

TrackGo: A Flexible and Efficient Method for Controllable Video Generation

Haitao Zhou^{*1,2}, Chuang Wang^{*1,2}, Rui Nie¹, Jinlin Liu², Dongdong Yu²,
Qian Yu^{1†}, Changhu Wang^{2†}

¹ Beihang University

² AISphere Tech

{zhouhaitao, chuangwang, qianyu, nierui_}@buaa.edu.cn

{yudongdong, linjinxiao, wangchanghu}@aishi.ai

Abstract

Recent years have seen substantial progress in diffusion-based controllable video generation. However, achieving precise control in complex scenarios, including fine-grained object parts, sophisticated motion trajectories, and coherent background movement, remains a challenge. In this paper, we introduce *TrackGo*, a novel approach that leverages free-form masks and arrows for conditional video generation. This method offers users with a flexible and precise mechanism for manipulating video content. We also propose the *TrackAdapter* for control implementation, an efficient and lightweight adapter designed to be seamlessly integrated into the temporal self-attention layers of a pretrained video generation model. This design leverages our observation that the attention map of these layers can accurately activate regions corresponding to motion in videos. Our experimental results demonstrate that our new approach, enhanced by the TrackAdapter, achieves state-of-the-art performance on key metrics such as FVD, FID, and ObjMC scores.

Project Page — <https://zhtjtcz.github.io/TrackGo-Page/>

Introduction

With the rapid development of diffusion models (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2020; Dhariwal and Nichol 2021; Song et al. 2020), video generation has witnessed significant progress, with the quality of generated videos continuously improving. Unlike text-based (Blattmann et al. 2023b; Zhang et al. 2023; Guo et al. 2023b) or image-based (Blattmann et al. 2023a) video generation, controllable video generation (Hu et al. 2023; Yin et al. 2023; Wu et al. 2024) focuses on achieving precise control over object movement and scene transformations in generated videos. This capability is particularly valuable in industry such as film production and cartoon creation.

Controllable video generation remains a highly challenging task. The primary challenge is precise control, which includes managing the target movement objects and their trajectories. Existing methods often struggle to achieve precise control over these elements. For instance, DragAny-




Figure 1: Attention map visualization of the last temporal self-attention layer in Stable Video Diffusion Model. The highlighted areas in the attention map correspond to the moving areas in the video. The video has a total of 25 frames, and we selected frames 1, 12, and 23 at equal intervals for visualization. And $Attn(i, j)$ denotes the temporal attention map between frame i and frame j .

thing (Wu et al. 2024) employs a center point and a Gaussian map to guide the target object along a predefined path. However, it fails to control the movement of partial or fine-grained objects effectively. Another approach, Boximator (Wang et al. 2024a), utilizes bounding boxes to dictate motion control. It uses a box to specify the target area, where the sequence of movements of the box guides the motion of the target. Unfortunately, bounding boxes often encompass redundant regions, which can interfere with the motion of the target and disrupt the coherence of the background in the generated videos. The second challenge is efficiency. Existing works often incorporate conditions in a way that significantly increases the number of model’s parameters. For instance, DragAnything utilizes the architecture of ControlNet (Zhang, Rao, and Agrawala 2023), and DragNUWA (Yin et al. 2023) employs heavy encoders to map guidance signals into the latent space of the pretrained model. These design choices inevitably lead to slower inference times, which can impede the practical deployment of these models in real-world applications.

In this work, we tackle the task of controllable video generation by addressing two crucial questions: First, *what* type of control should be employed to accurately describe the motion of the target? Second, *how* can this control be implemented efficiently?

For the *first* question regarding the type of control suitable for describing the motion of the target, we propose a novel combination of a free-form mask and an arrow to

*These authors contributed equally. Interns in AISphere Tech.

†Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 2: Example videos generated by our proposed *TrackGo*. Given an initial frame, users specify the target moving object(s) or part(s) using free-form masks and indicate the desired movement trajectory with arrows. *TrackGo* is capable of generating subsequent video frames with precise control. It can handle complex scenarios that involve multiple objects, fine-grained object parts, and sophisticated movement trajectories.

guide motion. Specifically, users can define the target area with a brush, allowing for precise specification ranging from entire objects to partial areas. The trajectory of the movement is indicated by an arrow, also drawn by the user, which provides clear directional guidance. For the *second* question concerning efficient control implementation, we introduce a novel approach that involves injecting conditions into the temporal self-attention layers. We observed that the attention maps generated by these layers effectively highlight areas of motion within a video, a finding also supported by previous research (Ma, Lewis, and Kleijn 2023). As demonstrated in Fig.1, the region of a moving train is distinctly activated in the attention map of the temporal self-attention layer. Building on this insight, we propose directly manipulating the attention map of the temporal self-attention layer to achieve precise motion control. This method not only enhances accuracy but also minimizes additional computational overhead.

We introduce **TrackGo**, a novel framework for controllable video generation that leverages user inputs to direct the generation of video sequences. *TrackGo* uses free-form masks and arrows provided by users to define target regions and movement trajectories, respectively, as illustrated in Fig. 2. This approach consists of two stages: Point Trajectories Generation and Conditional Video Generation. In the first stage, *TrackGo* automatically extracts point trajectories from the user-defined masks and arrows. These trajectories serve as precise blueprints for video generation. In the second stage, we use the Stable Video Diffusion Model

(SVD) (Blattmann et al. 2023a) as the base model, accompanied by an encoder that encodes the motion information. To ensure that the guidance is precisely interpreted by our model, we introduce the novel **TrackAdapter**. This adapter effectively modifies the existing temporal self-attention layers of a pre-trained video generation model to accommodate new conditions, enhancing the model’s control over the generated video.

