

CustomTTT: Motion and Appearance Customized Video Generation via Test-Time Training

Xiuli Bi¹, Jian Lu¹, Bo Liu¹, Xiaodong Cun^{2*}, Yong Zhang³, Weisheng Li¹, Bin Xiao^{1,4}

¹ School of Computer Science and Technology, Chongqing University of Posts and Telecommunications, Chongqing, China,

² GVC Lab, Great Bay University, Dongguan, China,

³ Meituan, China

⁴ Jinan Inspur Data Technology Co., Ltd., Jinan, China

{bixl, boliu, liws, xiaobin}@cqupt.edu.cn, s230201078@stu.cqupt.edu.cn, cun@gbu.edu.cn, zhangyong201303@gmail.com

Abstract

Benefiting from large-scale pre-training of text-video pairs, current text-to-video (T2V) diffusion models can generate high-quality videos from the text description. Besides, given some reference images or videos, the parameter-efficient fine-tuning method, *i.e.* LoRA, can generate high-quality customized concepts, *e.g.*, the specific subjects or the motions. However, combining the trained multiple concepts from different references into a single network shows obvious artifacts. To this end, we propose CustomTTT, where we can joint custom the appearance and the motion of the given video easily. In detail, we first analyze the prompt influence in the current video diffusion model and find the LoRAs are only needed for the specific layers for appearance and motion customization. Besides, since each LoRA is trained individually, we propose a novel test-time training technique to update parameters after combination utilizing the trained customized models. We conduct detailed experiments to verify the effectiveness of the proposed methods. Our method outperforms several state-of-the-art works in both qualitative and quantitative evaluations.

1 Introduction

In the realm of text-to-video (T2V) diffusion models (Guo et al. 2023; Chen et al. 2023, 2024a; Sterling 2023), significant progress has been achieved through large-scale pre-training of text-video pairs, enabling the generation of high-quality videos from text descriptions. However, these pre-trained models present difficulties in generating specific objects or motions that are hard to describe from text prompts only.

To solve this problem, early model customization methods, *e.g.*, Dreambooth (Ruiz et al. 2023), LoRA (Hu et al. 2021), are proposed to produce customized concepts, such as specific subjects or motions, when providing reference images or videos. However, these methods focus only on the specific *one* concept. There are also some works for the multi-concept generation. However, they are mostly works for the two different subjects (Yuan et al. 2023; Kumari et al. 2023). For the customized appearance and motion, the pio-

neer works (Wei et al. 2024) still face the problem of appearance overfitting and low-fidelity.

We aim for multiple customization methods including the subject and the motion. Inspired by the current LoRA merged methods (Yadav et al. 2024; Yu et al. 2024), we can train different LoRAs and merge them for this purpose. Differently, in our task, we can naturally select different layers (*i.e.*, spatial and temporal layers) in video diffusion models for motion and appearance customization. However, the resulting video shows heavy artifacts. Thus, we should fine-tune the model when merging two LoRAs to reduce the conflicts during the merging process further.

Therefore, we present CustomTTT for joint motion and appearance customization. In detail, to find the most criteria layers for appearance and motion modeling in the VDM, we perform the text-embedding replacement experiment inspired by the text-to-image customization method (Frenkel et al. 2025). On the other hand, after training and interpolating two LoRA adapters for motion and appearance customization, we utilize the pre-trained single LoRA model to produce a suitable reference latent for distillation in a certain step. Thus, the multi-customized model can learn to guarantee the motion and appearance from their single customized reference. We conduct the experiments over various motions and the subject reference. Compared with state-of-the-art methods, our method shows much better visual quality and text-video alignment under the multi-customization settings.

Overall, the contribution can be summarized as:

- We analyze the prompt embedding importance in the T2V model for the first time and train LoRA adapters on the specific layers for better multiple-concept customization.
- We design a novel test-time training process to improve the text-video alignment in multiple concept customization.
- The experiments show the advantage of the proposed method over current methods.

* Corresponding Author.

