

EVE: Efficient Zero-Shot Text-Based Video Editing With Depth Map Guidance and Temporal Consistency Constraints

Yutao Chen^{1*}, Xingning Dong^{2*}, Tian Gan^{1†}, Chunluan Zhou², Ming Yang², Qingpei Guo^{2†}

¹Shandong University

²Ant Group

yt-chen@mail.sdu.edu.cn, dongxingning.dxn@antgroup.com, gantian@sdu.edu.cn,
zhouchunluan.zcl@antgroup.com, m.yang@antgroup.com, qingpei.gqp@antgroup.com

Abstract

Motivated by the superior performance of image diffusion models, more and more researchers strive to extend these models to the text-based video editing task. Nevertheless, current video editing tasks mainly suffer from the dilemma between the high fine-tuning cost and the limited generation capacity. Compared with images, we conjecture that videos necessitate more constraints to preserve the temporal consistency during editing. Towards this end, we propose EVE, a robust and Efficient zero-shot Video Editing method. Under the guidance of depth maps and temporal consistency constraints, EVE derives satisfactory video editing results with an affordable computational and time cost. Moreover, recognizing the absence of a publicly available video editing dataset for fair comparisons, we construct a new benchmark named ZVE-50 dataset. Through comprehensive experimentation, we validate that EVE achieves a satisfactory trade-off between performance and efficiency. Codebase, datasets, and video editing demos are available at <https://github.com/alipay/Ant-Multi-Modal-Framework/blob/main/prj/EVE>.

1 Introduction

Owing to powerful diffusion models [Zhang *et al.*, 2023a; Li *et al.*, 2023b], recent years have witnessed dramatic progress in text-based image synthesis and editing tasks, igniting the soaring research interest in extending these methods to the video editing field. Nevertheless, current text-based video editing methods, which manipulate attributes or styles of videos under the guidance of the driven text, mainly suffer from the dilemma between the considerable fine-tuning cost and the unsatisfied generation performance.

Recent video editing methods could be roughly divided into two categories: tuning-based methods [Zhao *et al.*, 2023; Wu *et al.*, 2022] and zero-shot ones [Couairon *et al.*, 2023; Qi *et al.*, 2023]. The former approaches mainly rely on fine-tuning image diffusion models to derive strong generative priors. Nevertheless, they are usually costly as the fine-tuning step would consume substantial time and GPUs. Towards this end, zero-shot video editing methods aim to directly

edit real-world videos without time-consuming fine-tuning. Nevertheless, the edited videos in the zero-shot manner may suffer from the spatio-temporal distortion and inconsistency. Besides, some zero-shot methods are built upon diffusion models fine-tuned on video datasets, which may not be free of the high cost as the tuning-based ones.

In this paper, we attempt to achieve a trade-off between editing performance and efficiency. Specifically, we adopt the approach of zero-shot video editing, while improving editing performance based upon initial image diffusion models rather than video tuning-based ones. Consequently, the primary challenge is how to preserve and improve the temporal consistency of edited videos.

Let’s begin by considering human editing. When dealing with images, adjusting object appearances or attributes is relatively straightforward. However, when it comes to videos, a comprehensive evaluation of all edited frames becomes imperative to prevent the spatio-temporal distortion and inconsistency in edited videos. As a result, we conjecture that **videos necessitate more temporal constraints** to preserve the time consistency, whose editing process could not be as unconstrained as images. This hypothesis also interprets unsatisfied performance when directly extending image diffusion models to videos, as current image editing methods seldom enforce explicit constraints.

Given this argument, as illustrated in Figure 1, different from current methods that neither explicitly control over individual frame editing nor enforce additional constraints on inter-frame generation, we propose two strategies to reinforce temporal consistency constraints during zero-shot video editing: 1) **Depth Map Guidance**. Depth maps locate spatial layouts and motion trajectories of moving objects, providing robust prior cues for the given video. Therefore, we incorporate depth maps into video editing to improve the temporal consistency. And 2) **Frame-Align Attention**. We enhance the temporal encoding by forcing models to place their attentions on both previous and current frames.

Moreover, by narrowing the gap of whether introducing depth maps into the noise-to-image inference procedure, we design an efficient parameter optimization strategy that directly updates target latent features without fine-tuning the complex diffusion model. In this way, it takes about 83.1 seconds to edit a video with 8 frames on average.

