# RACCooN: Versatile Instructional Video Editing with Auto-Generated Narratives

**Jaehong Yoon**[1,2,*] **Shoubin Yu**[1,*] **Mohit Bansal**[1]
[1]UNC Chapel Hill [2]Nanyang Technological University
jaehong.yoon@ntu.edu.sg {shoubin, mbansal}@cs.unc.edu

https://raccoon-mllm-gen.github.io/

## Abstract

Recent video generative models primarily rely on detailed, labor-intensive text prompts for tasks, like inpainting or style editing, limiting adaptability for personal/raw videos. This paper proposes RACCooN, a versatile and user-friendly **video-to-paragraph-to-video** editing method, supporting diverse video editing capabilities, such as removal, addition, and modification, through a unified pipeline. RACCooN consists of two main stages: Video-to-Paragraph (V2P), which automatically generates structured descriptions of scene and object details, and Paragraph-to-Video (P2V), where users can refine these to guide a video diffusion model for flexible content edits, including removing, changing, or adding objects. Key contributions of RACCooN include: (1) A multi-granular spatiotemporal pooling strategy for structured video understanding, capturing both global context and fine-grained object details to enable precise text-based video editing without complex human annotations. (2) A video generative model fine-tuned on a curated video-paragraph-mask dataset for improved editing and inpainting. (3) The ability to generate new objects by forecasting motion via auto-generated mask planning. In the end, users can easily edit complex videos with RACCooN's automatic explanations and guidance. We demonstrate its versatile capabilities in video-to-paragraph generation (up to $9.4\%p \uparrow$ improvement in human evaluations), video content editing (relative $49.7\% \downarrow$ in FVD).

## 1 Introduction

Recent advances in video generative models (Yan et al., 2021; Hong et al., 2022; Esser et al., 2023; Mei and Patel, 2023; Chai et al., 2023; Blattmann et al., 2023; Wu et al., 2023a), including Sora (openai, 2024), have demonstrated remarkable capabilities in creating high-quality videos. Simultaneously, video editing models (Geyer et al., 2023; Qi et al., 2023; Wang et al., 2024b; Yu et al., 2023; Wu et al., 2024; Zhang et al., 2023b) have gained significant attention for enabling users to modify content using *user-written* instructions. However, building a versatile, user-friendly framework that facilitates easy video modification for personal use remains challenging. Key challenges are as follows:

1) Training a unified model for multiple editing tasks (e.g., add, remove, change objects) is difficult (Yu et al., 2025), and most methods focus on narrow tasks, such as background inpainting (Yu et al., 2023; Wu et al., 2024), or attribute editing (Geyer et al., 2023; Qi et al., 2023; Jeong and Ye, 2023).

2) The need for well-structured prompts that accurately describe videos and support diverse editing tasks is critical, as prompt quality directly impacts model performance. However, generating such prompts is costly and time-consuming, with quality varying by annotator expertise. While Multimodal Large Language Models (MLLMs) (Liu et al., 2023b; Yang et al., 2023; Yu et al., 2024) have been explored for automatic video description, they often miss key details in complex scenes, limiting seamless and user-friendly editing.

To tackle these limitations, we introduce RACCooN: *A Versatile Instructional Video Editing with Auto-Generated Narratives*, a novel **video-to-paragraph-to-video (V2P2V)** generative method that facilitates diverse video editing (Remove, Add, and Change) capabilities based on auto-generated video descriptions. RACCooN enables seamless removal or modification of subject attributes and addition of new objects in videos **without the need for densely annotated prompts or extensive user input**. RACCooN operates in two main stages: *video-to-paragraph* (V2P) and *paragraph-to-video* (P2V). In the V2P stage, we introduce a new video descriptive module built on a pre-trained Video-LLM backbone (PG-Video-LLaVA (Munasinghe et al., 2023)). We find that existing Video-LLMs

[*]Equal Contribution.



*Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 27972–28008
November 4-9, 2025 ©2025 Association for Computational Linguistics

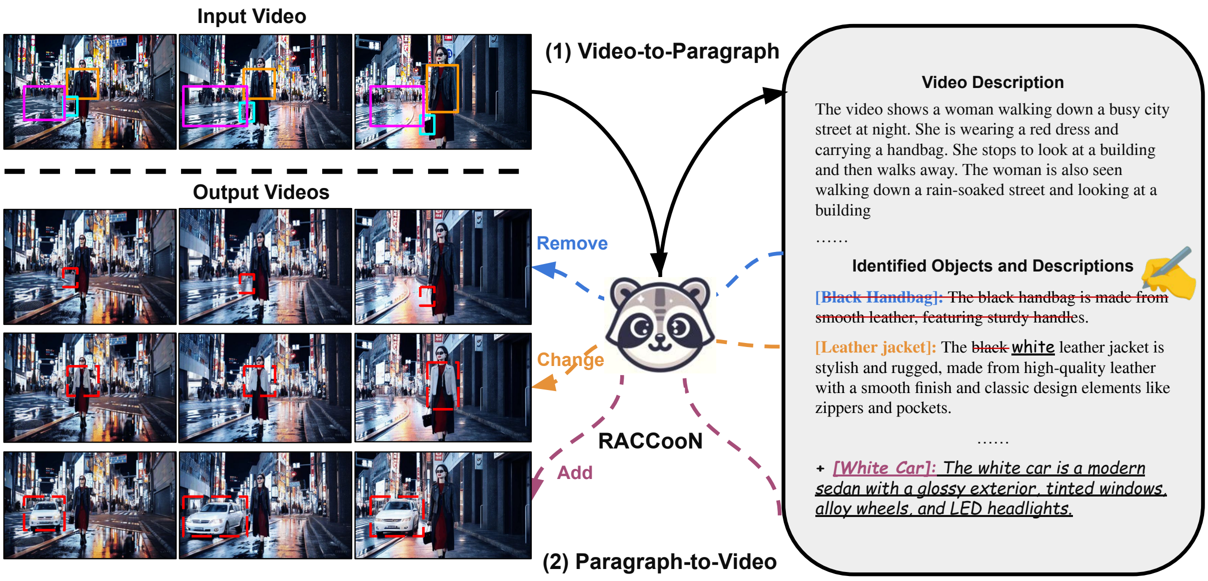


Figure 1: **Overview of RACCOON**, a versatile and user-friendly video-to-paragraph-to-video framework, enables users to remove, add, or change video content via updating auto-generated narratives.

effectively capture holistic video features, yet often overlook detailed cues that are critical for accurate video editing, as users may be interested in altering these missing contexts. To address this, we propose a novel multi-granular video perception strategy that leverages superpixels (Li and Chen, 2015; Ke et al., 2023) to capture diverse and informative localized contexts throughout a video. We first extract fine-grained superpixels using a lightweight predictor (Yang et al., 2020) and then apply overlapping k-means clustering (Cleuziou, 2007; Whang et al., 2015) to segment visual scenes into various levels of granularity. The suggested localized spatiotemporal segmentation assists the LLM's comprehension of objects, actions, and events within the video, enabling it to generate fluent and detailed natural language descriptions. Next, in the P2V stage, we fine-tune a video inpainting model to support multiple editing tasks using auto-generated detailed descriptions and object masks. Then, by leveraging user prompts derived from generated descriptions in the V2P stage, our video diffusion model accurately *paint* corresponding video regions, ensuring that textual edits are faithfully reflected across various editing tasks.

To further support model training, we introduce the **V**ideo **P**aragraph with **L**ocalized **M**ask (**VPLM**) dataset comprising over 7.2K quality video-paragraph data and 5.5K detailed object-level captions with masks, annotated from public datasets using GPT-4V (Achiam et al., 2023).

We emphasize that RACCOON enhances the quality and versatility of video editing by leveraging detailed, automatically generated prompts that minimize ambiguity and refine the scope of generation. We validate the extensive capabilities of the RACCOON in both V2P generation, text-based video content editing, and video generation on ActivityNet (Krishna et al., 2017), YouCook2 (Zhou et al., 2018a), UCF101 (Soomro et al., 2012), DAVIS (Pont-Tuset et al., 2017), and our proposed VPLM datasets. On the V2P side, RACCOON outperforms several strong video captioning baselines (Li et al., 2023; Munasinghe et al., 2023; Liu et al., 2023b), particularly improving by average **+9.1%p** on VPLM and up to **+9.4%p** on YouCook2 compared to PG-VL (Munasinghe et al., 2023), based on both automatic metrics and human evaluation. On the P2V side, RACCOON surpasses previous strong video editing/inpainting baselines (Geyer et al., 2023; Qi et al., 2023; Wang et al., 2024b; Yu et al., 2023; Wu et al., 2024) over three subtasks of video content editing (remove, add, and change video objects) over 9 metrics. Our contributions are as follows:

1. **Framework Contribution:** RACCooN offers a user-friendly, unified framework for diverse video editing tasks, enhancing interpretability and interaction by generating detailed, object-centric descriptions and layout plans tailored to editing goals, surpassing existing model combinations.

2. **Technical Contribution:** We present a novel **multi-granular pooling** strategy to capture local video contexts, enhancing video comprehension

by generating fluent and detailed descriptions in a zero-shot setting. This enables users to create new videos that retain the visual characteristics of the input and support targeted context editing.

3. **Training/Dataset Contribution:** To enable RACCooN to handle complex and diverse video editing requests, we present the **VPLM** dataset, comprising 7.2K high-quality video paragraphs and 5.5K object-level caption-mask pairs. These well-structured annotations enable accurate V2P and P2V stages at both video and object levels.

## 2 Related Work

**Video-to-Paragraph Generation.** Recent video-language tasks focus on generating comprehensive textual descriptions for long and complex video content (Shen et al., 2017; Krishna et al., 2017; Wang et al., 2018; Tewel et al., 2022; Wu et al., 2023b). Vid2Seq (Yang et al., 2023) introduces a novel dense event captioning approach for narrated videos, with time tokens and event boundaries. Video-LLaVA variants (Lin et al., 2023a; Munasinghe et al., 2023) present a large multi-modal model integrating text, video, and audio inputs for generative and question-answering tasks. Similarly, LLaVA-Next (Zhang et al., 2024) improves zero-shot video understanding by transferring multi-image knowledge through concatenated visual tokens. While these methods are effective in video description, they often miss key contextual details (Zhang et al., 2023a; Li et al., 2023). RACCooN captures both holistic and object-level details by leveraging localized spatiotemporal information, enhancing video editing and generation.

**Prompt-to-Video Editing.** Video editing (Ceylan et al., 2023; Liu et al., 2023c; Couairon et al., 2023; Kondratyuk et al., 2023; Wang et al., 2023; Zhang et al., 2023c) involves enhancing, modifying, or manipulating video content for desired effects. VideoComposer (Wang et al., 2024b) offers a multi-source controllable video generative framework. TokenFlow (Geyer et al., 2023) adapts text-to-image diffusion with flow matching for consistent text-driven video editing. LGVI (Wu et al., 2024) integrates an MLLM for complex language-based video inpainting. These methods often focus on specific tasks and may inadvertently alter unrelated regions due to limited contextual information. Our V2P2V framework overcomes these limitations by using auto-generated, detailed descriptions to integrate key contexts into diverse editing tasks.

## 3 A Versatile Instructional Video Editing with Auto-Generated Narratives

Conditional video generation and editing models struggle with complex scenes due to vague text descriptions and limited video understanding. To address this, we introduce RACCooN, a user-friendly, two-stage video-to-paragraph-to-video editing approach, with each stage detailed in Sec. 3.1 and Sec. 3.2. We also introduce the VPLM dataset, specifically curated to train models for detailed video editing, along with the training pipeline of RACCooN, detailed in Sec. 3.3.

### 3.1 V2P: Auto-Descriptive Module with Multi-Granular Video Perception

**Multimodal LLM for Video Paragraph Generation.** In the V2P stage, the RACCooN generates well-structured, detailed descriptions for both holistic videos and local objects. It employs a multimodal LLM with three main components: a visual encoder $E$, a multimodal projector, and an LLM. Given an input video $\boldsymbol{x} \in \mathbb{R}^{F\times C\times H\times W}$, where $F$, $C$, $H$, and $W$ represent the number of frames, channels, height, and width, respectively, we extract video features using the visual encoder: $\boldsymbol{e} = E(\boldsymbol{x}) \in \mathbb{R}^{t\times h\times w\times d}$. Here, $t$, $h$, $w$, and $d$ denote the encoded temporal dimension, the height and width of the tokens, and the feature dimension. To understand complex videos with multiple scenes, we use three pooling strategies: *spatial pooling*, *temporal pooling*, and *multi-granular spatiotemporal pooling*. Spatial pooling $\boldsymbol{e}^s = \text{Pooling}^s(\boldsymbol{e}) \in \mathbb{R}^{t\times d}$ aggregates tokens within the same frame, while temporal pooling $\boldsymbol{e}^t = \text{Pooling}^t(\boldsymbol{e}) \in \mathbb{R}^{(h\cdot w)\times d}$ averages features across the temporal dimension for the same region. Despite these strategies helping the LLM grasp the video holistically in space or time, they often overlook capturing key objects or actions localized throughout the video stream, especially in untrimmed, dynamic, multi-scene videos.

**Multi-Granular Spatiotemporal Pooling.** To address this issue, we introduce a novel superpixel-based spatiotemporal pooling strategy, coined *multi-granular spatiotemporal pooling* (MGS pooling). As illustrated in Fig. 2 left top, this strategy is designed to capture localized information via superpixels across spatial and temporal dimensions. Superpixels (Li and Chen, 2015; Yang et al., 2020; Ke et al., 2023) are small and coherent clusters of pixels that share similar characteristics, such as

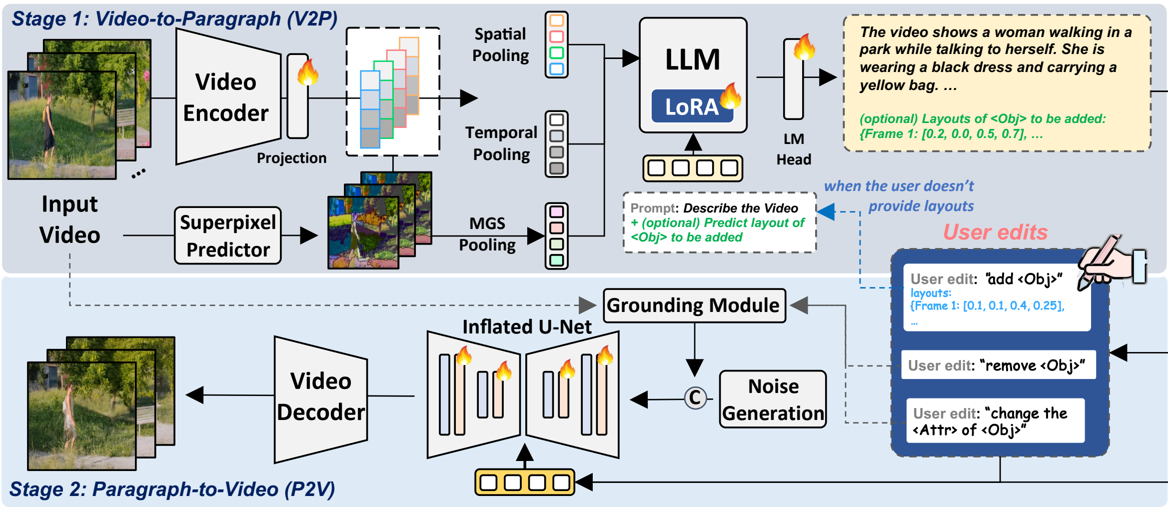


Figure 2: **Illustration of RACCooN.** RACCooN generates video descriptions with the three distinct pooled visual tokens, including Multi-Granular Spatiotemporal (MGS) Pooling. Next, users can edit the generated descriptions by adding, removing, or modifying words to create new videos. Note that for adding object tasks, if users do not provide layout information for the objects they want to add, RACCooN can predict the target layout in each frame.

color or texture. These clusters provide an efficient representation of visual scenes and are resilient to frame noise since they average out the pixel values within each cluster, effectively smoothing out variations induced by noise. As shown in Fig. 3, we use a lightweight superpixel predictor $\sigma(\cdot)$ (Yang et al., 2020) to generate superpixels across video frames, capturing the granular visuality of each local area. However, due to their limited coverage area, these fine-grained visual features often fail to capture attribute-level semantics, such as objects and actions (Zhang et al., 2023a; Li et al., 2023). Motivated by the importance of varying the compositions of multiple superpixels for different contexts in video understanding, we propose the use of overlapping k-means clustering (Cleuziou, 2007; Whang et al., 2015) for the obtained video superpixels, which improves the granularity from fine to coarse. This approach allows the LLM to gather informative cues about various objects and actions. We first obtain the pixel features and the superpixel index vector for the video pixels: $\boldsymbol{S}, \boldsymbol{g} = \sigma(\boldsymbol{x}, \boldsymbol{g}_{\text{init}})$, where $\boldsymbol{g}_{\text{init}}$ is the input superpixel indices, initialized by a region-based grid. Given the averaged pixel features of each superpixel, $\bar{\boldsymbol{S}} \in \mathbb{R}^{|\boldsymbol{g}| \times d_p}$, where $d_p$ denotes the pixel feature size, we generate the MGS tokens $\boldsymbol{e}^l$:

$$\begin{aligned} \boldsymbol{m} &= \text{OKM}\left(\bar{\boldsymbol{S}}, k, v\right) \in \{0,1\}^{k \times F \times H \times W}, \\ \boldsymbol{e}^l &= \text{AvgPool}(\boldsymbol{m}) \otimes \boldsymbol{e} \in \mathbb{R}^{k \times d}, \end{aligned} \tag{1}$$

where OKM represents the overlapping k-means algorithm with $k$ centroids and overlap scale $v$ for each cluster. $\boldsymbol{m}$ denotes the set of binary masks for superpixels. $\otimes$ denotes tensor multiplication. We describe the detailed MGS process and ablation of pooling strategies in the Appendix. Next, we concatenate the pooled video tokens and map them into the text embedding space using the multimodal linear projector. Combined with the embedding of the encoded text token $\boldsymbol{e}^p$ from the textual prompt, the LLM generates a well-structured and detailed description $\boldsymbol{a}$ of the video:

$$\begin{aligned} \hat{\boldsymbol{e}} &= \text{CONCAT}[\boldsymbol{e}^s; \boldsymbol{e}^l; \boldsymbol{e}^t] \cdot \boldsymbol{W}^\top, \\ \boldsymbol{a} &= LLM(\text{CONCAT}[\boldsymbol{e}^p; \hat{\boldsymbol{e}}]), \end{aligned} \tag{2}$$

where $\boldsymbol{W} \in \mathbb{R}^{d \times d'}$ is the weight matrix for linear projection into the text embedding dimension $d'$. We highlight that our video description module serves as an integrated, user-interactive tool for video-to-paragraph generation and video content editing. (Fig. 2 top right).