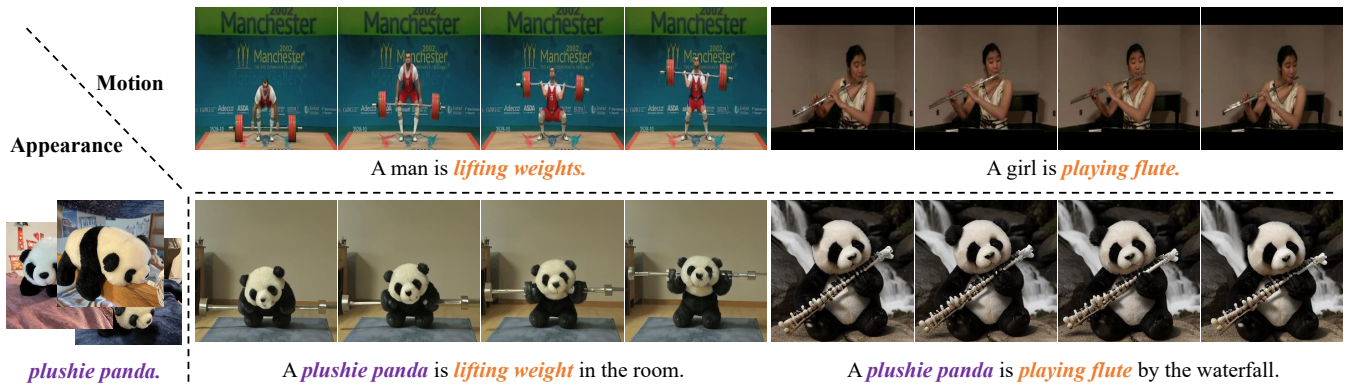


Figure 1: Given a *single* video for motion reference and *a few* images for appearance reference, our method can generate customized videos with multiple customized concepts in terms of the combinations of appearance and motion.

2 Related Work

2.1 Text-to-Video Diffusion Model

The Diffusion Model (Ho, Jain, and Abbeel 2020; Song, Meng, and Ermon 2021) has reinvented the research in video generation with the pre-trained text-to-image latent diffusion model (Rombach et al. 2022) as the basis. Different from previous methods which only can generate the video in the specific domain (*e.g.*, face (Siarohin et al. 2019; Zhang et al. 2023), body (Siarohin et al. 2021)), these text-to-video diffusion models (Guo et al. 2023; Chen et al. 2023, 2024a; Sterling 2023) are based on the pre-trained text-to-image model (Rombach et al. 2022) and train the temporal layers to model motion. After large-scale training, these methods can generate open-domain realistic videos from text prompts. However, the specific object or motion is hard to describe via text only. Thus, training the diffusion-based model for the specific appearance or motion is critical.

2.2 Model Customization

Fine-tuning the pre-trained large text-to-image/video models (Rombach et al. 2022; Chen et al. 2023, 2024a) to generate the customized image (or video) is naturally. Since the specific object/motion can not be described easily in text. Early works (Ruiz et al. 2023; Hu et al. 2021; Gal et al. 2022) focus on single subject customization via additional training, as well as multiple subjects (Yuan et al. 2023; Kumari et al. 2023) and also for video (Jiang et al. 2024). Recent works are to customize the motion of the video by decoupling the training of the spatial and temporal (Zhao et al. 2025; Wang et al. 2024; Ren et al. 2024). The most related work to our method is DreamVideo (Wei et al. 2024), they train additional adapters for motion and appearance customization. However, the performance is still restricted by the appearance overfit or the motion correctness.

2.3 Test-time Training

Test-time training (TTT) aims to improve the performance of the specific sample when testing without the label. Previous works have proved its efficiency on image classification (Wang et al. 2020; Gandelsman et al. 2022; Sun et al.