Specifically, the *TrackAdapter* introduces a dual-branch architecture within the existing temporal self-attention layers. It integrates an additional self-attention branch running parallel to the original. This new branch is specifically designed to focus on the motion within the target area, ensuring that the movement dynamics are captured with high fidelity. Meanwhile, the original branch continues to handle the rest areas. This architecture not only ensures accurate and cohesive generation of both the specific movements of the target and the overall video context but also modestly increases the computational cost. Furthermore, we introduce an attention loss to accelerate model convergence, thereby enhancing efficiency. This balance between control fidelity and efficiency is crucial for practical applications of video generation.

In summary, our contributions are three fold:

- We introduce a novel motion-controllable video generation approach named *TrackGo*. This method offers users a flexible mechanism for motion control, combining masks and arrows to achieve precise manipulation in complex scenarios, including those involving multi-

ple objects, fine-grained object parts, and sophisticated movement trajectories.

- A new component, the TrackAdapter, is developed to integrate motion control information into temporal self-attention layers effectively and efficiently.
- We conduct extensive experiments to validate our approach. The experimental results demonstrate that our model surpasses existing models in terms of video quality (FVD), image quality (FID), and motion faithfulness (ObjMC).

Related Work

Diffusion Based Image and Video Generation

Diffusion models (Ho, Jain, and Abbeel 2020) have made great progress in the field of text-to-image generation (Betker et al. 2023; Peebles and Xie 2023; Rombach et al. 2022; Saharia et al. 2022; Zhang, Rao, and Agrawala 2023; Xing et al. 2023), which directly promotes the progress of basic video diffusion, and many excellent works (Chen et al. 2024; Henschel et al. 2024; Ho et al. 2022; Zhang et al. 2023; Wang et al. 2023a; Zeng et al. 2023) have emerged. Early works such as VLDM (Blattmann et al. 2023b) and AnimateDiff (Guo et al. 2023b) try to insert temporal layers to complete the generation of text to video. Recent models benefit from the stability of diffusion-based trained models. I2VGen-XL (Zhang et al. 2023), Stable Video Diffusion (SVD) (Blattmann et al. 2023a) achieve surprising results on text-to-video generation and image-to-video, respectively, with large-scale high-quality data. While these models are capable of producing high-quality videos, they primarily depend on coarse-grained semantic guidance from text or image prompts, which can lead to actions that do not align with the user’s intentions.

Controllable Image and Video Generation

In pursuit of enhancing controllability in image and video content generation, numerous recent studies (Mou et al. 2024; Ye et al. 2023; Khachatryan et al. 2023; Ceylan, Huang, and Mitra 2023) have integrated diverse methodologies to incorporate additional forms of guidance. Notably, Disco (Wang et al. 2023b), MagicAnimate (Xu et al. 2023), DreamPose (Karras et al. 2023), and Animate Anyone (Hu et al. 2023) have each adopted pose-directed approaches, enabling the creation of videos featuring precisely prescribed poses. These advancements reflect a concerted effort towards achieving finer-grained manipulation and tailored video synthesis through pose-guided techniques. Despite demonstrating remarkable performance, these techniques are specifically tailored to and thus limited in application to videos depicting human subjects.

To broaden the scope of controllable video synthesis (Wang et al. 2024b; Guo et al. 2023a; Ma et al. 2024; Wang et al. 2023c; Yin et al. 2023; Wang et al. 2024a; Wu et al. 2024; Tu et al. 2023; Jiang et al. 2023; Qi et al. 2023) and enhance its general applicability, several approaches incorporate control signals into pre-trained video diffusion models. The work of AnimateDiff endeavors, to integrate

LoRA (Low-Rank Adaptation) (Hu et al. 2021) within temporal attention mechanisms to grasp camera motion learning. Nevertheless, LoRA’s potential shortcomings emerge when tasked with thoroughly manipulating and synthesizing sophisticated movements, as its learning capacity and predictive prowess might be constrained by repetitive motion patterns present in the training dataset. Innovations like MotionCtrl (Wang et al. 2023c) and DragNUWA (Yin et al. 2023) encode sparse optical flow into dense optical flow as guidance information to inject into the diffusion model to control object motion. Boximator (Wang et al. 2024a), on the other hand, pioneers motion regulation by correlating objects with bounding boxes, capitalizing on the model’s innate tracking capabilities. Fundamentally, this approach employs dual trajectory sets defined by the bounding box’s upper-left and lower-right vertices. A caveat arises from the boxes occasionally misinterpreting background details as part of the foreground, which can inadvertently taint output accuracy. DragAnything (Wu et al. 2024) has introduced an approach that employs masks to pinpoint a central point, subsequently generating a Gaussian map for tracking this center to produce a guiding trajectory for model synthesis. It’s crucial to note, however, that not all masked areas denote regions of motion. And relying on the ControlNet (Zhang, Rao, and Agrawala 2023) structure makes it difficult to achieve satisfactory results in terms of efficiency.