### 3.2 P2V: User-Interactive Video Editing with Auto-Generated Descriptions

With the well-structured, detailed, and object-centric video description generated from the *Video-to-Paragraph* stage, users can 'read' the video details and interactively modify the content by altering the model-generated description. This approach shifts users' focus from labor-intensive video observation to content editing. We categorize general video content editing into three important subtasks: **(1) Video Object Adding:** add extra

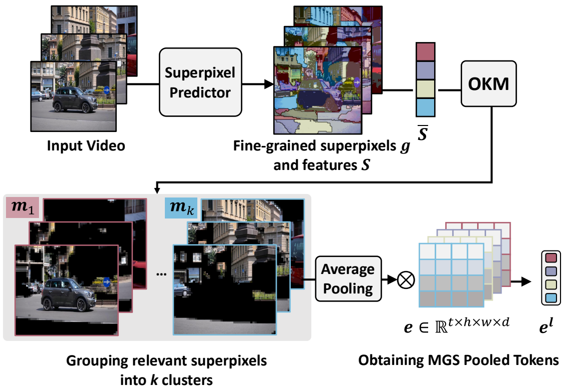


Figure 3: **Illustration of MGS pooling.** We obtain MGS pooling tokens using a spatiotemporal mask $\boldsymbol{m}$ via overlapping k-means clustering (OKM) of averaged superpixel features $\bar{S}$.

objects to a video. **(2) Video Object Removing:** delete target objects and re-generate the object region as the background. **(3) Video Object Changing:** change objects' attributes (e.g., color, textural, material). Many previous works have made great progress in video editing (Wu et al., 2024; Geyer et al., 2023; Qi et al., 2023; Zhang et al., 2023b; Fan et al., 2024), but usually focus on one of these subtasks. In this paper, we propose a unified generative model for video content editing that integrates all those crucial subtasks. Specifically, we formulate these subtasks as text-based video painting tasks and leverage a single video diffusion model for adding, removing, and changing video objects in the form of inpainting.

As shown in Fig. 2 bottom, our video diffusion model processes input video $\boldsymbol{x} \in \mathbb{R}^{F\times C\times H\times W}$ with a predicted binary mask $\boldsymbol{m}' \in \mathbb{R}^{F\times 1\times H\times W}$ targeting specific regions for modification. Following image inpainting techniques (Xie et al., 2023; Rombach et al., 2022), we apply the mask[1] to the video to designate the editing region. The masked video is then encoded using a Variational Autoencoder (VAE (Kingma and Welling, 2013)) to serve as the generation condition. The model can then be informed on which video region should be edited for localized editing. Driven by the detailed description, the diffusion model can conduct diverse video editing that reflects the text prompts.

In addition, we provide details regarding the process of adding objects in video editing. Indeed, adding objects can be considered a unique video editing task, distinct from removing objects or changing attributes. Unlike the latter scenarios, where the target objects are already present in the initial video, adding objects involves introducing entirely new elements, which necessitates a modified editing process. As illustrated in Fig. 2, the MLLM in V2P provides not only detailed descriptions but also frame-wise placement suggestions for new objects in the form of bounding box sequences. The object insertion process in RACCooN in inference is conducted through the following steps:

***1. User Edit:*** The user provides an instruction to add a specific object.

***2. MLLM Output:*** The finetuned MLLM in V2P generates fine-grained video descriptions along with frame-wise bounding box suggestions for new objects. For example, "Layouts of <Obj> to be added: {Frame 1: `[0.2, 0.0, 0.5, 0.7]`, Frame 2: `[0.2, 0.1, 0.4, 0.65]`, ...}" specifies the layout for each frame, where `[x1, y1, x2, y2]` represents the top-left and bottom-right corners of the bounding box, with coordinates normalized to the range $[0, 1]$ (yellow box in Fig. 2, top right).

***3. Video Editing:*** Generate videos based on the MLLM-generated output, including the frame-wise layout of the object to be added.

### 3.3 VPLM Dataset Collection and RACCooN Pipeline Training

**Dataset Collection.** We utilize video datasets (Majumdar et al., 2020; Gavrilyuk et al., 2018) from previous video inpainting work (Wu et al., 2024). Each raw video is accompanied by multiple inpainted versions with specific objects removed and includes binary masks of these objects. Although well-annotated with object masks and inpainted backgrounds, these datasets lack detailed descriptions of holistic video and specific local objects, hindering RACCooN's training for producing well-structured captions for video editing. To address this, we use GPT-4V (Achiam et al., 2023) to annotate detailed video descriptions. We first re-arrange uniformly sampled video frames into a grid-image (Fan et al., 2021) and add visual prompts by numbering each frame. We then ask GPT-4V to generate detailed captions for both the entire video and key objects, in a well-structured format. Next, we train V2P and P2V stages in our pipeline separately (Fig. 2). In the end, RACCooN can automatically generate detailed, well-structured descriptions for raw videos and adapt these descriptions based on user updates for various video content editing tasks.

**MLLM Instructional Fine-tuning.** To enable the

[1] We use image grounding (Liu et al., 2023d) and video tracking models (Cheng et al., 2023) as the off-the-shelf mask predictor in inference.

MLLM to output detailed video descriptions for content editing, we construct an instructional fine-tuning dataset based on VPLM with two video-instruction (Liu et al., 2023b) designs: (1) For object editing and removal, the MLLM generates structured video captions identifying key objects in the original video $\boldsymbol{x}$, using annotated descriptions as the learning objective. This allows users to edit videos directly from these descriptions without exhaustive analysis. (2) For object insertion, the MLLM provides not only detailed descriptions but also frame-wise placement suggestions for new objects, enhancing its utility in video editing by avoiding manual trajectory outlining. For training, we convert video object segmentation masks into bounding boxes by selecting maximal and minimal coordinates and follow the box planning strategy using LLMs (Lin et al., 2023b). We input box coordinates as a sequence of numbers and train RACCOON to predict these layouts given inpainted videos $\widehat{\boldsymbol{x}}$. We perform parameter-efficient fine-tuning with LoRA (Hu et al., 2022)[2] on these mixed datasets with CE loss. We freeze the visual encoder and LLM backbone, updating the projector, LoRA, and LLM head.

**Video Diffusion Model Fine-tuning.** Our video diffusion model builds on the prior image inpainting model (Rombach et al., 2022), enhanced with temporal attention layers to capture video dynamics. The model is designed to generate video that aligns with input prompts, focusing on object-centric video content editing. To support this, we develop a training dataset of mask-object-description triples. We use GPT-4 to produce single-object descriptions from long, detailed video narratives, framing this task as a multi-choice QA problem. Next, for the three video editing subtasks, we design specific input-output combinations: (1) **Video Object Addition:** Inputs: inpainted video $\widehat{\boldsymbol{x}}$, object bounding boxes from segmentation masks $m$, and detailed object description $\boldsymbol{p}$. Output: original video $\boldsymbol{x}$. (2) **Video Object Removal:** Inputs: original video $\boldsymbol{x}$, object segmentation masks $m$, and a fixed background prompt. Output: inpainted video $\widehat{\boldsymbol{x}}$. (3) **Video Object Change:** Inputs: original video $\boldsymbol{x}$, object segmentation masks $m$, and object description $\boldsymbol{p}$. Output: original video $\boldsymbol{x}$. The model is fine-tuned following the prior work (Wu et al., 2023a), updating only the temporal layers and the query projections within the self-attention and cross-attention modules. We employ the MSE loss between generated and random noise. See Appendix for more details on the dataset and training.

[2]We employ LoRA for query and value for each self-attention.

## 4 Experimental Results

**Tasks & Datasets**: We evaluate RACCOON on diverse video datasets across tasks, including video captioning (**YouCook2** (Zhou et al., 2018a), **VPLM**), text-based video content editing (**DAVIS** (Pont-Tuset et al., 2017), **VPLM**), and conditional video generation (**ActivityNet** (Krishna et al., 2017), **YouCook2** (Zhou et al., 2018a), **UCF101** (Soomro et al., 2012)).

**Metrics**: For each task, we evaluate our approach with various metrics. **(1) Video Caption:** following previous works (Yang et al., 2023; Zhou et al., 2018b), we conduct a comprehensive human evaluation and adopt general metrics for our long video descriptions, including SPICE (Anderson et al., 2016), BLEU-4 (Vedantam et al., 2015), and CIDEr (Vedantam et al., 2015). **(2) Video Object Layout Planning:** following the prior work (Lin et al., 2023b), we evaluate RACCOON for object layout planning by bounding box IoU, FVD (Unterthiner et al., 2019), and CLIP-score (Radford et al., 2021). **(3) Text-based Video Content Editing:** following prior works (Geyer et al., 2023; Ceylan et al., 2023; Yang et al., 2024), we evaluate RACCOON's video editing capabilities by CLIP-Text, CLIP-Frame, Qedit (Yang et al., 2024), and SSIM (Hore and Ziou, 2010). **(4) Conditional Video Generation:** we measure FVD (Unterthiner et al., 2019), CLIP-Score (Radford et al., 2021), and SSIM (Hore and Ziou, 2010). Implementation details are provided in the Appendix.

### 4.1 Video-to-Paragraph Generation

**Video-Paragraph Alignment.** We conducted a quantitative evaluation of our proposed RACCOON's video-to-paragraph generation capabilities, comparing it against strong baselines with a focus on object-centric captioning and object layout planning. The results, summarized in Tab. 1, show that open-source video-LLMs (e.g., PG-VL, Video-Chat), which have smaller LLMs ($< 13B$ parameters), struggle with object-centric captioning and usually fail to generate layout planning. This is primarily due to their lack of instructional fine-tuning and insufficient video detail modeling without multi-granular pooling. In contrast,

| Methods | S | B | C | IoU | FVD | CLIP |
|---|---|---|---|---|---|---|
| *open-source MLLMs* | | | | | | |
| LLaVA (Liu et al., 2023b) | 17.4 | 27.5 | 18.5 | - | - | - |
| Video-Chat (Li et al., 2023) | 18.2 | 25.3 | 19.1 | - | - | - |
| PG-VL (Munasinghe et al., 2023) | 18.2 | 27.4 | 14.6 | - | - | - |
| *proprietary MLLMs* | | | | | | |
| Gemini 1.5 Pro (Team et al., 2023) | 19.2 | 23.5 | 11.0 | 0.115 | **371.63** | 0.978 |
| GPT-4o (gpt 4o, 2024) | 20.6 | 28.0 | **37.4** | 0.179 | 447.67 | 0.977 |
| RACCooN | **23.1** | **31.0** | <u>33.5</u> | **0.218** | <u>432.42</u> | **0.983** |

Table 1: **Single Object Prediction** on VPLM test set. Metrics indicate: **S**: *SPICE*, **B**: *BLEU-4*, **C**: *CIDEr*.

| Methods | Logic | Lang. | Summ. | Details | Avg. |
|---|---|---|---|---|---|
| Ground Truth | 66.7 | 42.2 | 41.7 | **72.2** | 55.7 |
| PG-VL (Munasinghe et al., 2023) | 77.2 | 81.1 | 69.4 | 62.8 | 72.6 |
| **RACCooN** | **80.6** | **85.0** | **72.2** | **72.2** | **77.5** |

Table 2: **Results of Human Evaluation** on YouCook2. We measure the quality of the description through four metrics: Logic Fluency (Logic), Language Fluency (Lang.), Video Summary (Summ.), and Video Details (Details). We report the normalized score $s \in [0, 100]$.

| Model | Change Object | | | Remove Object | | | Add Object | | |
|---|---|---|---|---|---|---|---|---|---|
| | CLIP-T↑ | CLIP-F↑ | Qedit↑ | FVD↓ | SSIM↑ | PSNR↑ | FVD↓ | SSIM↑ | PSNR↑ |
| *Inversion-based Models* | | | | | | | | | |
| LOVECon (Liao and Deng, 2023) | 29.36 | 94.77 | 1.29 | 1319.51 | 60.40 | 17.78 | 1433.12 | 58.51 | 17.35 |
| FateZero (Qi et al., 2023) | 25.18 | 94.47 | 1.01 | 1037.05 | 47.35 | 15.16 | 1474.80 | 47.65 | 15.45 |
| TokenFlow (Geyer et al., 2023) | 29.25 | 96.23 | 1.31 | 1317.29 | 47.06 | 15.83 | 1373.20 | 49.95 | 15.95 |
| *Inpainting-Based Models* | | | | | | | | | |
| Inpaint Anything (Yu et al., 2023) | 24.86 | 92.01 | 1.01 | 383.81 | 82.33 | 27.69 | 712.59 | 77.75 | 22.41 |
| LGVI (Wu et al., 2024) | 23.82 | **95.33** | 1.04 | 915.24 | 56.16 | 19.14 | 1445.43 | 47.93 | 16.09 |
| VideoComposer (Wang et al., 2024b) | 27.61 | 94.18 | **1.25** | 827.04 | 47.34 | 17.55 | 1151.90 | 48.01 | 15.76 |
| PGVL + SD-v2.0-inpainting | 24.01 | 90.11 | 1.01 | 282.31 | 82.33 | 27.69 | 1579.65 | 43.21 | 15.76 |
| **RACCooN** | **27.85** | <u>94.78</u> | <u>1.15</u> | **162.03** | **84.38** | **30.34** | **415.82** | **77.81** | **23.38** |

Table 3: **Results of Video Content Editing on three sub-tasks** on VPLM test. We gray out models that conduct the DDIM inversion process and have a different focus on our inpainting-based model.

RACCooN demonstrates superior performance in both object-centric captioning and complex object layout planning, benefiting from the instructional tuning on our VPLM dataset. Additionally, our method achieves competitive performance with proprietary MLLMs (e.g., Gemini 1.5 Pro, GPT-4o) in key object captioning and layout planning, demonstrating its superior instruction following and generation quality.

**Human Evaluation & Qualitative Examples.** We conducted a human evaluation to compare our auto-generated captions with those from a strong baseline and human annotations on ten randomly selected YouCook2 videos (each three to five minutes long with multiple scenes and complex viewpoints). Five evaluators rated these based on Logic Fluency, Language Fluency, Video Summary, and Video Details (details in the supplementary material). The average scores for each criterion and their overall mean are illustrated in Tab. 2. Our method significantly surpassed both the PG-VL-generated and ground truth captions in all metrics, showing a **4.9%p** and **21.8%p absolute improvement** respectively, and matched the ground truth in capturing *Video Details* with a **9.4%p** enhancement over the baseline, highlighting RACCooN's superior capability in capturing video details. We additionally visualize descriptions generated by RACCooN using a well-known generated video from the Sora (openai, 2024) generated demo example. As shown in Fig. 4 (in the supplementary material), it demonstrates our robust capability to auto-describe complex video content without human textual input.

## 4.2 Instructional Video Editing with RACCooN

**Quantitative Evaluation.** As shown in Tab. 3, we quantitatively compare the video editing ability of RACCooN with strong video editing models based on inpainting or DDIM-inversion (Hertz et al., 2022) across three object-centric video content editing subtasks: *object changing*, *removal*, and *adding*. In general, RACCooN outperforms all baselines across 9 metrics. For object changing, RACCooN outperforms the best-performing baseline by 0.8% on CLIP-T, indicating better video-text alignment while maintaining temporal consistency, as demonstrated by comparable CLIP-F and Qedit scores. Note that LGVI is not designed to alter video attributes and tends to preserve video content with marginal change (i.e., identical input and output videos), resulting in im-

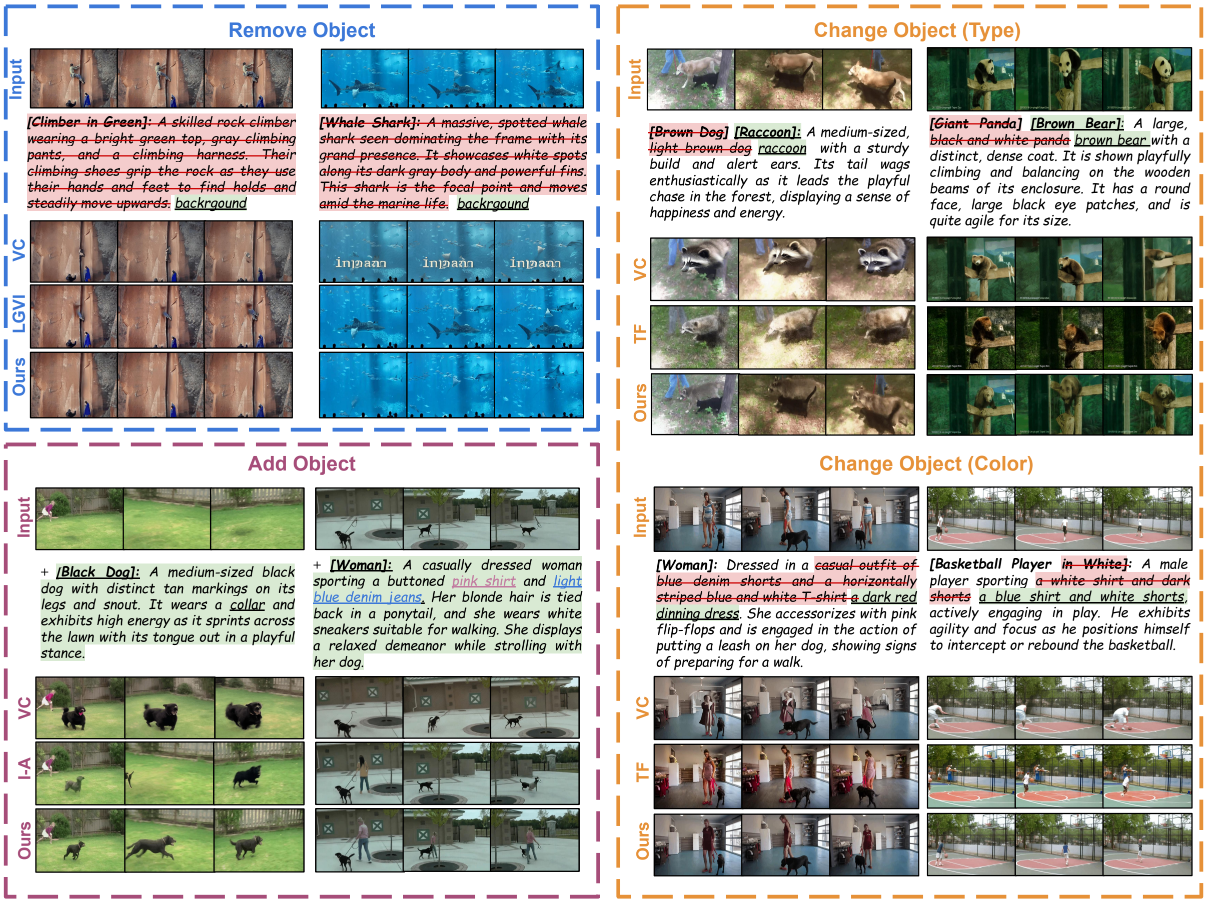


Figure 4: **Qualitative Comparison between RACCOON and other baselines.** Baseline names are abbreviated: **VC**: VideoComposer, **I-A**: Inpainting Anything, **TF**: TokenFlow. We underlined visual details in our caption. More visualizations are in the supplementary material.

proved CLIP-F scores. In the object removal task, RACCOON shows significant improvements over strong baselines (relatively +57.8% FVD, +2.5% SSIM, +9.6% PSNR). Such improvements are maintained in the addition task (relatively +41.6% FVD, +4.3% PSNR). Meanwhile, some DDIM inversion-based models (e.g., TokenFlow (Geyer et al., 2023)) work well for specific tasks (change objects), but do not handle other types of editing. In contrast, our method is an all-rounder player.

