

Exploring Pre-trained Text-to-Video Diffusion Models for Referring Video Object Segmentation

Zixin Zhu¹, Xuelu Feng¹, Dongdong Chen², Junsong Yuan¹, Chunming Qiao¹,
and Gang Hua³

¹University at Buffalo ²Microsoft GenAI ³Dolby Laboratories
{zixinzhu,xuelufen,jsyuan,qiao}@buffalo.edu, {cddlyf,ganghua}@gmail.com

Abstract. In this paper, we explore the visual representations produced from a pre-trained text-to-video (T2V) diffusion model for video understanding tasks. We hypothesize that the latent representation learned from a pretrained generative T2V model encapsulates rich semantics and coherent temporal correspondences, thereby naturally facilitating video understanding. Our hypothesis is validated through the classic referring video object segmentation (R-VOS) task. We introduce a novel framework, termed “VD-IT”, tailored with dedicatedly designed components built upon a fixed pretrained T2V model. Specifically, VD-IT uses textual information as a conditional input, ensuring semantic consistency across time for precise temporal instance matching. It further incorporates image tokens as supplementary textual inputs, enriching the feature set to generate detailed and nuanced masks. Besides, instead of using the standard Gaussian noise, we propose to predict the video-specific noise with an extra noise prediction module, which can help preserve the feature fidelity and elevates segmentation quality. Through extensive experiments, we surprisingly observe that fixed generative T2V diffusion models, unlike commonly used video backbones (e.g., Video Swin Transformer) pretrained with discriminative image/video pre-tasks, exhibit better potential to maintain semantic alignment and temporal consistency. On existing standard benchmarks, our VD-IT achieves highly competitive results, surpassing many existing state-of-the-art methods. The code is available at <https://github.com/buxiangzhiren/VD-IT>.

Keywords: Text-to-Video · Referring Segmentation · Diffusion

1 Introduction

Recent studies on pre-trained text-to-image diffusion models [34,36] have sparked growing interest in various image understanding tasks [20,32,48,49,53]. These studies focus on analyzing and understanding the internal representations for individual images. Specifically, The benefits of pre-trained diffusion models, particularly their ability to provide open vocabulary knowledge and improve semantic differentiation between objects in the same image, are well investigated.

However, the potential of pre-trained text-to-video diffusion models [2,30,42,56] in video understanding tasks remains much less explored. Unlike image

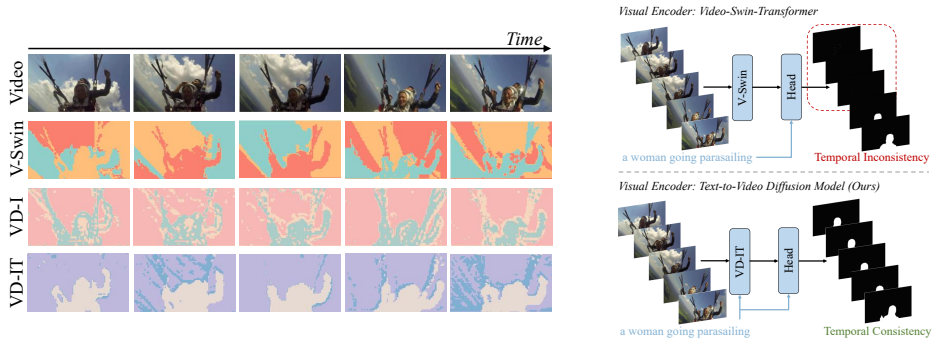


Fig. 1: Left: Analysis of learned features of existing methods that use discriminative backbone (Video Swin Transformer) and our methods (VD-I and VD-IT) that use fixed pretrained generative T2V model. Right: Temporal inconsistency in visual features will subsequently cause temporally inconsistent masks.

understanding, which primarily deals with single images, video comprehension necessitates the consideration of both spatial information within each frame and temporal information across all video frames, making it more challenging. Adhering to the principle of “*what I cannot create, I do not understand*”, we hypothesize that a pre-trained text-to-video diffusion model capable of generating coherent, high-quality video sequences based on text prompts inherently possesses sufficient requisite knowledge to help video understanding tasks.

We validate our hypothesis based on a popular text-to-video model ModelScopeT2V [42] for the classic task of Referring Video Object Segmentation (R-VOS). R-VOS, guided by descriptive natural language prompts, aiming to delineate a specific object within a video sequence in a temporally consistent way. It has a variety of applications, such as video Editing and Video Retrieval. Over recent years, many R-VOS methods [28, 29, 46, 47] have been proposed and achieved state-of-the-art performance. But most video backbone (e.g., Video Swin Transformer [27]) used in these methods are pretrained with discriminative (e.g., classification) pre-tasks, and then need to be jointly finetuned with the text-driven mask decoder. Different from these methods, we explore directly exploiting the latent feature representation of a fixed generative text-to-video diffusion model, offering a fresh perspective in this domain.

In this paper, by analyzing the task characteristic of R-VOS, we propose a new R-VOS framework termed “VD-IT,” built upon the open-sourced text-to-video diffusion model [42]. VD-IT comprises two core components: video feature extraction and a mask decoder responsible for generating final object masks based on the provided text query and extracted video features. To ensure temporal consistency and enhance feature detail for high-quality mask generation, we incorporate two key designs in feature extraction from the text-to-video diffusion model: Text-Guided Image Projection and Video-specific Noise Prediction. Specifically, for Text-Guided Image Projection, we propose leveraging both referring text and visual tokens from each frame to guide the T2V model in producing

the latent feature, rather than solely relying on visual tokens. This approach not only ensures visual feature consistency across time, aiding temporal instance matching but also enriches feature detail for better spatial differentiation. The latter, Noise Prediction, diverges from the standard Gaussian noise commonly applied in diffusion, predicting specific noise patterns during the diffusion forward process to further preserve the fidelity of extracted video features.