Methodology

Overview

Our task is motion-controllable video generation. For an input image $I \in \mathbb{R}^{H \times W \times 3}$, and point trajectories $P \in \mathbb{R}^{L \times H \times W \times 3}$ extracted from arrows describes the trajectories information, generating a video $V \in \mathbb{R}^{L \times H \times W \times 3}$ in line with the trajectories, where L is the length of video.

We use Stable Video Diffusion Model (SVD) (Blattmann et al. 2023a) as our base architecture. The SVD model, similar to most video latent diffusion models (Blattmann et al. 2023b), adds a series of temporal layers on the U-Net of image diffusion to form 3D U-Net. On this basis, We pass the point trajectories P through a trainable encoder \mathcal{E} to obtain a compressed representation f . This representation is then injected into each temporal self-attention module, using down-sampling to process and adapt it to the appropriate resolution. We introduce the **TrackAdapter** and at each temporal self-attention of SVD, a TrackAdapter is added to inject f , as shown in Fig. 3.

In the following sections, we will cover three main topics: (1) The advantages of point trajectories and how we obtain and use them. (2) The structure of TrackAdapter and how it helps SVD to understand complex motion patterns and complete the generation of complex actions. (3) The process of training and inferring our model.

Point Trajectories Generation

In inference, when the user provides the first frame picture, the masks of editing area, and corresponding arrows. Our approach can convert the masks and arrows into point trajectories P through preprocessing, as shown in Fig. 3.

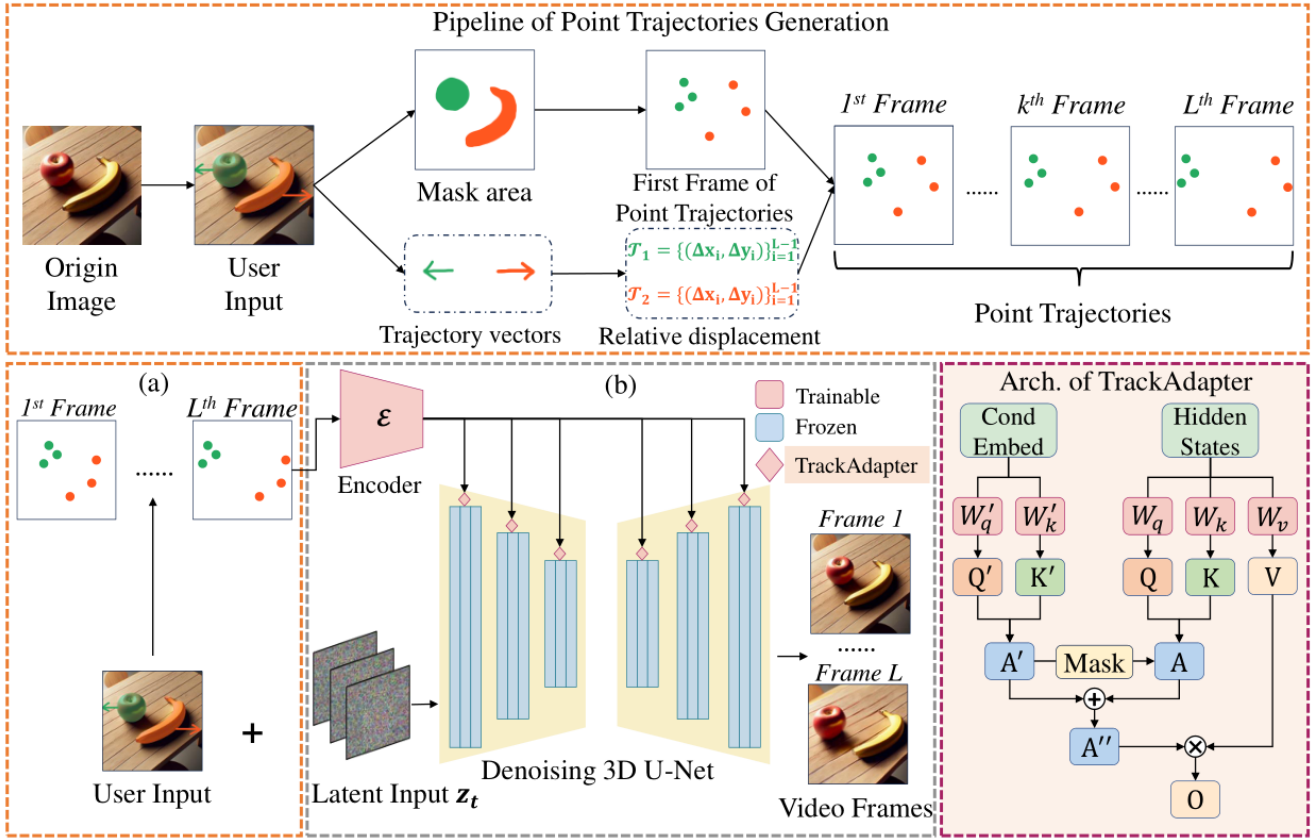


Figure 3: **Top: Pipeline of Point Trajectories Generation.** User’s inputs are divided into masks and trajectory vectors for processing. Each mask corresponds to a trajectory vector. For each mask area, $K * s$ points are randomly selected. The trajectory vector is then subdivided by the frame number to attain the relative displacement \mathcal{T} of each point between frames of adjacent frames. The final step is to combine this relevant data to construct point trajectories. **Bottom: Overview of TrackGo.** TrackGo generates videos by taking user input I and latent input z_t as inputs based on an image-to-video diffusion model. Through the pipeline of point trajectories generation(part (a)), point trajectories P can be obtained from I . Then the point trajectories P are passed through the Encoder \mathcal{E} and injected into the model via the TrackAdapter, as shown in part (b). **Architecture of TrackAdapter** describes the calculation process of TrackAdapter.