We further emphasize that **a simple combination of existing models cannot achieve an effective video editing pipeline**, leading to inferior instruction-following and editing abilities. This is evident in the degraded performance of open-source Video-LLM baselines (Tabs. 1 and 2) and other video editing models (Tab. 3). To address these limitations, we made unique novelty in technical and dataset contributions and achieved significant improvements in both video understanding and editing. This is evident in the comparison of RACCOON with a multi-agent baseline combining a powerful open-source video reasoning and video diffusion models (PG-Video-LLaMA + StableDiffusion 2.0-inpainting) in Tab. 3. RACCOON outperforms the *PGVL + SD 2.0-inpainting* by a significant margin across all metrics and editing tasks, highlighting the effectiveness of our proposed editing method.

**Visualization.** In Fig. 4, we compare videos generated by RACCOON with several SoTA baselines across three video content editing tasks. For object removal, RACCOON demonstrates superior results, naturally and smoothly inpainting the background, whereas VideoComposer generates unexpected content and LGVI fails to accurately remove objects across frames. For object addition, compared to Inpainting-Anything and VideoComposer, which often miss objects or produce distorted generations, RACCOON generates objects with more fluent and natural motion, accurately reflecting caption details (e.g., the *collar* of the dog, the *pink shirt*, and *blue jeans* for the woman). For changing objects, our method outperforms inpainting-

| Settings | FVD↓ | SSIM↑ | PSNR↑ |
|---|---|---|---|
| *add object* | | | |
| RACCooN | 415.80 | 77.81 | 23.38 |
| w/o detail caption | 476.01 | 76.80 | 23.14 |
| w/o oracle planning | 969.95 | 76.65 | 21.21 |
| *remove object* | | | |
| RACCooN | 162.03 | 84.38 | 30.34 |
| w/o oracle mask | 398.01 | 81.60 | 27.15 |
| **Setting** | **CLIP-T** | **CLIP-F** | **Qedit** |
| *change object* | | | |
| RACCooN | 27.85 | 94.78 | 1.15 |
| w/o oracle mask | 27.23 | 94.33 | 1.14 |

Table 4: **Ablation on video object changing, removing, and adding** with different inputs.

based VideoComposer and inversion-based TokenFlow. RACCooN accurately re-paints objects to achieve object editing for color (*white→blue*) and type (*dog→raccoon*), while others struggle to meet requirements.

**Ablation Studies.** As shown in Tab. 4, we further validate the effectiveness of components by replacing detailed descriptions with short captions, and oracle masks/planning boxes with model-generated ones. In adding objects, our detailed object descriptions can benefit generation by providing accurate details, leading to improved quantitative results (relatively $+14.4\%$ FVD). We further replace GT boxes with boxes predicted by RACCooN, and still show superior performance over other baseline methods with oracle boxes in Tab. 3. It demonstrates that our V2P stage can thus automatically generate planning from a given video to eliminate users' labor. Next, in object removal and changing, we replace the oracle masks with grounding (Liu et al., 2023d) and tracking (Cheng et al., 2023) tools generated mask, which shows a marginal decrement for changing objects, and RACCooN still shows strong results over other baselines in Tab. 3 with oracle masks. It suggests that RACCooN is effective and robust to handle diverse editing skills in a non-orcale setting (See the appendix).

## 5 Conclusion

We newly introduce an auto-descriptive video-to-paragraph-to-video generative approach. We automatically generate video descriptions by leveraging a multi-granular spatiotemporal pooling strategy, enhancing the model's ability to recognize detailed, localized video information. Our approach then uses these enriched descriptions to edit and generate video content, offering users the flexibility to modify content through textual updates, thus eliminating the need for detailed video annotations. Our video editing and generation abilities highlight notable effectiveness and enable a broader range of users to engage in video creation and editing tasks without the written textual prompts.

## Limitations

Our proposed RACCooN has shown a remarkable ability to interpret input videos, producing well-structured and detailed descriptions that outperform strong video captioning baselines and even ground truths. However, it has the potential to produce inaccuracies or hallucination (Liu et al., 2023a; Wang et al., 2024a; Zhou et al., 2024; Ma et al., 2023) in the generated text outputs. In addition, the performance of RACCooN in paragraph generation, video generation, and editing is influenced by the employed pre-trained backbones, including an LLM (Touvron et al., 2023), base Inpainting Model (Rombach et al., 2022), Video Diffusion Model (Xing et al., 2023), and Video Editor (Geyer et al., 2023). However, our key contributions are independent of these backbones, and we emphasize that RACCooN's capabilities can be further enhanced with future advancements in these generative model backbones.

LLM-empowered video description and photo-realistic video creation/editing inherit biases from their training data, leading to several broader impacts, including societal stereotypes, biased interpretation of actions, and privacy concerns. To mitigate these broader impacts, it is essential to carefully develop and implement generative and video description models, such as considering diversifying training datasets, implementing fairness and bias evaluation metrics, and engaging communities to understand and address their concerns.

## Acknowledgement

We thank the reviewers and Jaemin Cho, Han Lin, Jialu Li, and Elias Stengel-Eskin for the thoughtful discussion. This work was supported by DARPA ECOLE Program No. HR00112390060, NSF-AI Engage Institute DRL-2112635, DARPA Machine Commonsense (MCS) Grant N66001-19-2-4031, ARO Award W911NF2110220, ONR Grant N00014-23-1-2356, and Accelerate Foundation Models Research program. The views contained in this article are those of the authors and not of the funding agency.

## References


Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. arXiv preprint arXiv:2303.08774.

Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. Spice: Semantic propositional image caption evaluation. In Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V 14, pages 382–398. Springer.

Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. 2021. Frozen in time: A joint video and image encoder for end-to-end retrieval. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 1728–1738.

Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, and 1 others. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets. arXiv preprint arXiv:2311.15127.

Duygu Ceylan, Chun-Hao P Huang, and Niloy J Mitra. 2023. Pix2video: Video editing using image diffusion. In Proceedings of the International Conference on Computer Vision (ICCV).

Wenhao Chai, Xun Guo, Gaoang Wang, and Yan Lu. 2023. Stablevideo: Text-driven consistency-aware diffusion video editing. In Proceedings of the International Conference on Computer Vision (ICCV).

Haoxin Chen, Menghan Xia, Yingqing He, Yong Zhang, Xiaodong Cun, Shaoshu Yang, Jinbo Xing, Yaofang Liu, Qifeng Chen, Xintao Wang, and 1 others. 2023. Videocrafter1: Open diffusion models for high-quality video generation. arXiv preprint arXiv:2310.19512.

Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. 2023. Segment and track anything. arXiv preprint arXiv:2305.06558.

Guillaume Cleuziou. 2007. A generalization of k-means for overlapping clustering. Rapport technique.

Paul Couairon, Clément Rambour, Jean-Emmanuel Haugeard, and Nicolas Thome. 2023. Videdit: Zero-shot and spatially aware text-driven video editing. arXiv preprint arXiv:2306.08707.

Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. 2023. Structure and content-guided video synthesis with diffusion models. In Proceedings of the International Conference on Computer Vision (ICCV).

Quanfu Fan, Rameswar Panda, and 1 others. 2021. An image classifier can suffice for video understanding. arXiv preprint arXiv:2106.14104, 2.

Xiang Fan, Anand Bhattad, and Ranjay Krishna. 2024. Videoshop: Localized semantic video editing with noise-extrapolated diffusion inversion. arXiv preprint arXiv:2403.14617.

Kirill Gavrilyuk, Amir Ghodrati, Zhenyang Li, and Cees GM Snoek. 2018. Actor and action video segmentation from a sentence. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 5958–5966.

Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. 2023. Tokenflow: Consistent diffusion features for consistent video editing. arXiv preprint arXiv:2307.10373.

gpt 4o. 2024. https://openai.com/index/hello-gpt-4o/.

Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. 2022. Prompt-to-prompt image editing with cross attention control. arXiv preprint arXiv:2208.01626.

Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. 2022. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. arXiv preprint arXiv:2205.15868.

Alain Hore and Djemel Ziou. 2010. Image quality metrics: Psnr vs. ssim. In 2010 20th international conference on pattern recognition, pages 2366–2369. IEEE.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. Lora: Low-rank adaptation of large language models. In Proceedings of the International Conference on Learning Representations (ICLR).

Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. 2017. Flownet 2.0: Evolution of optical flow estimation with deep networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 2462–2470.

Hyeonho Jeong and Jong Chul Ye. 2023. Ground-a-video: Zero-shot grounded video editing using text-to-image diffusion models. arXiv preprint arXiv:2310.01107.

Yongrae Jo, Seongyun Lee, Aiden SJ Lee, Hyunji Lee, Hanseok Oh, and Minjoon Seo. 2023. Zero-shot dense video captioning by jointly optimizing text and moment. arXiv preprint arXiv:2307.02682.

Tsung-Wei Ke, Sangwoo Mo, and X Yu Stella. 2023. Learning hierarchical image segmentation for recognition and by recognition. In Proceedings of the International Conference on Learning Representations (ICLR).

Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114.

Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. 2023. Segment anything. arXiv:2304.02643.

Dan Kondratyuk, Lijun Yu, Xiuye Gu, José Lezama, Jonathan Huang, Rachel Hornung, Hartwig Adam, Hassan Akbari, Yair Alon, Vighnesh Birodkar, and 1 others. 2023. Videopoet: A large language model for zero-shot video generation. arXiv preprint arXiv:2312.14125.

Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In Proceedings of the IEEE international conference on computer vision, pages 706–715.

KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2023. Videochat: Chat-centric video understanding. arXiv preprint arXiv:2305.06355.

Zhengqin Li and Jiansheng Chen. 2015. Superpixel segmentation using linear spectral clustering. In Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR).

Zhenyi Liao and Zhijie Deng. 2023. Lovecon: Text-driven training-free long video editing with controlnet. arXiv preprint arXiv:2310.09711.

Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. 2023a. Video-llava: Learning united visual representation by alignment before projection. arXiv preprint arXiv:2311.10122.

Han Lin, Abhay Zala, Jaemin Cho, and Mohit Bansal. 2023b. Videodirectorgpt: Consistent multi-scene video generation via llm-guided planning. arXiv preprint arXiv:2309.15091.

Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2023a. Mitigating hallucination in large multi-modal models via robust instruction tuning. In Proceedings of the International Conference on Learning Representations (ICLR).

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. In Advances in Neural Information Processing Systems (NeurIPS).

Shaoteng Liu, Yuechen Zhang, Wenbo Li, Zhe Lin, and Jiaya Jia. 2023c. Video-p2p: Video editing with cross-attention control. arXiv preprint arXiv:2303.04761.

Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and 1 others. 2023d. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv preprint arXiv:2303.05499.

Ziqiao Ma, Jiayi Pan, and Joyce Chai. 2023. World-to-words: Grounded open vocabulary acquisition through fast mapping in vision-language models. arXiv preprint arXiv:2306.08685.

Arjun Majumdar, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, and Dhruv Batra. 2020. Improving vision-and-language navigation with image-text pairs from the web. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16, pages 259–274. Springer.

Kangfu Mei and Vishal Patel. 2023. Vidm: Video implicit diffusion models. In Proceedings of the AAAI National Conference on Artificial Intelligence (AAAI).

Shehan Munasinghe, Rusiru Thushara, Muhammad Maaz, Hanoona Abdul Rasheed, Salman Khan, Mubarak Shah, and Fahad Khan. 2023. Pg-video-llava: Pixel grounding large video-language models. arXiv preprint arXiv:2311.13435.

openai. 2024. https://openai.com/sora.

Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. 2017. The 2017 davis challenge on video object segmentation. arXiv preprint arXiv:1704.00675.

Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. 2023. Fatezero: Fusing attentions for zero-shot text-based video editing. arXiv preprint arXiv:2303.09535.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In Proceedings of the International Conference on Machine Learning (ICML).

Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang, Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. 2024. Grounded sam: Assembling open-world models for diverse visual tasks. arXiv preprint arXiv:2401.14159.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 10684–10695.

Zhiqiang Shen, Jianguo Li, Zhou Su, Minjun Li, Yurong Chen, Yu-Gang Jiang, and Xiangyang Xue. 2017. Weakly supervised dense video captioning. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1916–1924.

Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. 2012. Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402.

Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, and 1 others. 2023. Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805.

Yoad Tewel, Yoav Shalev, Roy Nadler, Idan Schwartz, and Lior Wolf. 2022. Zero-shot video captioning with evolving pseudo-tokens. arXiv preprint arXiv:2207.11100.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and 1 others. 2023. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288.

Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphaël Marinier, Marcin Michalski, and Sylvain Gelly. 2019. Fvd: A new metric for video generation. In ICLR workshop.

Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4566–4575.

Bairui Wang, Lin Ma, Wei Zhang, and Wei Liu. 2018. Reconstruction network for video captioning. In Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR).

Lei Wang, Jiabang He, Shenshen Li, Ning Liu, and Ee-Peng Lim. 2024a. Mitigating fine-grained hallucination by fine-tuning large vision-language models with caption rewrites. In International Conference on Multimedia Modeling. Springer.

Teng Wang, Ruimao Zhang, Zhichao Lu, Feng Zheng, Ran Cheng, and Ping Luo. 2021. End-to-end dense video captioning with parallel decoding. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 6847–6857.

Wen Wang, Yan Jiang, Kangyang Xie, Zide Liu, Hao Chen, Yue Cao, Xinlong Wang, and Chunhua Shen. 2023. Zero-shot video editing using off-the-shelf image diffusion models. arXiv preprint arXiv:2303.17599.

Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiuniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. 2024b. Videocomposer: Compositional video synthesis with motion controllability. Advances in Neural Information Processing Systems, 36.

Joyce Jiyoung Whang, Inderjit S Dhillon, and David F Gleich. 2015. Non-exhaustive, overlapping k-means. In Proceedings of the 2015 SIAM international conference on data mining, pages 936–944. SIAM.

Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. 2023a. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In Proceedings of the International Conference on Computer Vision (ICCV).

Jianzong Wu, Xiangtai Li, Chenyang Si, Shangchen Zhou, Jingkang Yang, Jiangning Zhang, Yining Li, Kai Chen, Yunhai Tong, Ziwei Liu, and 1 others. 2024. Towards language-driven video inpainting via multimodal large language models. arXiv preprint arXiv:2401.10226.

Wenhao Wu, Haipeng Luo, Bo Fang, Jingdong Wang, and Wanli Ouyang. 2023b. Cap4video: What can auxiliary captions do for text-video retrieval? In Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR).

Shaoan Xie, Zhifei Zhang, Zhe Lin, Tobias Hinz, and Kun Zhang. 2023. Smartbrush: Text and shape guided object inpainting with diffusion model. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 22428–22437.

Jinbo Xing, Menghan Xia, Yong Zhang, Haoxin Chen, Xintao Wang, Tien-Tsin Wong, and Ying Shan. 2023. Dynamicrafter: Animating open-domain images with video diffusion priors. arXiv preprint arXiv:2310.12190.

Wilson Yan, Yunzhi Zhang, Pieter Abbeel, and Aravind Srinivas. 2021. Videogpt: Video generation using vq-vae and transformers. arXiv preprint arXiv:2104.10157.

Antoine Yang, Arsha Nagrani, Paul Hongsuck Seo, Antoine Miech, Jordi Pont-Tuset, Ivan Laptev, Josef Sivic, and Cordelia Schmid. 2023. Vid2seq: Large-scale pretraining of a visual language model for dense video captioning. In Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR).

Fengting Yang, Qian Sun, Hailin Jin, and Zihan Zhou. 2020. Superpixel segmentation with fully convolutional networks. In Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR).

Xiangpeng Yang, Linchao Zhu, Hehe Fan, and Yi Yang. 2024. Eva: Zero-shot accurate attributes and multi-object video editing. arXiv preprint arXiv:2403.16111.

Zongxin Yang, Yunchao Wei, and Yi Yang. 2021. Associating objects with transformers for video object segmentation. In Advances in Neural Information Processing Systems (NeurIPS).

Zongxin Yang and Yi Yang. 2022. Decoupling features in hierarchical propagation for video object segmentation. In Advances in Neural Information Processing Systems (NeurIPS).

Shoubin Yu, Difan Liu, Ziqiao Ma, Yicong Hong, Yang Zhou, Hao Tan, Joyce Chai, and Mohit Bansal. 2025. Veggie: Instructional editing and reasoning of video concepts with grounded generation. arXiv preprint arXiv:2503.14350.

Shoubin Yu, Jaehong Yoon, and Mohit Bansal. 2024. Crema: Multimodal compositional video reasoning via efficient modular adaptation and fusion. arXiv preprint arXiv:2402.05889.

Tao Yu, Runseng Feng, Ruoyu Feng, Jinming Liu, Xin Jin, Wenjun Zeng, and Zhibo Chen. 2023. Inpaint anything: Segment anything meets image inpainting. arXiv preprint arXiv:2304.06790.

Hang Zhang, Xin Li, and Lidong Bing. 2023a. Video-llama: An instruction-tuned audio-visual language model for video understanding. In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP).

Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. 2024. Llava-next: A strong zero-shot video understanding model. https://llava-vl.github.io/blog/2024-04-30-llava-next-video/.

Zhixing Zhang, Bichen Wu, Xiaoyan Wang, Yaqiao Luo, Luxin Zhang, Yinan Zhao, Peter Vajda, Dimitris Metaxas, and Licheng Yu. 2023b. Avid: Any-length video inpainting with diffusion model. arXiv preprint arXiv:2312.03816.