We evaluate the effectiveness of VD-IT on four standard R-VOS benchmarks: Ref-Youtube-VOS [38], Ref-DAVIS17 [15], A2D-Sentences [12], and JHMDB-Sentences [12]. The quantitative and qualitative results demonstrate that our VD-IT achieves highly competitive performances and outperforms previous state-of-the-art methods. More interestingly, through in-depth analysis, we observe that the visual features learned in VD-IT exhibit significantly better temporal semantic consistency and spatial smoothness than existing methods [29] that employ discriminatively finetuned video backbone (e.g., Video Swin Transformer). In Figure 1, we present the K-Means clustering results of the visual features coming from our methods (VD-IT and another variant VD-I that only uses visual tokens as guide the latent feature extraction) and one existing method [29] that uses fine-tuned Video Swin Transformer (V-Swin). It can be seen that the V-Swin feature changes significantly due to variations in lighting, affecting feature consistency across frames and complicating segmentation tasks. In contrast, VD-I and VD-IT maintain strong temporal consistency, with VD-IT further reducing internal noise through text guidance.

We believe that the superior temporal consistency observed in the T2V diffusion models can be attributed to the use of global text prompts as conditional inputs, which guides the generation of semantically consistent image frames throughout a video. This conditioning ensures that the semantic attributes across frames are aligned, enhancing temporal coherence. Additionally, the intrinsic denoising capability of the diffusion model bolsters its robustness against external disturbances like illumination changes and camera movement. By leveraging these mechanisms, the diffusion features exhibit enhanced temporal semantic consistency, proving advantageous for video understanding tasks. This idea is supported by analysis in following Section 4.1 and Section 5.2. It is crucial to clarify that our contribution involves more than merely substituting a visual backbone from V-Swin to a text-to-video diffusion model for enhanced performance. In fact, in the subsequent analysis section, we demonstrate that without our dedicated designs, simply utilizing the initial T2V diffusion model cannot achieve superior performances compared to existing methods. The main contributions of this work are:

- To our knowledge, we are the first to explore the video prior (temporal consistency) acquired by pre-trained text-to-video diffusion models for video understanding task. Our findings reveal that, compared to conventional discriminatively fine-tuned video encoders, pretrained text-to-video diffusion models exhibit superior temporal consistency.

- We propose a new R-VOS framework VD-IT with several innovative designs, which significantly improves the extracted feature quality and boost the final performances.
- Extensive experiments and analysis on Ref-Youtube-VOS [38], Ref-DAVIS17 [15], A2D-Sentences and JHMDB-Sentences [12] show that VD-IT can achieve super competitive results and outperform previous methods.

2 Related Work

Diffusion Model. In recent years, diffusion models have exhibited remarkable efficacy across various generative tasks, markedly enhancing generation quality, particularly in text-to-image [11, 13, 34, 36, 37, 43, 55] and text-to-video generation [1, 2, 30, 56]. Moreover, there have been recent endeavors to design specialized diffusion models tailored for diverse image and video understanding tasks [4, 6, 39, 54]. While these methods approach understanding tasks from a generative perspective, other works [20, 32, 48, 49, 53] directly leverage feature priors learned in pretrained text-to-image models for image understanding, showcasing the discriminative capabilities of latent feature representations learned in pretrained text-to-image models. However, to the best of our knowledge, our study marks the first exploration of temporal-consistency priors in pre-trained text-to-video diffusion models for video understanding tasks.

Referring Video Object Segmentation. In contrast to traditional semi-supervised video object segmentation methods [31, 57], relying on ground-truth mask annotations in the initial frame for target identification, R-VOS aims to delineate the object based on textual queries, demanding a nuanced understanding of cross-modal sources—vision and language. Previous R-VOS approaches [23, 29, 46] leverage deep neural networks for vision-and-language interaction, enhancing visual features with linguistic information for pixel-level segmentation. For instance, [23] introduces a top-down R-VOS approach addressing limitations of bottom-up strategies, while SgMg [29] performs direct segmentation on encoded features, refining masks to combat drift. Diverging from discriminative feature representation learning in these methods, this paper studies a new perspective, i.e., directly leverage the feature representation learned in generative text-to-video model for video understanding. And the technique advancements in previous methods (e.g., better mask decoder design) are complementary to our method.

3 Preliminary of Text-to-Video Diffusion Model

In this section, we provide a brief overview of T2V diffusion models and describe how we extract features for video understanding tasks.

Overview. Given a text prompt, T2V diffusion models is designed to generate a coherent video based on the semantic meaning of the prompt. The overall architecture often consists of three parts: an encoder \mathcal{E} , a decoder \mathcal{D} that is

derived from VQGAN/VQVAE [10], translating image pixels into/back from latent spaces respectively, and a denoising U-Net \mathcal{U} functioning in the latent space. The training procedure of T2V diffusion models is a forward-backward process. Specifically, the input video is first mapped to a latent vector. Then the latent vector is distorted by adding a small level of Gaussian noise. Meanwhile, the text prompt is encoded into a text embedding with a pre-trained text encoder [33]. The denoising U-Net is optimised to predict the noise and recover the initial video signal, conditioned on the text embedding and the noised video. In the inference stage, the denoising U-Net takes video-shaped pure Gaussian noise and the text embedding of a user provided description as input, and progressively de-noises it to a realistic video through many iterations.

Visual Representation Extraction. The denoising U-Net architecture is structured with downsampling, spatio-temporal, and upsampling blocks. It integrates cross-attention [40] mechanisms that align text embeddings with U-Net features across spatial dimensions, alongside self-attention [40] mechanisms that process these features through the temporal dimension for enhanced context comprehension. Thus the feature maps produced by each U-Net block encapsulate semantically aligned and temporally consistent attributes. To generate the latent features of the denoising U-Net for video understanding, it requires two inputs: the noise-added video and the text embedding. Following [49, 53], we introduce a single step of low-intensity noise to the video, aiming to preserve the original video signal as much as possible. However, our approach diverges from them [49, 53] for image segmentation, where image projection serves as an implicit text embedding to circumvent the necessity for explicitly captioned image data. In our case, we employ both referring text and image projection, where referring text is to ensure temporal consistency across visual features while image projection is utilized as a supplementary condition to enhance the richness and detail of these features. Furthermore, in contrast to the referenced methods [49, 53] that add pure Gaussian noise into the image, we leverage an extra noise prediction module that learns video-specific noise, which can help retain the fidelity in the generated video feature.

