

VideoCoF: Unified Video Editing with Temporal Reasoner

Xiangpeng Yang¹ Ji Xie² Yiyuan Yang¹ Yue Ma³ Yan Huang^{1✉} Min Xu¹ Qiang Wu¹
¹University of Technology Sydney ²Zhejiang University ³HKUST

<https://videocof.github.io/>

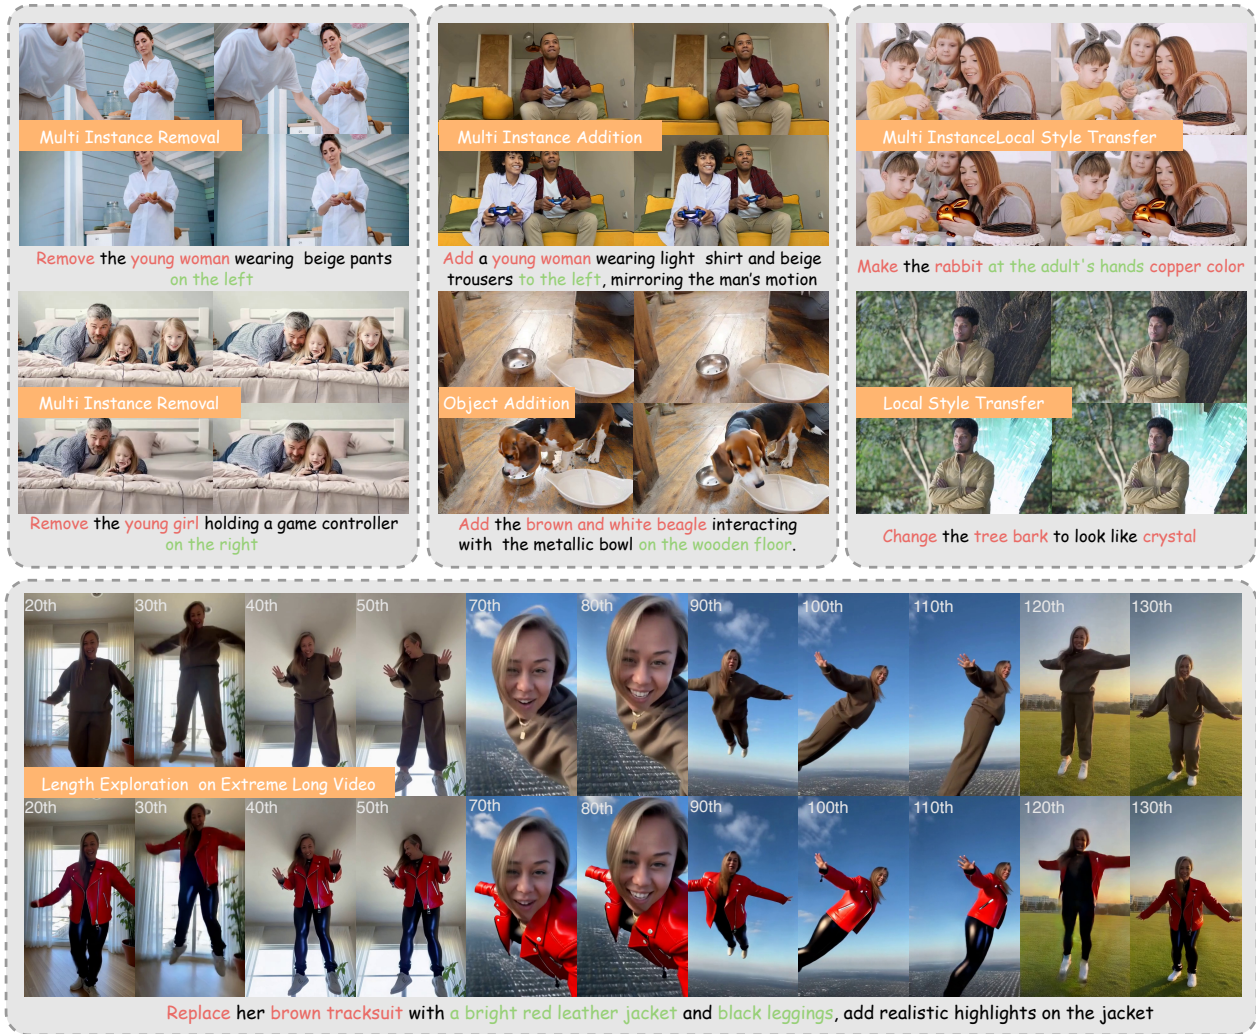


Figure 1. VideoCoF’s video editing capabilities emerge from its **seeing, reasoning, then editing** framework. Trained on only **50k** data (33 frames), this teaser shows multi-instance editing and robust $4\times$ length generalization.

Abstract

Existing video editing methods face a critical trade-off: expert models offer precision but rely on task-specific priors like masks, hindering unification; conversely, unified

temporal in-context learning models are mask-free but lack explicit spatial cues, leading to weak instruction-to-region mapping and imprecise localization. To resolve this conflict, we propose **VideoCoF**, a novel **Chain-of-Frames** approach inspired by Chain-of-Thought reasoning. VideoCoF enforces a “see → reason → edit” procedure by compelling

✉ Corresponding author.

the video diffusion model to first predict **reasoning tokens** (edit-region latents) before generating the target video tokens. This explicit reasoning step removes the need for user-provided masks while achieving precise instruction-to-region alignment and fine-grained video editing. Furthermore, we introduce a RoPE alignment strategy that leverages these reasoning tokens to ensure motion alignment and enable length extrapolation beyond the training duration. We demonstrate that with a minimal data cost of only 50k video pairs, VideoCoF achieves state-of-the-art performance on VideoCoF-Bench, validating the efficiency and effectiveness of our approach. Our code, weight, data are available at <https://github.com/knightyxp/VideoCoF>.

1. Introduction

The development of Video Diffusion Models (VDM) [16, 42, 50, 57] has enabled high-fidelity video generation across a wide range of concepts. Building on these advances, video editing methods support users in designing video by adding [41], removing [21, 70], swapping [9, 56] visual concepts, and performing global style transformation [60].

Current video editing methods mainly follow two strategies: (i) **expert models** [1, 21, 41, 56, 64], which use adapter-based modules to feed *external masks* into the video generation model, yielding precise, localized edits but requiring additional inputs and per-task overhead; and (ii) **unified temporal in-context learning models** [13, 22, 59], which concatenate source tokens with noised edit tokens along the temporal dimension and use self-attention mechanism to guide the edit. However, without explicit spatial cues, these models often exhibit weak accuracy, especially in cases that need multi-instance recognition or spatial reasoning (Fig. 2, left). In short, there is a *trade-off*: expert models are accurate but mask-dependent, while unified in-context models are mask-free but less precise; This raises a critical question: **Can we maintain former’s precision and latter’s unification without the mask dependency?**

Inspired by Chain-of-Thought (CoT) multi-step reasoning [46], we *compel* the video diffusion model to first predict the edit region and then perform the edit, enforcing a “**see** → **reason** → **edit**” procedure. Accordingly, we propose **VideoCoF**, a Chain-of-Frames approach that predicts *reasoning tokens* (edit-region latents) before generating the target video tokens, thereby removing the need for user-provided masks while achieving precise instruction-to-region alignment. To explicitly model the reasoning process, we leverage visual grounding, which is naturally suited to simulating reasoning about the edit region. Empirically, we find a soft, gradually highlighted grayscale region is the most effective reasoning format. Additionally, we introduce a RoPE alignment strategy. By explicitly accounting for the reasoning latent, we reset the temporal indices

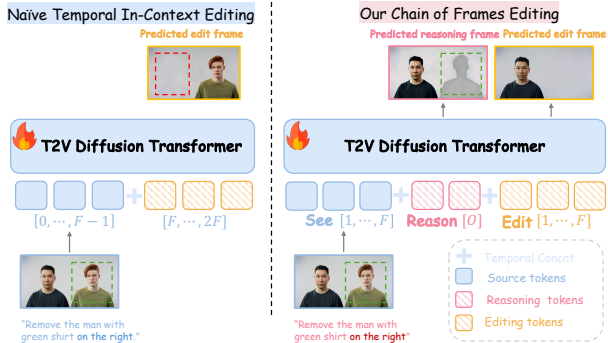


Figure 2. Illustration of the difference between previous methods and our VideoCoF. We enhance the editing accuracy by forcing the video diffusion model to first predict the editing area, and then perform the editing.

of the edited video’s rotary position embeddings to match those of the source video, ensuring motion alignment and length extrapolation.

To holistically evaluate fine-grained video editing, we further construct VideoCoF-Bench. VideoCoF trained on only 50k video pairs, outperforms a strong baseline ICVE [22] that uses ~1M pretraining videos plus 150k for fine-tuning. Specifically, we improve the instruction-following score by +15.14% and the success ratio by +18.6%. Our contributions can be summarized as follows:

- We propose VideoCoF, the first framework to introduce a Chain of Frames approach to video editing, enabling temporal reasoning for fine-grained video editing.
- Building on VideoCoF, we explore an effective reasoning format for video diffusion models, and introduce a RoPE alignment strategy that allows generalization to longer frames beyond the training duration.
- We demonstrate that with a minimal data cost (only 50k video pairs), we achieve state-of-the-art quantitative and qualitative performance on VideoCoF-Bench, validating the efficiency and effectiveness of our approach.

2. Related Work

Video Editing Methods. Early training-free video editing methods [6, 8, 26, 32, 50, 55] rely on inversion and consistency techniques, but often lack precise control and struggle with complex edits. Recent training-based methods [2, 3, 7, 24, 27, 35] have become the dominant paradigm, offering higher quality and edit diversity. Meanwhile, concurrent works [23, 25, 29, 45, 62] integrate MLLMs to guide the editing process, though this adds significant training and inference cost, which our pure VDM approach avoids.