Zicheng Zhang, Bonan Li, Xuecheng Nie, Congying Han, Tiande Guo, and Luoqi Liu. 2023c. Towards consistent video editing with text-to-image diffusion models. In Advances in Neural Information Processing Systems (NeurIPS).

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. Advances in Neural Information Processing Systems, 36.

Luowei Zhou, Chenliang Xu, and Jason Corso. 2018a. Towards automatic learning of procedures from web instructional videos. In Proceedings of the AAAI Conference on Artificial Intelligence.

Luowei Zhou, Yingbo Zhou, Jason J Corso, Richard Socher, and Caiming Xiong. 2018b. End-to-end dense video captioning with masked transformer. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 8739–8748.

Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, and Huaxiu Yao. 2024. Analyzing and mitigating object hallucination in large vision-language models. In Proceedings of the International Conference on Learning Representations (ICLR).

# Appendix

In this appendix, we present the following:

- More details about VPLM dataset collection (Sec. A.1), experimental setups (Sec. A.2), more implementation details (Sec. A.3).
- Additional analysis including ablations (Sec. B.2, Sec. B.1).
- Additional qualitative examples with RACCOON on video content editing (Sec. B.3).

## A Experimental Setup

### A.1 VPLM data collection.

As we mention in Sec. 3.3, to facilitate our model training, we start from open-source video inpainting data (Wu et al., 2024)[3] to build a high-quality dataset that includes the well-structured, detailed caption for both video and each object in the video. Specifically, we leverage GPT-4V [4] to annotate each video. As shown in Fig. 6, we first convert a video to a superimage (Fan et al., 2021) by concatenating uniformly sampled frames, and we also draw frame IDs on each frame as a visual prompt to present temporal order. Then we prompt the GPT-4V by providing a detailed prompt with in-context examples (left of Fig. 6). In this case, we obtained well-structured, detailed captions that contain both holistic video and local objects (bottom right of Fig. 6). To ensure the annotation quality, we sampled 1 annotated video from each 100 batches and then did a human cross-check, and refined the batch annotations according to the sampled example. Through this pipeline, we obtained 7.2K high-quality quality well-structured, detailed video descriptions with an average of 238.0 words for each video.

Next, to obtain paired object-mask-description triplets for video inpainting model training, we build an automatic detailed object caption and object name matching pipeline using GPT4. As in our base dataset (Wu et al., 2024), we already have class labels for each object mask, we framed this matching as a multi-choice QA to ask GPT4 which object caption can in Fig. 6 matched to the given object classes. We further filtered out the triplets with too small masks ($< 1\%$ mask areas.) In this case, we obtained 5.5K object-mask-description triplets with an average of 37.2 words for each object to support our video diffusion model training.

[3] MIT License: https://choosealicense.com/licenses/mit/

[4] version 1106

### A.2 Benchmarks and Datasets Details

As mentioned in the main paper (Sec. 4), we evaluate our proposed RACCOON on various tasks. For video-to-paragraph generation, we test our model on the standard video caption dataset YouCook2 (Zhou et al., 2018a) (validation set) as well as our VPLM dataset. We next test video content editing with three subtasks on our VPLM dataset. Regarding the experiments of incorporating RACCOON with other conditional video generation models, we test RACCOON on diverse videos from ActivityNet, YouCook2, and UCF101. We uniformly selected 100 videos from those 3 datasets to build the test bed. For the experiments of incorporating RACCOON with other video editing models, we follow the previous work (Geyer et al., 2023), and select 30 unique videos from the DAVIS dataset. For each video, we annotate two different types of editing, attribute editing and instance editing. It leads to 60 text-video pairs in our video editing evaluation. We choose object captions that contain the same keywords for editing in human captions to represent the model-generated caption.

### A.3 Implementation Details

In V2P generation, we set $k = [20, 25]$ and $v = [5, 6]$ for superpixel clustering. We use CLIP-L/14@336 (Radford et al., 2021) as the image encoder and Vicuna-1.5 (Zheng et al., 2024) as the LLM. Our P2V model starts from StableDiffusion-2.0-Inpainting (Rombach et al., 2022). We split the VPLM datasets into train and test sets, with the test set containing 50 unique video-paragraph pairs (for V2P) and 180 mask-object-description triples (for P2V). We manually annotate the editing prompts for the object-changing subtask. We quantitatively compare RACCOON and other baselines on the VLPM test set. To focus on generation results rather than grounding ability, we apply the same ground truth masks and captions to all methods for P2V evaluation. See the supplementary material for more details on datasets, metrics, implementations, ablations, and qualitative analysis.

**Metrics**: We provide more details about our metrics. CLIP-Text measures the similarity between the edit prompt and the embedding of each frame in the edited video. CLIP-Frame computes the

**Input Video**

**RACCooN Generated Paragraphs**

The video shows a woman walking down a busy city street at night. She is wearing a red dress and carrying a handbag. She stops to look at a building and then walks away. The woman is also seen walking down a rain-soaked street and looking at a building. She then walks away again. The video also shows a man walking down the street and looking at a building. The woman is also seen walking down a street and looking at a building. The video ends with a woman walking down a street and looking at a building. Overall, the video captures the busy and bustling atmosphere of a city at night. The woman's red dress and handbag stand out as she walks down the street, and her attention to the buildings she looks at adds depth to the scene. The man walking down the street also adds to the overall sense of a busy urban environment.

Object Catalog:
1. Woman: A stylish woman wearing a black leather jacket, a long red dress, and black boots. She carries a black purse and has sunglasses on.
2. Traffic lights: There are traffic lights visible in the background, indicating that the street is a busy intersection.
3. Building: A large, multi-story structure with a prominent presence in the background. It might be a commercial or residential building.
4. Handbag: A small, handheld accessory carried by the woman. It could be used for holding personal belongings or as a fashionable accessory.

Figure 5: **Qualitative V2P example of our RACCOON** on Sora video.

average CLIP similarity between the image embeddings of consecutive frames to measure the temporal coherence. SSIM measures the structural similarity between the original and edited video for evaluating localized editing. $Q_{edit} = CLIP - T/Wrap - Err$, it is a comprehensive score for video editing quality, where $Warp - Err$ calculates the pixel-level difference by warping the edited video frames according to the estimated optical flow of the source video, extracted by FlowNet2.0 (Ilg et al., 2017). For layout planning, we compute the FVD and CLIP-Image similarity between the ground truth and the predicted bounding box.

**More Details about Multi-granular Spatiotemporal Pooling.** As mentioned Sec. 3.1, we proposed a novel Multi-granular Spatiotemporal Pooling (MGS Pooling) to address the lack of complex spatial-temporal modeling in video. We further provide a more intuitive visualization for our proposed MGS Pooling in Fig. 3. We first adopt a lightweight 165 superpixel predictor that generates superpixels across video frames, then use the overlapping k-means clustering for the obtained video superpixels. In this case, we gather informative cues about various objects and actions for LLM.

**Human Evaluation on Video-to-Paragraph Generation.** We conduct a human evaluation on nine randomly selected videos from the YouCook2 dataset. Videos are three- to five-minute-long and contain multiple successive scenes with complex viewpoints. We provide these videos to four different annotators with ground truth captions, descriptions generated by PG-Video-LLAVA, and our method, RACCOON, where the captions/descriptions for each video are randomly shuffled. We leverage four distinct human evaluation metrics: Logic Fluency (Logic), Language Fluency (Language), Video Summary (Summary) and Video Details (Details). To avoid a misinterpretation of methods' capabilities due to relative evaluation, we instruct the annotators to independently rate the quality of each set of captions based on these four different criteria, by giving a score from 1 to 5 (i.e., choice: [1, 2, 3, 4, 5]).

**Off-shelf Video Editing Models.** We utilize TokenFlow (Geyer et al., 2023) and FateZero (Qi et al., 2023) as our video editing tools. TokenFlow generates a high-quality video corresponding to the target text, while preserving the spatial layout and motion of the input video. For SSIM computation, we compute SSIM for the region-of-no-interest since we want to keep those regions unchanged. We first mask out regions of interest with the ground truth mask provided by the DAVIS dataset, then we compute SSIM on masked images and conduct mean pooling as the video-level metrics.

**Off-shelf Conditional Video Generation Models.** We leverage both VideoCrafter (Chen et al., 2023) and DynamiCrafter (Xing et al., 2023) as our video generation backbone. VideoCrafter is one of the SoTA video generation models that can handle different input conditions (image, text). DynamiCrafter is based on the open-source VideoCrafter and T2I Latent Diffusion model (Rombach et al., 2022), and was trained on WebVid10M (Bain et al., 2021), it provides bet-

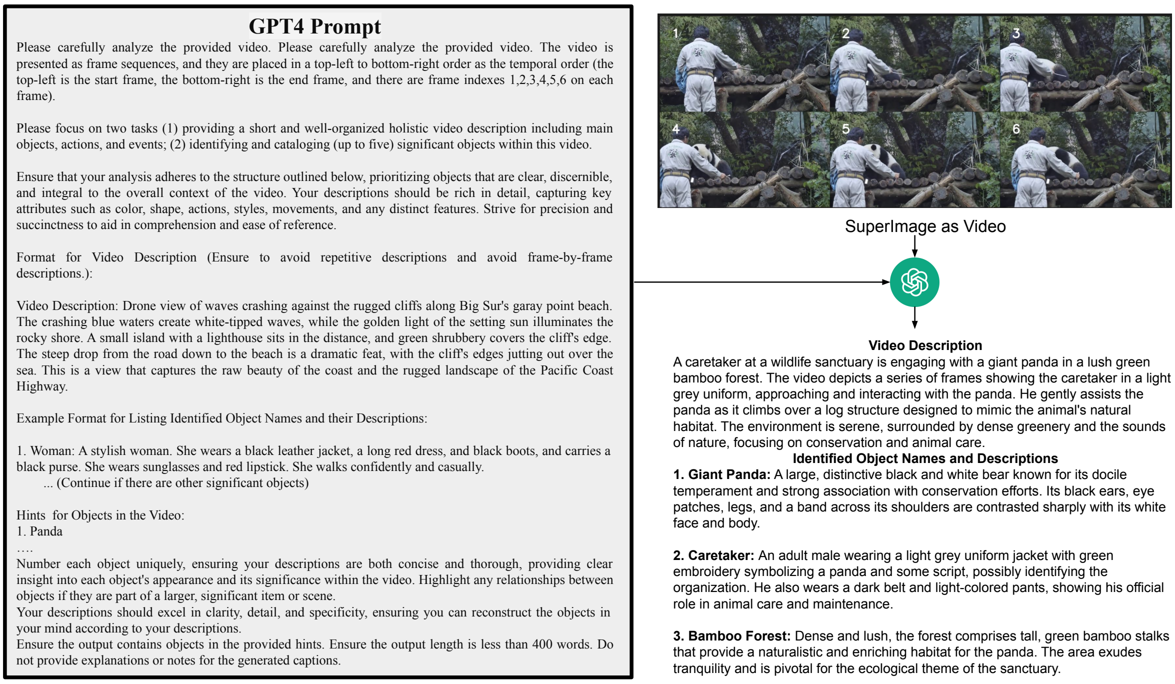


Figure 6: **Pipeline of our VPLM dataset annotation with GPT4V**. We first convert a video as a superimage and then give some in-context examples to prompt GPT-4V to annotate detailed and well-structured video descriptions.

ter dynamics and stronger coherence. We adopt VideoCrafter-512 and DynamiCrafter-512 variants. For each video, we use CLIP similarity to retrieve multiple keyframes corresponding to each caption. Those keyframes result in multiple generated video clips via the video generation model. For FVD computation, we conduct mean pooling over those clips to represent a video. We use $k = 25$ and $v = 6$ for generated captions in all experiments. The experiments are conducted on the 4 × 48GB A6000 GPUs machine.

# B Additional Analysis

## B.1 Ablation study

**The effect of $k$ and $v$** As shown in Tab. 5, we did initialized hyperparameter probing experiments on ActivityNet-Cap and YouCook2 datasets. we observe that all variants of our approach with varying $k$ and $v$ generally achieve improved performance compared to baselines in terms of multiple video captioning metrics: *SPICE*, *BLEU-4*, *METEOR*, and *ROUGE*. This result demonstrates the efficacy of our multi-granular spatiotemporal pooling approach with a fine-to-coarse search of video contexts based on superpixels. In addition, we observe that RACCOON shows a small gap between variants in each dataset, highlighting the robustness of our approach to the hyperparameter setups and datasets.

**The effect of Superpixel Overlap** We introduce overlapping k-means clustering to aggregate video superpixels, capturing a variety of visual contexts while allowing for partial spatiotemporal overlap. To investigate the effect of our suggested overlapping approach, we also evaluate the variant of our approach without overlap (i.e., $v = 1$) on video-to-paragraph generation tasks in Tab. 5. As shown, our approach with overlap (i.e., $v > 1$) surpasses the non-overlapping version of RACCOON across various scales of visual contexts $k$, as indicated by the video captioning metrics we evaluated. This emphasizes the advantage of permitting overlap in understanding video contexts, which enhances the input video's comprehension by allowing for diverse and fluent interpretations of local visual regions with surroundings associated at the same time.

For simplicity, we use $k = 25$ and $v = 6$ for all experiments on conditional video generation and video editing tasks, demonstrating the robustness of RACCOON for hyperparameters.

## B.2 Comparison with Pre-trained Grounding Models

We further investigate the applicability of recent powerful pre-trained visual grounding models (Kirillov et al., 2023; Cheng et al., 2023; Ren et al., 2024). Segment-Anything (Kirillov et al., 2023) and Grounded SAM (Ren et al., 2024) are strong

Table 5: **Ablation of RACCooN** for Video-to-Paragraph Generation on ActivityNet and YouCook2. Metrics are abbreviated: **M**: *METEOR*, **B**: *BLEU-4*, **S**: *SPICE*, **R**: *ROUGE*. $v = 1$ indicates the version without superpixel overlap. We highlight the hyperparameter setup used in the main experiment.

| Models | k | v | ActivityNet | | | | YouCook2 | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | | S | B | M | R | S | B | M | R |
| PDVC (Wang et al., 2021) | - | - | - | 2.6 | 10.5 | - | - | 0.8 | 4.7 | - |
| Vid2Seq (Yang et al., 2023) | - | - | 5.4 | - | 7.1 | - | 4.0 | - | 4.6 | - |
| ZeroTA (Jo et al., 2023) | - | - | 2.6 | - | 2.7 | - | 1.6 | - | 2.1 | - |
| PG-VL (Munasinghe et al., 2023) | - | - | 13.6 | 13.9 | 14.2 | 18.1 | 6.2 | 16.5 | 8.6 | 15.8 |
| **RACCooN (Ours)** | 20 | 1 | 13.5 | 13.9 | 14.2 | 18.1 | 6.3 | 16.9 | 8.7 | 15.9 |
| | | 2 | 13.7 | 14.6 | 14.4 | 18.2 | 6.4 | 17.5 | 8.7 | 16.1 |
| | | 4 | 13.6 | 14.3 | 14.3 | 18.2 | 6.6 | 16.2 | 8.8 | 16.0 |
| | | 5 | **13.8** | **15.0** | **14.5** | **18.4** | 6.4 | 17.9 | 8.8 | **16.2** |
| | 25 | 1 | 13.6 | 14.1 | 14.3 | 18.0 | 6.1 | 16.9 | 8.6 | 15.9 |
| | | 2 | **13.8** | 14.4 | 14.3 | 18.3 | 6.4 | 16.3 | **9.0** | 16.0 |
| | | 4 | 13.6 | 14.3 | 14.3 | 18.2 | 6.6 | 17.1 | 8.9 | 16.1 |
| | | 6 | 13.7 | 14.5 | 14.4 | 18.2 | **6.9** | 18.0 | **9.0** | 16.1 |
| | | 10 | - | - | - | - | 6.4 | 16.5 | 8.7 | 16.1 |
| | 30 | 1 | 13.6 | 14.1 | 14.3 | 18.0 | 6.3 | 16.5 | 8.8 | 16.0 |
| | | 2 | 13.7 | 14.2 | 14.2 | 18.1 | 6.6 | 17.1 | **9.0** | **16.2** |
| | | 4 | 13.5 | 14.5 | 14.3 | 18.2 | 6.6 | **18.1** | 8.8 | 16.1 |
| | | 6 | 13.6 | 14.4 | 14.4 | 18.2 | 6.4 | 17.2 | 8.8 | **16.2** |

Table 6: **RACCooN variants with different grounding methods** for Video-to-Paragraph Generation on YouCook2.

| Method | Localization | Clustering | SPICE | BLEU-4 | METEOR | ROUGE |
|---|---|---|---|---|---|---|
| PG-VL (Munasinghe et al., 2023) | - | - | 6.2 | 16.5 | 8.6 | 15.8 |
| Ours | SAM (Kirillov et al., 2023) | - | 6.4 | 16.9 | 8.7 | 15.9 |
| | Grounded SAM (Ren et al., 2024) | - | 6.5 | 16.5 | 8.7 | **16.1** |
| | SAM-Track (Cheng et al., 2023) | k-means | 6.2 | 16.5 | 8.8 | **16.1** |
| | SAM-Track (Cheng et al., 2023) | overlapping k-means | 6.5 | 17.4 | **9.0** | 16.1 |
| | Superpixel | overlapping k-means | **6.9** | **18.0** | **9.0** | **16.1** |

open-ended object segmentation models for images, and we directly compute our localized granular tokens based on their segmentation masks. We select 25 segmentation masks in total, from uniformly sampled frames in each video for fair comparison with RACCooN ($k = 25$). As shown in Tab. 6, these variants of RACCooN achieve improved performance against the best-performing baseline, PG-VL, but are often suboptimal since they focus on regional information and cannot contain the temporal information of the videos. Unlike these image-based segmentation methods, SAM-Track (Cheng et al., 2023) generates coherent masks of observed objects over successive frames in videos, by adopting multiple additional pre-trained modules, including GroundingDino (Liu et al., 2023d) and AOT (Yang et al., 2021; Yang and Yang, 2022). We adopt SAM-Track to initialize superpixels in videos and conduct overlapping k-means clustering ($k = 25$). Here, we observe that RACCooN with SAM-Track superpixel initialization achieves reasonable performance, and is beneficial for video editing tasks. It enables the model to coherently edit targeted regions in videos with edited keywords.

### B.3 Additional Visualizations

In this section, we provide more qualitative examples of various tasks, including three types of video content editing, ablation on removing oracle planning and GT masks, enhanced video editing, and conditional video generation.