2020; Schneider et al. 2020), video object segmentation (Liu et al. 2024a; Wang et al. 2023). The core of the TTT is to design a self-supervised proxy function, *e.g.*, the rotation (Sun et al. 2020), the masked image reconstruction (Gandelsman et al. 2022), so as the pre-trained knowledge can be updated via this proxy gradient. In our multiple customizations of appearance and motion, we use this method in diffusion-based customization to update the parameters after the LoRA combination. We construct the relationship between the multiple and the single customization model.

3 Method

We aim to generate a highly customized video from text, where the motion and appearance are from *different* given references. These references may include a video for motion guidance and 3-5 images for appearance guidance. To achieve this goal, we utilize the pre-trained Text-to-Video Diffusion Models (Guo et al. 2023) as the base model and train the LoRAs (Hu et al. 2021) for motion and appearance customization, respectively. However, directly incorporating individually trained temporal and spatial LoRAs into a single pre-trained model for multiple customizations will typically fail. To this end, we first perform an in-depth prompt importance analysis in the pre-trained video diffusion model and show that LoRAs should be added in the correct layers only to influence the motion and appearance. To further improve the combined results, we introduce a novel test-time training phase when interpolating two LoRAs. Below, we first introduce the preliminaries in the video diffusion model and LoRA as the background knowledge in Sec. 3.1. Then, we give the details of LoRA analysis and the LoRA customization in Sec. 3.2 and Sec. 3.3, respectively. Finally, we give the details of the proposed test-time training process after LoRA combinations in Sec. 3.4.

3.1 Preliminary

Text-to-Video Diffusion Models Text-to-video diffusion models aim to generate realistic videos from text prompts. Based on the text-to-image latent diffusion model, *i.e.*, stable diffusion model (Rombach et al. 2022), these methods




Figure 2: The overall pipeline. We first train the LoRAs on the specific layers for appearance (a) and motion (b) customization individually. Then, we design a test-time training method to further improve the results when combining.

add the Gaussian noise in the encoded video, *i.e.* latent, in training, where a denoising UNet is used to predict the added step noise. Given a text condition c , the parameters θ of the denoising U-Net ϵ_θ are optimized using the following reconstruction loss:

$$E_{z_0, c, \epsilon \sim \mathcal{N}(0, I), t \sim \mathcal{U}(0, T)} [\|\epsilon - \epsilon_\theta(z_t, t, \tau_c)\|_2^2], \quad (1)$$

where z_0 represents the latent of the encoded training video, τ_c refers to the pre-trained text encoder τ with prompt c , ϵ is the Gaussian noise added to the latent code, and t denotes the time step. The latent variable z_t at time t can be expressed as a function of the initial latent variable z_0 , the noise ϵ , and the time step t . Typically, this process is formulated as:

$$z_t = \sqrt{\alpha_t} z_0 + \sqrt{1 - \alpha_t} \epsilon, \quad (2)$$

where α_t is a time-dependent scaling factor. After training, this diffusion model can generate the video from text and noise with a T step progressive sampling process.

LoRA Low-Rank Adaptation (LoRA, (Hu et al. 2021)) is an effective technique for adapting large pre-trained models to down-streaming tasks with few training-able parameters. To achieve this goal, LoRA introduces a low-rank decomposition-based method to the model’s weight matrix, enabling efficient adaptation to new tasks while maintaining the model’s original capabilities.

Given the weight matrix $W_0 \in \mathbb{R}^{m \times n}$ of the original pre-trained model, LoRA uses two low-rank matrices

$A \in \mathbb{R}^{m \times r}$ ($r \ll m$) and $B \in \mathbb{R}^{r \times n}$ ($r \ll n$) to shift the trained distribution according to the new data training. Thanks to the low-rank matrix in A and B , LoRA updates the model more efficiently than the full rank matrix and shows comparable results with full training. Specifically, the new weight matrix W can be represented as:

$$W = W_0 + \Delta W = W_0 + AB. \quad (3)$$

3.2 Prompt Influence Analysis in VDM

As introduced previously, directly combining the trained spatial and temporal LoRA results in low-fidelity videos. To address this problem, instead of adding LoRAs in each layer, we aim to find *the most crucial layers* for appearance (or motion) modeling so that we can customize different layers to decouple training and reduce the risk of overfitting.