In-Context Video Editing. Recently, in-context learning (ICL) has emerged as a promising paradigm for unified editing [14, 53, 66]. Methods like UNIC [59] and ICVE [22] concatenate video conditions along the temporal axis

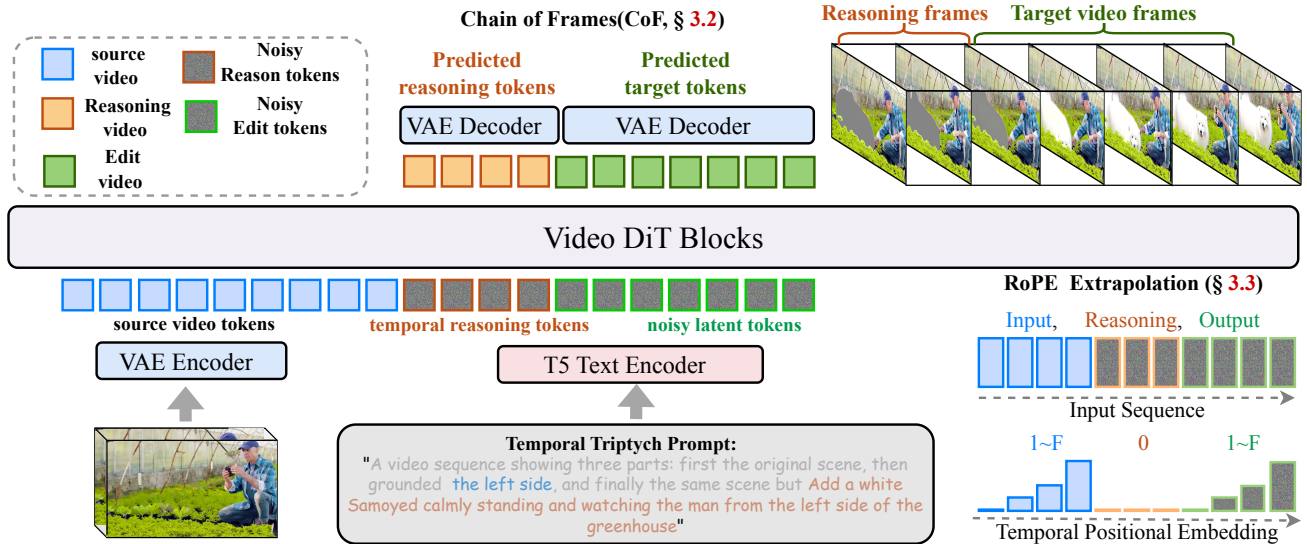


Figure 3. Overview of VideoCoF framework. Our model processes source (blue), reasoning (orange), and target (green) tokens in a unified sequence to “reason” then “edit”. **Bottom right:** Our RoPE design enables length extrapolation.

to perform ICL. However, these methods are often limited by mask requirements [59] or, as we identify, suffer from fundamental issues with editing accuracy and a lack of length extrapolation due to their naive temporal concatenation. While EditVerse [13] also explored unified in-context learning, it was built on a LLaMA-style DiT backbone, whereas our work explores these capabilities within a standard video diffusion transformer.

Chain of Thought in Vision. Chain-of-Thought (CoT) prompting [15, 46] elicits multi-step reasoning in LLMs by having them “think step-by-step.” This concept of emergent reasoning has also been identified in large video generative models [5, 47] that can solve visual puzzles. However, how to leverage visual reasoning for the task of unified video editing remains unexplored. In this work, we investigate whether generative video models can perform a “chain of frames” reasoning to achieve this.

3. Methods

3.1. VideoCoF Framework

As illustrated in Figure 3, VideoCoF employs a VideoDiT [42] for unified video editing. We model editing as a reasoning-then-generation process: the model first reasons where to edit, then generates the intended content in that area. We call this process “**Chain of Frames (CoF)**” (Sec 3.2). All visual inputs (source, reasoning, and target frames) are encoded separately by a Video VAE and then concatenated temporally. The unified frame sequence is then fed into the model, performing unified in-context learning via self-attention and language control via cross-attention. To enable video alignment and variable-length inference, we revisit the design of positional encoding. We adapt the tem-

poral RoPE for source-to-target alignment and reasoning tokens’ RoPE for explicit spatial guidance (Sec 3.3). Subsequent sections detail the training and inference paradigm (Sec 3.4), and the data curation pipeline (Sec 3.5).

3.2. Chain of Frames

Seeing, Reasoning, then Editing. Previous video in-context editing methods, such as UNIC [59], ICVE [22], or EditVerse [13], perform in-context learning by temporally concatenating clean source video tokens with noised editing video tokens. However, this approach lacks an explicit constraint mapping the editing instruction to the specific editing region, leading to editing accuracy problems, as shown in Fig 2. Recently, VDM have been shown to possess reasoning capabilities, as demonstrated in [47]. Inspired by this, we explicitly model the reasoning tokens, forcing the model to actively learn the relationship between the editing instruction and the target edit region first. The edit is then executed *after* reasoning, following a “seeing, reasoning, then editing” process.

Inspired by Chain of Thought prompting in Large Language Models (LLMs) [46], we argue that a video generative model should also have an analogous chain-reasoning ability. Given the generative priors in video editing, the visual-chain should be progressive, moving from the original video to a visual reference of the editing region, and finally to the edited video. Visual grounding is naturally suitable for this representation. Since video diffusion models are often insensitive to grounding masks (e.g., black or white pixels), we instead use a gray highlight to delineate the “grounding region,” which is also evidenced in [10]. Finally, the gray-highlighted area is used as the ground truth for the reasoning frames, teaching the diffusion model to

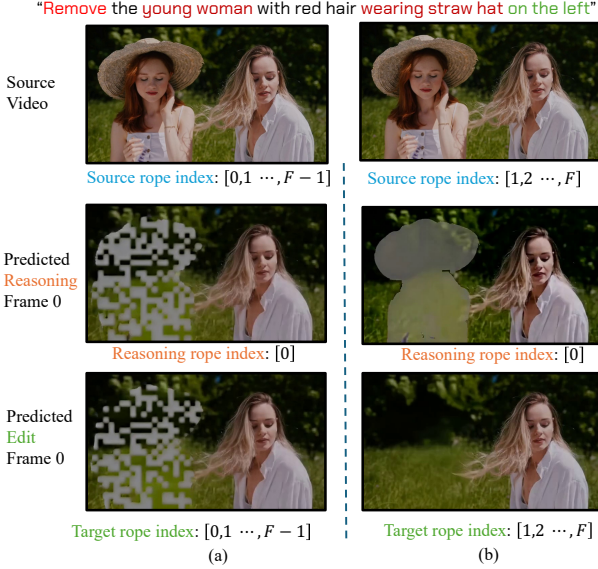


Figure 4. How our RoPE design avoids index collision.

reason about where the edit should occur.

Consequently, the entire video editing task is reformulated as a chained process: first “seeing” the original video, then “reasoning” by predicting the grounding region, and finally “editing” to generate the new video content within that specified area. We call this **Chain of Frames (CoF)**.

Let $\mathcal{E}(\cdot)$ denote the video VAE encoder. We use F and L for frames in the source/target and reasoning latent space, respectively, and denote channel, height, and width by C , H , and W . Given a triplet source-reasoning-target video pair $\{s, r, e\}$, we first encode them into latent representations. The source s and target video e yield latent $z_s = \mathcal{E}(s)$ and $z_e = \mathcal{E}(e)$, both with shape $\mathbb{R}^{F \times C \times H \times W}$. The reasoning video r yields a latent $z_r = \mathcal{E}(r)$ with shape $\mathbb{R}^{L \times C \times H \times W}$. This separate encoding ensures intra-causal relations and inter-video independence. Then, we perform temporal concatenation to get the unified representation:

$$\mathbf{z}_{full}^{(t)} = \underbrace{z_s^{(0)}}_{\text{seeing}} \parallel \underbrace{z_r^{(t)}}_{\text{reasoning}} \parallel \underbrace{z_e^{(t)}}_{\text{editing}} \in \mathbb{R}^{(F+L+F) \times C \times H \times W}, \quad (1)$$

where the $z_s = \mathbf{z}_{0:F-1}^{(0)}$ denotes anchoring the source video latent at timestep 0. $z_r = \mathbf{z}_{F:F+L-1}^{(t)}$ and $z_e = \mathbf{z}_{F+L:2F+L-1}^{(t)}$ mean the reasoning and target noised video latents at timestep t . At each denoising step, only the $L + F$ reasoning and target frames are denoised, and the source video latents are kept clean.

3.3. RoPE Design for Length Extrapolation

In VideoDiT, 3D factorized RoPE [38] provides spatio-temporal positions. A naive in-context learning approach applies sequential temporal indices (e.g., 0 to $2F - 1$) across

concatenated source and target videos. However, this hinders video length extrapolation, as the model overfits to a static $[0, F - 1] \rightarrow [F, 2F - 1]$ mapping and fails to generalize to videos longer than F frames.

A better strategy is to repeat the temporal indices. For our CoF triplet (consider $L = 1$ for reasoning frame), a straightforward reset configuration is to assign temporal indices: $[0, F - 1]$ to the source, “0” to the reasoning frame, and $[0, F - 1]$ to the target.

However, as illustrated in Figure 4 (a), this naive reset leads to index collisions at temporal position 0, shared by the source, reasoning, and target frames. This overlap introduces visual artifacts that propagate from the reasoning tokens into the first target frame.