4 Text-to-Video Diffusion Model for R-VOS

In this paper, we use the referring video segmentation task as an example to explore the potential of learned representation of generative text-to-video models for video understanding tasks. Considering the pretrained text-to-video model ModelScopeT2V [42] was open-sourced at the beginning of this project, we use it by default in all following experiments. Given a video clip $\mathcal{I} = \{I_t\}_{t=1}^T$ comprising T frames, R-VOS task aims to generate T -frame binary segmentation masks for the referred object $\mathcal{S} = \{s_t\}_{t=1}^T, s_t \in \mathbb{R}^{H \times W}$, guided by a referring expression $\mathcal{E} = \{e_l\}_{l=1}^L$ consisting of L words, in an temporally coherent way.

By analyzing the task characteristics of R-VOS, we propose a new framework named VD-IT, constructed upon ModelScopeT2V. Illustrated in Figure 2, VD-IT initially generates visual features utilizing the fixed pretrained text-to-

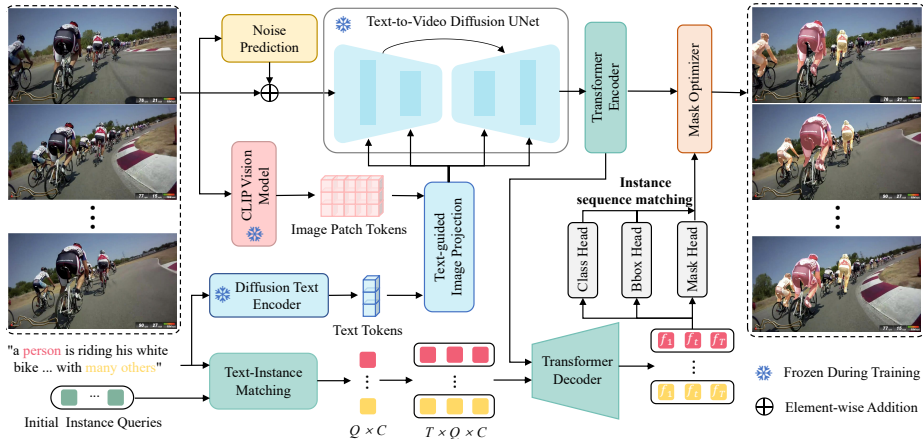


Fig. 2: This framework comprises two core components: visual feature extraction and the mask segmentation head. The feature extraction progresses through three phases: (1) prompt generation via text-guided image projection conditions the text-to-video diffusion model; (2), predicted noise is applied to the video; (3) this noisy video and prompt are processed by a diffusion U-Net for visual feature extraction. The segmentation head generates instance queries from the text and merges them with the U-Net’s features to create the final masks.

video diffusion model, which are subsequently fed into the segmentation head for generating final object masks. As detailed earlier, VD-IT operates as a T2V diffusion model and requires two inputs for extracting the visual representation: a prompt as condition embedding and a noise-added video clip. In the below section, we will describe the corresponding designs in details.

4.1 Video Feature Extraction

As introduced in Section 3, extracting diffusion features requires two inputs: the text prompt embedding as condition and the noise-added video. We use Text-Guided Image Projection to generate condition prompt embedding and Noise Prediction to predict the added noise. The methodologies for these two components are elaborated below.

Text-Guided Image Projection. We investigate three distinct approaches for condition prompt embedding. The first straightforward way is directly using the referring text as the conditional prompt, denoted as “VD-T”. However, we observed that the resulting segmentation masks often lacked fine-grained details. Illustrated in Figure 3, despite identifying the target object, the mask quality around the object boundary is not precise enough. We hypothesize that the text prompt serving as a generalized condition and fails to furnish detailed instance-specific information of the U-Net feature for each frame. To validate our guess, we visualize the low-level features of VD-T in Figure 3 and show that these features indeed lack detail. Inspired by [49, 53], we further try using the CLIP

vision model to extract the visual tokens of each frame followed by a learnable projection layer as an implicit text embedding, termed ‘‘VD-I’’. Through further visualization, we find it can preserve the fine-grained details in low-level visual features. However, it will cause some confusion in differentiating object instances in high-level features. As demonstrated in Figure 3, the final mask erroneously incorporates elements of the background larger dog. We guess this confusion stems from the ignorance of referring text in feature extraction, thus introducing semantic noise into the features and increasing the risk of inaccurately including multiple instances during segmentation.

Based on the above analysis, we propose to use both referring text tokens and visual tokens as prompts (our ‘‘VD-IT’’). As shown in Figure 3, VD-IT exhibits richer image details in its low-level features compared to VD-T, which leads to more precise mask boundaries. For high-level features, VD-IT demonstrates better temporally aligned semantic visual features with the query text compared to VD-I, enabling accurate temporal matching of subjects, such as a puppy.

More specifically, to extract visual tokens from video frames, we extract vectors using the CLIP [33] vision model to get embedding vectors $\{p_{v,t}\}_{t=1}^T, p_{v,t} \in \mathbb{R}^{N_p \times C}$, where N_p denotes the number of tokens and C denotes the channel dimension. For referring text tokens, we utilize the text encoder provided by the text-to-video diffusion model to get the vectors $p_e \in \mathbb{R}^{L \times C}$. We use text token vectors guided image token projection as final prompt, which is formulated by

$$p_{ve,t} = \text{MLP}\left(p_e + \text{SoftMax}\left(\frac{p_e W^Q (p_{v,t} W^K)^T}{\sqrt{d_k}}\right) p_{v,t} W^V\right), \quad (1)$$

where $W^Q, W^K, W^V \in \mathbb{R}^{C \times d_k}$ are learnable parameters, d_k is the number of channels and MLP denotes Multi-Layer Perception. The text tokens play the role of *query* in attention mechanism, and image tokens play the role of *key* and *value*. Therefore, the final prompt combining image and text tokens ensures that the text-to-video diffusion model generates features with both temporal semantic consistency and detailed accuracy.