Before introducing our analysis, we first show the overview architecture and some definitions of the current T2V diffusion model to clarify. As shown in Fig. 3, take AnimateDiff (Guo et al. 2023) as an example, the denoising UNet often comprises four down-sampling modules, one middle module, and four up-sampling modules. Each layer consists of a spatial module and a temporal module, respectively. The spatial module contains several res-blocks (He et al. 2016), self-attention (Vaswani et al. 2017) and cross-attention layers while every motion module contains temporal attention layers to model the temporal information. To

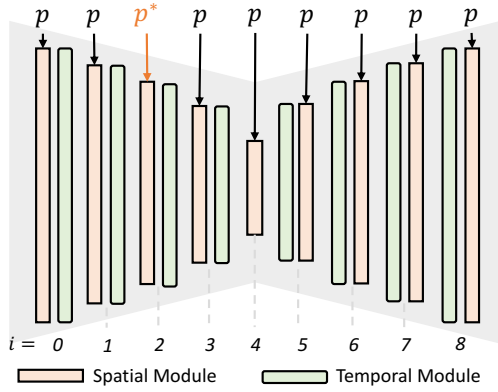


Figure 3: Examines the influence of the i -th layer on the appearance and motion in video generation. The text prompt p^* is injected into the i -th layer, while the text prompt p is injected into all other layers.




Figure 4: The effect of prompt injection on appearance. Injecting p^* into both $i = 2, 6$ shows comparable results with p^* all injections.

generate a video, a single text prompt p is embedded into all the cross-attention layers of the pre-trained spatial module. Considering each spatial module as a whole layer, we can mark the video generation process as $V_{p \rightarrow i \in [0, \dots, 8]}$, where i is the index of the layers using p .

Inspired by the customized text-to-image method, *i.e.*, B-LoRA (Frenkel et al. 2025), to find the importance of the prompt embedding in each layer, we analyze by replacing one of the text embeddings in all the pre-trained video diffusion model with the new one and keep other text embeddings unchanged. So, we can observe the relationships between the given hybrid text prompts and the generated video. Formally, considering the original prompt p and the new prompt p^* , we can generate the video via $V_{p^* \rightarrow I, p \rightarrow J}$, where I and J are the module indexes of the prompt p^* and p , respectively.

As shown in Fig. 4, for appearance, when we only insert p^* in the layer 6, $V_{p^* \rightarrow i=6, p \rightarrow j \neq i}$, the video predominantly depicts a tiger instead of a koala. However, if we modify the embedding in other locations, the results might not show the correct changes. Besides, if we change the prompt in layer $i = 6$ only, the tiger’s behavior still resembles that of a koala. Consequently, we undertake an additional round of the experiment by incorporating a new text embedding replacement within the remaining embeddings. As shown in Fig. 4, by injecting the new prompt p^* at both $i = 2, 6$, while keeping p unchanged for other layers, the resulting video ac-

$p = \text{“A dog is swimming”}$ $p^* = \text{“A dog is running”}$

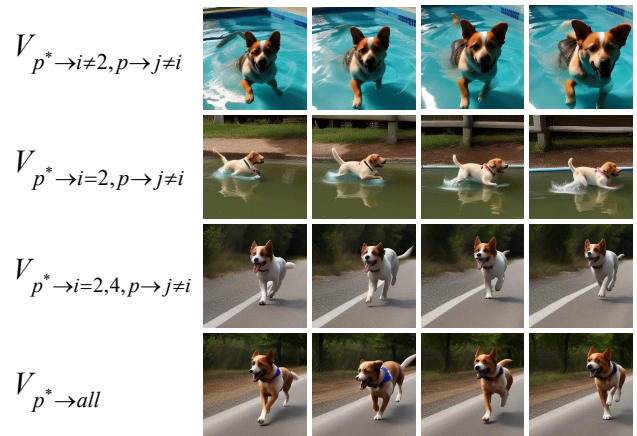


Figure 5: The effect of prompt injection on motion. Injecting p^* into both $i = 2, 4$ can remove the influence of prompt p .

curately depicts tiger characteristics without koala interference. This outcome is comparable to injecting “a tiger” into all layers, indicating that $i = 2, 6$ are the most important layers for controlling appearance.