To resolve this index collision, we set the temporal indices for both the source video and the target video to the range $[1, F]$, while keeping the reasoning frame’s temporal index at 0. This isolates the reasoning token and prevents artifact leakage while maintaining length generalization.

3.4. Training and Inference Paradigm

Algorithm 1: Chain of Frame (CoF) Training

Input: Dataset \mathcal{D} with tuples $(\mathbf{z}_s^{(0)}, \mathbf{z}_r^{(0)}, \mathbf{z}_e^{(0)}, c)$

Output: Fine-tuned parameters θ

foreach minibatch $(\mathbf{z}_s^{(0)}, \mathbf{z}_r^{(0)}, \mathbf{z}_e^{(0)}, c) \sim \mathcal{D}$ **do**

foreach sample in minibatch **do**

$\mathbf{z}_{full}^{(0)} \leftarrow \mathbf{z}_s^{(0)} \parallel \mathbf{z}_r^{(0)} \parallel \mathbf{z}_e^{(0)}$;

Sample $t \sim \mathcal{U}[0, 1]$;

Sample $\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ with the same shape $\mathbf{z}_{full}^{(0)}$;

$\mathbf{v} \leftarrow (\varepsilon - \mathbf{z}_{full}^{(0)})$;

$\mathbf{z}_{r,e}^{(t)} \leftarrow (1 - t)(\mathbf{z}_r^{(0)} \parallel \mathbf{z}_e^{(0)}) + t(\varepsilon_{F:2F+L-1})$;

$\mathbf{z}^{(t)} \leftarrow \mathbf{z}_s^{(0)} \parallel \mathbf{z}_{r,e}^{(t)}$;

$\hat{\mathbf{v}} \leftarrow \mathbf{F}_\theta(\mathbf{z}^{(t)}, t, c)$;

$\mathcal{L} \leftarrow \frac{1}{L+F} \sum_{i=F}^{2F+L-1} \|\mathbf{v}_i - \hat{\mathbf{v}}_i\|_2^2$;

Update θ using gradients of \mathcal{L} ;

Given a concatenated full latent sequence $\mathbf{z}_{full}^{(0)} = \text{TemporalConcat}(\mathbf{z}_s^{(0)}, \mathbf{z}_r^{(0)}, \mathbf{z}_e^{(0)})$, we treat the reasoning+editing block as the generation target during training.

Given timestep $t \in [0, 1]$ and Gaussian noise $\varepsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, we only progressively noise the reasoning and editing parts, $\mathbf{z}_{r,e}^{(t)} = (1 - t)(\mathbf{z}_r^{(0)} \parallel \mathbf{z}_e^{(0)}) + t\varepsilon_{F:2F+L-1}$, and form the model input $\mathbf{z}^{(t)} = \mathbf{z}_s^{(0)} \parallel \mathbf{z}_{r,e}^{(t)}$. The target velocity field is $\mathbf{v} = \varepsilon - \mathbf{z}_{full}^{(0)}$. Our model $\mathbf{F}_\theta(\cdot)$ predicts this velocity field from the partially noised input, and we train it by minimizing the mean squared error between predicted and true velocities. Concretely, we only supervise the reasoning and target frames, so the training loss can be written

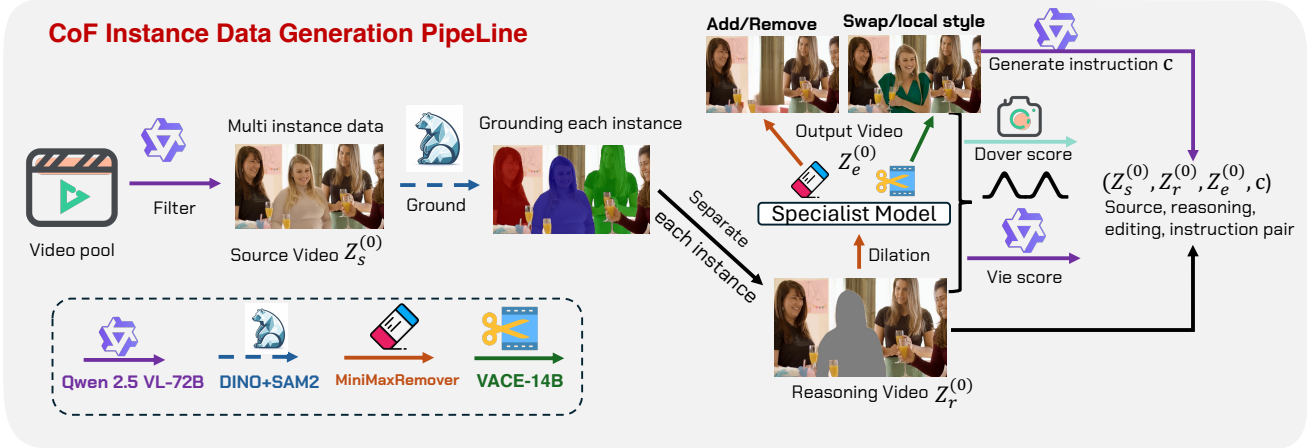


Figure 5. Our data curation pipeline for multi-instance data.

in per-frame form as

$$\mathcal{L} = \frac{1}{L+F} \sum_{i=F}^{2F+L-1} \left\| \mathbf{v}_i - [\mathbf{F}_\theta(\mathbf{z}^{(t)}, t, \mathbf{c})]_i \right\|_2^2, \quad (2)$$

where $[\mathbf{F}_\theta(\mathbf{z}^{(t)}, t, \mathbf{c})]_i$ denotes the model’s prediction for frame i and \mathbf{c} is the text condition. The model parameters $\mathbf{F}_\theta(\cdot)$ are updated via a gradient step computed from this loss. The full training procedure is summarized in Algorithm 1.

During inference we initialize the reasoning+editing block from Gaussian noise, $\mathbf{z}_{r,e}^{(1)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and form the full latent at $t = 1$ by temporal concatenation with the clean source $\mathbf{z}_{\text{full}}^{(1)} = \text{TemporalConcat}(\mathbf{z}_s^{(0)}, \mathbf{z}_{r,e}^{(1)})$. An ODE solver guided by our model \mathbf{F}_θ evolves $\mathbf{z}_{\text{full}}^{(t)}$ to $\mathbf{z}_{\text{full}}^{(0)}$. The source latents $\mathbf{z}_s^{(0)}$ are held fixed during inference, so only the reasoning/editing parts change. We then extract the edited-target latent using the same slicing index as in training: $\mathbf{z}_{\text{edit}}^{(0)} = (\mathbf{z}_{\text{full}}^{(0)})_{F+L:2F+L-1}$ and decode the final edited video: $\mathbf{x}_{\text{edit}} = \mathcal{D}(\mathbf{z}_{\text{edit}}^{(0)})$.

3.5. Video Data Curation

The training of our VideoCoF requires a large and diverse dataset structured as source, reasoning, and edited video triplets. However, existing video editing datasets and methods predominantly focus on single-instance-level object manipulation. This limitation is a significant barrier, as real-world videos contain complex visual cues, multiple interacting instances, and intricate spatial relationships (e.g., physical left/right, object-to-object interactions). Enabling a generative model to comprehend these complex, instance-level dynamics is a critical step toward true reasoning-based video editing. Therefore, we develop a comprehensive data curation pipeline, illustrated in Figure 5, to specifically generate and process complex, instance-level video data.

Instance-Level Curation Pipeline. Our pipeline begins with a large pool of diverse videos sourced from Pexels

[31]. First, we employ the Qwen-VL 72B [43] to perform multi-instance identification, scanning the videos to find scenes that contain multiple, distinct objects. Once these videos are identified, we use Grounding-SAM2 [34] to perform precise segmentation, generating distinct segmentation masks for each individual instance. With these instance-specific masks, we generate triplets for a variety of editing tasks:

- **Object Addition/Removal:** We utilize the Minimaxremover [70] to erase a specific instance from the video. The data for object addition is then created by simply reversing this process.
- **Object Swap and Local Style Transfer:** For these tasks, we leverage the VACE-14B [12] in its inpainting mode to fill the specified masked regions. Critically, the creative prompts for these inpainting edits are generated by GPT-4o[30], as we found Qwen-VL 72B’s imaginative capabilities for this specific task to be limited.

Filtering and Final Dataset. All generated video pairs are rigorously evaluated to ensure quality. We use the Dover Score [49] to assess aesthetic quality and the VIE Score [17] to measure editing fidelity and coherence. A weighted combination of these scores is used to filter for high-quality, successful edits. Finally, we use this pipeline to filter from the large-scale open-source Señorita 2M [71] dataset, and distill a high-quality subset of **50k** videos to supplement our training data. This multi-pronged approach yields our final large-scale dataset, rich in the instance-level complexity required for reasoning-based video editing.

4. Experiments

4.1. Implementation Details.

VideoCoF is trained on WAN-14B [42]. We employ a resolution-bucketing strategy to support multiple aspect ratios, using spatial resolutions of 336×592, 400×704, 400×752, and 400×944 (and the corresponding vertical variants, e.g., 592×336). Training videos are sourced from

Table 1. We compare VideoCoF with SOTA baselines on VideoCoF-Bench: TokenFlow [8], InsV2V [7], Señorita [71] (an I2V model guided by InsP2P [2]), VACE-14B [12] (using GPT-4o-generated captions), the concurrent ICVE [22] (pretrained on 1M videos and fine-tuned on 150k), and LucyEdit [40]. Despite the extensive training data used by baselines, VideoCoF is fine-tuned on only 50k video pairs and achieves superior instruction-following and success ratio.