Currently, there lack public video editing datasets for fair

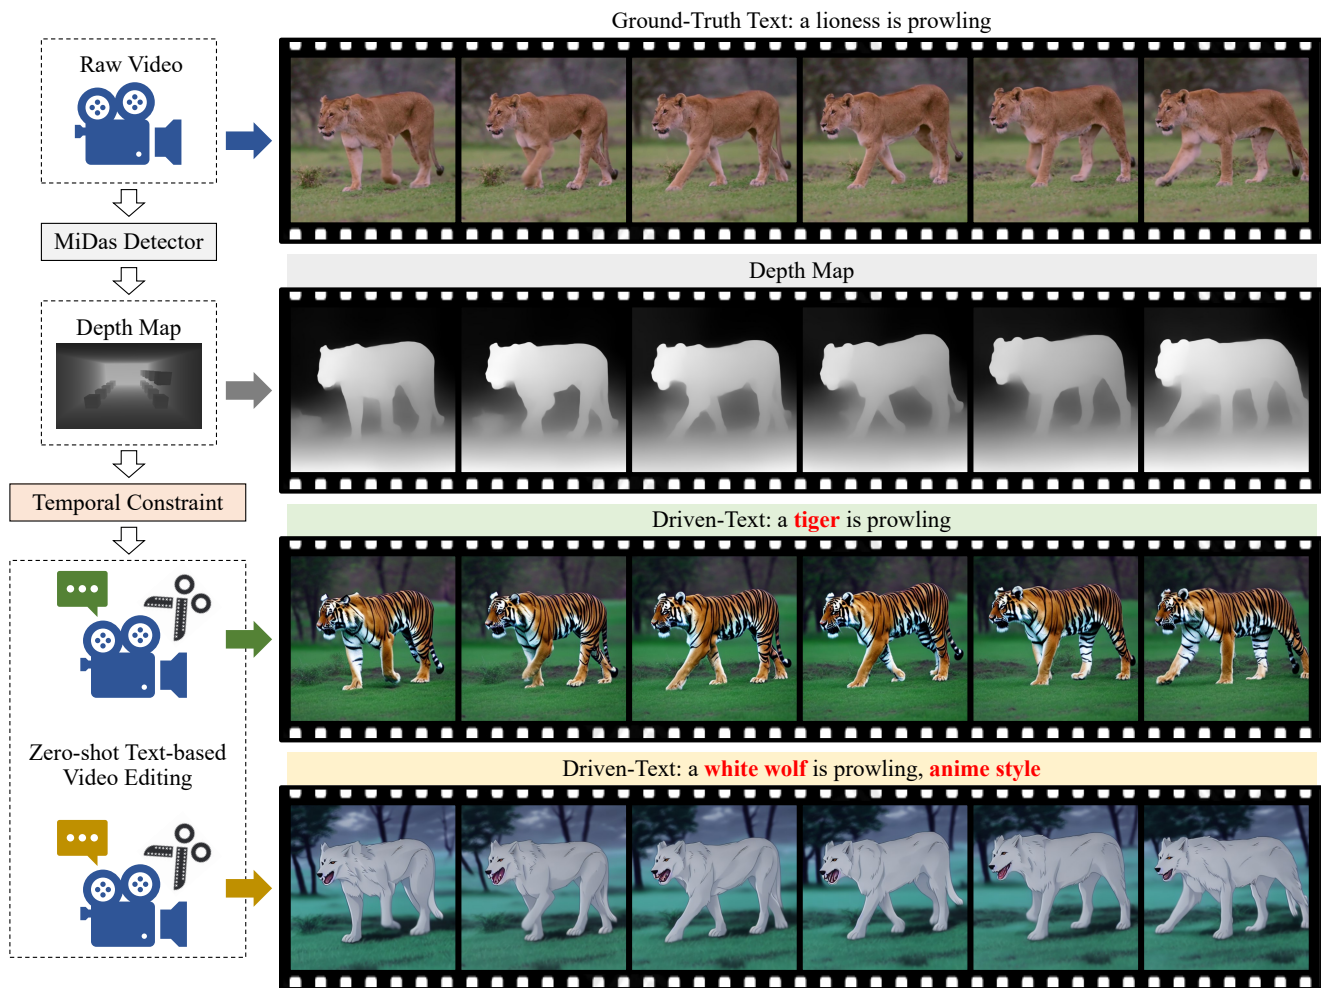


Figure 1: We present EVE, a robust and efficient zero-shot text-based video editor, which trades off editing performance and efficiency.

performance comparisons. Towards this end, we construct a new ZVE-50 dataset, where each collected video is associated with four corresponding driven text. We conduct extensive experiments to benchmark our ZVE-50 dataset.

Our contributions are summarized in four folds:

- We propose EVE, a robust and efficient video editor.
- We propose two strategies to improve the temporal consistency, achieving robust editing performance.
- We construct a new benchmark ZVE-50 dataset. To the best of our knowledge, ZVE-50 is the first dataset for zero-shot text-based video editing, which facilitates future researchers to perform a fair comparison.
- We conduct extensive experiments to affirm the robustness of EVE as a sound method for video editing.