**Remove, Add, and Change Object the videos.** We provide more qualitative examples in this Appendix across different types of video content editing ( Fig. 7), including removing (Fig. 8, Fig. 9, Fig. 10), adding (Fig. 11, Fig. 12, Fig. 13), and changing/editing (Fig. 14, Fig. 15, Fig. 16). According to the visualization, our RACCooN generally outperforms other strong baselines on all three subtasks. RACCooN can reflect the updated text description more accurately, thus aiding in user-friendly video editing. For example, our method can accurately change the color of the hat (*red→blue*), which is a very small area in the video, while other methods struggle to meet the requirement.

**Ablation Study Visualization.** We illustrate extra visualizations for replacing orecle mask with grounding&tracking tools generated ones for video object removal ( Fig. 17 and  Fig. 18) and changing ( Fig. 19 and  Fig. 20), as well as replacing oracle object boxes with our model-predicted one ( Fig. 21 and  Fig. 22). RACCOON shows robust results with LLM planning and off-shelf segmentation tools. We further show that the failure cases of removing and changing objects mainly come from the missing mask prediction of the video segmentation masks.


Figure 7: More visualization of diverse editing skills on Sora video and comparison with other methods

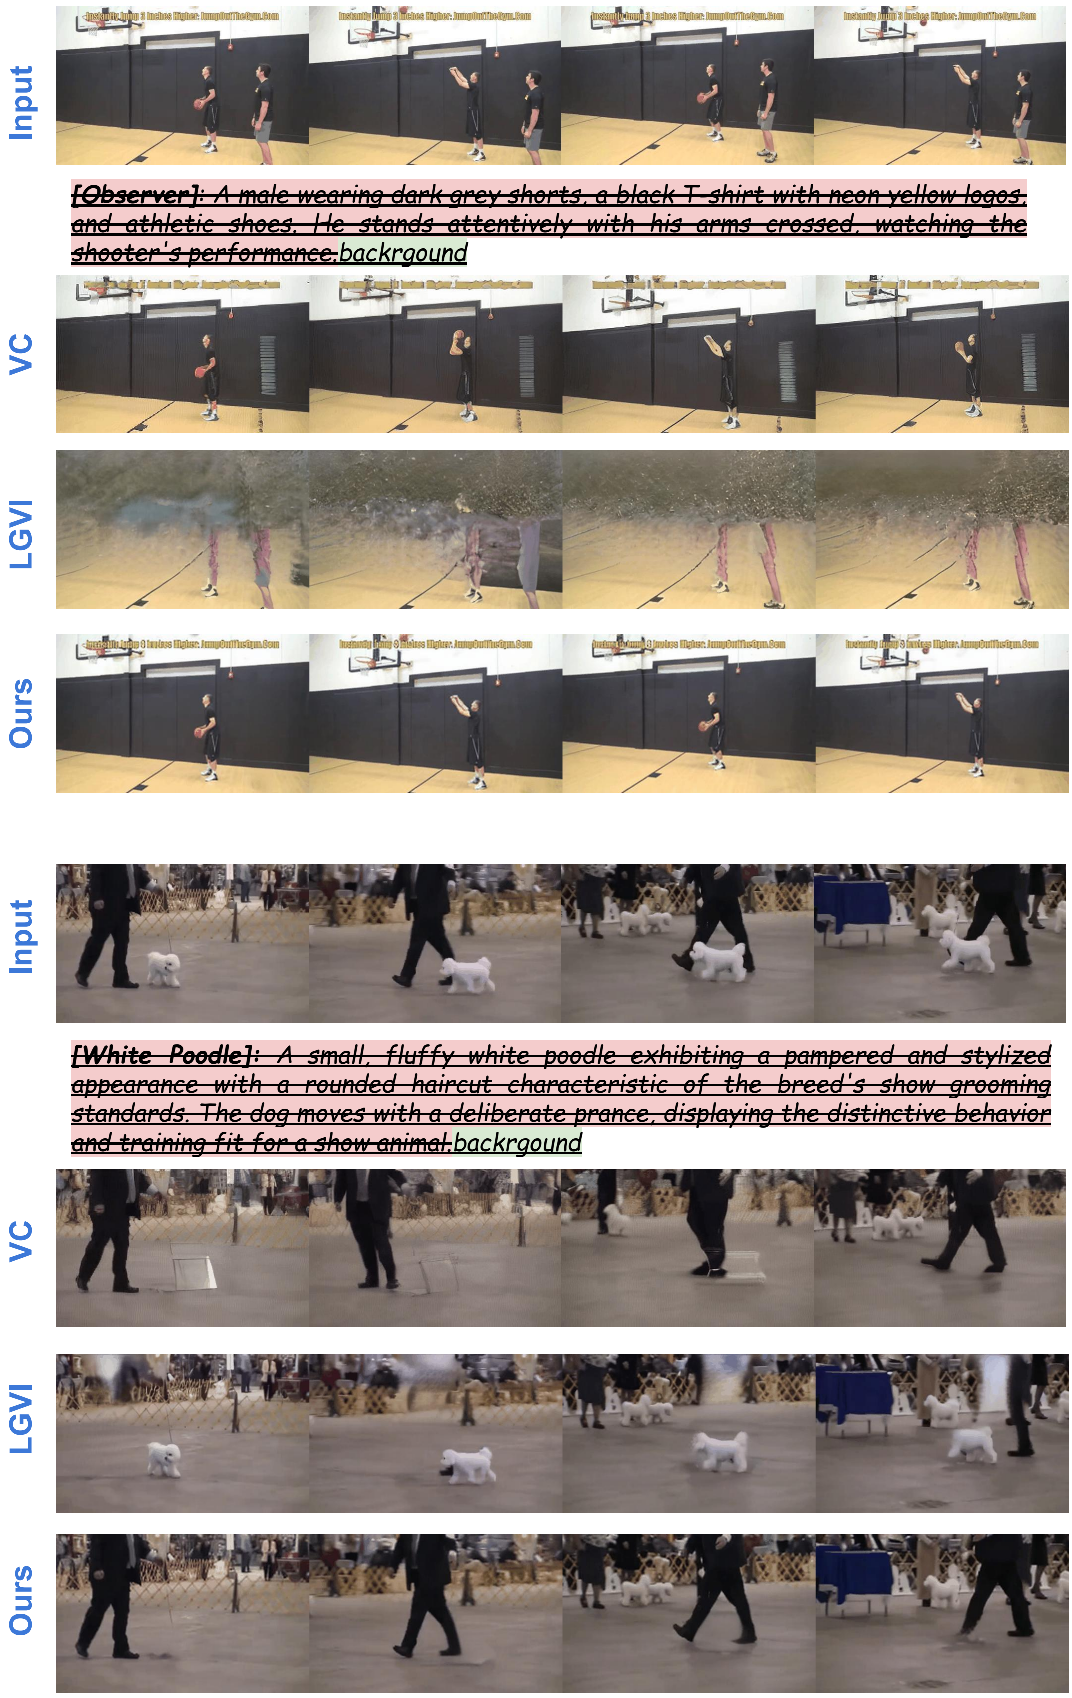


Figure 8: More visualization of **removing** video objects


~~**[Seagull]:** A white seagull with wings spread wide, showcasing gray tips. As it flies, its yellow beak and black eye markings are noticeable, adding to its distinctive features. The bird's underbelly is white, reflecting sunlight, while the light and smooth lines of its body exemplify a strong yet graceful figure in mid-flight.~~ backrgound



~~**[Large Shark]:** A sizable and powerful shark, possibly a Great White, with a grey and white body. Its dorsal fin is tall and prominent, and it has a distinctly streamlined shape. The shark conveys a sense of calm authority as it moves effortlessly through the water.~~backrgound