Noise Prediction. As described in Section 3, for a video clip $\mathcal{I} = \{I_t\}_{t=1}^T$, a text-to-video diffusion model first projects the clip into a latent space using an autoencoder. Then, it proceeds with the diffusion forward pass based on the latent vectors $\mathcal{F}_o = \{f_{o,t}\}_{t=1}^T, f_o \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times 4}$. During the forward process, the diffusion model will gradually introduce Gaussian noise to the original video signals \mathcal{F}_o based on the *step*. The higher the *step*, the more intense the noise. To maximize the preservation of the original video’s visual signals in VD-IT, we set *step* = 0, thereby adding the least amount of noise. Moreover, to avoid blurring key details by using Gaussian noise like [49, 53], we propose to predict video-correlated noise in place of independent Gaussian noise to retain more details in the features. Specifically, we first send latent vectors \mathcal{F}_o into a convolution layer and get $\mathcal{F}_n = \{f_{n,t}\}_{t=1}^T, f_n \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times 4}$. The prediction of normalized noise $\mathcal{N} = \{n_t\}_{t=1}^T, n_t \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times 4}$ is formulated by

$$n_t = (f_{n,t} W^N - \mu(f_{n,t} W^N)) / (\sigma(f_{n,t} W^N) + \epsilon) \quad (2)$$




Fig. 3: “VD-I” denotes Image-conditioned Video Diffusion based on the video clip. “VD-T” denotes Text-conditioned Video Diffusion based on the referring text, *i.e.*, expression. “VD-IT” denotes Image-Text-conditioned Video Diffusion based on both video clip and referring text.

where μ and σ represent the mean and standard deviation, respectively, and $W^N \in \mathbb{R}^{4 \times 4}$ is the trainable weight. ϵ is a small constant to prevent division by zero and set as $1e-5$ in our experiments by default.

With the prompts $p_{ve,t}$ and predicted noise n_t fed into the pretrained text-to-video model, we extract side outputs from each intermediate layer of the diffusion U-Net as visual features \mathcal{F}_v^D . The formulation is

$$\mathcal{F}_v^D = \text{U-Net}(\alpha_0 \mathcal{F}_o + (1 - \alpha_0) \mathcal{N}, \{p_{ve,t}\}_{t=1}^T), \quad (3)$$

where $1 - \alpha_0$ denotes the noise strength of the 0-th step, determined by the noise schedule inherent to the diffusion model. \mathcal{F}_v^D includes 3-level visual features downsampled by factors of $8 \times$ to $32 \times$. By incorporating features downsampled by $4 \times$ from the intermediate layer of the autoencoder, we construct a 4-level hierarchy of visual features \mathcal{F}_v , with downsampling rates from $4 \times$ to $32 \times$.

4.2 Mask Prediction Head

After obtaining the video feature representations, we follow the widely used query-based segmentation model design and propose a simple mask prediction head that first extracts instance queries from referring text tokens and then fuse these queries with visual features to produce the final segmentation masks.

Text-Instance Matching. The mask prediction head first extracts instance queries from referring text tokens. Here, we first have Q learnable query instance vectors $q_o \in \mathbb{R}^{Q \times C}$ which are randomly initialized, where Q represents the number of instance queries. These instance vectors undergo cross-attention with word vectors $\mathcal{F}_e = \{f_{e,l}\}_{l=1}^L, f_{e,l} \in \mathbb{R}^C$ extracted from the referring text by RoBERTa [26]. In detail, q_o will serve as the *query*, while \mathcal{F}_e serves as both the *key* and *value*. The updated query instance vectors $q_e \in \mathbb{R}^{Q \times C}$ then represent the different objects described within the referring expression.

Vision-Language Fusion. Then the segmentation head fuses instance queries and visual features to output final masks. Following the common practice [3, 29,

47], we feed the visual features \mathcal{F}_v into a Deformable Transformer [58] encoder to deal with multi-scale features. Next, we use them as the *key* and *value* in a Deformable Transformer decoder. Concurrently, instance query vectors $q_e \in \mathbb{R}^{Q \times C}$ are replicated T times along the temporal dimension and introduced into the transformer decoder as *query*, resulting in the extraction of cross-modality features $\mathcal{F}_{ve} \in \mathbb{R}^{T \times \frac{H}{8} \times \frac{W}{8} \times C_d}$.

Segmentation Head. With the cross-modality features \mathcal{F}_{ve} , the segmentation head sends it into three sub-heads: bounding box head, classification head and mask head. We configure the bounding box head as a 3-layer feed-forward network (FFN) with ReLU activation for all except the final layer, while the classification head consists of a linear layer. These two heads process \mathcal{F}_{ve} , to produce final bounding box predictions, $\mathcal{B} \in \mathbb{R}^{T \times Q \times 4}$ and classification predictions, $\mathcal{C} \in \mathbb{R}^{T \times Q}$. The classification head functions as a binary classifier, determining whether each frame contains the object of interest. Concurrently, dynamic convolution is performed between \mathcal{F}_{ve} and \mathcal{F}_v to produce the segmented masks $\mathcal{M}_o \in \mathbb{R}^{T \times Q \times \frac{H}{8} \times \frac{W}{8}}$. Finally, the predicted masks are enlarged $8 \times$ through four times of upsampling to reach final masks $\mathcal{M} \in \mathbb{R}^{T \times Q \times H \times W}$. Before the first and third rounds of upsampling, they are combined with visual features \mathcal{F}_v to achieve high-quality masks.

To accurately identify each referred object in videos, we use Hungarian algorithm [16] to select the most suitable match from Q instance queries. For the matched instance, We adopt the same training losses and weights as used in [3, 29, 47] for a fair comparison. Specifically, we use Dice loss [19] and Focal loss [24] for mask \mathcal{M} , Focal loss [24] for confidence scores \mathcal{S} , and L1 and GIoU [35] loss for bounding boxes \mathcal{B} . In inference, we choose the instance query which has the high confidence scores \mathcal{S} as our final results.