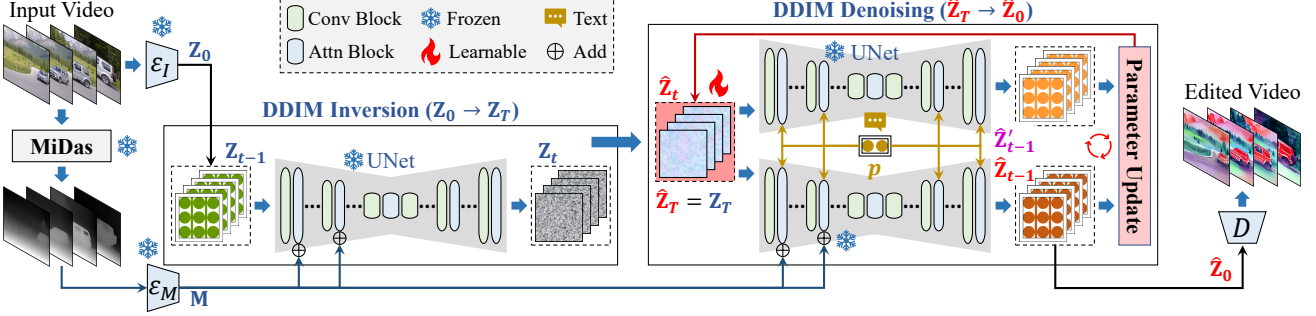
2 Related Work

Diffusion Models. Large-scale diffusion models [Rombach *et al.*, 2022; Croitoru *et al.*, 2023; Dhariwal and Nichol, 2021] have achieved start-of-the-art performance in image synthesis and translation. Diffusion models, in essence, are generative

probabilistic models that approximate a data distribution $p(x)$ by gradually denoising a normally distributed variable. Nevertheless, training a diffusion model from scratch is often expensive and time-consuming.

Text-based Image Generation and Editing. Text-based image editing [Mokady *et al.*, 2023; Zhang *et al.*, 2023b; Choi *et al.*, 2023] aims to manipulate the attributes or styles of one image with the guidance of the driven text. Based on powerful diffusion models, researchers have proposed various methods. *E.g.*, DreamBooth [Ruiz *et al.*, 2023] proposes a subject-driven generation technology by fine-tuning diffusion models, while T2I-Adapters [Mou *et al.*, 2023] provides an efficient image editing approach with a low training cost.

Text-based Video Generation and Editing. Motivated by text-based image editing, video editing [Molad *et al.*, 2023; Lee *et al.*, 2023] has attracted increasing research interest recently, which could be roughly divided into two categories: tuning-based ones [Esser *et al.*, 2023] and zero-shot ones [Qi *et al.*, 2023]. The former approaches mainly edit video attributes by fine-tuning powerful image diffusion


 Figure 2: The **simplified** version of the proposed EVE, presenting the overall video editing pipeline.

 Figure 3: The **elaborated** version of the proposed EVE, detailing the DDIM inversion and denoising procedures.

models, whose training cost is inevitably expensive. Alternatively, FateZero [Qi *et al.*, 2023] proposes the zero-shot video editing task, attempting to generate a text-driven video without extra optimization on complicated generative priors. Nevertheless, FateZero suffers from the limited video editing performance due to weak constraints on the temporal consistency. Moreover, FateZero still heavily relies on Tune-A-Video [Wu *et al.*, 2022], which is a tuning-based diffusion model and is still costly and time-consuming.

3 Methodology

3.1 Preliminary: DMs, LDMs, and DDIMs

DMs. Diffusion Models (DMs) [Sohl-Dickstein *et al.*, 2015] are essentially generative probabilistic models that approximate a data distribution $p(x)$ by gradually denoising a normally distributed variable. Specifically, diffusion models learn to reconstruct the reverse process of a fixed forward Markov chain x_1, x_2, \dots, x_T , where T is the length. The forward Markov chain ($1 \rightarrow T$) could be treated as an image-to-noise procedure, where each Markov transition step $q(x_t|x_{t-1})$ is usually formulated as a Gaussian distribution (\mathcal{N}) with a variance schedule $\beta_t \in (0, 1)$, that is:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}). \quad (1)$$

The reverse Markov chain ($T \rightarrow 1$) could be treated as a noise-to-image procedure, where each reverse Markov transition step $p(x_{t-1}|x_t)$ is formulated as:

$$p(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)), \quad (2)$$

where θ denotes learnable parameters to guarantee that the reverse process is close to the forward one.

Empirically, current diffusion models could be interpreted as an equally weighted sequence of denoising auto-encoders $\epsilon_\theta(x_t, t)$, which is utilized to recover a denoised variant of

their input x_t , and x_t is a noisy version of the input x . The optimization objective could be simplified as:

$$\mathbb{E}_{x, \epsilon \sim \mathcal{N}(0, 1), t} [\|\epsilon - \epsilon_\theta(x_t, t)\|_2^2]. \quad (3)$$

LDMs. Latent Diffusion Models (LDMs) [Rombach *et al.*, 2022] are trained in the learned latent space z_t rather than redundant spatial dimensionality x_t , aiming to remove the noise added to latent image features ϵ_x . LDMs are generally composed of an encoder \mathcal{E} , a time-conditional UNet \mathcal{U} , and a decoder \mathcal{D} , where $z = \mathcal{E}(x)$ and $x \approx \mathcal{D}(\mathcal{E}(x))$. The optimization objective could be formulated as:

$$\mathbb{E}_{\epsilon_x, \epsilon \sim \mathcal{N}(0, 1), t} [\|\epsilon - \epsilon_\theta(z_t, t)\|_2^2]. \quad (4)$$