Model	GPT-4o Score (avg.)				Perceptual Quality (avg.)		
	Instruct Follow \uparrow	Preservation \uparrow	Quality \uparrow	Success Ratio \uparrow	CLIP-T \uparrow	CLIP-F \uparrow	DINO \uparrow
TokenFlow [8]	3.12	5.85	5.10	4.25%	25.42	0.982	0.970
InsV2V [7]	3.41	6.15	5.51	6.39%	26.19	0.988	0.978
Señorita [71]	3.26	6.30	5.48	10.35%	26.04	0.994	0.988
VACE [12]	7.47	5.82	7.61	26.60%	27.02	0.994	0.990
ICVE [22]	7.79	8.06	8.14	57.76%	27.49	0.992	0.986
Lucy Edit [40]	5.24	6.50	6.37	29.64%	26.98	0.991	0.986
VideoCoF (Ours)	8.97	8.20	<u>7.77</u>	76.36%	28.00	0.992	0.991

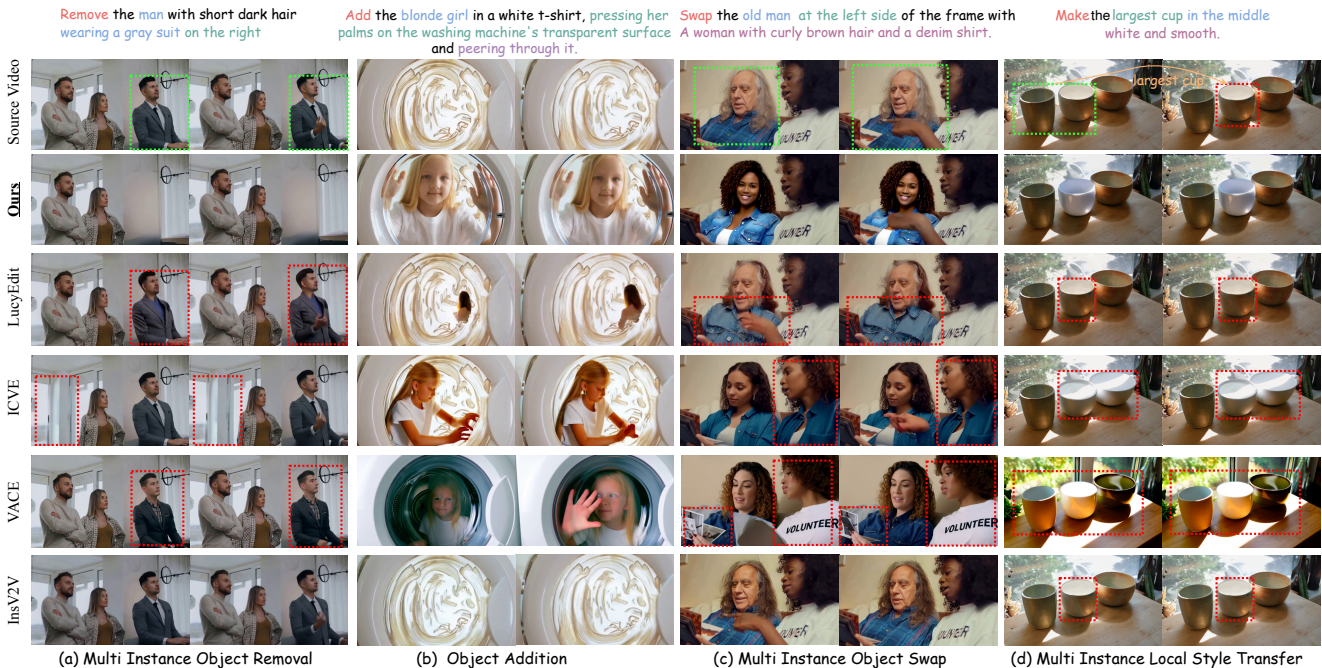


Figure 6. Visual comparison between our VideoCoF and other methods on diverse video editing tasks.

Señorita [71] and are 33 frames long. We ultimately train on only 50k curated video samples. Thanks to our RoPE alignment design, the model generalizes to longer sequences at inference (e.g., 513+ frames). By default, we use 33-frame source and edited videos, together with a 4-frame reasoning clip. We train for 8k iterations on 16 H100 GPUs. With DMD-LoRA [61], inference takes only 4 steps and about 10 seconds to edit 33 frames on a single H100 GPU.

4.2. VideoCoF-Bench and Experimental Setting

VideoCoF-Bench. Previous video-editing benchmarks such as V2VBench [39], TGVE [51], and FIVE-Bench [19] focus on target-prompt edits and are mostly limited to

class-level object swaps. They were mainly designed for training-free methods and are not suitable for instruction-guided or instance-level video editing. Real-world editing requires precise instruction understanding, including instance- and part-level control (e.g., distinguishing multiple people or left vs. right), and complex reasoning. To address these gaps, we introduce VideoCoF-Bench. It contains 200 high-quality videos collected from Pexels [31], covering diverse scenes and both landscape and portrait aspect ratios. VideoCoF-Bench includes four tasks: Object Removal, Object Addition, Object Swap, and Local Style Transfer, each with 50 samples. Half of these samples per task are instance-level cases with instance-focused editing

Table 2. Ablation on Chain of frames and RoPE design.

Ablation on Chain of frames and RoPE design			
CoF	Naive Temporal in Context		VideoCoF
RoPE Design	✗	✗	✓
	0-2F-1	0-F-1, 0-F-1	1-F, 0, 1-F
<i>GPT-4o Score</i>			
Instruct Follow↑	8.109	8.064	8.973
Preservation↑	7.930	7.793	8.203
Quality↑	7.394	7.217	7.765
Success Ratio↑*	72.41%	65.52%	76.36%
<i>Perceptual Quality</i>			
CLIP-T↑	26.880	27.088	28.000
CLIP-F↑	0.9907	0.9905	0.9915
DINO↑	0.9857	0.9826	0.9913

prompts.

Evaluation Metrics. To evaluate editing performance on VideoCoF-Bench, we employ MLLM-as-a-Judge to provide a holistic evaluation score. This is achieved by prompting **GPT-4o** [30] to assess multiple criteria given the original video, edited video, and user instruction: (1) Instruction Following (editing accuracy), (2) Preservation (unedited regions), (3) Video Quality. (4) Success ratio: we prompt the GPT-4o to provide a binary Success Ratio (Yes/No) to judge the overall success of the edit. We report three perceptual quality metrics to quantify low- and high-level visual similarity between source and target frames: CLIP-T for image-text alignment, CLIP-F for temporal consistency, and DINO for structural consistency.

4.3. Comparison on VideoCoF-Bench

We show qualitative and quantitative comparisons of VideoCoF-Bench in this section. As shown in Table 1, we evaluate VideoCoF against five baseline methods on the VideoCoF-Bench benchmark, which spans four distinct video editing tasks: multi-instance removal, object addition, multi-instance swap, and multi-instance local style transfer.

Overall, VideoCoF demonstrates the best performance in **Instruct Follow** and **Success Ratio** across all categories. Compared to naive temporal in-context editing approaches like ICVE [22], our method achieves significantly higher success rates and better instruction adherence using only **50k** reasoning pairs, whereas ICVE is pre-trained on 1M samples and fine-tuned on 150k data.

Qualitatively (see Figure 6), our method also shows clearer, more faithful edits at the instance level: (a) Multi-instance removal: we precisely remove the right instance while ICVE[22] incorrectly removes the left instance. (b) Object addition: the added girl is correctly placed inside the washing machine, matching the instruction. (c) Object swap: we replace the elderly person’s face and update clothing; Lucy Edit [40] changes only clothing, ICVE fails to disambiguate instances, and VACE often alters non-target peo-



Figure 7. Length extrapolation beyond the training length.

ple. (d) Local style (multi-instance): our model correctly identifies and edits the largest cup among several similar objects; other methods either fail to edit or mistakenly edit a bowl. These qualitative examples demonstrate VideoCoF’s stronger instance-level reasoning and higher editing fidelity.

4.4. Ablation Study

To verify our novel Chain of Frames (CoF) design, particularly its “reasoning frames” and the RoPE design for length extrapolation, we conduct an ablation study on the reasoning frames, RoPE alignment strategy and reasoning format. **Naive Temporal Incontext VS. CoF.** As shown in Table 2, we compare VideoCoF against a “Naive Temporal in-context” baseline. This applies temporal in-context learning by using the source video as a condition through temporal concatenation, an approach similar to ICVE [22].

In contrast, our approach introduces **reasoning frames** as a core component of the (CoF) design. This ensures the video editing follows a reasoning process, i.e., forcing the model to predict the editing region first and then execute the versatile edit within that specific area.

The efficacy of this design is evident when comparing the first ($[0, 2F - 1]$) and third (VideoCoF) columns in Table 2. The inclusion of CoF brings substantial gains: the instruct follow score increases by 10.65% and the success ratio improves by 5.46%. Furthermore, the 4.16% increase in CLIP-T confirms that our reasoning frames effectively enhance the model’s editing accuracy and precision.

RoPE Design for Length Extrapolation. As illustrated in Fig 7, the naive approach ($[0, 2F - 1]$) only learns a fixed temporal mapping (e.g., mapping frame 0_{th} to frame 33_{th}). This prevents length extrapolation, causing severe degradation (blurriness, motion misalignment, and artifacts) when a 33-frame trained model is tested on 81 frames (second row).

In contrast, our RoPE alignment design ($[1 - F, 0, 1 - F]$) generalizes to unseen lengths without quality degradation (third row). As demonstrated in Fig 1, our model extrapolates to 141 frames ($4\times$ training length) and even 513 frames ($16\times$ training length), supporting strong generalization to much longer video sequences.