For the training process, we first use DEVA (Cheng et al. 2023) to segment the main components in the ground truth video P and obtain the corresponding segmentation sequence M_i^j , where i denotes the i -th frame and j denotes the j -th component. Then, for the mask sequence $\{M_1^j\}_{j=1}^s$ of the first frame, we need to select K points in each mask area as control points, where s denotes the number of components. For mask M_i^j , we randomly select $3K$ points in the white area, and then use the K-means (MacQueen et al. 1967) to obtain K points. This guarantees the uniformity of the chosen k points without incurring too much time overhead. We will have a total of $s * K$ control points after this stage. After obtaining the control points, we use the Co-Tracker (Karaev et al. 2023) to track these points and obtain the corresponding motion trajectories $\mathcal{T} = \{\{(x_i^1, y_i^1)\}_{i=1}^L, \{(x_i^2, y_i^2)\}_{i=1}^L, \dots, \{(x_i^{s*K}, y_i^{s*K})\}_{i=1}^L\}$, where L is the length of video. Finally, we assign a color to the control points corresponding to the same component, plot the trajectory \mathcal{T} and get point trajectories

$P \in \mathbb{R}^{L \times H \times W \times 3}$. We built the training dataset using this method, and after data cleaning and other operations, we finally got 110k $(V, P, \{M_1^j\}_{j=1}^s)$ triple pairs.

Injecting Motion Conditions via TrackAdapter

Motion Conditions Extraction. We use the same encoder structure from Animate Anyone (Hu et al. 2023) to extract the timed features. This process can be obtained from Eq. 1, where f is the compressed temporal representation of P .

$$f = \mathcal{E}(P) \quad (1)$$

The f will be down-sampled according to the resolution of different temporal self-attention layers and aligned with its input size.

TrackAdapter Design. In order to use the compressed temporal representation f to guide the model to generate a video corresponding to this action, a straightforward approach is to construct the attention map shown in Fig. 1 us-

ing f . Therefore, we propose a lightweight and simple structure called TrackAdapter. The function of TrackAdapter is to activate a motion region corresponding to a specified object, thus guiding the model generation process. When injecting point trajectories, TrackAdapter is responsible for activating the motion region of the specified object. We first compute the attention map \mathbf{A}' for the TrackAdapter:

$$\mathbf{A}' = \text{softmax}\left(\frac{\mathbf{Q}'(\mathbf{K}')^T}{\sqrt{d}}\right) \quad (2)$$

where $\mathbf{Q}' = f\mathbf{W}'_q$, $\mathbf{K}' = f\mathbf{W}'_k$ are the query, key matrices of the TrackAdapter, f is the compressed representation of \mathbf{P} , obtained from Eq. 1.

To avoid the impact of the origin temporal self-attention branch on the final active region, we obtain an attention mask according to the attention map to suppress the areas activated by the origin temporal self-attention branch.

We transform the attention map \mathbf{A}' into the corresponding attention mask \mathbf{A}_M by setting a threshold α :

$$\mathbf{A}_{Mij} = \begin{cases} -\text{inf} & \text{if } \mathbf{A}'_{ij} \geq \alpha \\ 0, & \text{if } \mathbf{A}'_{ij} < \alpha \end{cases} \quad (3)$$

The motion region attended to by the TrackAdapter is the part of the attention map \mathbf{A}' that exceeds α . By setting the equivalent area in the attention mask \mathbf{A}_M to $-\text{inf}$, the original temporal self-attention will no longer pay attention to this part and produce the separation effect. Then, the attention map of the original temporal self-attention can be rewritten as,

$$\mathbf{A} = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T}{\sqrt{d}} + \mathbf{A}_M\right) \quad (4)$$

Finally, we get the temporal self-attention output of the current block: $\mathbf{O} = (\mathbf{A} + \mathbf{A}')\mathbf{V} = \mathbf{A}''\mathbf{V}$, where $\mathbf{Q} = \mathbf{X}\mathbf{W}_q$, $\mathbf{K} = \mathbf{X}\mathbf{W}_k$, $\mathbf{V} = \mathbf{X}\mathbf{W}_v$ are the query, key, and value matrices of the temporal self-attention operation, \mathbf{X} is input feature and \mathbf{A}'' is final attention map. We complete the separation of the corresponding area and the unspecified area during the calculation of attention.