Figure 9: More visualization of **removing** video objects

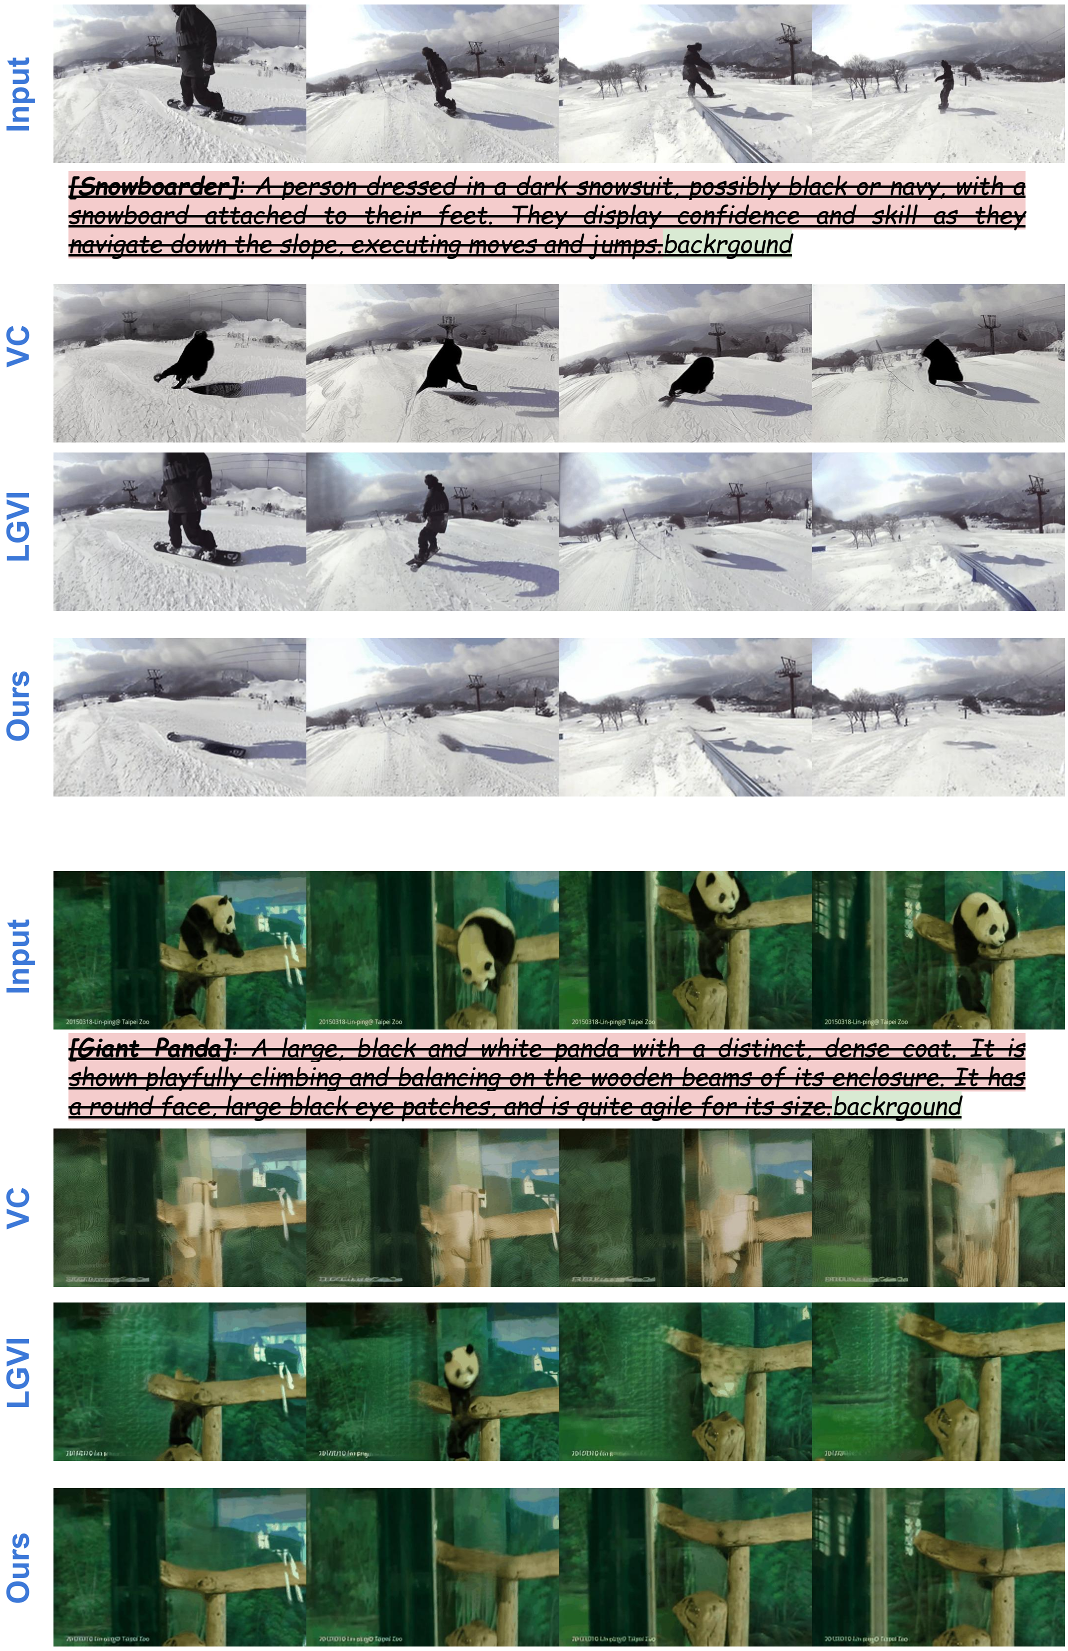


Figure 10: More visualization of **removing** video objects

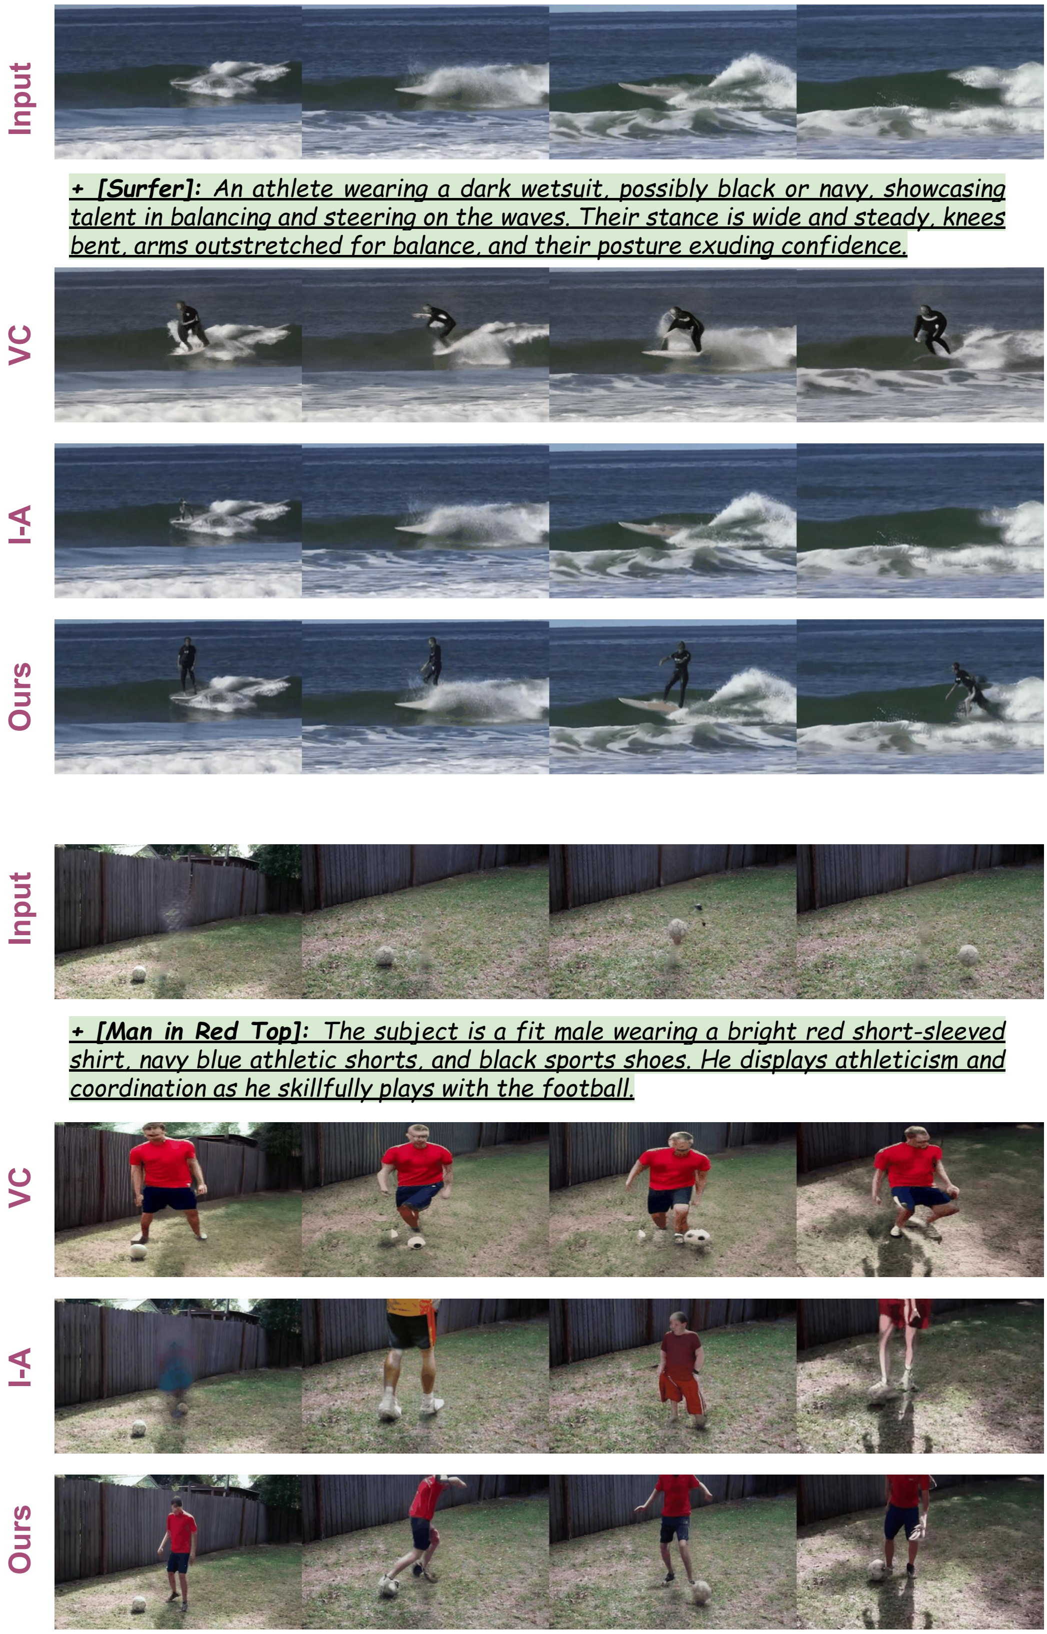


Figure 11: More visualization of **adding** video objects

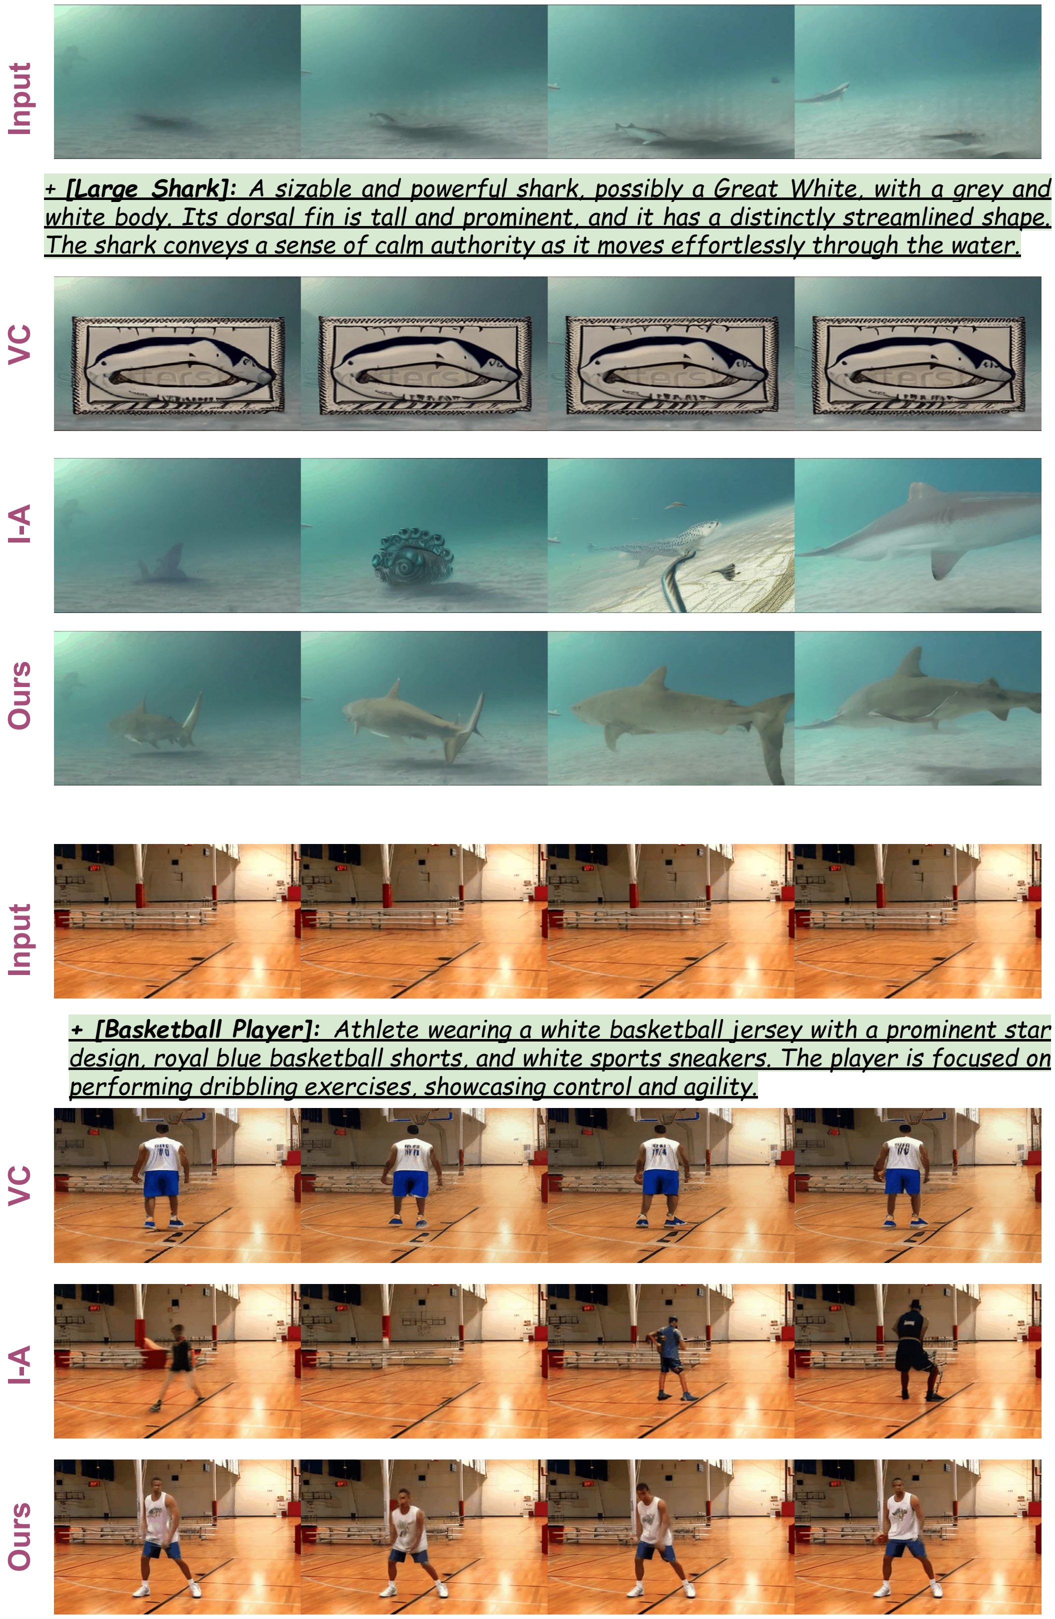


Figure 12: More visualization of **adding** video objects

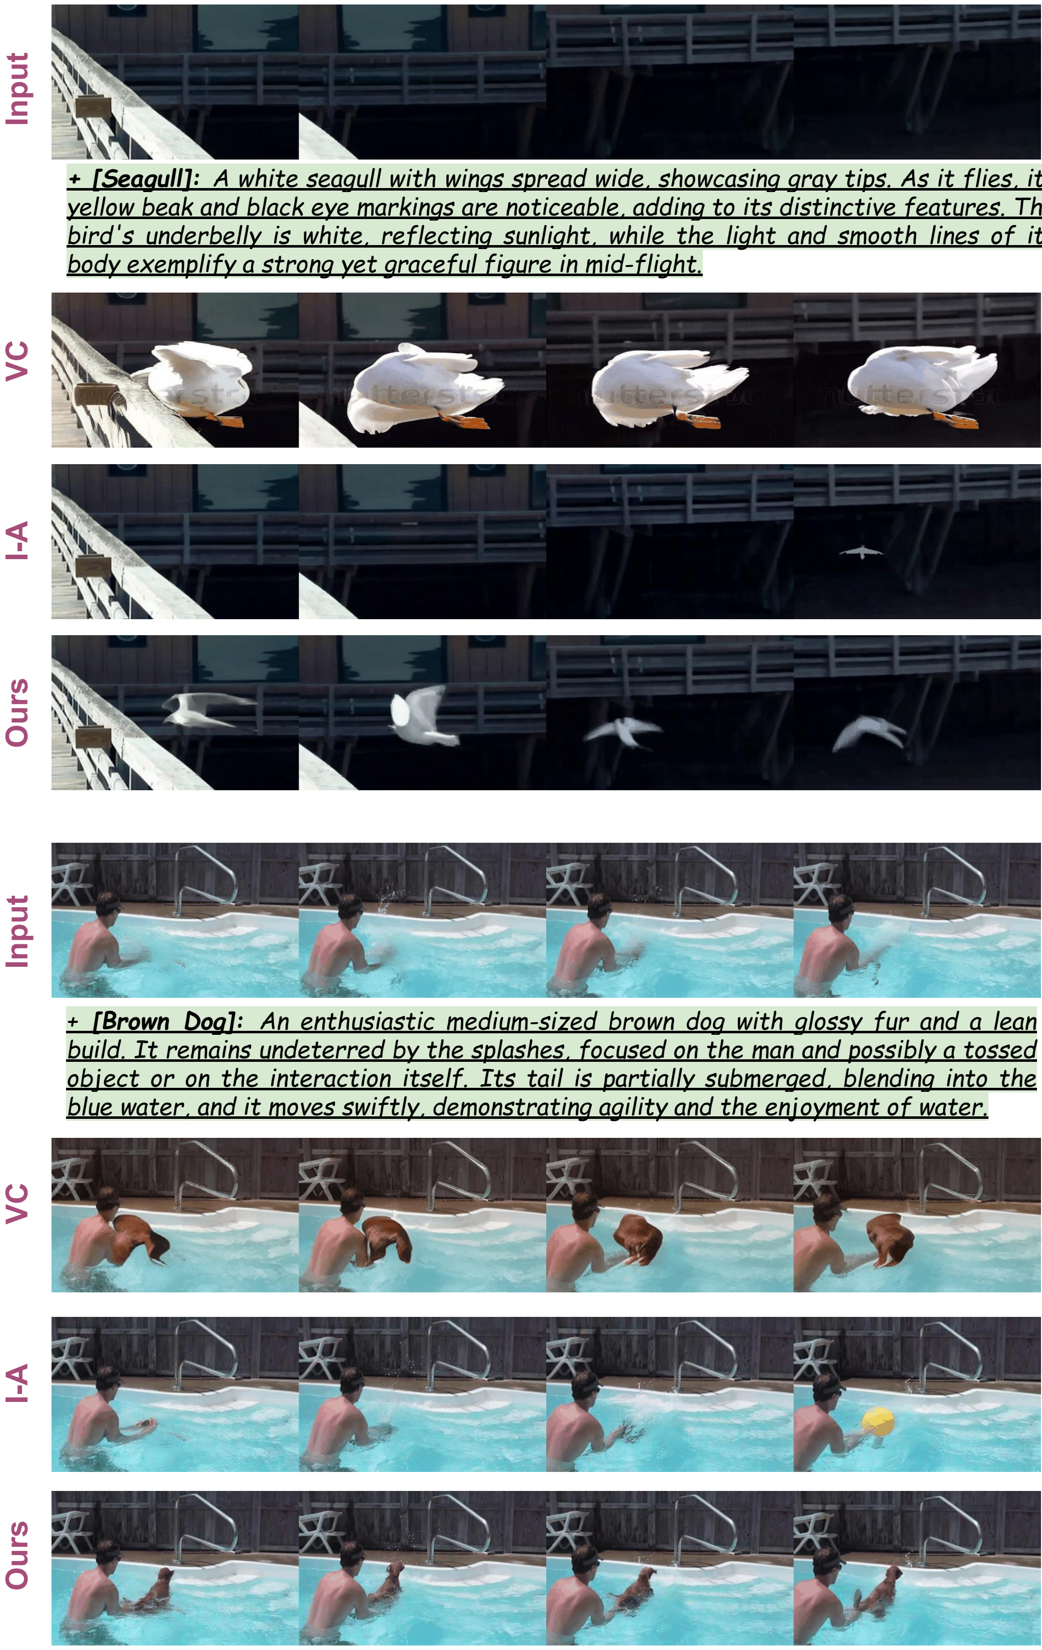


Figure 13: More visualization of **adding** video objects

Input

~~[Large Shark]~~ **[Orange Shark]:** A sizable and powerful shark, possibly a Great White, with a ~~grey and white~~ orange body. Its dorsal fin is tall and prominent, and it has a distinctly streamlined shape. The shark conveys a sense of calm authority as it moves effortlessly through the water.

VC

TF

Ours

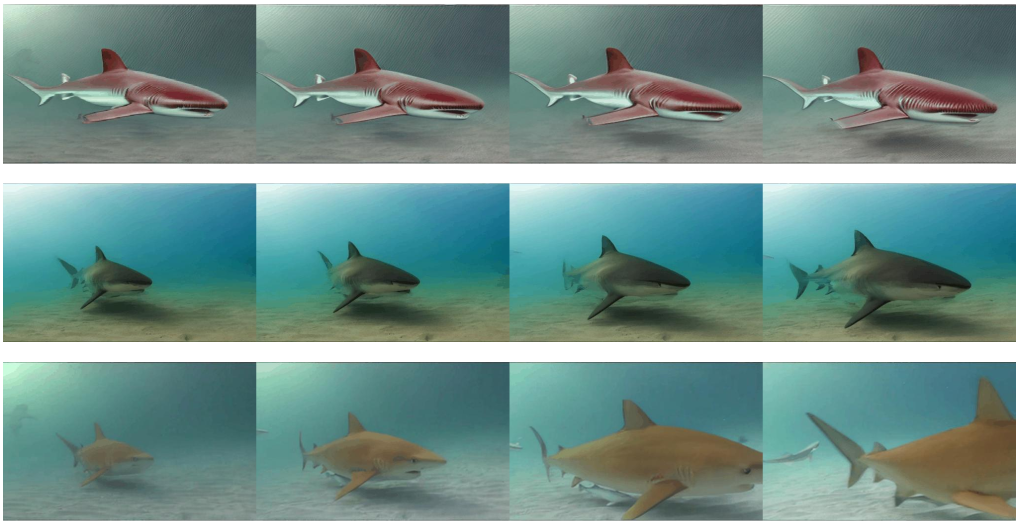

Input

~~[Man]~~ **[Man in orange]:** An athletic man wearing an orange shirt, ~~black~~ white shorts, and sports shoes. He is running on a well-maintained playing area characterized by manicured green grass, a smooth outfield, and a perimeter wall adorned with sponsorship banners and a scoreboard.