DDIMs. Denoising Diffusion Implicit Models (DDIMs) [Mokady *et al.*, 2023] could accelerate the sampling from the distribution of images/videos at the denoising step. During inference, deterministic DDIM sampling ($T \rightarrow 1$) aims to recover a clean latent z_0 from a random noise z_T with a noise schedule α_t , which could be formulated as:

$$z_{t-1} = \sqrt{\frac{\alpha_{t-1}}{\alpha_t}} z_t + \left(\sqrt{1 - \alpha_{t-1}} - \sqrt{\frac{1}{\alpha_t} - 1} \right) \cdot \epsilon_\theta. \quad (5)$$

On the contrary, DDIM inversion ($1 \rightarrow T$) aims to process a clean latent z_0 into a noise one \hat{z}_T , which could be simplified as:

$$\hat{z}_t = \sqrt{\frac{\alpha_t}{\alpha_{t-1}}} \hat{z}_{t-1} + \left(\sqrt{1 - \alpha_t} - \sqrt{\frac{1}{\alpha_{t-1}} - 1} \right) \cdot \epsilon_\theta. \quad (6)$$

Compared with conventional DMs that directly employ random noise as inputs and attempt to map each noise vector to a specific image, we exploit DDIM inversion to produce a T steps trajectory between the clean latent z_0 to a Gaussian noise vector z_T . Then we treat z_T as the start vector of the denoising step. This configuration seems appropriate for our

video editing task, since it ensures that the generated video would be close to the original one.

Note that we employ **LDMs** and **DDIM inversion/denoising** in zero-shot text-based video editing. Readers can refer to [Rombach *et al.*, 2022] (LDM) and [Song *et al.*, 2020] (DDIMs) for more details of formulation derivations if necessary.

3.2 Problem Formulation

Given a video V and a prompt text P , zero-shot text-based video editing aims to generate an edited video \hat{V} , which aligns with the description outlined in the prompt P and looks similar to the original video V .

3.3 Overall Framework

As shown in Figure 2 and 3, we present both simplified and elaborated versions of the overall framework. The simplified version could be treated as a flow chart that reveals the whole processing pipeline of our EVE. While the complex one presents detailed information mainly on the iterative DDIM inversion and denoising procedures.

As shown in Figure 2, our EVE is built upon the pre-trained latent diffusion model (LDM), which is composed of a UNet for T-timestep DDIM inversion and denoising. To enforce the temporal consistency of the generated video, we introduce depth maps and exploit them to guide the editing process. Moreover, we propose two consistency constraints to prevent edited videos from spatial or temporal distortion.

We first present the overall pipeline of our EVE based upon Figure 3, including the following five steps.

1. Frozen Features Extraction. Given a video V , we first derive K frames from V , and utilize an image encoder \mathcal{E}_I to obtain **frozen** latent features $\mathbf{Z}_0 = \mathcal{E}(V)$, where $\mathbf{Z}_0 = \{z_0^i\}_{i=1}^K$. Meanwhile, we employ the MiDas Detector [Ranftl *et al.*, 2020] to generate K depth maps from V , and utilize another visual encoder \mathcal{E}_M to obtain **frozen** depth-map features $\mathbf{M} = \{m^i\}_{i=1}^K$. Moreover, we utilize a text encoder \mathcal{E}_p to process the prompt P into **frozen** text features p .

2. DDIM Inversion. Then we repeat DDIM inversion for T steps to derive Gaussian noise vectors \mathbf{Z}_T from video latent features \mathbf{Z}_0 . Each DDIM inversion at timestep t could be formulated as:

$$\mathbf{Z}_t = \text{DDIM}_{\text{inv}}(\mathbf{Z}_{t-1} \mid \mathbf{M}, t) \quad t = 1 \rightarrow T, \quad (7)$$

where DDIM_{inv} denotes DDIM inversion shown in Eq. 6.

To prevent the edited video from temporal distortion and inconsistency, we improve the image-based DDIM inversion operation by introducing depth-map features into the down-sampling pass of the **frozen** UNet, which could rectify the discrepancies among neighboring frames at each inversion step. In this way, we ensure that the generated noise vectors \mathbf{Z}_T would not severely spoil the temporal consistency.

3. DDIM Denoising. Afterward, we repeat DDIM denoising for T steps to obtain edited video features $\hat{\mathbf{Z}}_0$ from DDIM inverted noise $\hat{\mathbf{Z}}_T$, where $\hat{\mathbf{Z}}_T = \mathbf{Z}_T$. Each DDIM denoising at timestep t could be formulated as:

$$\hat{\mathbf{Z}}_{t-1} = \text{DDIM}_{\text{den}}(\hat{\mathbf{Z}}_t \mid p, \mathbf{M}, t), \quad t = T \rightarrow 1, \quad (8)$$

Algorithm 1: DDIM Denoising Procedure.