Training and Inference of TrackAdapter

Training The video diffusion model iteratively predicts noise ϵ in the noisy input, gradually transforming Gaussian noise into meaningful video frames. The optimization of the model ϕ is achieved through noise prediction loss,

$$\mathcal{L} = \mathbb{E}_{t \sim \mathcal{U}(0,1), \epsilon \sim \mathcal{N}(0, \mathbf{I})} \|\epsilon_\theta(\mathbf{z}_t; c) - \epsilon\|_2^2 \quad (5)$$

where t represents the timesteps, θ represents the U-Net's parameters, c represents conditions and \mathbf{z}_t is a noisily transformed version of ground truth video \mathbf{z}_0 :

$$\mathbf{z}_t = \alpha_t \mathbf{z}_0 + \sigma_t \epsilon \quad (6)$$

Here, α_t and σ_t denotes a predefined constant sequence. On the basis of Eq 5, we add point trajectories \mathbf{P} and an image \mathbf{I} as conditions, and the optimization objective can be written as,

$$\mathcal{L}_{\text{mse}} = \|\epsilon_\theta(\mathbf{z}_t; \mathcal{E}(\mathbf{P}), \mathbf{I}, t) - \epsilon\|_2^2 \quad (7)$$

In order to adapt the original temporal self-attention to the new input mode quickly and to accelerate the model's convergence, we design an attention map based loss function. We gather attention maps from different temporal self-attention layers to get a set \mathcal{C} , which contains 16 different attention maps. For $\mathbf{A}_q \in \mathcal{C}$, $\mathbf{A}_q(x, y)$ denotes the temporal self-attention map between frame x and frame y at block q . The purpose of the attention loss is to suppress the area corresponding to the mask in the original branch's attention map, *i.e.* the motion area,

$$\mathcal{L}_{\text{attn}} = \sum_{\mathbf{A}_q \in \mathcal{C}} \sum_{i=1}^l (\mathbf{A}_q(i, i) * \phi(\mathbf{M}_i))^2 \quad (8)$$

where ϕ denotes the down-sampling operation, and \mathbf{M}_i denotes the mask of all moving components in i -th frame. In total, our final loss function is then defined by the weighted average of the two terms,

$$\min_{\theta} \nabla_{\theta} \mathcal{L}_{\text{mse}} + \lambda \mathcal{L}_{\text{attn}} \quad (9)$$

where λ is a hyperparameter.

Inference. During inference, we set the intensity of the unspecified area to τ , that is, we set the part of the Eq 3 less than α to τ . Users can adjust τ to control the movement of the unspecified area in cases where it needs to move synchronously with the foreground movement or when sensory interference from the unspecified area needs to be mitigated. This feature will greatly enhance the creation of fluid and highly synchronized motion videos.

Experiment Settings

Implementation Details. We employ SVD (Blattmann et al. 2023a) as our base model. All experiments were conducted using PyTorch with 8 NVIDIA A100-80G GPUs. AdamW (Loshchilov and Hutter 2017) is configured as our optimizer, running for a total of 18,000 training steps with a learning rate of $3e-5$ and a batch size of 8. Following the method proposed in Animate Anyone (Hu et al. 2023), we have developed a lightweight encoder \mathcal{E} . This encoder employs a total of six convolution layers, two pooling layers, and a final fully-connected layer. Its primary function is to align the point trajectories \mathbf{P} to the appropriate resolution. The query and key matrices of the TrackAdapter are initialized from the original temporal attention branch.

Dataset. For our experiments, we utilized an internal dataset characterized by superior video quality, comprising about 200K video clips. Following the experimental design, we further filtered the data to obtain a subset of about 110K videos as our final training dataset. During the training process, each video was resized to a resolution of 1024×576 and standardized to 25 frames per clip.

Our test set comprised the VIPSeg validation set along with an additional 300-video subset from our internal validation dataset. Notably, all videos in the VIPSeg dataset are formatted in a 16:9 aspect ratio. To maintain consistency, we adjusted the resolution of all videos in the validation sets to 1024×576 . For evaluation purposes, we extracted the trajectories from first 14 frames of each video in the test set.