Since motion is critical for video, we also conduct a similar experiment replacing the motion-related words in the prompt. *e.g.*, in Fig. 5, we replace one of the text embeddings of the text prompt p (“a dog is swimming”) with p^* (“a dog is running”). Similar to appearance analysis, initially, injecting p^* into a single layer reveals that $i = 2$ could generate the action described by p^* as shown in the second row of the Fig. 5. However, using only a single p^* injection still results in motion influenced by p as the dog is running in water, which corresponds to p ’s description of “a dog is swimming”. This indicates that other layers also significantly impact motion. Building on $(V_{p^* \rightarrow i=2, p \rightarrow j \neq i})$, we add p^* to other layers and discovered that injecting p^* into $i = 4$ resulted in a realistic running motion. Therefore, we conclude that the most critical layers for motion are $i = 2$ and $i = 4$.

3.3 Train LoRAs Individually for Customization

As shown in the Fig. 2 (a), (b), based on the above observations, we train the LoRA weight $\Delta W_s^{2,6}$ on the spatial modules’ 2, 6-th layers for appearance customization. As for the motion customization, current video diffusion models do not have a temporal layer at $i = 4$ and do not have cross-attention at $i = 3$ and $i = 5$, making it impossible to determine their importance through the injection of prompts p and p^* . Thus, we train the LoRAs $\Delta W_t^{2,5}$ on the neighbor temporal layers at $i = 2, 5$ for motion customization. All the experiments are done using the basic training method as introduced in Sec 3.1. Since spatial and temporal LoRAs are placed in different layers, this combination more effectively mitigates conflict in merging weights.