5 Experiments

Datasets and Evaluation Metrics. We conduct evaluations on four popular R-VOS benchmarks. Ref-Youtube-VOS [38] is the largest, with 3,978 videos and 15,000 language descriptions. Ref-DAVIS17 [15] includes descriptions for specific objects in 90 videos. A2D-Sentences [12] has 3,782 videos and each video is annotated with 3-5 frames, totaling over 6,600 annotations. JHMDB-Sentences [12] has 928 videos and corresponding descriptions, covering 21 action classes. On Ref-Youtube-VOS and Ref-DAVIS17, the performance is measured using standard metrics: region similarity (\mathcal{J}), contour accuracy (\mathcal{F}), and their mean ($\mathcal{J\&F}$). We also report the frames per second (FPS) achieved on Ref-Youtube-VOS. For A2D-Sentences and JHMDB-Sentences, we employ Overall IoU, Mean IoU, and mAP within interval of 0.50 to 0.95, incremented by 0.05.

Training Details. Following common practices, we train our models on the Ref-YouTube-VOS training set and test them on the validation splits of Ref-YouTube-VOS and Ref-DAVIS17. Since some methods [29, 47] report their results after pre-training on Ref-COCO [52], we also report our results on this setting followed by fine-tuning on Ref-YouTube-VOS. The pre-trained model is also fine-

Table 1: Comparison with the state-of-the-art methods on Ref-YouTube-VOS and Ref-DAVIS17. * denotes that we run the official codes to get the results.

Method	Backbone	Ref-YouTube-VOS				Ref-DAVIS17		
		$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	FPS	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
CMSA [51]	ResNet-50	36.4	34.8	38.1	-	40.2	36.9	43.5
URVOS [38]	ResNet-50	47.2	45.3	49.2	-	51.5	47.3	56.0
CMPC-V [25]	I3D	47.5	45.6	49.3	-	-	-	-
PMINet [9]	ResNeSt-101	53.0	51.5	54.5	-	-	-	-
YOFO [17]	ResNet-50	48.6	47.5	49.7	10	53.3	48.8	57.8
LBDT [8]	ResNet-50	49.4	48.2	50.6	-	54.3	-	-
MLRL [45]	ResNet-50	49.7	48.4	51.0	-	52.8	50.0	55.4
MTTR [3]	Video-Swin-T	55.3	54.0	56.6	-	-	-	-
MANet [5]	Video-Swin-T	55.6	54.8	56.5	-	-	-	-
ReferFormer [47]	Video-Swin-T	56.0	54.8	57.3	50	-	-	-
SgMg [29]	Video-Swin-T	58.9	57.7	60.0	65	56.7	53.3	60.0
SgMg* [29]	Video-Swin-B	61.6	59.7	63.5	40	-	-	-
VD-IT (Ours)	Video Diffusion	64.8	63.1	66.6	21	63.0	59.9	66.1
Pre-training with RefCOCO+/g								
ReferFormer [47]	Video-Swin-T	59.4	58.0	60.9	50	59.6	56.5	62.7
SgMg [29]	Video-Swin-T	62.0	60.4	63.5	65	61.9	59.0	64.8
ReferFormer [47]	Video-Swin-B	62.9	61.3	64.6	33	61.1	58.1	64.1
OnlineRefer [46]	Video-Swin-B	62.9	61.0	64.7	-	62.4	59.1	65.6
SgMg [29]	Video-Swin-B	65.7	63.9	67.4	40	63.3	60.6	66.0
VD-IT (Ours)	Video Diffusion	66.5	64.4	68.5	21	69.4	66.2	72.6

tuned on the A2D-Sentences training set, and we then evaluate its performance on the validation splits of A2D-Sentences and JHMDB-Sentences. The number Q of queries across all datasets is set to 5. The training of the models utilizes two NVIDIA A100 GPUs, processing 5 frames per clip over the course of 9 epochs. To accommodate model requirements, all frames are adjusted to ensure the longest side to be 640 pixels. Additional training details can be found in the supplementary materials.

5.1 Comparison with State-of-the-Art Methods

Ref-YouTube-VOS and Ref-DAVIS17. We compare our method with other state-of-the-art methods in Table 1. Specifically, on Ref-YouTube-VOS and Ref-DAVIS17 datasets, VD-IT outperforms state-of-the-art methods.

A2D-Sentences and JHMDB-Sentences. We further evaluate our method on the A2D-Sentences and JHMDB-Sentences datasets and compare its performance with other state-of-the-art methods in Table 2. In alignment with methodologies from recent studies [29, 47], the approach involves pre-training on the RefCOCO+/g datasets before fine-tuning on A2D-Sentences, while JHMDB-Sentences is exclusively used for evaluation purposes. It shows that VD-IT outperforms state-of-the-art methods by 2.9 and 1.5 in mAP for A2D-Sentences and JHMDB-Sentences, respectively.

Table 2: Quantitative comparison to state-of-the-art R-VOS methods on A2D-Sentences and JHMDB-Sentences.

Method	Backbone	A2D-Sentences			JHMDB-Sentences		
		mAP	Overall	Mean	mAP	Overall	Mean
Hu et al. [14]	VGG-16	13.2	47.4	35.0	17.8	54.6	52.8
Gavrilyuk et al. [12]	I3D	19.8	53.6	42.1	23.3	54.1	54.2
ACAN [41]	I3D	27.4	60.1	49.0	28.9	57.6	58.4
CMPC-V [25]	I3D	40.4	65.3	57.3	34.2	61.6	61.7
ClawCraneNet [22]	ResNet-50/101	-	63.1	59.9	-	64.4	65.6
MTTR [3]	Video-Swin-T	46.1	72.0	64.0	39.2	70.1	69.8
ReferFormer [47]	Video-Swin-T	52.8	77.6	69.6	42.2	71.9	71.0
SgMg [29]	Video-Swin-T	56.1	78.0	70.4	44.4	72.8	71.7
ReferFormer [47]	Video-Swin-B	55.0	78.6	70.3	43.7	73.0	71.8
SgMg [29]	Video-Swin-B	58.5	79.9	72.0	45.0	73.7	72.5
VD-IT (Ours)	Video Diffusion	61.4	81.5	73.2	46.5	74.4	73.4

Table 3: Ablation study on the components of VD-IT on Ref-YouTube-VOS. “NP” denotes Noise Prediction.