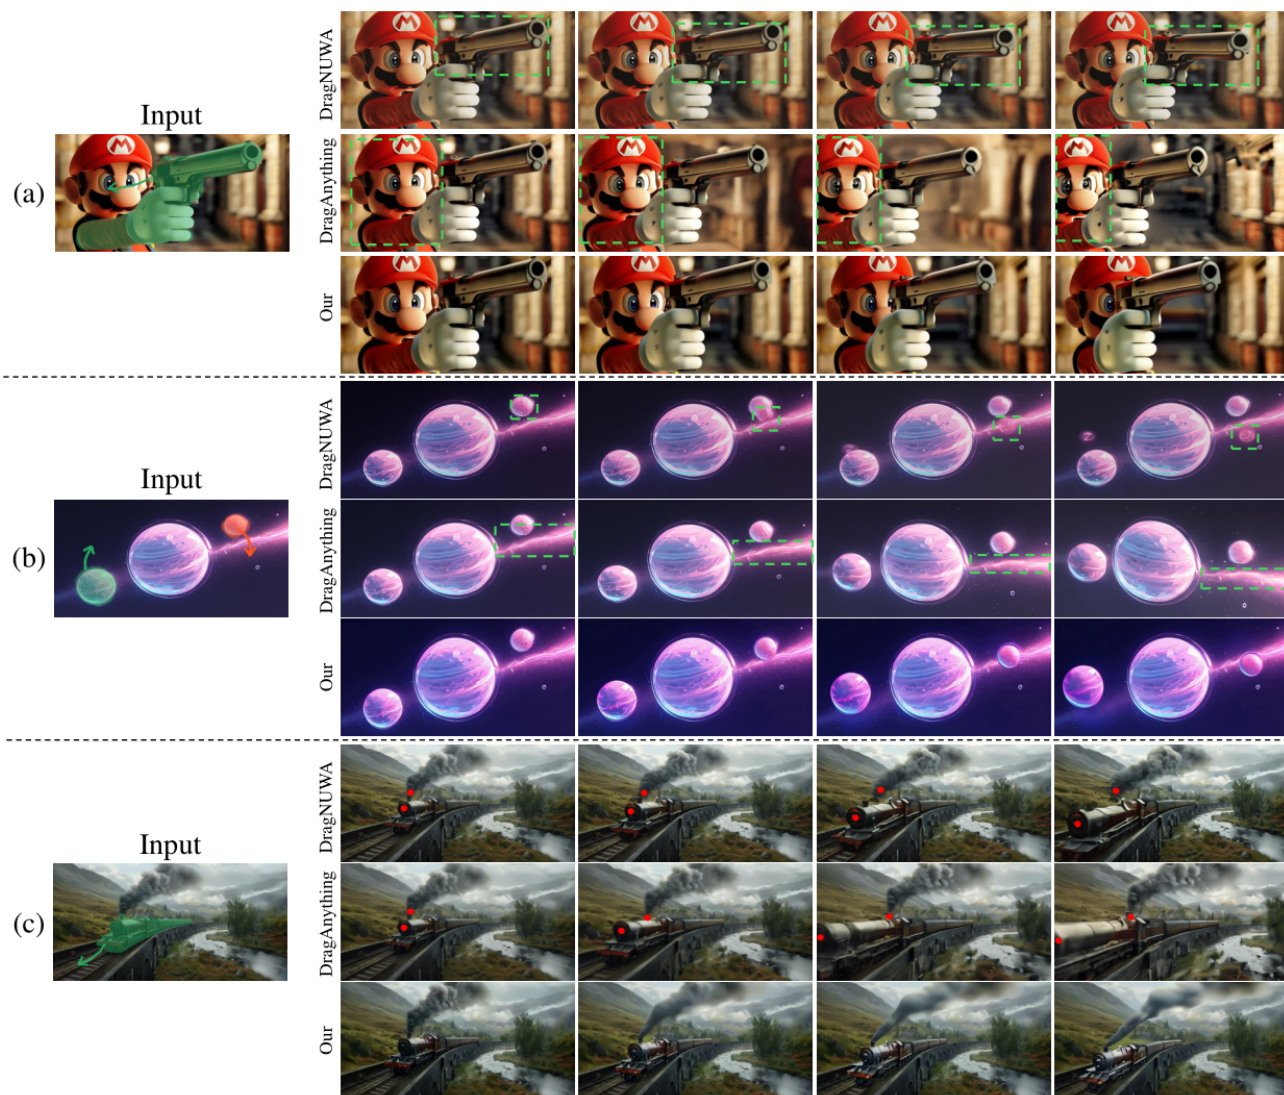


Figure 4: Qualitative comparisons between our method and baseline methods, DragAnything and DragNUWA. We use colorful symbols to highlight the undesired parts of the results generated by the other two approaches.

Evaluation Metrics and Baseline Methods. We measure video quality using FVD (Unterthiner et al. 2018) and image quality using FID (Heusel et al. 2017). We compare our methods with DragNUWA and DragAnything, which can also use trajectory information as conditional input. Following DragAnything, ObjMC is used to evaluate the motion control performance by computing the Euclidean distance between the predicted and ground truth trajectories.

Quantitative Evaluation

Quantitative comparisons of our method with baseline methods are shown in Table 1. We test all models on VIPSeg validation set and internal validation set. From the results, TrackGo outperforms other approaches across all metrics, producing videos with higher visual quality and better fidelity to input motion control. We also compared the model

parameters and inference speed of the three approaches. Since all use the same base model, our comparison focused on the weights of the newly added modules. To assess the model’s inference speed, we conducted 100 inference tests for each approach using identical input data on an NVIDIA A100 GPU. The results demonstrate that our approach not only delivers the best visual quality but also achieves the fastest inference speed, all while requiring the fewest additional parameters.

Qualitative Evaluation

Visualizations. We present a visualization comparison with DragAnything and DragNUWA, as shown in Fig. 4. We can have the following observations: **First**, DragNUWA struggles with perceiving the control area, which may lead to incomplete or inaccurate optical flow. In case (b), the planet

Method	VIPSeg			Internal validation dataset			Performance Metrics	
	FVD ↓	FID ↓	ObjMC ↓	FVD ↓	FID ↓	ObjMC ↓	Parameters	Inference Time
DragNUWA	321.31	30.15	298.98	178.37	38.07	129.80	160.38M	58.12s
DragAnything	294.91	28.16	236.02	169.73	32.85	133.89	685.06M	152.98s
TrackGo	248.27	25.60	191.15	136.11	29.19	79.52	29.36M	33.94s

Table 1: Quantitative comparisons of our approach with DragAnything and DragNUWA. All three methods are based on the same basic model called SVD. **Bolded values** indicate the best scores in each column.

is not correctly perceived, and in case (a), the movement of the gun is also incorrect. In case (c), while the optical flow of the train is successfully predicted, the absence of flow in the smoke creates a jarring effect. Second, DragAnything struggles with partial or fine-grained object movement. In case (a), only the gun and Mario’s hand should move, but Mario’s entire position shifts unexpectedly, with a similar issue in case (b). DragAnything also fails to produce a harmonious background; in case (c), the smoke doesn’t follow the moving train. In contrast, TrackGo generates videos where target region movement aligns precisely with user input while maintaining background consistency and harmony. This significantly enhances the visual quality and coherence of the videos, showcasing TrackGo’s effectiveness. Additional examples are shown in Fig. 2.