This effectiveness is also quantified in Table 2 (third vs. first column). We observe a 3.4% relative increase in the

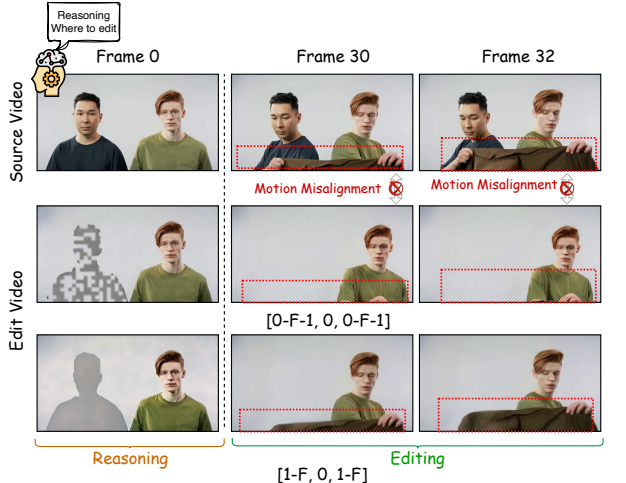


Figure 8. Motion alignment benefit by our RoPE design.

preservation score. Furthermore, the improved DINO score confirms that our RoPE design better preserves the original video’s spatio-temporal structure during editing.

RoPE Design for Motion Alignment. Setting the temporal index for the reasoning frame latent is a critical design choice. A naive approach is to set its index to 0, aligning it with the first video frame. This causes two severe issues.

First, it leads to significant motion misalignment (e.g., the subject fails to perform the “lifting clothes” motion in Fig 8, second row). Second, this “0-index” design causes interference with the first editing video frame (also index 0), leading to artifacts where the model incorrectly predicts the first frame as the reasoning frame (Fig 4).

Therefore, we fix the reasoning latent’s index to 0, while the source and edited video indices range from 1 to F (denoted as $[1 - F, 0, 1 - F]$). This strategy allows the reasoning frame to provide clear spatial guidance on **where** to edit, without disrupting the video’s temporal structure and motion alignment. The improvements across all metrics in Tab 2 (column 3 vs. column 2) validate this design.

Reasoning Frame Format. First, we explore the most suitable color for the reasoning frame mask. As shown in Table 3, we compare three formats: (1) A black mask over the unedited region; (2) A red, 50% transparent highlight, same as Veggie [62]; and (3) A gray, 50% transparent mask. The quantitative results show that using a gray mask (column 3) for the edit region yields the best performance.

Furthermore, we argue that the reasoning frame should act as a gradual transition from the source video to the edited video. Therefore, we test a progressive gray mask. Instead of a single static mask, we interpolate the gray reasoning frame with the editing frame, with the transparency progressively increased (e.g., 0%, 25%, 50%, 75%). As shown by comparing column 4 and column 3 in Table 3, this progressive gray reasoning frame approach works best.

Table 3. Ablation on the reasoning frame format.

Ablation on Reasoning Frame Format				
Color	Black (bg)	Red	Gray	Gray
Transparency	(0%)	(50%)	(50%)	(0-75%)
<i>GPT-4o Score</i>				
Instruct Follow↑	7.512	7.805	8.150	8.973
Preservation↑	7.034	7.350	7.443	8.203
Quality↑	6.155	6.501	6.645	7.765
Success Ratio↑*	52.17%	60.33%	68.45%	76.36%
<i>Perceptual Quality</i>				
CLIP-T↑	26.550	26.810	27.220	28.000
CLIP-F↑	0.9810	0.9855	0.9889	0.9915
DINO↑	0.9750	0.9790	0.9803	0.9913

Qualitatively, as shown in Figure 9, the mask format is critical. The black mask fails the deletion task, while the red mask incorrectly deletes content on the right side. In contrast, our progressive gray mask accurately performs the intended deletion on the left. We conclude from these experiments that the optimal reasoning format is a gray mask with progressive transparency.




Figure 9. Ablation on reasoning frame format.

5. Conclusion

In this paper, we introduced VideoCoF, a unified model for universal video editing via temporal reasoning. We identified that existing temporal in-context learning approaches often fail due to a lack of explicit spatial cues, leading to weak instruction-to-region mapping and imprecise localization. To address these issues, we proposed the innovative Chain of Frames. VideoCoF compels the video diffusion model to follow a “see, reason, then edit” process by first predicting the editing region before executing the versatile edit. Furthermore, to solve the length generalization challenge, we developed a novel RoPE alignment paradigm that accounts for the reasoning latent. This design enables up to $16\times$ length extrapolation during the inference. Experimental results show that VideoCoF achieves SOTA performance using only 50k video pairs, validating the efficiency and effectiveness of our temporal reasoning design.

References

- [1] Yuxuan Bian, Zhaoyang Zhang, Xuan Ju, Mingdeng Cao, Liangbin Xie, Ying Shan, and Qiang Xu. Videopainter: Any-length video inpainting and editing with plug-and-play context control. *arXiv preprint arXiv:2503.05639*, 2025. **2**
- [2] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18392–18402, 2023. **2, 6, 5**
- [3] Yuanhao Cai, He Zhang, Xi Chen, Jinbo Xing, Yiwei Hu, Yuqian Zhou, Kai Zhang, Zhifei Zhang, Soo Ye Kim, Tianyu Wang, et al. Omnivcus: Feedforward subject-driven video customization with multimodal control conditions. *arXiv preprint arXiv:2506.23361*, 2025. **2**
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. **4**
- [5] Lan Chen, Yuchao Gu, and Qi Mao. Univid: Unifying vision tasks with pre-trained video generation models. *arXiv preprint arXiv:2509.21760*, 2025. **3**
- [6] Yiyang Chen, Xuanhua He, Xiujun Ma, and Yue Ma. Contextflow: Training-free video object editing via adaptive context enrichment. *arXiv preprint arXiv:2509.17818*, 2025. **2**
- [7] Jiaxin Cheng, Tianjun Xiao, and Tong He. Consistent video-to-video transfer using synthetic dataset. *arXiv preprint arXiv:2311.00213*, 2023. **2, 6, 1, 5**
- [8] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Tokenflow: Consistent diffusion features for consistent video editing. *arXiv preprint arXiv:2307.10373*, 2023. **2, 6**
- [9] Yuchao Gu, Yipin Zhou, Bichen Wu, Licheng Yu, Jia-Wei Liu, Rui Zhao, Jay Zhangjie Wu, David Junhao Zhang, Mike Zheng Shou, and Kevin Tang. Videoswap: Customized video subject swapping with interactive semantic point correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7621–7630, 2024. **2**
- [10] Nicholas Guttenberg. Diffusion with offset noise. <https://www.crosslabs.org/blog/diffusion-with-offset-noise>, 2023. **3**
- [11] Lianghua Huang, Wei Wang, Zhi-Fan Wu, Yupeng Shi, Huanzhang Dou, Chen Liang, Yutong Feng, Yu Liu, and Jingren Zhou. In-context lora for diffusion transformers. *arXiv preprint arxiv:2410.23775*, 2024. **2**
- [12] Zeyinzi Jiang, Zhen Han, Chaojie Mao, Jingfeng Zhang, Yulin Pan, and Yu Liu. Vace: All-in-one video creation and editing. *arXiv preprint arXiv:2503.07598*, 2025. **5, 6, 3**
- [13] Xuan Ju, Tianyu Wang, Yuqian Zhou, He Zhang, Qing Liu, Nanxuan Zhao, Zhifei Zhang, Yijun Li, Yuanhao Cai, Shaoteng Liu, et al. Editverse: Unifying image and video editing and generation with in-context learning. *arXiv preprint arXiv:2509.20360*, 2025. **2, 3, 4**
- [14] Xuan Ju, Weicai Ye, Quande Liu, Qiulin Wang, Xintao Wang, Pengfei Wan, Di Zhang, Kun Gai, and Qiang Xu. Fulldit: Multi-task video generative foundation model with full attention. *arXiv preprint arXiv:2503.19907*, 2025. **2**
- [15] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213, 2022. **3**
- [16] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*, 2024. **2**
- [17] Max Ku, Dongfu Jiang, Cong Wei, Xiang Yue, and Wenhua Chen. ViScore: Towards explainable metrics for conditional image synthesis evaluation, 2023. **5**
- [18] Maksim Kuprashevich, Grigori Alekseenko, Irina Tolstykh, Georgii Fedorov, Bulat Suleimanov, Vladimir Dokholyan, and Aleksandr Gordeev. Nohumansrequired: Autonomous high-quality image editing triplet mining. *arXiv preprint arXiv:2507.14119*, 2025. **4**
- [19] Minghan Li, Chenxi Xie, Yichen Wu, Lei Zhang, and Mengyu Wang. Five: A fine-grained video editing benchmark for evaluating emerging diffusion and rectified flow models. *arXiv preprint arXiv:2503.13684*, 2025. **6**
- [20] Xirui Li, Chao Ma, Xiaokang Yang, and Ming-Hsuan Yang. VidtoMe: Video token merging for zero-shot video editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7486–7495, 2024. **1**
- [21] Xiaowen Li, Haolan Xue, Peiran Ren, and Liefeng Bo. Diffuseraser: A diffusion model for video inpainting. *arXiv preprint arXiv:2501.10018*, 2025. **2**
- [22] Xinyao Liao, Xianfang Zeng, Ziyue Song, Zhoujie Fu, Gang Yu, and Guosheng Lin. In-context learning with unpaired clips for instruction-based video editing. *arXiv preprint arXiv:2510.14648*, 2025. **2, 3, 6, 7, 1, 4, 5**
- [23] Han Lin, Xichen Pan, Ziqi Huang, Ji Hou, Jialiang Wang, Weifeng Chen, Zecheng He, Felix Juefei-Xu, Junzhe Sun, Zhipeng Fan, et al. Exploring mllm-diffusion information transfer with metacanvas. *arXiv preprint arXiv:2512.11464*, 2025. **2**
- [24] Shaoteng Liu, Tianyu Wang, Jui-Hsien Wang, Qing Liu, Zhifei Zhang, Joon-Young Lee, Yijun Li, Bei Yu, Zhe Lin, Soo Ye Kim, et al. Generative video propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17712–17722, 2025. **2**
- [25] Zhiheng Liu, Weiming Ren, Haozhe Liu, Zijian Zhou, Shoufa Chen, Haonan Qiu, Xiaoke Huang, Zhaochong An, Fanny Yang, Aditya Patel, et al. Tuna: Taming unified visual representations for native unified multimodal models. *arXiv preprint arXiv:2512.02014*, 2025. **2**
- [26] Zeqian Long, Mingzhe Zheng, Kunyu Feng, Xinhua Zhang, Hongyu Liu, Harry Yang, Linfeng Zhang, Qifeng Chen, and Yue Ma. Follow-your-shape: Shape-aware image editing via trajectory-guided region control. *arXiv preprint arXiv:2508.08134*, 2025. **2**
- [27] Yue Ma, Yulong Liu, Qiyuan Zhu, Ayden Yang, Kunyu Feng, Xinhua Zhang, Zhifeng Li, Sirui Han, Chenyang Qi, and Qifeng Chen. Follow-your-motion: Video motion transfer