Input: DDIM inverted noise $\hat{\mathbf{Z}}_T$, text prompt features p , depth-map features \mathbf{M} , learning rate λ

Output: edited video features $\hat{\mathbf{Z}}_0$

```

1 for  $i \leftarrow T$  to 1 do
2    $\hat{\mathbf{Z}}_{t-1} = \text{DDIM}_{\text{den}}(\hat{\mathbf{Z}}_t \mid p, \mathbf{M}, t)$ ;
3    $\hat{\mathbf{Z}}'_{t-1} = \text{DDIM}_{\text{den}}(\hat{\mathbf{Z}}_t \mid p, t)$ ;
4    $\mathcal{L} = 1 - \text{cosin}(\hat{\mathbf{Z}}_{t-1}, \hat{\mathbf{Z}}'_{t-1})$ 
5    $\hat{\mathbf{Z}}_{t-1} = \hat{\mathbf{Z}}_{t-1} - \lambda \Delta_{\hat{\mathbf{Z}}_{t-1}}(\mathcal{L})$ 
6 end
    
```

where DDIM_{den} denotes DDIM denoising shown in Eq. 5.

To prevent the edited video from temporal distortion and inconsistency, we improve the image-based DDIM denoising operation from two aspects: 1) We introduce depth-map features into the down-sampling pass of the **frozen** UNet as DDIM inversion. And 2) we propose the frame-aligned attention to place explicit temporal constraints on the edited video, which will be discussed in the following subsection.

4. Parameter Optimization. To reduce the computation cost and make the generation process more efficient, we freeze all feature extractors (*i.e.*, \mathcal{E}_I , \mathcal{E}_M , and \mathcal{E}_p) and Unets, and only set noise vectors $\hat{\mathbf{Z}}_t$ in DDIM denoising to be trainable. In another words, different from conventional editing methods that update “neural networks”, we directly update “latent noise” to obtain edited videos.

Specifically, at each timestep t in DDIM denoising, except for $\hat{\mathbf{Z}}_{t-1}$, we also derive auxiliary vectors $\hat{\mathbf{Z}}'_{t-1}$ as:

$$\hat{\mathbf{Z}}'_{t-1} = \text{DDIM}_{\text{den}}(\hat{\mathbf{Z}}_t \mid p, t), \quad t = T \rightarrow 1. \quad (9)$$

Compared with $\hat{\mathbf{Z}}_{t-1}$ (Eq. 8), $\hat{\mathbf{Z}}'_{t-1}$ is obtained without strict depth map constraints, which could be treated as free image editing that could unleash the generation capacity of powerful image-based diffusion models. In brief, $\hat{\mathbf{Z}}_{t-1}$ sacrifices the creativity to preserve the temporal consistency, while $\hat{\mathbf{Z}}'_{t-1}$ is just the opposite. Therefore, we leverage the more creative $\hat{\mathbf{Z}}'_{t-1}$ and more temporal consistent $\hat{\mathbf{Z}}_{t-1}$, pursuing to achieve a trade-off between diversity and quality.

The detailed DDIM denoising procedure is illustrated in Algorithm 1, including the parameter optimization step (Lines 4-5). $\Delta_x(\mathcal{L})$ denotes updating trainable x by the gradient descent procedure according to the loss \mathcal{L} , and $\text{cosin}(\cdot, \cdot)$ denotes the cosine similarity computation.

5. Edited Video Decoding. Ultimately, we feed the **frozen** visual decoder \mathcal{D} with generated latent features $\hat{\mathbf{Z}}_0$, obtaining the edited video $\hat{V} = \mathcal{D}(\hat{\mathbf{Z}}_0)$.

3.4 Temporal Consistency Constraints

As aforementioned, we assume that videos necessitate more temporal constraints to preserve the time consistency. Therefore, we propose two strategies to alleviate temporal distortion and inconsistency problems.




Figure 4: The architecture of the attention block within the UNet. Note that we propose the Frame-Aligned Attention to improve the temporal consistency.

1. Depth Map Guidance. Depth maps record visual representations of the distance information, revealing spatial layouts and motion trails of all objects within a video. Therefore, depth maps could be treated as strong prior cues to guide the video editing procedure close to the initial version. Nevertheless, recent video editing methods seldom take advantage of depth maps and neglect the significance of explicitly intervening in the video editing procedure, resulting in intractable temporal distortion and inconsistency problems. Towards this end, we introduce depth maps into the down-sampling pass of the frozen UNet for both DDIM inversion and denoising procedures, forcing the editing process to imitate motion trails and scene transformations of the origin video. In this way, the stability and consistency of the edited video would be improved.

2. Frame-Aligned Attention. We propose the frame-align attention (FAA) to explicitly introduce the temporal information during video editing. As illustrated in Figure 4, a typical UNet comprises a series of “Conv-Attn” blocks to conduct the down-sampling and up-sampling calculation. The conventional attention block (Attn) contains a self-attention (SA) module [Vaswani *et al.*, 2017], a cross-attention (CA) module [Zheng *et al.*, 2021], and a feed-forward network (FFN). The computation of SA(Q, K, V) and CA(Q, K, V) could be formulated as:

$$\begin{cases} \text{SA} : Q = W^Q z^i, K = W^K z^i, V = W^V z^i, \\ \text{CA} : Q = W^Q z^i, K = W^K p, V = W^V p, \end{cases} \quad (10)$$

where W denotes **frozen** projection matrices, z^i is the latent features of the i^{th} frame within the video, and p is the latent features of the text prompt.