VC

TF

Ours

Figure 14: More visualization of **editing** video objects

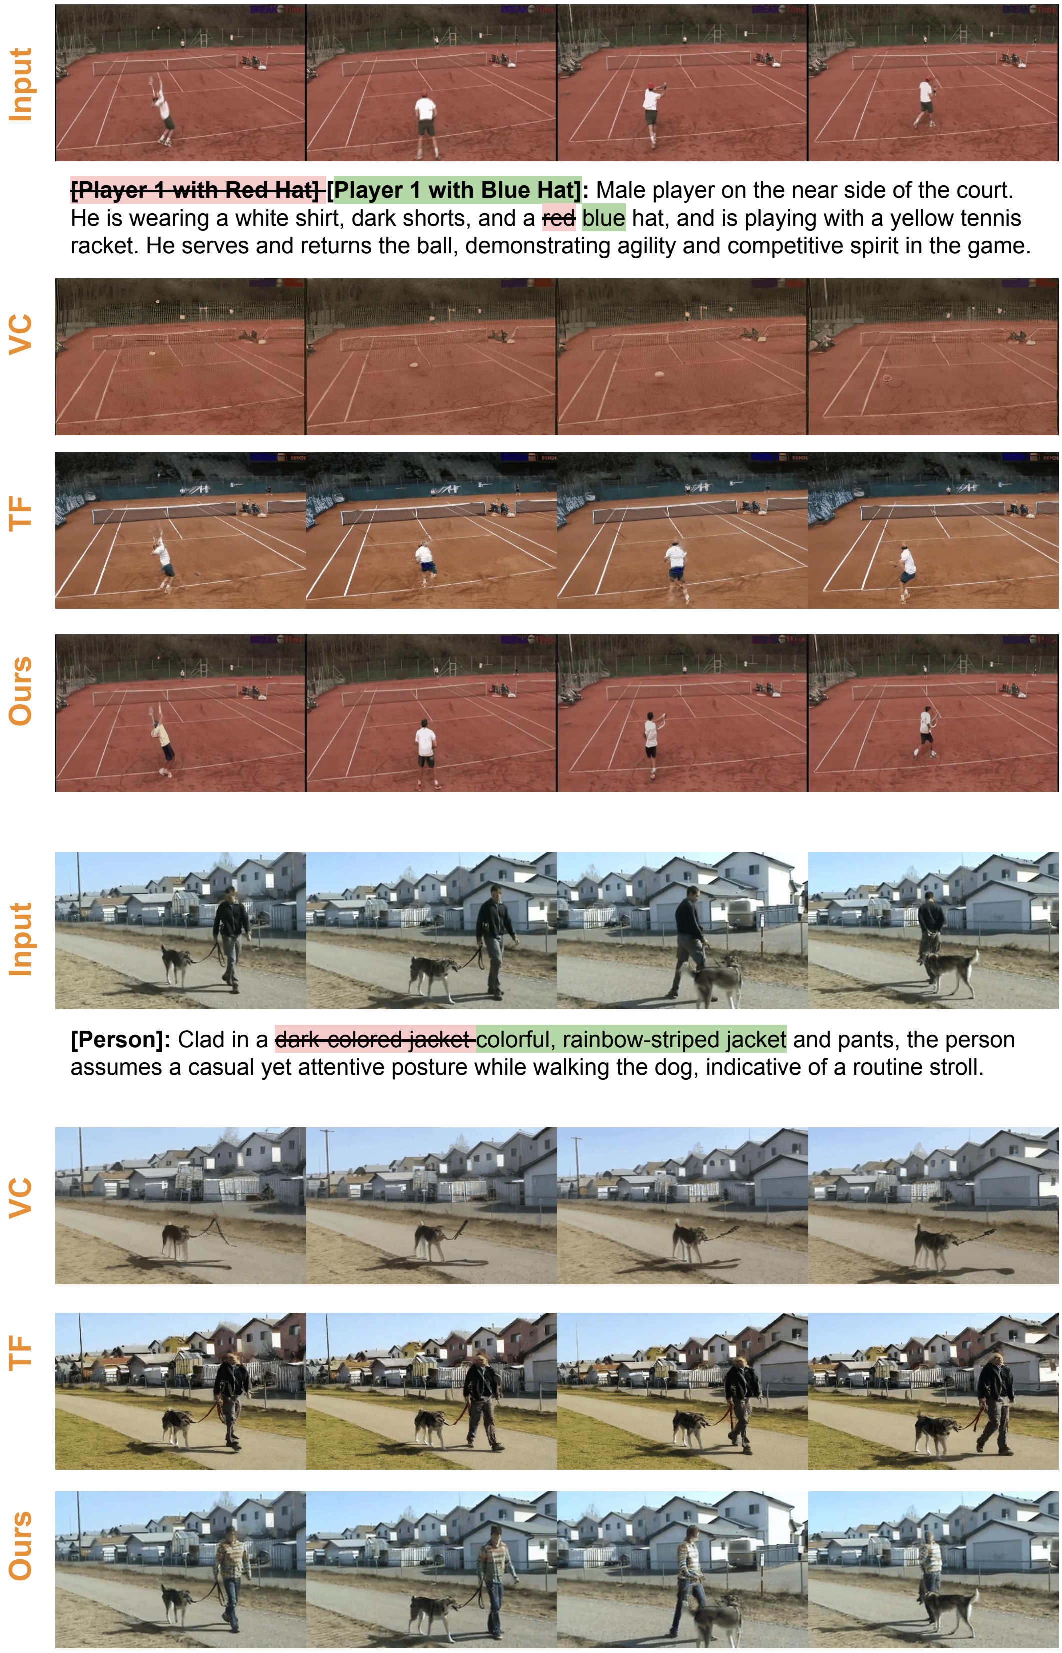


Figure 15: More visualization of **editing** video objects

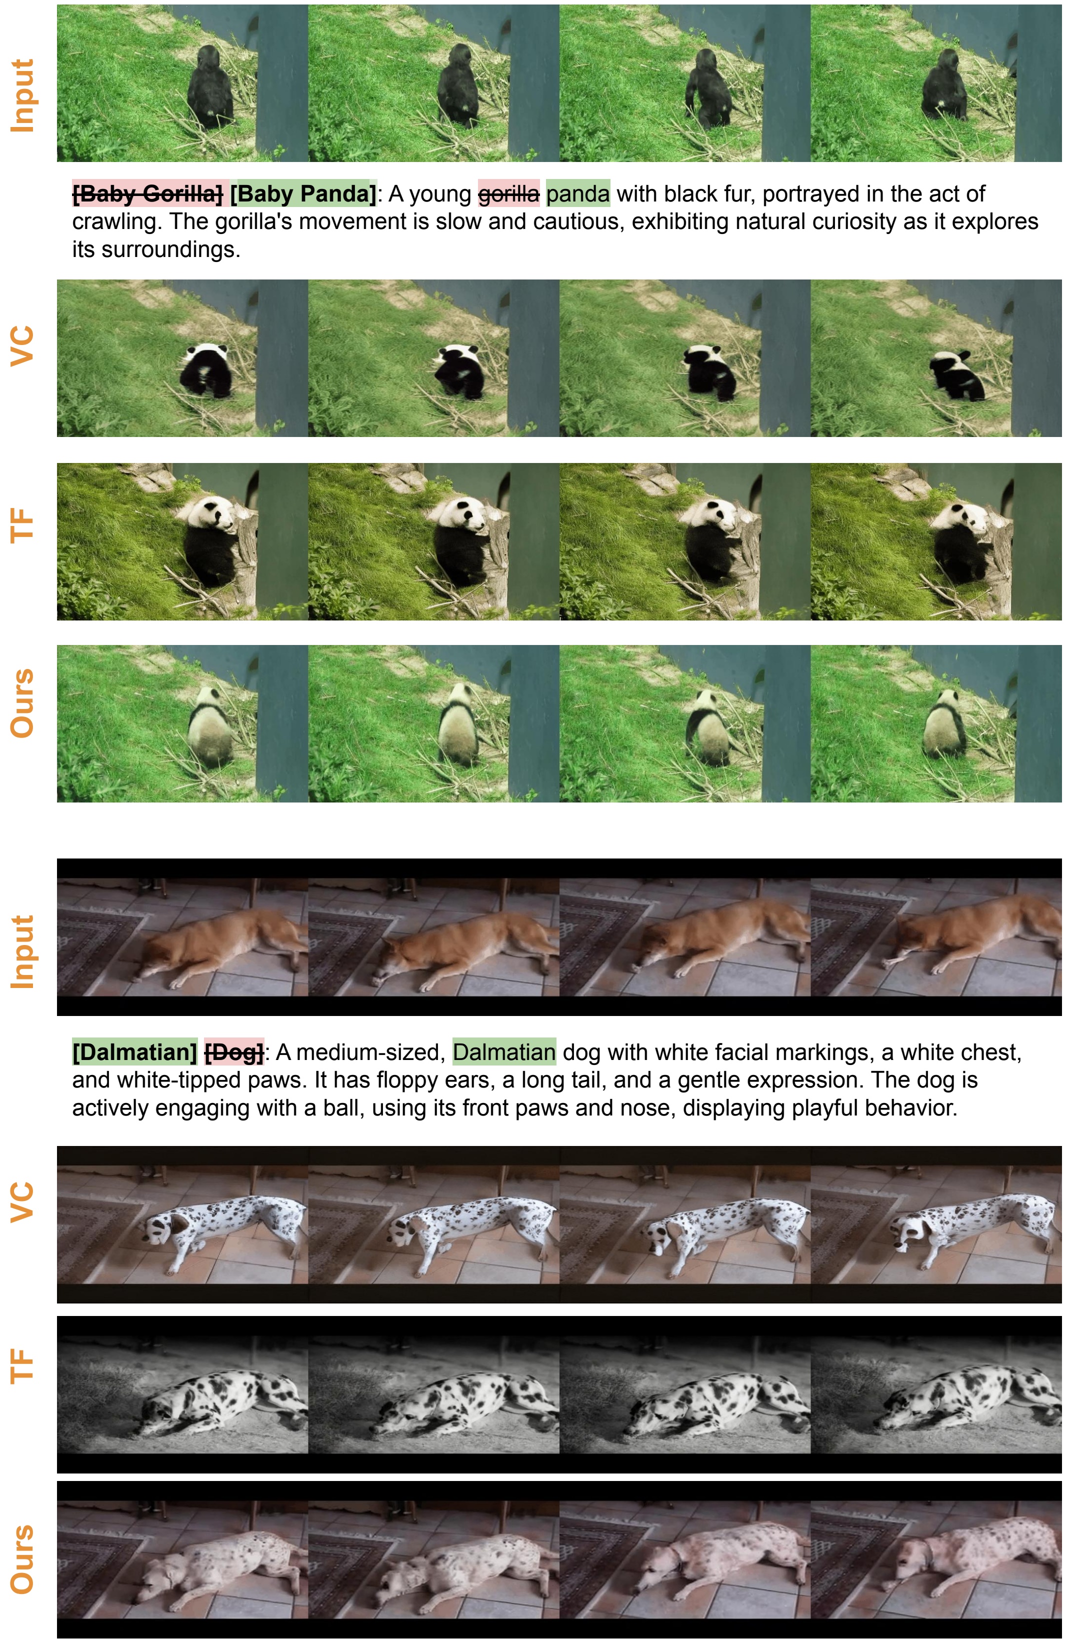


Figure 16: More visualization of **editing** video objects

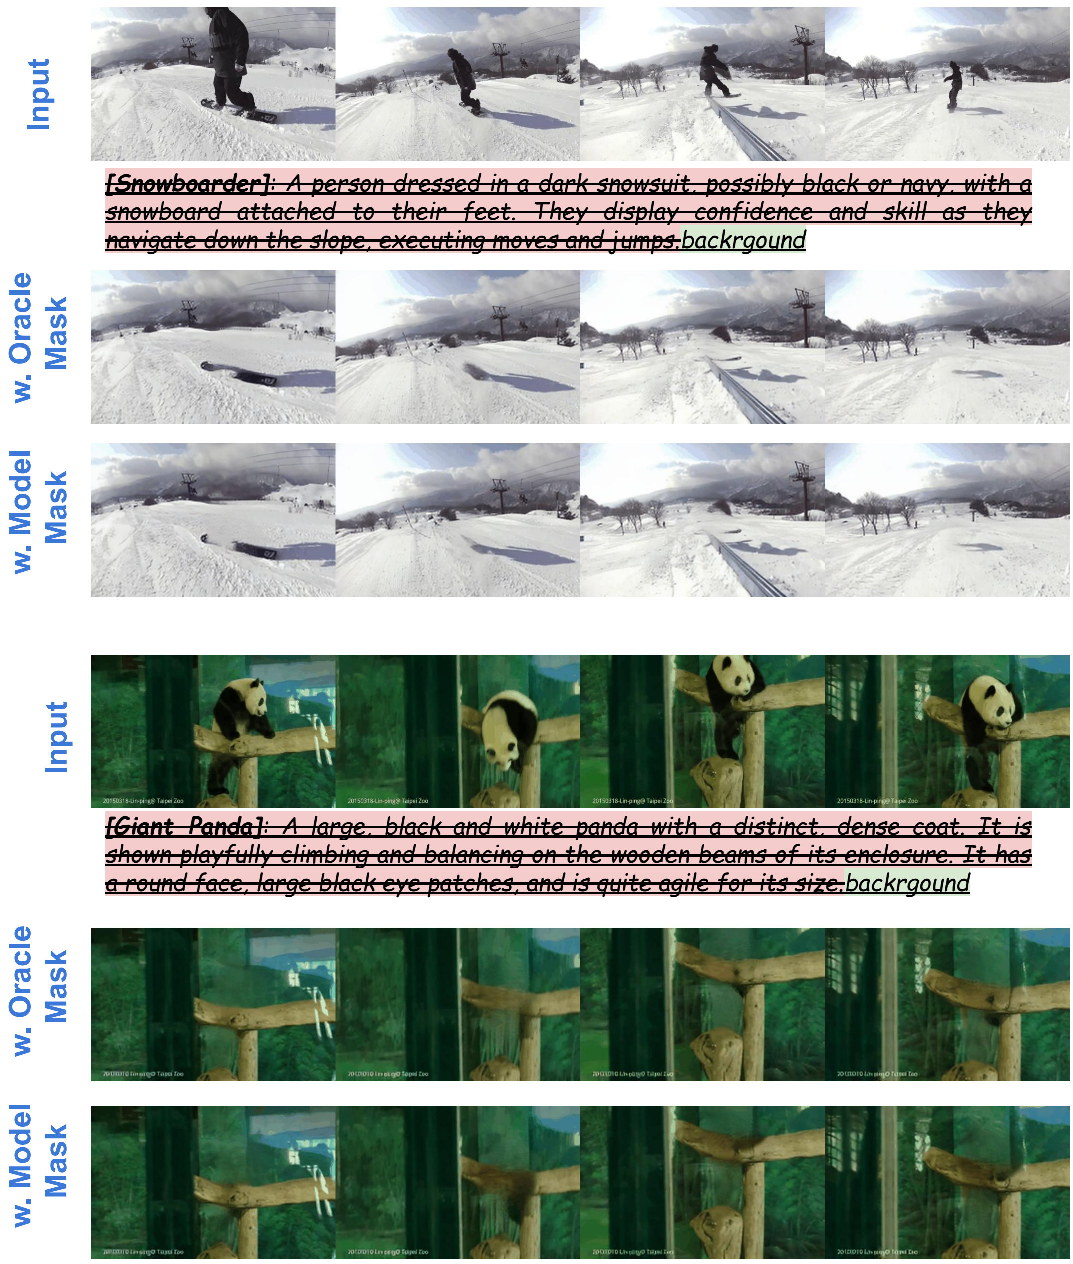


Figure 17: More visualization of **removing** video objects

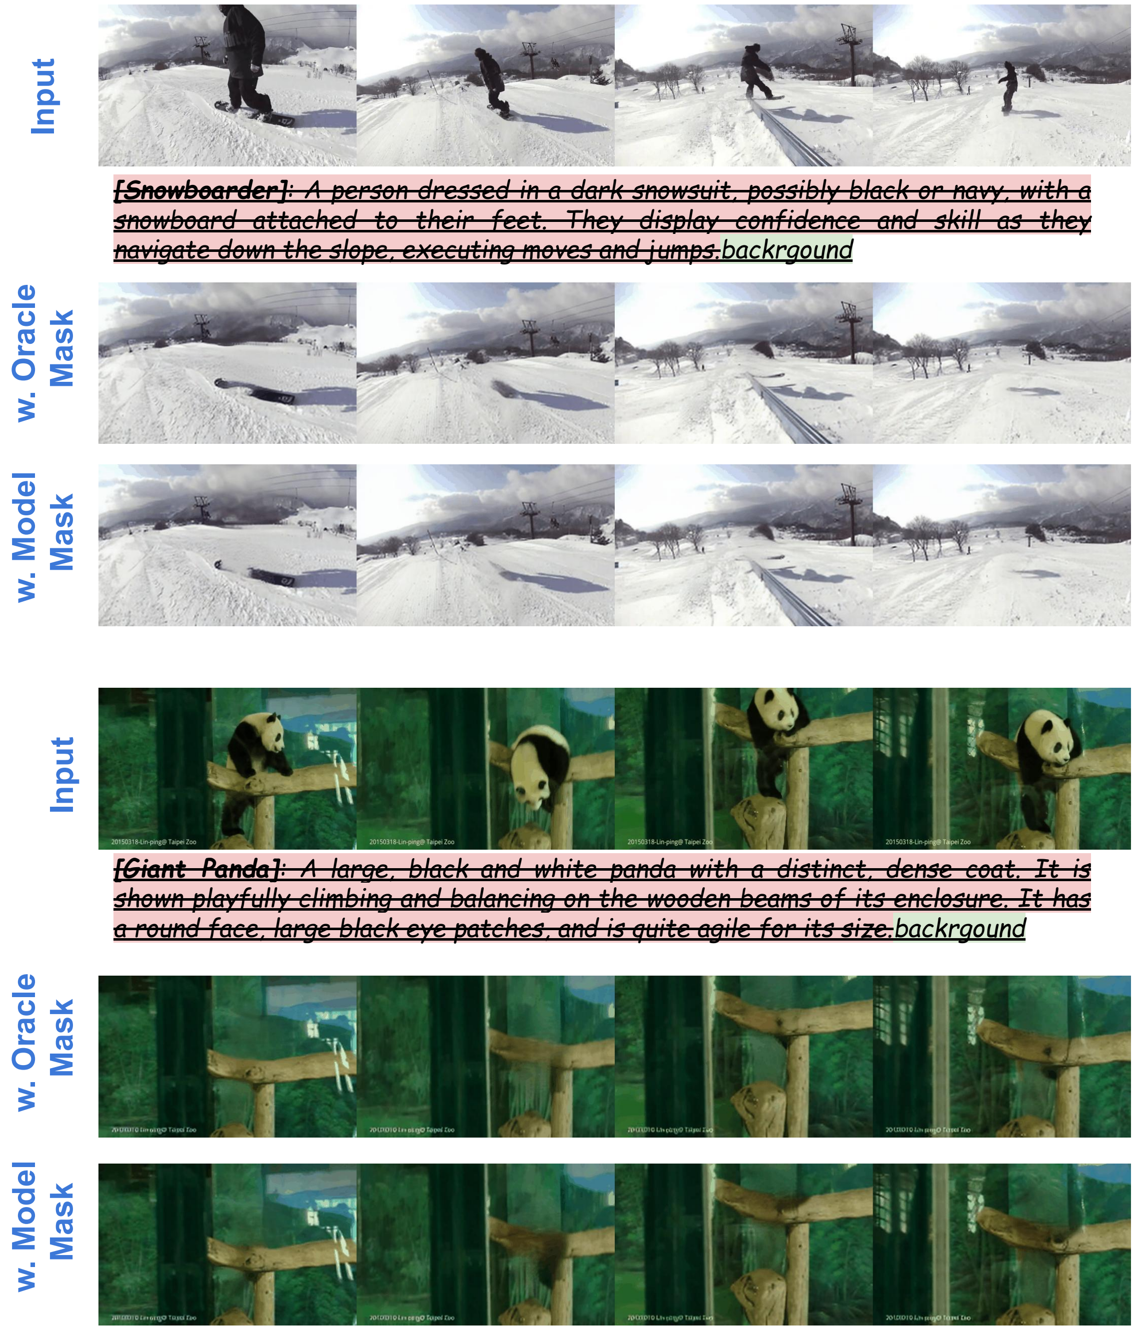


Figure 18: More visualization of **removing** video objects


~~[Baby Gorilla]~~ **[Baby Panda]**: A young ~~gorilla~~ panda with black fur, portrayed in the act of crawling. The gorilla's movement is slow and cautious, exhibiting natural curiosity as it explores its surroundings.



~~*[Brown Dog]*~~ *[Raccoon]:* *A medium-sized,* ~~*light brown dog*~~ *raccoon with a sturdy build and alert ears. Its tail wags enthusiastically as it leads the playful chase in the forest, displaying a sense of happiness and energy.*


Figure 19: More visualization of **editing** video objects


~~**[Large Shark]**~~ **[Orange Shark]:** A sizable and powerful shark, possibly a Great White, with a ~~grey and white~~ orange body. Its dorsal fin is tall and prominent, and it has a distinctly streamlined shape. The shark conveys a sense of calm authority as it moves effortlessly through the water.