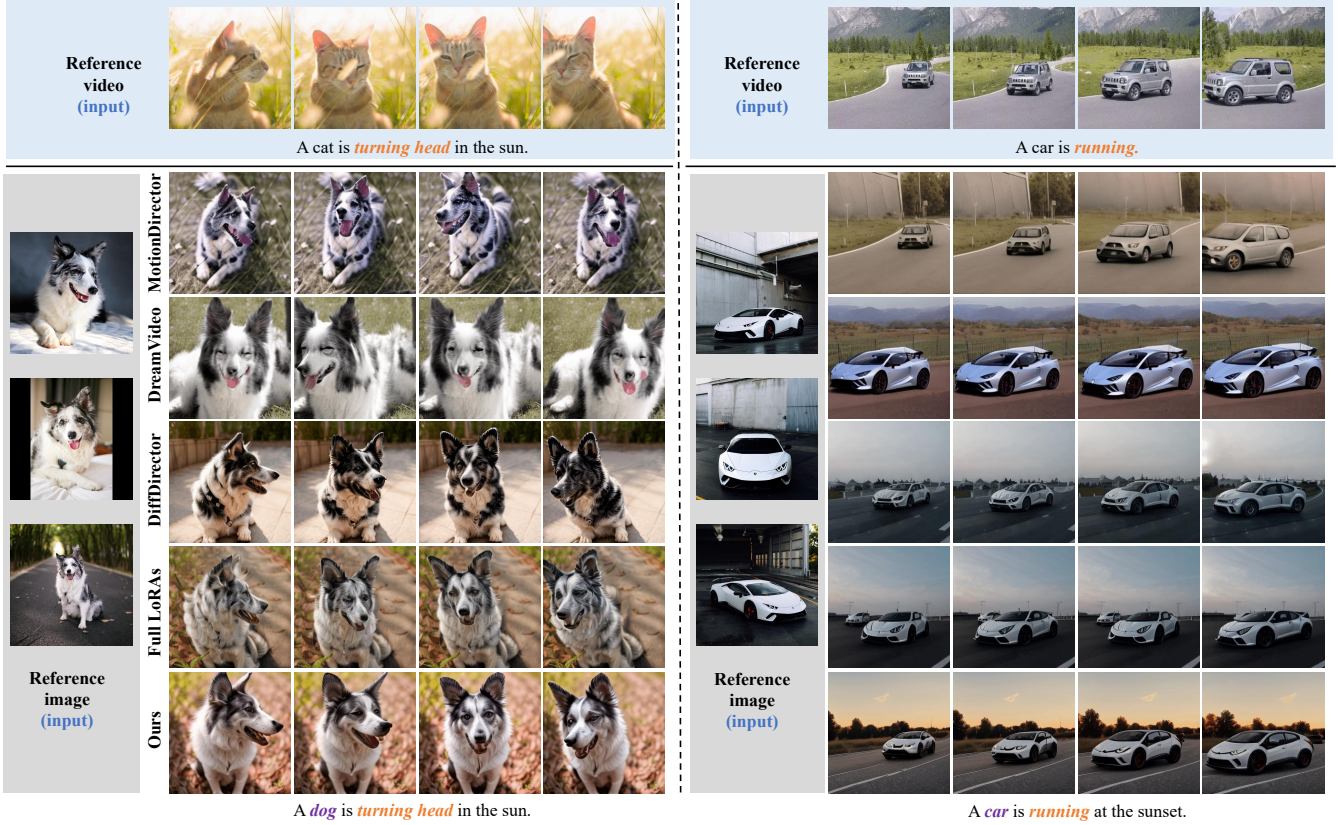


Figure 6: Visual comparison with other state-of-the-art methods.

3.4 Test-time Training for LoRA Combination

Directly combining our individually learned LoRA $\Delta W_t^{2,5}$ and $\Delta W_s^{2,6}$ into a single base model already shows better performance than the fully added ones. To further improve the text-video alignment, we propose a novel test-time training process. This is achieved by distilling the knowledge from the single LoRA model as the teacher model.

As shown in Fig. 2 (c), formally, given the pre-trained T2V diffusion model ϵ_θ , trained LoRA weights of $\Delta W_s^{2,6}$ and $\Delta W_t^{2,5}$ individually, we first inference f step using DDIM (Song, Meng, and Ermon 2021) to generate a reference latent z_f^s via $DDIM(x, \epsilon(c'; \Delta W_s^{2,6}))$, where $\epsilon(\cdot; \Delta W_s^{2,6})$ is the solo appearance customized model and c' are the related prompt to increase the diversity of the reference model. Then, we consider it as the target and add the Gaussian noise ϵ' to z_f^s to obtain z_{f+1}^s as the input for finetune. We call it appearance preservation loss \mathcal{L}_{ap} since it measures the difference between the multi-customized model, which can be written as follows:

$$\mathcal{L}_{ap} = \|\epsilon_\theta(z_{f+1}^s, \tau_c; \Delta W_s^{2,6}, \Delta W_t^{2,5}) - \epsilon'\|^2, \quad (4)$$

where $\Delta W_s^{2,6}$ and $\Delta W_t^{2,5}$ are the training-able copy of $\Delta W_s^{2,6}$ and $\Delta W_t^{2,5}$, respectively.

Meanwhile, we design a temporal preservation loss to have a similar temporal motion as the videos of the original model. Inspired by the temporal debiased loss (Zhao

et al. 2025), for the latent ϵ_i in each frame i , we calculate the added noise of each latent as:

$$\phi(\epsilon_i) = \sqrt{\beta^2 + 1}\epsilon_i - \beta_{anchor}, \quad (5)$$

where β_{anchor} is the selected frame as the anchor frame to remove the influence of the appearance. Similar to appearance finetune, we generate a reference motion latent of z_f^t via $DDIM(f, \epsilon(c''; \Delta W_t^{2,5}))$. We then calculate the temporal preservation loss as follows:

$$\mathcal{L}_{tp} = \|\phi(\epsilon'') - \phi(\epsilon_\theta(z_{f+1}^t, \tau_c; \Delta W_s^{2,6}, \Delta W_t^{2,5}))\|^2, \quad (6)$$

where ϵ'' is the added noise for training.