Conventional self-attention modules are inherited from image diffusion models, which encodes each frame separately and seem insufficient in preserving the temporal consistency for the video editing. Therefore, we propose the frame-align attention (FAA) to replace K and V with the first frame features z^1 , forcing models to emphasize both previous and current frames for better temporal encoding. The computation of FAA(Q, K, V) could be formulated as:

$$\text{FAA} : Q = W^Q z^i, K = W^K z^1, V = W^V z^1. \quad (11)$$

4 Experiments

4.1 Dataset Construction

Since zero-shot text-based video editing is a novel task, to the best of our knowledge, there lacks a public dataset to perform fair performance and efficiency comparisons. Towards this end, we construct *Zero-shot Video Editing 50* (dubbed as **ZVE-50**) to fulfill this job.

Data Collection. We collect videos from two resources: DAVIS-2017 [Pont-Tuset *et al.*, 2017] and stock-video-footage¹. DAVIS-2017 is a competition dataset for the video object segmentation task [Yao *et al.*, 2020], while stock-video-footage is a public website for free stock video clips and motion graphics. After filtering out videos with similar scenes and styles to avoid the repeatability and promote the diversity, we collect 14 short videos from DAVIS-2017 and 36 ones from stock-video-footage, resulting in the ZVE-50 dataset.

Caption Generation. Then we feed collected videos into BLIP2 [Li *et al.*, 2023a] to obtain the corresponding captions. Specifically, we generate several candidate captions and select the longest one as the ground-truth text.

Prompt Generation. Afterward, we employ GPT-4² to generate the driven text derived from video captions and our manually made prompts. There are four types of driven text, requiring models to edit the given video by 1) Object Replacement (OR), 2) Object Adding (OA), 3) Style Transfer (ST), and 4) Background Changing (BC). Here we present an example of feeding GPT-4 with the manually written prompt and ground-truth caption to obtain the driven text:

Q (human): *Here is a sentence. Please replace the object with another object with a similar shape: “a pink lotus flower in the water with green leaves”*

A (GPT-4): *a pink lotus flower floating in a tranquil koi pond with lily pads*

Ultimately, we manually check all videos, captions, and prompt text to ensure the correctness of the ZVE-50 dataset.

4.2 Experimental Settings

Implementation Details. Zero-shot text-based video editing directly takes a given video and outputs its edited version, which differs from previous methods with explicit training or testing procedure. Specifically, we freeze the pre-trained Latent Diffusion Model as our basic model, where the visual encoder \mathcal{E}_I , the UNet, and the visual decoder \mathcal{D} are inherited from [Rombach *et al.*, 2022] with the version of v1.5. We employ MiDas [Ranftl *et al.*, 2020] to derive depth maps, and utilize frozen Resnet blocks [He *et al.*, 2016] to extract depth map features M . The text encoder \mathcal{E}_p is the pre-trained CLIP text encoder [Radford *et al.*, 2021].

During video editing, following [Wu *et al.*, 2022] and [Qi *et al.*, 2023], we uniformly sample 8 frames at the resolution of $512 * 512$ from each video, and conduct DDIM inversion and denoising steps 50 (T) times. The learning rate λ is 0.8. It takes about **83 seconds** to edit a video on an A40 GPU.

¹<https://www.videvo.net/stock-video-footage/>

²<https://openai.com/gpt-4>

No.	Model	DMG	Attn	Temporal Consistency					Prompt Consistency				
				OR	OA	ST	BC	AVG	OR	OA	ST	BC	AVG
<i>Performance comparisons between the proposed EVE and FateZero.</i>													
A1	FateZero	-	STSA	95.53	95.64	95.84	96.62	<u>95.91</u>	28.92	29.43	29.24	28.90	29.12
A2	VidToMe	-	-	95.07	95.43	95.40	95.60	<u>95.37</u>	29.77	30.78	29.80	30.04	<u>30.10</u>
A3	EVE	✓	FAA	96.41	96.65	96.43	96.70	96.54	30.39	31.01	31.27	28.94	30.40
<i>Ablation study of two proposed temporal consistency constraints within EVE.</i>													
B1	EVE	-	SA	92.59	92.94	92.07	96.30	<u>93.48</u>	25.39	26.43	27.38	28.88	<u>27.02</u>
B2		-	FAA	95.26	95.20	95.16	95.68	<u>95.33</u>	27.77	29.37	29.11	29.71	<u>28.99</u>
B3		✓	SA	94.57	94.02	94.88	95.45	<u>94.73</u>	30.58	31.12	31.16	28.45	<u>30.33</u>
B4		✓	SCA	95.74	96.18	96.13	96.62	<u>96.17</u>	30.26	31.01	31.00	28.91	<u>30.29</u>
B5		✓	FAA	96.41	96.65	96.43	96.70	96.54	30.39	31.01	31.27	28.94	30.40

Table 1: Performance comparisons between our EVE with FateZero and VidToMe, and ablation study of two proposed temporal consistency constraints within EVE. We report the detailed results of four video editing missions and their average performance (underlined), where *OR* = *Object Replacement*, *OA* = *Object Adding*, *ST* = *Style Transfer*, and *BC* = *Background Changing*. All experiments are conducted on one A40 GPU under the same setting. “DMG” denotes with/without the depth map guidance.