- via efficient spatial-temporal decoupled finetuning. *arXiv preprint arXiv:2506.05207*, 2025. 2
- [28] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 1
- [29] Chong Mou, Qichao Sun, Yanze Wu, Pengze Zhang, Xinghui Li, Fulong Ye, Songtao Zhao, and Qian He. Instructx: Towards unified visual editing with mllm guidance. *arXiv preprint arXiv:2510.08485*, 2025. 2, 3, 4
- [30] OpenAI. Hello gpt-4o. Blog post, 2024. 5, 7, 3
- [31] Pexels. Pexels: Free stock photos, royalty free stock images & videos. <https://www.pexels.com/>, 2025. Accessed: 2025-11-06. 5, 6, 3
- [32] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15932–15942, 2023. 2
- [33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 3
- [34] 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. Grounded sam: Assembling open-world models for diverse visual tasks, 2024. 5
- [35] Yutao Shen, Junkun Yuan, Toru Aonishi, Hideki Nakayama, and Yue Ma. Follow-your-preference: Towards preference-aligned image inpainting. *arXiv preprint arXiv:2509.23082*, 2025. 2
- [36] Chaehun Shin, Jooyoung Choi, Heeseung Kim, and Sungroh Yoon. Large-scale text-to-image model with inpainting is a zero-shot subject-driven image generator. 2024. 2
- [37] Uriel Singer, Amit Zohar, Yuval Kirstain, Shelly Sheynin, Adam Polyak, Devi Parikh, and Yaniv Taigman. Video editing via factorized diffusion distillation. In *European Conference on Computer Vision*, pages 450–466. Springer, 2024. 1
- [38] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 4
- [39] Wenhao Sun, Rong-Cheng Tu, Jingyi Liao, and Dacheng Tao. Diffusion model-based video editing: A survey. *arXiv preprint arXiv:2407.07111*, 2024. 6, 1
- [40] DecartAI Team. Lucy edit: Open-weight text-guided video editing. 2025. 6, 7, 5
- [41] Yuanpeng Tu, Hao Luo, Xi Chen, Sihui Ji, Xiang Bai, and Hengshuang Zhao. Videoanydoor: High-fidelity video object insertion with precise motion control. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers*, pages 1–11, 2025. 2
- [42] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Fei Wu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025. 2, 3, 5
- [43] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 5
- [44] Xiang Wang, Hangjie Yuan, Shiwei Zhang, Dayou Chen, Jiniu Wang, Yingya Zhang, Yujun Shen, Deli Zhao, and Jingren Zhou. Videocomposer: Compositional video synthesis with motion controllability. *Advances in Neural Information Processing Systems*, 36:7594–7611, 2023. 1
- [45] Cong Wei, Quande Liu, Zixuan Ye, Qiulin Wang, Xintao Wang, Pengfei Wan, Kun Gai, and Wenhua Chen. Univideo: Unified understanding, generation, and editing for videos. *arXiv preprint arXiv:2510.08377*, 2025. 2
- [46] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022. 2, 3
- [47] Thaddäus Wiedemer, Yuxuan Li, Paul Vicol, Shixiang Shane Gu, Nick Matarese, Kevin Swersky, Been Kim, Priyank Jaini, and Robert Geirhos. Video models are zero-shot learners and reasoners. *arXiv preprint arXiv:2509.20328*, 2025. 3
- [48] Bichen Wu, Ching-Yao Chuang, Xiaoyan Wang, Yichen Jia, Kapil Krishnakumar, Tong Xiao, Feng Liang, Licheng Yu, and Peter Vajda. Fairy: Fast parallelized instruction-guided video-to-video synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8261–8270, 2024. 1
- [49] Haoning Wu, Erli Zhang, Liang Liao, Chaofeng Chen, Jingwen Hou Hou, Annan Wang, Wenxiu Sun Sun, Qiong Yan, and Weisi Lin. Exploring video quality assessment on user generated contents from aesthetic and technical perspectives. In *International Conference on Computer Vision (ICCV)*, 2023. 5
- [50] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7623–7633, 2023. 2, 1
- [51] Jay Zhangjie Wu, Xiuyu Li, Difei Gao, Zhen Dong, Jinbin Bai, Aishani Singh, Xiaoyu Xiang, Youzeng Li, Zuwei Huang, Yuanxi Sun, et al. Cvpr 2023 text guided video editing competition. *arXiv preprint arXiv:2310.16003*, 2023. 6
- [52] Jay Zhangjie Wu, Xiuyu Li, Difei Gao, Zhen Dong, Jinbin Bai, Aishani Singh, Xiaoyu Xiang, Youzeng Li, Zuwei Huang, Yuanxi Sun, et al. Cvpr 2023 text guided video editing competition. *arXiv preprint arXiv:2310.16003*, 2023. 1

- [53] Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xin-grun Xing, Ruiran Yan, Chaofan Li, Shuting Wang, Tiejun Huang, and Zheng Liu. Omnigen: Unified image generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13294–13304, 2025. 2
- [54] Zhen Xing, Qi Dai, Han Hu, Zuxuan Wu, and Yu-Gang Jiang. Simda: Simple diffusion adapter for efficient video generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7827–7839, 2024. 1
- [55] Xiangpeng Yang, Linchao Zhu, Hehe Fan, and Yi Yang. Eva: Zero-shot accurate attributes and multi-object video editing. *arXiv preprint arXiv:2403.16111*, 2024. 2
- [56] Xiangpeng Yang, Linchao Zhu, Hehe Fan, and Yi Yang. Videograin: Modulating space-time attention for multi-grained video editing. In *The Thirteenth International Conference on Learning Representations*, 2025. 2
- [57] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2
- [58] Danah Yatim, Rafail Fridman, Omer Bar-Tal, Yoni Kasten, and Tali Dekel. Space-time diffusion features for zero-shot text-driven motion transfer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8466–8476, 2024. 1
- [59] Zixuan Ye, Xuanhua He, Quande Liu, Qiulin Wang, Xintao Wang, Pengfei Wan, Di Zhang, Kun Gai, Qifeng Chen, and Wenhan Luo. Unic: Unified in-context video editing. *arXiv preprint arXiv:2506.04216*, 2025. 2, 3
- [60] Zixuan Ye, Huijuan Huang, Xintao Wang, Pengfei Wan, Di Zhang, and Wenhan Luo. Stylemaster: Stylize your video with artistic generation and translation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2630–2640, 2025. 2
- [61] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Frédo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *CVPR*, 2024. 6
- [62] Shoubin Yu, Difan Liu, Ziqiao Ma, Yicong Hong, Yang Zhou, Hao Tan, Joyce Chai, and Mohit Bansal. Veggie: Instructional editing and reasoning video concepts with grounded generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15147–15158, 2025. 2, 8
- [63] Kai Zhang, Lingbo Mo, Wenhui Chen, Huan Sun, and Yu Su. Magicbrush: A manually annotated dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*, 36:31428–31449, 2023. 4
- [64] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3836–3847, 2023. 2
- [65] Zechuan Zhang, Ji Xie, Yu Lu, Zongxin Yang, and Yi Yang. In-context edit: Enabling instructional image editing with in-context generation in large scale diffusion transformer. *arXiv preprint arXiv:2504.20690*, 2025. 2
- [66] Zechuan Zhang, Ji Xie, Yu Lu, Zongxin Yang, and Yi Yang. In-context edit: Enabling instructional image editing with in-context generation in large scale diffusion transformer. *arXiv preprint arXiv:2504.20690*, 2025. 2
- [67] Min Zhao, Guande He, Yixiao Chen, Hongzhou Zhu, Chongxuan Li, and Jun Zhu. Riflex: A free lunch for length extrapolation in video diffusion transformers. *arXiv preprint arXiv:2502.15894*, 2025. 1
- [68] Min Zhao, Hongzhou Zhu, Yingze Wang, Bokai Yan, Jintao Zhang, Guande He, Ling Yang, Chongxuan Li, and Jun Zhu. Ultravico: Breaking extrapolation limits in video diffusion transformers. *arXiv preprint arXiv:2511.20123*, 2025. 1
- [69] Rui Zhao, Yuchao Gu, Jay Zhangjie Wu, David Junhao Zhang, Jia-Wei Liu, Weijia Wu, Jussi Keppo, and Mike Zheng Shou. Motiondirector: Motion customization of text-to-video diffusion models. In *European Conference on Computer Vision*, pages 273–290. Springer, 2024. 1
- [70] Bojia Zi, Weixuan Peng, Xianbiao Qi, Jianan Wang, Shihao Zhao, Rong Xiao, and Kam-Fai Wong. Minimax-remover: Taming bad noise helps video object removal. *arXiv preprint arXiv:2505.24873*, 2025. 2, 5, 3
- [71] Bojia Zi, Penghui Ruan, Marco Chen, Xianbiao Qi, Shaozhe Hao, Shihao Zhao, Youze Huang, Bin Liang, Rong Xiao, and Kam-Fai Wong. Señorita-2M: A high-quality instruction-based dataset for general video editing by video specialists. *arXiv preprint arXiv:2502.06734*, 2025. 5, 6