Figure 5: Comparison results of unspecified area suppression intensity τ . The top row shows the results of the last frame generated for various τ . The bottom row provides a magnified view, highlighting the differences more clearly with red boxes, to better observe the variations.

Attention Mask Control of Background Movement.

Our model possesses the capability to adjust the intensity of motion in unspecified areas through specific parameters, which is shown in the Fig. 5. We define the motor inhibition intensity in unspecified areas with the parameter τ . As shown in Fig. 5, when the hands move, it is necessary for the steering wheel to move correspondingly to enhance the realism of the output. When $\tau = -10^4$, the motion in unspecified areas is significantly suppressed, allowing only the hands to follow their trajectory while other parts remain static. This often results in distortion. When $\tau = -10^2$, the unspecified area is less suppressed, allowing hands and the steering wheel to move simultaneously. However, this setting can still produce disharmonious outcomes, such as the car logo on the steering wheel remaining static. When $\tau = 0$, the unspecified areas can move freely, which typically results in a more cohesive and harmonious video. Neverthe-

less, not all movements in the unspecified areas are desirable, and excessive movement can damage the video quality. Therefore, it is crucial to carefully manage the suppression level to balance realism with artistic control.

Train Step	14k	16k	18k
w/o Attn Mask and Loss	219.15	208.31	218.50
w/o Attn Loss	216.54	191.54	165.12
Full Method	204.02	184.03	136.11

Table 2: FVD tested on the validation set of different variants after 14k, 16k, and 18k training steps.

Ablation Study

To validate the effectiveness of the attention mask and attention loss, we report the FVD metrics on the internal validation set at various training steps, as shown in Table 2. Under the same number of training steps, the model without attention loss shows a slightly higher FVD compared to the model with attention loss. When attention loss is not utilized, the FVD is higher compared to when attention loss is applied. This discrepancy becomes particularly pronounced at the 18K training mark. This demonstrates that using attention loss can accelerate model training and aid convergence. Without both the attention mask and attention loss, the FVD stabilizes around 16K steps but remains significantly higher than the FVD under the full setting.

Conclusion

In this paper, we introduce point trajectories to capture complex temporal information in videos. We propose the Track-Adapter to process these point trajectories, focusing on the motion of specified targets, and employ an attention mask to mitigate the influence of original temporal self-attention on specified regions. During inference, the attention mask can regulate the movement of unspecified areas, resulting in video output that aligns more closely with user input. Extensive experiments demonstrate that our TrackGo achieves state-of-the-art FVD, FID, and ObjMC scores.

Acknowledgments

This work was supported by the National Science and Technology Major Project (No. 2022ZD0117800), Young Elite

Scientists Sponsorship Program by CAST, and the Fundamental Research Funds for the Central Universities.

References

- Betker, J.; Goh, G.; Jing, L.; Brooks, T.; Wang, J.; Li, L.; Ouyang, L.; Zhuang, J.; Lee, J.; Guo, Y.; et al. 2023. Improving image generation with better captions. *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2(3): 8.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendeleevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; et al. 2023a. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*.
- Blattmann, A.; Rombach, R.; Ling, H.; Dockhorn, T.; Kim, S. W.; Fidler, S.; and Kreis, K. 2023b. Align your latents: High-resolution video synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22563–22575.
- Ceylan, D.; Huang, C.-H. P.; and Mitra, N. J. 2023. Pix2video: Video editing using image diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23206–23217.
- Chen, J.; Ge, C.; Xie, E.; Wu, Y.; Yao, L.; Ren, X.; Wang, Z.; Luo, P.; Lu, H.; and Li, Z. 2024. PixArt- Σ : Weak-to-Strong Training of Diffusion Transformer for 4K Text-to-Image Generation. *arXiv preprint arXiv:2403.04692*.
- Cheng, H. K.; Oh, S. W.; Price, B.; Schwing, A.; and Lee, J.-Y. 2023. Tracking anything with decoupled video segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1316–1326.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Guo, Y.; Yang, C.; Rao, A.; Agrawala, M.; Lin, D.; and Dai, B. 2023a. Sparsctrl: Adding sparse controls to text-to-video diffusion models. *arXiv preprint arXiv:2311.16933*.
- Guo, Y.; Yang, C.; Rao, A.; Wang, Y.; Qiao, Y.; Lin, D.; and Dai, B. 2023b. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*.
- Henschel, R.; Khachatryan, L.; Hayrapetyan, D.; Poghosyan, H.; Tadevosyan, V.; Wang, Z.; Navasardyan, S.; and Shi, H. 2024. Streamingt2v: Consistent, dynamic, and extendable long video generation from text. *arXiv preprint arXiv:2403.14773*.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. volume 33, 6840–6851.
- Ho, J.; Salimans, T.; Gritsenko, A.; Chan, W.; Norouzi, M.; and Fleet, D. J. 2022. Video diffusion models. *Advances in Neural Information Processing Systems*, 35: 8633–8646.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Hu, L.; Gao, X.; Zhang, P.; Sun, K.; Zhang, B.; and Bo, L. 2023. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. *arXiv preprint arXiv:2311.17117*.
- Jiang, Y.; Wu, T.; Yang, S.; Si, C.; Lin, D.; Qiao, Y.; Loy, C. C.; and Liu, Z. 2023. VideoBooth: Diffusion-based Video Generation with Image Prompts. *arXiv preprint arXiv:2312.00777*.
- Karaev, N.; Rocco, I.; Graham, B.; Neverova, N.; Vedaldi, A.; and Rupprecht, C. 2023. Cotracker: It is better to track together. *arXiv preprint arXiv:2307.07635*.
- Karras, J.; Holynski, A.; Wang, T.-C.; and Kemelmacher-Shlizerman, I. 2023. Dreampose: Fashion image-to-video synthesis via stable diffusion. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 22623–22633. IEEE.
- Khachatryan, L.; Movsisyan, A.; Tadevosyan, V.; Henschel, R.; Wang, Z.; Navasardyan, S.; and Shi, H. 2023. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15954–15964.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Ma, W.-D. K.; Lewis, J.; and Kleijn, W. B. 2023. TrailBlazer: Trajectory Control for Diffusion-Based Video Generation. *arXiv preprint arXiv:2401.00896*.
- Ma, Y.; He, Y.; Wang, H.; Wang, A.; Qi, C.; Cai, C.; Li, X.; Li, Z.; Shum, H.-Y.; Liu, W.; et al. 2024. Follow-Your-Click: Open-domain Regional Image Animation via Short Prompts. *arXiv preprint arXiv:2403.08268*.
- MacQueen, J.; et al. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, 281–297. Oakland, CA, USA.
- Mou, C.; Wang, X.; Xie, L.; Wu, Y.; Zhang, J.; Qi, Z.; and Shan, Y. 2024. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4296–4304.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4195–4205.
- Qi, C.; Cun, X.; Zhang, Y.; Lei, C.; Wang, X.; Shan, Y.; and Chen, Q. 2023. Fatezero: Fusing attentions for zero-shot text-based video editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15932–15942.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.

- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Tu, S.; Dai, Q.; Cheng, Z.-Q.; Hu, H.; Han, X.; Wu, Z.; and Jiang, Y.-G. 2023. MotionEditor: Editing Video Motion via Content-Aware Diffusion. *arXiv preprint arXiv:2311.18830*.
- Unterthiner, T.; Van Steenkiste, S.; Kurach, K.; Marinier, R.; Michalski, M.; and Gelly, S. 2018. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*.
- Wang, J.; Yuan, H.; Chen, D.; Zhang, Y.; Wang, X.; and Zhang, S. 2023a. Modelscope text-to-video technical report. *arXiv preprint arXiv:2308.06571*.
- Wang, J.; Zhang, Y.; Zou, J.; Zeng, Y.; Wei, G.; Yuan, L.; and Li, H. 2024a. Boximator: Generating Rich and Controllable Motions for Video Synthesis. *arXiv preprint arXiv:2402.01566*.
- Wang, T.; Li, L.; Lin, K.; Zhai, Y.; Lin, C.-C.; Yang, Z.; Zhang, H.; Liu, Z.; and Wang, L. 2023b. Disco: Disentangled control for realistic human dance generation. *arXiv preprint arXiv:2307.00040*.
- Wang, X.; Yuan, H.; Zhang, S.; Chen, D.; Wang, J.; Zhang, Y.; Shen, Y.; Zhao, D.; and Zhou, J. 2024b. Videocomposer: Compositional video synthesis with motion controllability. *Advances in Neural Information Processing Systems*, 36.
- Wang, Z.; Yuan, Z.; Wang, X.; Chen, T.; Xia, M.; Luo, P.; and Shan, Y. 2023c. Motionctrl: A unified and flexible motion controller for video generation. *arXiv preprint arXiv:2312.03641*.
- Wu, W.; Li, Z.; Gu, Y.; Zhao, R.; He, Y.; Zhang, D. J.; Shou, M. Z.; Li, Y.; Gao, T.; and Zhang, D. 2024. DragAnything: Motion Control for Anything using Entity Representation.
- Xing, J.; Xia, M.; Zhang, Y.; Chen, H.; Wang, X.; Wong, T.-T.; and Shan, Y. 2023. Dynamicrafter: Animating open-domain images with video diffusion priors. *arXiv preprint arXiv:2310.12190*.
- Xu, Z.; Zhang, J.; Liew, J. H.; Yan, H.; Liu, J.-W.; Zhang, C.; Feng, J.; and Shou, M. Z. 2023. Magicanimate: Temporally consistent human image animation using diffusion model. *arXiv preprint arXiv:2311.16498*.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Yin, S.; Wu, C.; Liang, J.; Shi, J.; Li, H.; Ming, G.; and Duan, N. 2023. Dragnuwa: Fine-grained control in video generation by integrating text, image, and trajectory. *arXiv preprint arXiv:2308.08089*.
- Zeng, Y.; Wei, G.; Zheng, J.; Zou, J.; Wei, Y.; Zhang, Y.; and Li, H. 2023. Make pixels dance: High-dynamic video generation. *arXiv preprint arXiv:2311.10982*.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhang, S.; Wang, J.; Zhang, Y.; Zhao, K.; Yuan, H.; Qin, Z.; Wang, X.; Zhao, D.; and Zhou, J. 2023. I2vgen-xl: High-quality image-to-video synthesis via cascaded diffusion models. *arXiv preprint arXiv:2311.04145*.