Model	Tune-A-Video	FateZero	VidToMe	EVE
Time	~ 30 min	247.6 sec	159.4 sec	83.1 sec

Table 2: Efficiency comparisons between EVE with tuning-based Tune-A-Video, zero-shot FateZero and VidToMe on one A40 GPU.

Evaluation Metrics. We employ two metrics, *i.e.*, Temporal Consistency (TC) and Prompt Consistency (PC), to thoroughly evaluate the quality of edited videos: 1) Following [Esser *et al.*, 2023], we first extract CLIP visual embeddings of all frames within the edited video, and calculate the average cosine similarity between all pairs of neighborhood frames to derive the Temporal Consistency score. 2) Following [Qi *et al.*, 2023], we utilize Text-Video CLIP Score to evaluate the Prompt Consistency between the edited video \hat{V} (K frames) and the driven text p , which is formulated as:

$$\text{CLIP}(\hat{V}, p) = \frac{1}{K} \sum_{k=1}^K \text{CLIP}(\hat{v}^k, p). \quad (12)$$

4.3 Performance and Efficiency Comparisons

As aforementioned, zero-shot text-based video editing is a novel task without public datasets and widely-employed baselines. Moreover, some up-to-date video editing methods [Couairon *et al.*, 2023] do not release their codebase for users to reproduce. Thus, it is intractable to conduct a fair performance comparison with these methods. Therefore, we compare the video editing efficiency between our EVE with the tuning-based Tune-A-Video [Wu *et al.*, 2022], the zero-shot Fatezero [Qi *et al.*, 2023], and the zero-shot VidToMe [Li *et al.*, 2024]. We also compare editing performance among our EVE, Fatezero, and VidToMe in two quantitative metrics.

Regarding efficiency comparisons, as illustrated in Table 2, compared with the tuning-based Tune-A-Video, zero-shot video editing methods are much more efficient (30min \rightarrow <5min). Moreover, compared with the baseline FateZero and VidToMe, the proposed EVE only costs about 1/3 and 1/2 of the total time (247sec / 159sec \rightarrow 83sec) to edit a video, which is more time-efficient and user-friendly.

Regarding performance comparisons, as illustrated in Table 1 (A1 / A2 vs. A3), we observe that our proposed EVE

outperforms the baseline FateZero and VidToMe in all four tasks on the constructed ZVE-50 dataset. It indicates that EVE is a robust and efficient video editing method, which improves the temporal consistency of the generated video.

4.4 Ablation Study

Based on the argument that video editing necessitates more temporal constraints to preserve the time consistency, we propose two constraints to alleviate temporal distortion and inconsistency problems. We conduct several ablation study to verify their effectiveness on our ZVE-50 dataset.

As illustrated in Table 1, we have three observations:

1) Depth maps are strong generative priors that prevent the edited video from temporal distortion and inconsistency. Compared with B2 (without DMG) and B5, we witness an obvious performance decay on both temporal and prompt consistency, indicating the indispensability of the proposed depth map guidance strategy.

2) The proposed Frame-Align Attention (FAA) reinforce the temporal encoding to improve the consistency of edited videos. Compared with B3 (without FAA) and B5, methods equipped with FAA would outperform conventional ones with SA by a large margin, especially on the metric of the temporal consistency.

3) We also compare our FAA with the Sparse-Causal Attention (SCA) mechanism proposed by Tune-A-Video. SCA calculates attentions among current frames and the previous neighborhood ones, which could be formulated as:

$$\text{SCA} : Q = W^Q \dot{z}^i, K = W^K [\dot{z}^1; \dot{z}^{i-1}], V = W^V [\dot{z}^1; \dot{z}^{i-1}], \quad (13)$$

where $[\cdot]$ denotes the concatenation operation. We implement SCA with the same experimental setting. Compared with B4 and B5, we outperform SCA on both temporal and prompt consistency in all four tasks. One possible reason is that SCA employs previous frames, which may result in error propagation. Since each editing step may bring visual distortions, using previous frames will magnify this temporal inconsistency.

As illustrated in Figure 5, our EVE supports four types of applications towards zero-shot text-based video editing.



Figure 5: Results of EVE on four applications: Object Replacement, Object Adding, Style Transfer, and Background Changing.

5 Conclusion

We present EVE, a robust and efficient zero-shot text-based video editing method. Motivated by the observation that videos necessitate more constraints to preserve the time consistency, we introduce depth maps and two temporal consistency constraints to guide the video editing procedure. In this way, the proposed EVE achieves a satisfactory trade-off between performance and efficiency.