Components		Performance			
Image-Cond	Text-Cond	NP	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
✓			59.7	57.9	61.6
	✓		61.9	60.1	63.7
✓	✓		63.8	62.0	65.5
✓	✓	✓	64.8	63.1	66.6

Table 4: Quantitative evaluation on the validation split of RefCOCO+/g. Overall IoU is adopted as the evaluation metric.

Method	Ref.	Ref.+	Ref.g
MaIL [21]	70.1	62.2	62.5
CRIS [44]	70.5	62.3	59.9
RefTR [18]	70.6	-	-
LAVT [50]	72.7	62.1	61.2
VLT [7]	73.0	63.5	63.5
SgMg [29]	76.3	66.4	70.0
VD-IT (Ours)	76.7	66.5	70.3

5.2 Ablation Study

In this section, we perform extensive ablation studies on Ref-YouTube-VOS to study the effect of different components in our model. “Image-Cond” and “Text-Cond” shown in Table 3 denotes only using visual token or text token in VD-I and VD-T, while “NP” denotes Noise Prediction.

Component Analysis. As described in Section 4.1, we have tried three different settings for the prompt embedding to extract the feature representation and the corresponding ablation results are shown in Table 3. By comparing the results with the state-of-the-art results shown in Table 1, we can observe that just using visual tokens as the condition embedding like [49, 53] cannot perform as well as SgMg [29] that uses the discriminative Video Swin model as the backbone. This also indicates that our state-of-the-art (SOTA) performance stems neither from increased capacity nor from more extensive data exposure in the text-to-video diffusion model. By considering temporal semantics provided by the referring text, VD-T achieves a stronger performance with 2.2 point improvement. This supports our initial hypothesis that referring text plays a crucial role in en-

Table 5: Temporal consistency ablation. "IoU diff." denotes the IoU variation between frames and "Inter." denotes the frame interval. Averages are calculated on Ref-YouTube-VOS.

Method	IoU Diff.	
	Inter. 1	Inter. 5
SgMg [29]	7.24	11.15
VD-I	6.52	9.43
VD-T	5.41	8.25
VD-IT	5.19	7.89

Table 6: Ablation study shows the impact of fine-grained features on mask quality in Ref-YouTube-VOS, reporting the ratio of high-quality (IoU > 0.9) to moderate-quality (IoU > 0.5) masks.

Components			Ratio = $\frac{\text{IoU}>0.9}{\text{IoU}>0.5}$
Image-Cond	Text-Cond	NP	
✓			47.26%
	✓		44.54%
✓	✓		46.71%
✓	✓	✓	47.36%

hancing visual features. The oversight of referring text in the visual encoders of previous methods is not optimal. However, relying solely on referring text for conditioning captures only temporal semantic consistency without extracting detailed visual features. By leveraging image tokens for detail alongside text tokens for semantic guidance, our final model sees a large performance boost. Besides the prompt embedding choice, in the forward process, we choose to predict video-specific noise instead of using normal Gaussian noise, further boosting the performance by 1.0 point.

Temporal Consistency Analysis. To demonstrate that our VD-IT offer better temporal consistency than V-Swin as backbones, Table 5 presents the difference in IoU between two frames, *i.e.*, frame interval is 1 and 5. A smaller difference in IoU indicates improved temporal coherence. When the visual encoder does not incorporate referring text, the distinction between VD-I and SgMg becomes less significant. The minor advantage observed with VD-I is primarily due to its enhanced denoising capabilities, which improve resilience to various types of noise such as changes in lighting and camera movements. Consequently, the fluctuation in IoU scores for VD-I shows a reduction of about 0.7 compared to SgMg, reflecting a steadier performance across diverse scenarios. By incorporating referring text within the visual encoder to guide feature extraction, VD-T and VD-IT exhibit much improved temporal consistency. Their fluctuation in performance measures is around 2 point less than that observed with SgMg. This reduction in variability demonstrates that VD-IT significantly enhances performance through better maintenance of temporal consistency.

Enhancing Mask Quality. To further study our design in enhancing the model to generate high-quality masks, we use the ratio of masks with IoU greater than 0.9 to those with IoU above 0.5 as the metric. A higher ratio reflects a better capacity for generating more high-quality masks. As shown in Table 6, VD-I, when utilizing visual tokens as prompts, provides additional details to visual features compared to VD-T, which relies solely on referring text for prompts. Compared to only using visual tokens, combining referring text tokens with visual tokens will slightly worsen the high-quality mask generation. But by using our predicting noise module, our VD-IT can extract finer features and produce




Fig. 4: Temporal Semantic Consistency. Averaged over 1,000 samples from Ref-Youtube-VOS, the cosine similarity between the Region of Interest (RoI) features of the initial frame and the following seven frames is reported.




Fig. 5: Robustness against light noise. We modify the brightness of various frames randomly and compare the IoU of segmentation results under changing lighting conditions. The results are reported on Ref-Youtube-VOS.

slightly higher-quality masks. This ablation analysis highlights the importance of integrating detailed visual information via image tokens and the improvements from noise prediction in refining feature extraction for superior mask generation.

Discussion on Feature Behavior. Previous experiments indicate the good temporal consistency of the T2V model through the segmentation results. However, we did not confirm whether it was indeed the prompt tuning that enabled the model to learn features with strong temporal semantic consistency, leading to temporally coherent outcomes. Therefore, we further test the temporal semantic similarity of the features. Figure 4 shows that both VD-I and VD-IT have a slower semantic decline, a phenomenon that visual features would gradually drift over consecutive video frames, compared to the V-Swin, indicating better temporal stability. Importantly, VD-IT benefits from using referring text for semantic guidance, which results in higher semantic similarity than VD-I. We hypothesize that the slower decline of VD-I and VD-IT compared to V-Swin is due to their stronger robustness to environmental variations in videos. Therefore, we evaluated the model performance under varying lighting conditions, as depicted in Figure 5. To make the results more convincing, we select SgMg [29], a state-of-the-art method that previously utilized V-Swin as the visual encoder, for comparison with VD-I and VD-IT. Figure 5 shows that VD-I and VD-IT are more resistant to lighting noise, explaining their slower rate of semantic decline.