Finally, the motion LoRAs $\Delta W_t^{2,5}$ and appearance LoRAs $\Delta W_s^{2,6}$ will be jointly trained via optimizing \mathcal{L}_{ap} and \mathcal{L}_{tp} , iteratively. After training, they can be directly used for inference.

4 Experiments

4.1 Settings and Implement Details

Dataset Following previous methods (Kumari et al. 2023; Wei et al. 2024), for appearance customization, we collect a total of 13 objects from Dreambooth (Ruiz et al. 2023) and CustomDiffusion (Kumari et al. 2023), including pets, vehicles, toys, *etc.* For motion customization, we source 18 sets of videos from the UCF101 (Soomro, Zamir, and Shah

Method	Train-able parameters	Objective evaluation			User study			
		CLIP-T	CLIP-I	Temporal consistency	Motion similarity	Appearance similarity	Prompt alignment	Video quality
Full LoRA	28.26M	0.294	0.687	0.977	3.791	3.618	3.945	3.627
DreamVideo	85.00M	0.271	0.681	0.969	2.758	3.518	2.718	2.891
MotionDirector	21.26M	0.269	0.690	0.965	3.164	3.264	2.836	3.055
DiffDirector	30.46M	0.287	0.685	0.971	3.773	3.491	3.727	3.427
Ours	12.12M	0.301	0.712	0.978	3.873	4.045	4.102	3.954

Table 1: Quantitative experimental results for different methods under the numerical evaluation metrics.

2012), the UCF (Soomro and Zamir 2015) Sports Action datasets, and the DAVIS dataset (Pont-Tuset et al. 2017). Each set includes 5 prompts for appearance and motion, covering various scenarios.

Implementation details Our approach uses AnimatedDiff (Guo et al. 2023) as the base T2V model, trained with the LION optimizer (Chen et al. 2024b). The learning rate for the spatial LoRA is set at $1e-5$, while the temporal LoRA is trained with a learning rate of $5e-5$. Both types of LoRA are trained for 500 steps with a rank set to 32. In the test-time training phase, we set the learning rate to $1e-6$ and train for 30 steps. During inference, we employ DDIM (Song, Meng, and Ermon 2021) sampling with 25 steps and a classifier-free guidance (Ho and Salimans 2022) scale of 9. We generate 16-frame videos at a resolution of 256×256 and 8 fps. All experiments are performed on a single A6000 GPU.

4.2 Comparison with Other SOTA Methods

To demonstrate the effectiveness of our method in customizing video appearance and motion, we compare it with several current open-source methods. (1) **DreamVideo** (Wei et al. 2024) achieves appearance and motion customization by separately adding an ID adapter and a Motion Adapter. (2) **MotionDirector** (Zhao et al. 2025) is based on the ZeroScope (Sterling 2023) T2V model, with a dual-path LoRAs architecture that decouples the learning of appearance and motion. (3) **DiffDirector** (MotionDirector-Animatediff) (Zhao et al. 2025) replaces the original ZeroScope in MotionDirector with the pre-trained Animatediff for a fair comparison. (4) **Full LoRA** fine-tunes *all* the spatial and temporal attention blocks using LoRA based on the pre-trained Animatediff individually.

Visual Comparison As shown in Fig. 6, given reference video and images for motion and appearance customization, directly training using full LoRAs shows the visual artifacts. MotionDirector and DiffDirector use the same method with different backbones, they perform well in motion customization. However, they have a very different appearance from the reference images. Besides, DreamVideo seems to overfit the given reference, which has weak connections with the text prompt. Differently, the proposed method achieves a similar behavior to the reference video and also shows high-quality subject customization and text-video alignment.