VideoCoF: Unified Video Editing with Temporal Reasoner

Supplementary Material

This document provides more details of our approach and additional experimental results, which are organized as follows:

- Discussion on RoPE Design (§6)
- Editing Length Upper Bound (§7)
- Full Comparison (§8)
- TGVE+ and V2VBench (§9)
- More Ablation studies (§10)
- Implementation Details (§11)
- Metrics (§12)
- Future Directions (§13)

6. Discussion on RoPE Design

Difference from UNIC. VideoCoF targets *length extrapolation* in V2V editing, whereas UNIC focuses on *source-target alignment*. To enable extrapolation, we **pre-position** the ID/reasoning frames so that their indices never overlap with the target video. In contrast, UNIC **post-positions** the ID/reasoning tokens (e.g., shifting them to a later index range). As the video becomes longer and reaches the same index range, overlap becomes unavoidable, leading to temporal index collisions and content perturbations.

Difference from T2V extrapolation methods. RIFLEX [67] and UltraViCo [68] mainly address motion or content repetition in T2V length extrapolation, whereas our V2V extrapolation is primarily limited by source-target alignment and temporal index collision. We apply UltraViCo to the baseline, and Fig. 10 shows that it still fails.



Figure 10. Comparison between VideoCoF and UltraViCo

7. Editing Length Upper Bound

Systematic Temporal Extrapolation Failure Cases. We test $16\times$ extrapolation in both single-shot and multi-shot settings. Single-shot setting supports $16\times$ extrapolation. For multi-shot, as shown in Fig. 11, it remains stable for the first 500 frames but shows slight degradation at 512 frames.

8. Full Comparison

As shown in Tab. 9, we provide a detailed breakdown of the results across four distinct tasks: Object Removal,



Figure 11. Temporal extrapolation failure case.

Object Addition, Object Swap, and Local Style Transfer. Our VideoCoF consistently achieves the highest scores in **instruction following** and **success ratio** across all tasks, demonstrating superior capability in understanding and executing editing requests. We note that our scores in Preservation and Quality are slightly lower than the concurrent work ICVE [22]. This performance gap is reasonable given that ICVE benefits from large-scale pre-training on 1M pairs, and its supervised fine-tuning (SFT) dataset scale (150K) is three times larger than ours (50k). Furthermore, in terms of perceptual quality, VideoCoF achieves the highest **CLIP-T** scores across all tasks. This further demonstrates superior video-text alignment, consistent with our leading performance in GPT-4o score.

9. TGVE+ and V2VBench

Since TGVE[52] has only 304 samples, we instead evaluate on the comprehensive TGVE+ subset used in EVE [37] (1,417 samples across 7 instruction editing tasks), and report V2VBench [39] results. As shown in the tables, VideoCoF achieves the best performance on both benchmarks.

Table 4. Comparison on the TGVE+ benchmark.

Dataset	Methods	PickScore \uparrow	CLIP-F \uparrow	ViCLIP _{dir} \uparrow	ViCLIP _{out} \uparrow
TGVE+	Tune-A-Video [50]	20.47	0.933	0.131	0.242
	SDEdit [28]	20.35	0.899	0.131	0.241
	STDF [58]	20.60	0.933	0.093	0.227
	Fairy [48]	19.81	<u>0.933</u>	0.140	0.197
	InsV2V [7]	20.37	0.925	0.174	0.236
	EVE [37]	<u>20.88</u>	0.926	<u>0.198</u>	<u>0.251</u>
	VideoCoF	20.90	0.956	0.213	0.257

Table 5. Comparison on the V2VBench benchmark.

V2VBench	Frames Quality \uparrow	Semantic Consis. \uparrow	Object Consis. \uparrow	Frames-Text Align. \uparrow	Frames Pick. \uparrow	Video-Text Align. \uparrow	Motion Align. \uparrow
Tune-A-Video [50]	5.001	0.934	0.917	27.513	20.701	0.254	-5.599
SimDA [54]	4.988	0.940	0.929	26.773	20.512	0.248	-4.756
VidToMe [20]	4.988	0.949	0.945	26.813	20.546	0.240	-3.203
VideoComposer [44]	4.429	0.914	0.905	<u>28.001</u>	20.272	<u>0.262</u>	-8.095
MotionDirector [69]	4.984	<u>0.949</u>	<u>0.951</u>	27.845	<u>20.923</u>	<u>0.262</u>	-3.088
VideoCoF (Ours)	5.024	0.954	0.956	28.336	21.154	0.275	-2.620

10. More Ablation Studies

In this section, we validate key design choices of VideoCoF: the length of reasoning frames and the dispatch prompt.

10.1. Ablation on Reasoning Frames

Table 6. Ablation on the number of Reasoning Frames. We investigate the impact of varying the number of reasoning frames from 1 to 5. Our default setting (4 frames) achieves the best balance.

Frames	Ablation on Reasoning Frames				
	1	2	3	4 (Ours)	5
<i>GPT-4o Score</i>					
Instruct Follow \uparrow	8.219	8.312	8.281	8.973	7.915
Preservation \uparrow	8.115	8.150	8.191	8.203	6.542
Quality \uparrow	7.692	7.752	7.735	7.765	5.274
Success Ratio \uparrow	68.47%	69.39%	68.32%	76.36%	29.06%
<i>Perceptual Quality</i>					
CLIP-T \uparrow	27.092	27.148	27.136	28.000	26.997
CLIP-F \uparrow	0.9892	0.9893	0.9899	0.9915	0.9849
DINO \uparrow	0.9815	0.9827	0.9836	0.9913	0.9719

Tab. 6 investigates the optimal number of reasoning frames (F) for spatial guidance. Considering the VideoVAE temporal compression formula $L = (F - 1) // 4 + 1$, frames $1 \sim 4$ map to a single latent frame ($L = 1$), while $F = 5$ introduces latent frames $L = 2$. Results show that $F = 4$ achieves the best performance. This indicates maximizing spatial information within a single latent frame is more effective than expanding to a second latent, which introduces unnecessary temporal complexity and degradation.

10.2. Ablation on Temporal Triptych Prompt

Table 7. Ablation on Temporal Triptych Prompt. We compare the performance of our model with and without the triptych patch prompt mechanism. The inclusion of the triptych prompt significantly enhances instruction following and overall success rates.

	Ablation on Temporal Triptych Prompt	
	w/o Triptych	w/ Triptych (Ours)
<i>GPT-4o Score</i>		
Instruct Follow \uparrow	8.064	8.973
Preservation \uparrow	8.094	8.203
Quality \uparrow	7.360	7.765
Success Ratio \uparrow	71.43%	76.36%
<i>Perceptual Quality</i>		
CLIP-T \uparrow	27.07	28.00
CLIP-F \uparrow	0.989	0.992
DINO \uparrow	0.980	0.991

To adapt a standard T2V model for instruction-based editing tasks, we draw inspiration from in-context image



Figure 12. Input Prompt Variants for In-Context Video Editing. We evaluate two prompt formats: (a) Temporal Triptych Prompt - instructions embedded in a structure “A video sequence showing three parts: first the original scene, then grounded {ground instruction}, and finally the same scene but {edit instruction}.” (b) Direct Instruction - explicit editing commands provided directly.

editing approaches [11, 36, 65]. Specifically, we implement a **temporal triptych prompt** mechanism in VideoCoF to describe the evolution of video content along the temporal dimension. As illustrated in Fig. 12, our prompt template is structured as follows: “A video sequence showing three parts: first the original scene, then grounded {ground instruction}, and finally the same scene but {edit instruction}.”

As evidenced in Tab. 7, this mechanism brings significant performance gains across all metrics. Crucially, unlike the concurrent work ICVE [22], which requires computationally expensive pre-training on 1M video pairs to align the T2V model with an instruction mode, our “temporal triptych prompt” approach offers a practically zero-cost solution to effectively bridge the gap between generation and editing without the need for massive instruction tuning.

11. Implementation Details

11.1. Training Dataset

To equip our model with robust instruction-following capabilities, we constructed a unified chain-of-frames video editing dataset comprising 50k video pairs. As detailed in Table 8, the dataset is strategically balanced across four core editing tasks: object addition, removal, swapping, and local stylization. The data construction pipeline integrates both filtered open-source data and high-quality synthetic data. Object Addition and Removal: These subsets (25k samples total) are derived from the Señorita dataset. We employ MiniMax-Remover [70] to synthesize paired data. Specifically, for the removal task (15k), we treat the original video

Table 8. Statistics of the VideoCoF Training Data. The dataset consists of 50k samples balanced across four tasks.