Referring Image Segmentation Results. We implemented VD-IT for referring image segmentation directly, comparing it to the current state-of-the-art methods on the RefCOCO/+g datasets, without making any changes to the architecture. According to the results shown in Table 4, VD-IT achieves strong performance on all three datasets. Nonetheless, its performance is quite similar to that of SgMg, which highlights our progress on video R-VOS datasets due to our focus on maintaining temporal consistency, not just on improving single-frame image mask quality. If improving single-frame mask quality were



Fig. 6: Qualitative comparison of our method with others. VD-IT has better temporal consistency.

our main achievement, we would expect much higher IoU scores on RefCOCO datasets compared to SgMg. This suggests our method’s strength lies in enhancing performance through attention to time-based continuity rather than simply boosting the quality of individual image masks.

Visualization Results. In Figure 6, we further provide some visual comparison results with the state-of-the-art method SgMg [29]. From the visual results, we can clearly see that our proposed VD-IT can segment and track the referred object in a more temporally consistent way than the baseline method.

6 Conclusion

This paper presents a pioneering exploration into leveraging feature representations learned in pretrained generative text-to-video diffusion models for video understanding tasks. We take the classical referring video object segmentation task as the testbed, and design a new R-VOS framework with some dedicated designs. Through comprehensive experimentation, we validate the effectiveness of our proposed R-VOS framework, showcasing superior performance over existing state-of-the-art methods. Our findings underscore the capacity of latent representations learned in generative text-to-video models to encapsulate rich semantics and coherent temporal correspondence. Consequently, the segmentation results based on these latent features exhibit markedly improved temporal consistency compared to methods utilizing feature representations of video backbones fine-tuned in a discriminative manner. We envision that our pioneering exploration will inspire further insightful research in this direction, propelling the unification of generative and discriminative models toward new horizons.

References

1. Blattmann, A., Dockhorn, T., Kulal, S., Mendeleevitch, D., Kilian, M., Lorenz, D., Levi, Y., English, Z., Voleti, V., Letts, A., et al.: Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127 (2023)
2. Blattmann, A., Rombach, R., Ling, H., Dockhorn, T., Kim, S.W., Fidler, S., Kreis, K.: Align your latents: High-resolution video synthesis with latent diffusion models. In: CVPR. pp. 22563–22575 (2023)
3. Botach, A., Zheltonozhskii, E., Baskin, C.: End-to-end referring video object segmentation with multimodal transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4985–4995 (2022)
4. Chen, S., Sun, P., Song, Y., Luo, P.: Diffusiondet: Diffusion model for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19830–19843 (2023)
5. Chen, W., Hong, D., Qi, Y., Han, Z., Wang, S., Qing, L., Huang, Q., Li, G.: Multi-attention network for compressed video referring object segmentation. In: ACM MM. pp. 4416–4425 (2022)
6. Chen, Z., Gao, R., Xiang, T.Z., Lin, F.: Diffusion model for camouflaged object detection. arXiv preprint arXiv:2308.00303 (2023)
7. Ding, H., Liu, C., Wang, S., Jiang, X.: Vlt: Vision-language transformer and query generation for referring segmentation. IEEE TPAMI (2022)
8. Ding, Z., Hui, T., Huang, J., Wei, X., Han, J., Liu, S.: Language-bridged spatial-temporal interaction for referring video object segmentation. In: CVPR. pp. 4964–4973 (2022)
9. Ding, Z., Hui, T., Huang, S., Liu, S., Luo, X., Huang, J., Wei, X.: Progressive multimodal interaction network for referring video object segmentation. The 3rd Large-scale Video Object Segmentation Challenge **8** (2021)
10. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: CVPR. pp. 12873–12883 (2021)
11. Fan, W.C., Chen, Y.C., Chen, D., Cheng, Y., Yuan, L., Wang, Y.C.F.: Frido: Feature pyramid diffusion for complex scene image synthesis. In: Thirty-Seventh AAAI Conference on Artificial Intelligence (AAAI) (2023)
12. Gavriluk, K., Ghodrati, A., Li, Z., Snoek, C.G.: Actor and action video segmentation from a sentence. In: CVPR. pp. 5958–5966 (2018)
13. Gu, S., Chen, D., Bao, J., Wen, F., Zhang, B., Chen, D., Yuan, L., Guo, B.: Vector quantized diffusion model for text-to-image synthesis. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2022) (2022)
14. Hu, R., Rohrbach, M., Darrell, T.: Segmentation from natural language expressions. In: ECCV. pp. 108–124. Springer (2016)
15. Khoreva, A., Rohrbach, A., Schiele, B.: Video object segmentation with language referring expressions. In: ACCV. pp. 123–141. Springer (2019)
16. Kuhn, H.W.: The hungarian method for the assignment problem. Naval Research Logistics **52** (1955)
17. Li, D., Li, R., Wang, L., Wang, Y., Qi, J., Zhang, L., Liu, T., Xu, Q., Lu, H.: You only infer once: Cross-modal meta-transfer for referring video object segmentation. In: AAAI. pp. 1297–1305 (2022)
18. Li, M., Sigal, L.: Referring transformer: A one-step approach to multi-task visual grounding. NeurIPS **34**, 19652–19664 (2021)