***[Woman]:*** *Dressed in a ~~casual outfit of blue denim shorts and a horizontally striped blue and white T-shirt~~ a dark red dinning dress. She accessorizes with pink flip-flops and is engaged in the action of putting a leash on her dog, showing signs of preparing for a walk.*

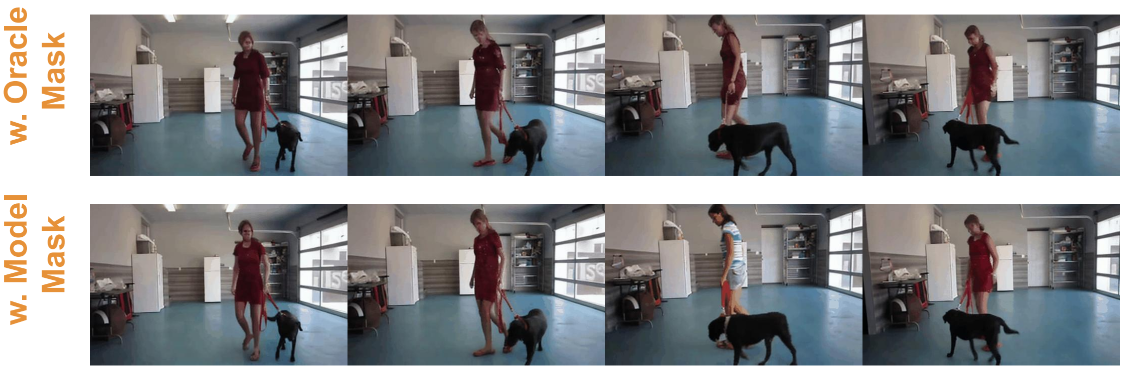


Figure 20: More visualization of **editing** video objects

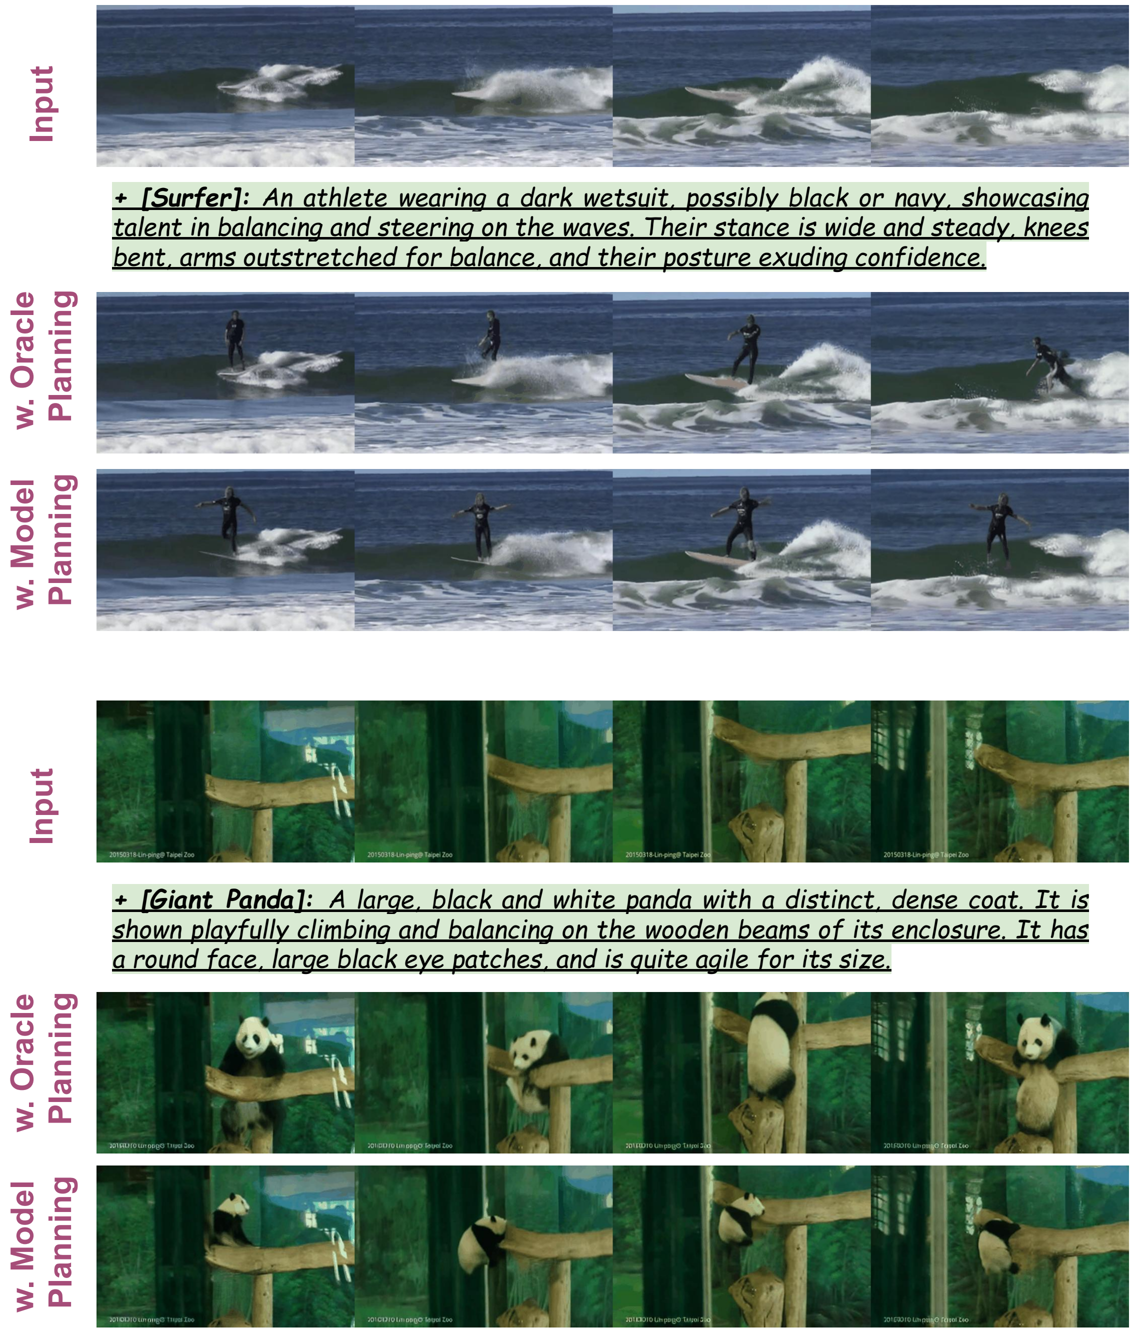


Figure 21: More visualization of **adding** video objects

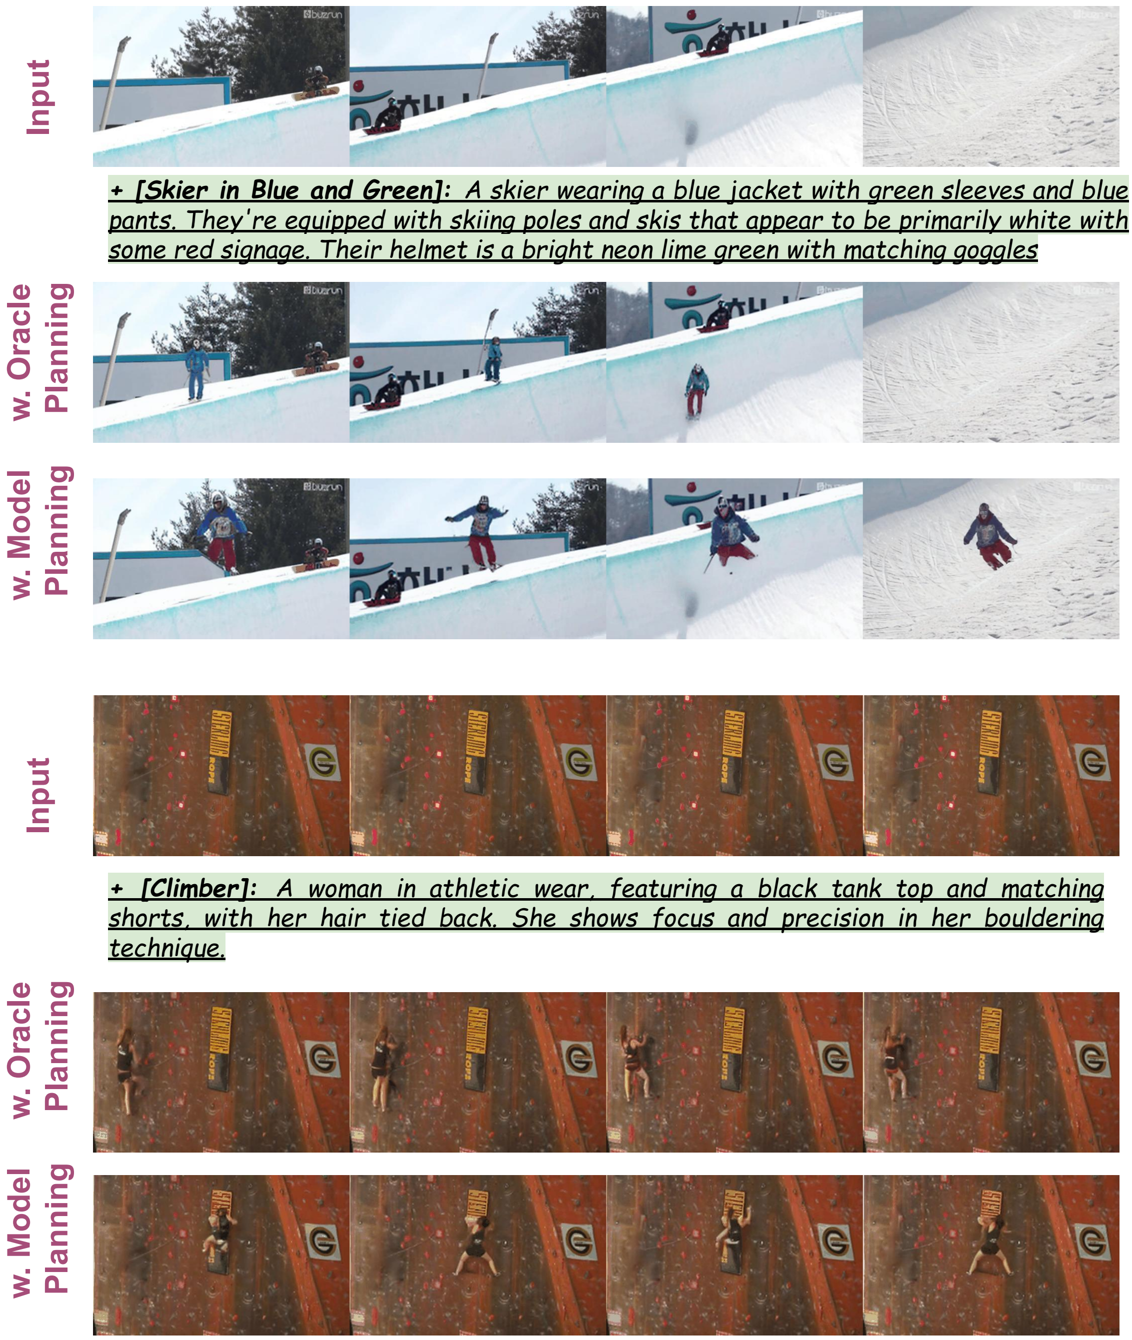


Figure 22: More visualization of **adding** video objects

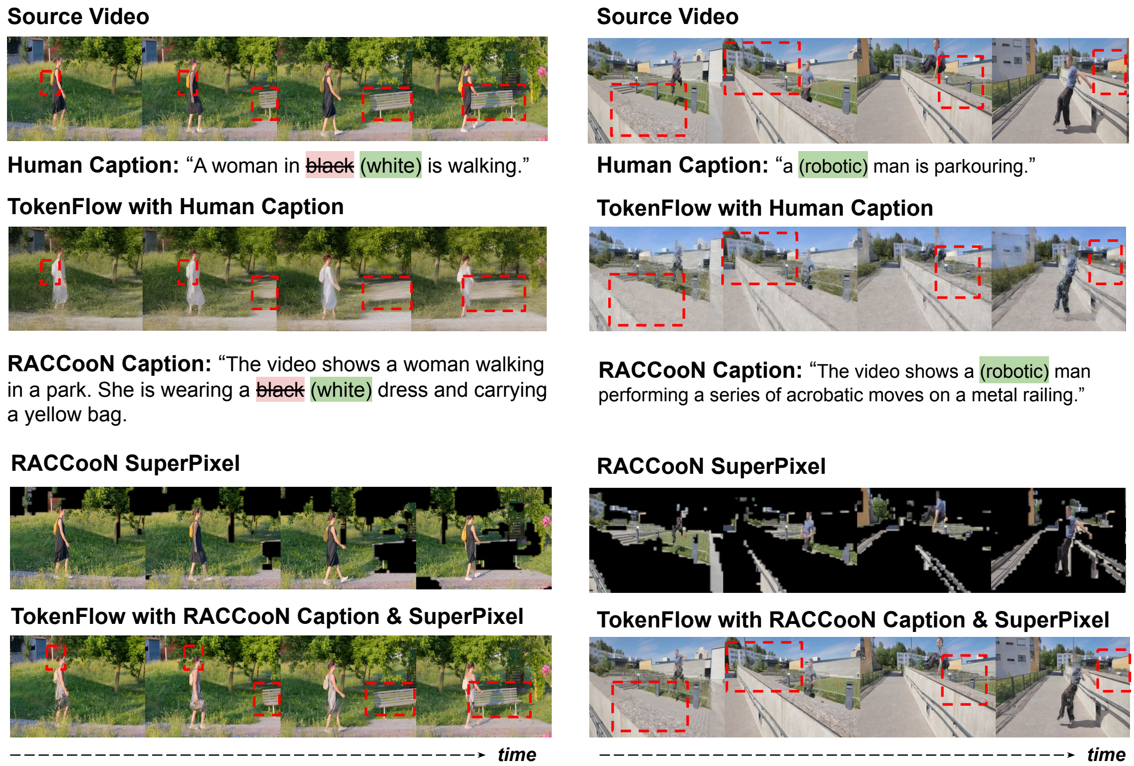


Figure 23: **Visualization of text-based video editing.** The edited words are marked with **Red**, and the target words are marked with **Green**. The RACCOON caption is selected from predicted dense captions. We highlight the region of interest with red dashed-line boxes for comparison.

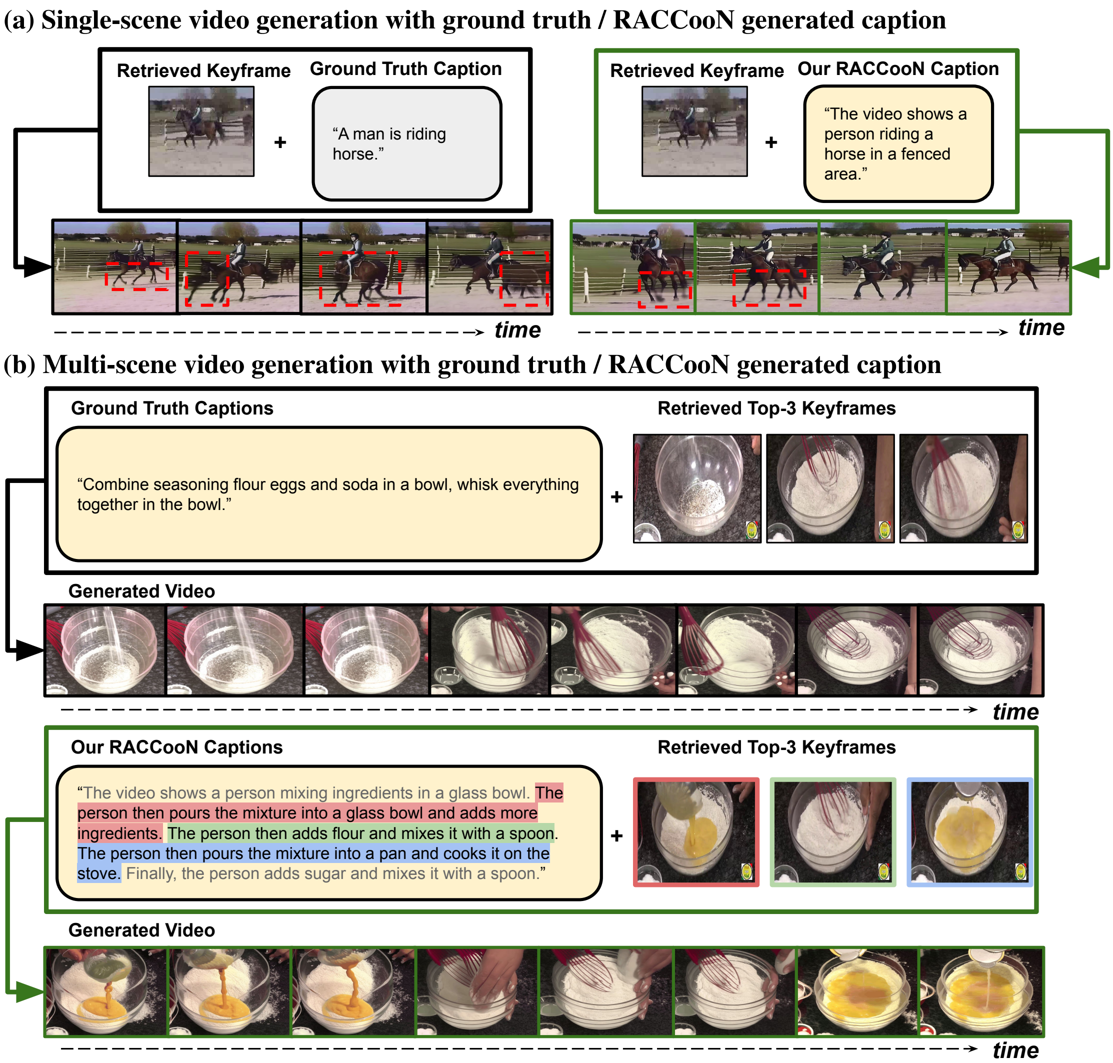


Figure 24: **Visualization of conditional video generation with VideoCrafter.** Top (a): we compare generation results conditional on different captions and with the **same** keyframe. Bottom (b): We leverage multiple keyframes retrieved by different captions to generate multi-scene video. We gray out captions that are not used for retrieval, and highlight captions used for keyframe retrieval. We highlight the region of distortion with red dashed-line boxes for detailed comparison.

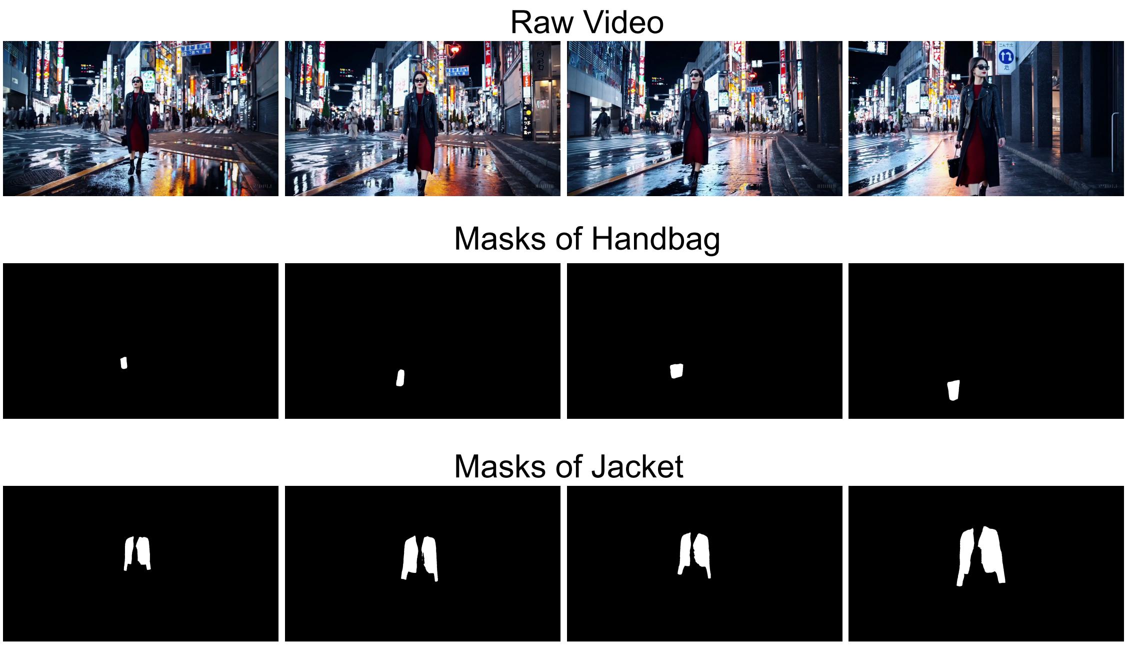


Figure 25: **Visualization of Masks.**