Dataset	#Samples	Information
Video Editing Tasks		
Obj. Addition	10,000	Derived from filtered Señorita. Source generated by removing objects from target via MiniMax-Remover [70]. (absent → present).
Obj. Removal	15,000	Derived from filtered Señorita. Target generated via MiniMax-Remover [70]. Includes 5k multi-instance samples. (present → absent).
Obj. Swap	15000	Generated via VACE-14B [12] using GPT-4o prompts and Grounding DINO masks. Covers rigid & non-rigid swaps and 5k multi-instance object swap samples.
Local Style	10000	Generated via VACE-14B [12] using GPT-4o prompts and Grounding DINO masks. Focuses on texture & stylization.
Total	50,000	Unified Dataset

as the source and the object-erased version as the target. Conversely, for the addition task (10k), we invert this pair (absent → present). Notably, the removal subset includes 5,000 samples featuring multi-instance objects to enhance model robustness in complex scenes.

Object Swap and Local Style: To capture fine-grained structural and textural changes, we generated 25,000 samples (15k for swap, 10k for style) utilizing VACE-14B [12]. The generation process is guided by GPT-4o for diverse prompt synthesis and Grounding DINO for precise mask extraction. The swap subset encompasses both rigid and non-rigid object replacements, while the local style subset focuses on texture modification and artistic stylization.

11.2. VideoCoF-Bench

Benchmark Construction. To strictly evaluate the generalization capability, we introduce VideoCoF-Bench, a diverse evaluation set specifically curated to have no overlap with the training domain. The benchmark is constructed from three distinct sources to ensure comprehensive coverage:

- **Pexels [31] Subset:** We manually curated a collection of high-quality videos from Pexels, comprising 50 samples for each editing task. These samples are balanced across the four core editing tasks (Addition, Removal, Swap, and Local Style) to test resolution adaptability and instruction following in varied scenes.
- **Standard Benchmark Integration:** To ensure a fair comparison with existing methods, we incorporated rep-

resentative samples from established benchmarks, including EditVerse [13] and UNIC-Bench [59].

- **Adaptation for Fairness:** Notably, for samples sourced from UNIC-Bench (which typically involves ID-driven editing), we removed the reference identity images. This adaptation unifies the evaluation protocol, focusing purely on text-driven editing capabilities.

This combination results in a highly diverse benchmark that challenges models with unseen content and complex editing instructions.

12. Metrics

GPT Evaluation. To comprehensively assess the editing performance, we employ the state-of-the-art Vision-Language Model, GPT-4o [30], serving as an automated judge. Following the protocol of InstructX [29], we sample three frames from each video pair and utilize structured prompts in [29] to evaluate the results across the following dimensions:

- **Instruction Following (Score 1-10):** This metric measures the precision with which the edit adheres to the user’s specific command. Higher scores indicate that the editing result strictly follows the prompt instructions without ambiguity.
- **Visual Quality (Score 1-10):** This evaluates whether the edited video is visually seamless, natural-looking, and aesthetically pleasing. It penalizes artifacts, distortions, or unnatural transitions introduced during the editing process.
- **Preservation (Score 1-10):** This assesses the coherence with the original video context. It strictly penalizes unintended changes to non-edited regions, ensuring the background and non-target objects remain intact.
- **Success Rate (Binary Yes/No):** To mitigate scoring variance, we incorporate a stricter discrete metric inspired by [22]. GPT-4o performs a binary judgment based on a rigorous three-step verification logic: (1) *Target Identification* (confirming the target matches the descriptor/position); (2) *Modification Accuracy* (verifying the specific edit is applied); and (3) *Strict Preservation* (ensuring no other instances are altered).

As presented in Table 9, VideoCoF achieves superior performance across all these metrics, validating the effectiveness of our reasoning-driven approach.

Perception Quality. In addition to semantic evaluation, we report quantitative metrics to measure the visual alignment and temporal consistency:

- **CLIP-T (Text-Image Alignment):** This metric assesses the semantic alignment between the editing instruction and the output video. We compute the cosine similarity between the CLIP [33] text embedding of the instruction

and the CLIP vision embedding of each output frame, reporting the average score across all frames.

- **CLIP-F (Frame-wise Consistency):** To evaluate temporal stability, we utilize the ViT-L/14 vision encoder from CLIP to extract features for each frame. The consistency score is calculated as the average cosine similarity between feature vectors of adjacent frames.
- **DINO (Structure Consistency):** While CLIP focuses on semantics, we aim to capture more fine-grained structural and textural consistency. We repeat the temporal consistency calculation using features extracted from a pre-trained DINOv2 [4] model. DINO’s self-supervised training enables it to capture object-level details that might be overlooked by CLIP.

13. Future Directions

Scaling up Chain-of-Frames. Currently, VideoCoF achieves SOTA performance in instruction following and success rate using only 50k source-reasoning-editing pairs. This demonstrates remarkable data efficiency compared to existing large-scale baselines. For instance, EditVerse [13] utilizes 4M videos and 8M images, ICVE [22] leverages 2M pre-training data with 150k SFT samples, and InstructX [29] employs 200k SFT samples with joint training. Despite the significant gap in data scale, our method’s superior performance suggests that the “reasoning-then-editing” paradigm is highly effective for Video Diffusion Models (VDMs). A promising future direction is to explore the performance ceiling of VideoCoF by scaling the dataset to 200k or even millions of samples. Investigating how the reasoning capabilities evolve with larger-scale data could reveal new upper limits for precise video editing.

Joint Image-Video Editing and Efficient Architectures. While our current work focuses on video data, integrating high-quality image editing datasets (e.g., MagicBrush [63], NHR-Edit[18]) presents a valuable opportunity. Many recent studies have shown that joint training can enhance visual quality and concept understanding. Future work could investigate the optimal mixture ratios between image and video datasets to maximize performance. Furthermore, designing unified and efficient attention mechanisms is crucial for handling the varying temporal dimensions of images and videos within a single model. Such advancements would likely improve the model’s cross-modal learning capabilities, allowing it to transfer fine-grained editing skills from images to complex video dynamics.

Generalizing VideoCoF to Broader Tasks. VideoCoF has demonstrated exceptional performance in local editing tasks. However, the underlying reasoning framework is inherently flexible and can be extended to a wider range of applications. For *Global Editing* (e.g., style transfer), the reasoning frame could employ a full-frame gray mask to guide

global transformations. For *ID-Driven Editing*, reference identity images could be integrated as “reasoning frames” to guide specific character insertions or swaps. Unifying these diverse tasks—ranging from local modifications to global stylization and ID injection—under the VideoCoF paradigm represents an exciting avenue for future exploration.

Table 9. Quantitative full comparison over 4 video editing tasks on VideoCoF-Bench. We compare VideoCoF with SOTA baselines: InsV2V [7]; Señorita [71] (an I2V model guided by an InstructPix2Pix [2] first frame); VACE-14B [12] (using GPT-4o generated captions); the concurrent work ICVE [22] (pre-trained 1M, fine-tuned 150k); and Lucy Edit Dev [40]. Despite extensive baseline training data, our VideoCoF is fine-tuned on only 50k source-reasoning-editing triplets and shows superior instruction following and success ratio.

Model	GPT-4o Score				Perceptual Quality		
	Instruct Follow \uparrow	Preservation \uparrow	Quality \uparrow	Success Ratio \uparrow	CLIP-T \uparrow	CLIP-F \uparrow	DINO \uparrow
Object Removal							
InsV2V [7]	3.11	4.02	3.77	3.92%	26.85	0.984	0.973
Señorita [71]	3.11	4.68	4.38	9.80%	26.96	<u>0.995</u>	0.990
VACE [12]	N/A	N/A	N/A	0.00%	25.57	0.996	<u>0.995</u>
ICVE [22]	<u>5.38</u>	<u>7.30</u>	7.68	<u>25.49%</u>	26.64	0.994	0.989
Lucy Edit [40]	2.06	4.09	4.45	1.96%	<u>27.37</u>	0.992	0.988
VideoCoF (Ours)	9.65	7.35	<u>6.94</u>	86.27%	27.50	0.988	0.996
Object Addition							
InsV2V [7]	2.71	5.31	4.84	2.04%	25.50	0.985	0.966
Señorita [71]	2.63	5.43	4.80	6.12%	25.26	0.990	<u>0.981</u>
VACE [12]	7.12	5.40	7.38	30.61%	28.01	0.990	0.980
ICVE [22]	<u>8.95</u>	<u>8.65</u>	8.33	<u>77.55%</u>	<u>29.13</u>	0.987	0.974
Lucy Edit [40]	6.96	7.29	6.78	44.90%	27.39	0.987	0.978
VideoCoF (Ours)	9.12	8.78	<u>8.27</u>	79.59%	29.60	0.988	0.982
Object Swap							
InsV2V [7]	1.52	7.37	6.54	0.00%	26.22	0.991	0.984
Señorita [71]	1.69	7.39	6.40	0.00%	25.97	<u>0.994</u>	0.990
VACE [12]	8.11	6.53	7.79	34.62%	<u>26.93</u>	0.995	<u>0.992</u>
ICVE [22]	<u>9.08</u>	8.40	8.57	<u>73.08%</u>	26.54	0.993	0.989
Lucy Edit [40]	6.81	7.58	7.50	44.23%	26.46	0.992	0.988
VideoCoF (Ours)	9.10	<u>8.39</u>	<u>8.14</u>	80.77%	27.10	0.993	0.996
Local Style Transfer							
InsV2V [7]	6.29	7.89	6.89	19.61%	26.19	0.992	0.987
Señorita [71]	5.60	7.69	6.33	25.49%	25.97	0.995	0.992
VACE [12]	7.18	5.53	7.65	41.18%	27.56	<u>0.996</u>	0.994
ICVE [22]	<u>7.75</u>	7.89	7.98	<u>54.90%</u>	<u>27.64</u>	0.994	0.991
Lucy Edit [40]	5.12	7.05	6.73	27.45%	26.71	0.993	<u>0.992</u>
VideoCoF (Ours)	8.02	8.29	<u>7.71</u>	58.82%	27.80	0.997	0.991