19. Li, X., Sun, X., Meng, Y., Liang, J., Wu, F., Li, J.: Dice loss for data-imbalanced nlp tasks. arXiv preprint arXiv:1911.02855 (2019)
20. Li, Z., Zhou, Q., Zhang, X., Zhang, Y., Wang, Y., Xie, W.: Open-vocabulary object segmentation with diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7667–7676 (2023)
21. Li, Z., Wang, M., Mei, J., Liu, Y.: Mail: A unified mask-image-language trimodal network for referring image segmentation. arXiv preprint arXiv:2111.10747 (2021)
22. Liang, C., Wu, Y., Luo, Y., Yang, Y.: Clawcranenet: Leveraging object-level relation for text-based video segmentation. arXiv preprint arXiv:2103.10702 (2021)
23. Liang, C., Wu, Y., Zhou, T., Wang, W., Yang, Z., Wei, Y., Yang, Y.: Rethinking cross-modal interaction from a top-down perspective for referring video object segmentation. arXiv preprint arXiv:2106.01061 (2021)
24. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV. pp. 2980–2988 (2017)
25. Liu, S., Hui, T., Huang, S., Wei, Y., Li, B., Li, G.: Cross-modal progressive comprehension for referring segmentation. IEEE TPAMI **44**(9), 4761–4775 (2021)
26. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
27. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: CVPR. pp. 3202–3211 (2022)
28. Mei, J., Piergiovanni, A., Hwang, J.N., Li, W.: Slvp: Self-supervised language-video pre-training for referring video object segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 507–517 (2024)
29. Miao, B., Bennamoun, M., Gao, Y., Mian, A.: Spectrum-guided multi-granularity referring video object segmentation. In: ICCV. pp. 920–930 (2023)
30. Molad, E., Horwitz, E., Valevski, D., Acha, A.R., Matias, Y., Pritch, Y., Leviathan, Y., Hoshen, Y.: Dreamix: Video diffusion models are general video editors. arXiv preprint arXiv:2302.01329 (2023)
31. Park, H., Yoo, J., Jeong, S., Venkatesh, G., Kwak, N.: Learning dynamic network using a reuse gate function in semi-supervised video object segmentation. In: CVPR. pp. 8405–8414 (2021)
32. Pnvr, K., Singh, B., Ghosh, P., Siddiquie, B., Jacobs, D.: Ld-znet: A latent diffusion approach for text-based image segmentation. In: ICCV. pp. 4157–4168 (2023)
33. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML. pp. 8748–8763. PMLR (2021)
34. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with clip latents. arXiv preprint arXiv:2204.06125 **1**(2), 3 (2022)
35. Rezatofghi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. In: CVPR. pp. 658–666 (2019)
36. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR. pp. 10684–10695 (2022)
37. Saharia, C., Chan, W., Chang, H., Lee, C., Ho, J., Salimans, T., Fleet, D., Norouzi, M.: Palette: Image-to-image diffusion models. In: ACM SIGGRAPH. pp. 1–10 (2022)
38. Seo, S., Lee, J.Y., Han, B.: Urvos: Unified referring video object segmentation network with a large-scale benchmark. In: ECCV. pp. 208–223. Springer (2020)

39. Tur, A.O., Dall’Asen, N., Beyan, C., Ricci, E.: Exploring diffusion models for unsupervised video anomaly detection. In: ICIP. pp. 2540–2544. IEEE (2023)
40. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *NeurIPS* **30** (2017)
41. Wang, H., Deng, C., Yan, J., Tao, D.: Asymmetric cross-guided attention network for actor and action video segmentation from natural language query. In: ICCV. pp. 3939–3948 (2019)
42. Wang, J., Yuan, H., Chen, D., Zhang, Y., Wang, X., Zhang, S.: Modelscope text-to-video technical report. arXiv preprint arXiv:2308.06571 (2023)
43. Wang, W., Bao, J., Zhou, W., Chen, D., Chen, D., Yuan, L., Li, H.: Semantic image synthesis via diffusion models. arXiv preprint arXiv:2207.00050 (2022)
44. Wang, Z., Lu, Y., Li, Q., Tao, X., Guo, Y., Gong, M., Liu, T.: Cris: Clip-driven referring image segmentation. In: CVPR. pp. 11686–11695 (2022)
45. Wu, D., Dong, X., Shao, L., Shen, J.: Multi-level representation learning with semantic alignment for referring video object segmentation. In: CVPR. pp. 4996–5005 (2022)
46. Wu, D., Wang, T., Zhang, Y., Zhang, X., Shen, J.: Onlinerefer: A simple online baseline for referring video object segmentation. In: ICCV. pp. 2761–2770 (2023)
47. Wu, J., Jiang, Y., Sun, P., Yuan, Z., Luo, P.: Language as queries for referring video object segmentation. In: CVPR. pp. 4974–4984 (2022)
48. Wu, W., Zhao, Y., Shou, M.Z., Zhou, H., Shen, C.: Diffumask: Synthesizing images with pixel-level annotations for semantic segmentation using diffusion models. arXiv preprint arXiv:2303.11681 (2023)
49. Xu, J., Liu, S., Vahdat, A., Byeon, W., Wang, X., De Mello, S.: Open-vocabulary panoptic segmentation with text-to-image diffusion models. In: CVPR. pp. 2955–2966 (2023)
50. Yang, Z., Wang, J., Tang, Y., Chen, K., Zhao, H., Torr, P.H.: Lavt: Language-aware vision transformer for referring image segmentation. In: CVPR. pp. 18155–18165 (2022)
51. Ye, L., Rochan, M., Liu, Z., Wang, Y.: Cross-modal self-attention network for referring image segmentation. In: CVPR. pp. 10502–10511 (2019)
52. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: ECCV. pp. 69–85. Springer (2016)
53. Zhang, J., Herrmann, C., Hur, J., Polania Cabrera, L., Jampani, V., Sun, D., Yang, M.H.: A tale of two features: Stable diffusion complements dino for zero-shot semantic correspondence. *NeurIPS* **36** (2024)
54. Zhao, H., Lin, K.Q., Yan, R., Li, Z.: Diffusionvmr: Diffusion model for video moment retrieval. arXiv preprint arXiv:2308.15109 (2023)
55. Zhao, S., Chen, D., Chen, Y.C., Bao, J., Hao, S., Yuan, L., Wong, K.Y.K.: Uni-controlnet: All-in-one control to text-to-image diffusion models. In: Thirty-Seventh Conference on Neural Information Processing Systems (NeurIPS 2023) (2023)
56. Zhou, D., Wang, W., Yan, H., Lv, W., Zhu, Y., Feng, J.: Magicvideo: Efficient video generation with latent diffusion models. arXiv preprint arXiv:2211.11018 (2022)
57. Zhu, W., Li, J., Lu, J., Zhou, J.: Separable structure modeling for semi-supervised video object segmentation. *IEEE TCSVT* **32**(1), 330–344 (2021)
58. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)