer look at referring expressions for video object segmentation. *arXiv preprint arXiv:2010.00263*, 2020. (cited on p. 6)
2. Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 06 2017. (cited on p. 7)
3. Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D. Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. (cited on pp. 1 and 2)
4. Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *European Conference on Computer Vision (ECCV)*, pages 213–229. Springer, August 2020. (cited on pp. 1, 2, 3, 4, and 5)
5. Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017. (cited on pp. 2 and 4)
6. Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. DeepLab: semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected CRFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4):834–848, 2017. (cited on p. 7)
7. Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. (cited on p. 2)
8. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. (cited on pp. 1, 2, and 7)
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. *The 3rd Large-scale Video Object Segmentation Challenge*, June 2021. (cited on p. 7)
10. Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. (cited on pp. 1, 2, and 3)
11. Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The PASCAL visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. (cited on p. 7)
12. Kirill Gavriluk, Amir Ghodrati, Zhenyang Li, and Cees G. M. Snoek. Actor and action video segmentation from a sentence. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. (cited on pp. 2, 4, 5, 6, 7, and 8)
13. Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from natural language expressions. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *European Conference on Computer Vision (ECCV)*, pages 108–124, Cham, 2016. Springer, Springer International Publishing. (cited on p. 6)
14. Tianrui Hui, Shaofei Huang, Si Liu, Zihan Ding, Guanbin Li, Wenguan Wang, Jizhong Han, and Fei Wang. Collaborative spatial-temporal modeling for language-queried video actor segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4187–4196, June 2021. (cited on pp. 2, 4, 6, and 7)
15. Hueihan Jhuang, Juergen Gall, Silvia Zuffi, Cordelia Schmid, and Michael J. Black. Towards understanding action recognition. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3192–3199, December 2013. (cited on pp. 5 and 7)
16. Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. MDETR – modulated detection for end-to-end multi-modal understanding. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1780–1790, October 2021. (cited on pp. 3 and 5)
17. Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, Mustafa Suleyman, and Andrew Zisserman. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. (cited on pp. 4 and 6)
18. Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017. (cited on p. 2)
19. Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected CRFs with gaussian edge potentials. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems (NIPS)*, volume 24. Curran Associates, Inc., 2011. (cited on p. 4)
20. Harold W. Kuhn. The Hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. (cited on p. 5)
21. 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. (cited on p. 7)
22. Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *IEEE Conference on Computer*

- Vision and Pattern Recognition (CVPR)*, July 2017. (cited on p. 4)
23. Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. (cited on p. 5)
 24. Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision (ECCV)*, pages 740–755. Springer, 2014. (cited on pp. 6 and 7)
 25. 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. (cited on pp. 2, 4, 6, 7, and 8)
 26. 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. (cited on pp. 2 and 7)
 27. 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. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10012–10022, October 2021. (cited on pp. 1, 2, 3, and 4)
 28. Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. (cited on pp. 2, 3, 4, and 6)
 29. Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. (cited on pp. 1 and 7)
 30. Bruce McIntosh, Kevin Duarte, Yogesh S. Rawat, and Mubarak Shah. Visual-textual capsule routing for text-based video segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. (cited on p. 2)
 31. Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-Net: fully convolutional neural networks for volumetric medical image segmentation. In *Fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016. (cited on p. 5)
 32. Ke Ning, Lingxi Xie, Fei Wu, and Qi Tian. Polar relative positional encoding for video-language segmentation. In Christian Bessiere, editor, *Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 948–954. International Joint Conferences on Artificial Intelligence Organization, 7 2020. Main track. (cited on pp. 2 and 4)
 33. Jeffrey Pennington, Richard Socher, and Christopher D. Manning. GloVe: Global vectors for word representation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, Oct. 2014. Association for Computational Linguistics. (cited on p. 7)
 34. 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 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. (cited on p. 6)
 35. Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. (cited on pp. 1 and 2)
 36. Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. (cited on p. 2)
 37. Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. Dynamic routing between capsules. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. (cited on p. 2)
 38. Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019. (cited on p. 7)
 39. Seonguk Seo, Joon-Young Lee, and Bohyung Han. URVOS: Unified referring video object segmentation network with a large-scale benchmark. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *European Conference on Computer Vision (ECCV)*, pages 208–223. Springer, August 2020. (cited on pp. 2, 6, 7, and 8)
 40. Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *European Conference on Computer Vision (ECCV)*, pages 213–229. Springer, August 2020. (cited on p. 4)
 41. Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. (cited on pp. 1, 2, and 3)
 42. Hao Wang, Cheng Deng, Fan Ma, and Yi Yang. Context modulated dynamic networks for actor and action video segmentation with language queries. *AAAI Conference on Artificial Intelligence*, 34(07):12152–12159, Apr. 2020. (cited on p. 2)
 43. 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 *IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. (cited on p. 2)
 44. Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. MaX-DeepLab: end-to-end panoptic segmentation with mask transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5463–5474, 2021. (cited on p. 4)
 45. Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8741–8750, June 2021. (cited on pp. 2, 4, 5, and 7)
 46. Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Péric Cistac, Clara Ma,

- Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online, Oct. 2020. Association for Computational Linguistics. (cited on p. 7)
47. Chenliang Xu, Shao-Hang Hsieh, Caiming Xiong, and Jason J. Corso. Can humans fly? Action understanding with multiple classes of actors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2264–2273, 2015. (cited on p. 5)
 48. Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. YouTube-VOS: sequence-to-sequence video object segmentation. In *European Conference on Computer Vision (ECCV)*, pages 585–601, September 2018. (cited on pp. 6 and 7)
 49. Jianhua Yang, Yan Huang, Kai Niu, Zhanyu Ma, and Liang Wang. Actor and action modular network for text-based video segmentation. *arXiv preprint arXiv:2011.00786*, 2020. (cited on pp. 2 and 6)
 50. Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R. Salakhutdinov, and Quoc V. Le. XLNet: generalized autoregressive pretraining for language understanding. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. (cited on pp. 1 and 2)
 51. Linwei Ye, Mrigank Rochan, Zhi Liu, Xiaoqin Zhang, and Yang Wang. Referring segmentation in images and videos with cross-modal self-attention network. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (01):1–1, jan 2021. (cited on p. 6)
 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. (cited on pp. 1 and 7)