6 Future Work

In the future, we aim to further improve the quality of edited videos, narrowing the performance gap between tuning-based video editing methods and zero-shot ones. *E.g.*, introducing the triplet attention mechanism [Zhou *et al.*, 2021] to promote the temporal stability; and generating pseudo labels by recording attention maps of Unets to build a knowledge distillation mechanism in the denoising step.

Acknowledgements

This work is supported by the National Natural Science Foundation of China under Grants 62176137, Ant Group.

Contribution Statement

Equal contribution: Yutao Chen and Xingning Dong.
Corresponding authors: Qingpei Guo and Tian Gan.

References

- [Choi *et al.*, 2023] Jooyoung Choi, Yunjey Choi, Yunji Kim, Junho Kim, and Sungroh Yoon. Custom-edit: Text-guided image editing with customized diffusion models. *arXiv preprint arXiv:2305.15779*, 2023.
- [Couairon *et al.*, 2023] Paul Couairon, Clément Rambour, Jean-Emmanuel Haugeard, and Nicolas Thome. Videdit: Zero-shot and spatially aware text-driven video editing. *arXiv preprint arXiv:2306.08707*, 2023.
- [Croitoru *et al.*, 2023] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [Dhariwal and Nichol, 2021] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34:8780–8794, 2021.
- [Esser *et al.*, 2023] Patrick Esser, Johnathan Chiu, Parmida Atighehchian, Jonathan Granskog, and Anastasis Germanidis. Structure and content-guided video synthesis with diffusion models. *arXiv preprint arXiv:2302.03011*, 2023.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [Lee *et al.*, 2023] Yao-Chih Lee, Ji-Ze Genevieve Jang, Yi-Ting Chen, Elizabeth Qiu, and Jia-Bin Huang. Shape-aware text-driven layered video editing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 14317–14326, 2023.
- [Li *et al.*, 2023a] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [Li *et al.*, 2023b] Yifan Li, Kun Zhou, Wayne Xin Zhao, and Ji-Rong Wen. Diffusion models for non-autoregressive text generation: A survey. *arXiv preprint arXiv:2303.06574*, 2023.
- [Li *et al.*, 2024] Xirui Li, Chao Ma, Xiaokang Yang, and Ming-Hsuan Yang. Vidtome: Video token merging for zero-shot video editing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2024.
- [Mokady *et al.*, 2023] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6038–6047, 2023.
- [Molad *et al.*, 2023] Eyal Molad, Eliahu Horwitz, Dani Valevski, Alex Rav Acha, Yossi Matias, Yael Pritch, Yaniv Leviathan, and Yedid Hoshen. Dreamix: Video diffusion models are general video editors. *arXiv preprint arXiv:2302.01329*, 2023.
- [Mou *et al.*, 2023] Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhongang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023.
- [Pont-Tuset *et al.*, 2017] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017.
- [Qi *et al.*, 2023] 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. *arXiv preprint arXiv:2303.09535*, 2023.
- [Radford *et al.*, 2021] 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.
- [Ranftl *et al.*, 2020] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(3):1623–1637, 2020.
- [Rombach *et al.*, 2022] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [Ruiz *et al.*, 2023] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023.
- [Sohl-Dickstein *et al.*, 2015] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [Song *et al.*, 2020] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.

- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- [Wu *et al.*, 2022] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Weixian Lei, Yuchao Gu, 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. *arXiv preprint arXiv:2212.11565*, 2022.
- [Yao *et al.*, 2020] Rui Yao, Guosheng Lin, Shixiong Xia, Jiaqi Zhao, and Yong Zhou. Video object segmentation and tracking: A survey. *ACM Transactions on Intelligent Systems and Technology*, 11(4):1–47, 2020.
- [Zhang *et al.*, 2023a] Chenshuang Zhang, Chaoning Zhang, Mengchun Zhang, and In So Kweon. Text-to-image diffusion model in generative ai: A survey. *arXiv preprint arXiv:2303.07909*, 2023.
- [Zhang *et al.*, 2023b] Zhixing Zhang, Ligong Han, Arnab Ghosh, Dimitris N Metaxas, and Jian Ren. Sine: Single image editing with text-to-image diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6027–6037, 2023.
- [Zhao *et al.*, 2023] Min Zhao, Rongzhen Wang, Fan Bao, Chongxuan Li, and Jun Zhu. Controlvideo: Adding conditional control for one shot text-to-video editing. *arXiv preprint arXiv:2305.17098*, 2023.
- [Zheng *et al.*, 2021] Lecheng Zheng, Yu Cheng, Hongxia Yang, Nan Cao, and Jingrui He. Deep co-attention network for multi-view subspace learning. In *Proceedings of the Web Conference*, pages 1528–1539, 2021.
- [Zhou *et al.*, 2021] Haoyi Zhou, Jianxin Li, Jieqi Peng, Shuai Zhang, and Shanghang Zhang. Triplet attention: Rethinking the similarity in transformers. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2378–2388, 2021.