Figure 7: We finetune LoRAs in the specific layers as discussed in Sec. 3.2. The selected layers show the best performance in terms of motion and appearance customization.

Numerical Comparison We employ three metrics to evaluate our method following previous works (Liu et al. 2024b; Wei et al. 2024). (1) **CLIP-T** measures the alignment between the text and the video by encoding both into embeddings using a clip encoder and calculating their average cosine similarity. (2) **CLIP-I** assesses the visual alignment between the generated video and reference images by computing the average similarity of their respective embeddings. (3) **Temporal Consistency** evaluates the continuity between consecutive video frames by determining the average cosine similarity of their clip embeddings. As shown in Tab. 1, the proposed method shows the best performance in terms of visual quality, temporal consistency, and the alignment between the reference image and video. Besides, since we only train the LoRAs on the specific layers, our method has fewer trainable parameters than others.

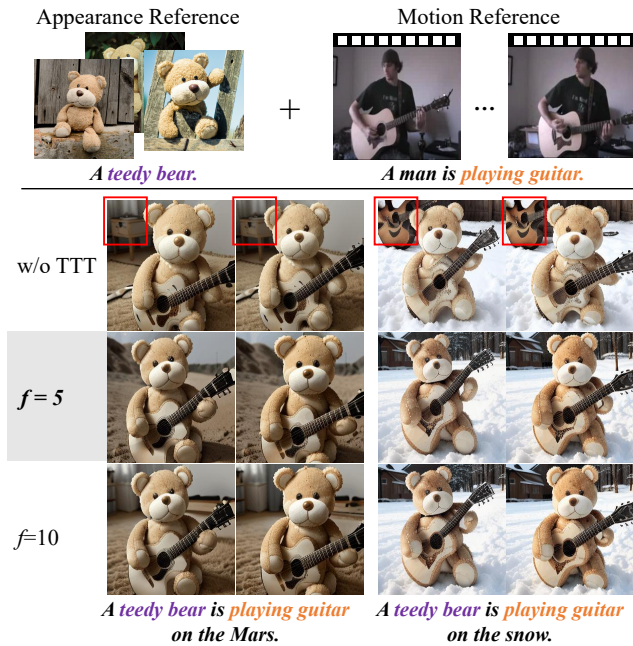


Figure 8: The influence of the sampling steps of the reference latent in our test-time training stage, using 5 steps of DDIM sampling improves the text-video alignment.

User Studies We also conduct user studies to show the effectiveness of the proposed methods. In detail, we invert 11 participants to vote for 10 generated videos in terms of motion similarity, appearance similarity, prompt alignment, and video quality. Each video will be ranked from different aspects, scoring from 1 to 5. Overall, we get 440 opinions. As shown in Table. 1, our methods achieve the best score over other state-of-the-art methods.

4.3 Ablation Studies

The effectiveness of training LoRA in crucial layers. In Sec. 3.2, we have analyzed the effectiveness of the prompt replacement. Here, we show that training LoRAs on the produced layers improves the performance. As shown in Fig. 7, training LoRA in other location, e.g., $\Delta W_s^{0,1}$ and $\Delta W_t^{1,7}$ can not customize the appearance and the motion. Fine-tuning the LoRAs on the $\Delta W_t^{2,5}$ or $\Delta W_s^{2,6}$ solely performs motion or appearance customization. Ideally, utilizing the $\Delta W_s^{2,6}$ and $\Delta W_t^{2,5}$ show the better performance on the multiple customization.

The effectiveness of reference sampling steps f . We validate the influence of our reference latent. As shown in Fig. 8, directly combining the LoRAs w/o TTT shows wrong backgrounds. However, training the model using reference latents will improve the text-video alignment. We evaluate the steps f of the reference latent, where it performs the best at $f = 5$ using DDIM sampling.

The effectiveness of test-time training steps. The number of training steps for TTT significantly affects the results, as shown in Fig. 9. We only require 30 steps to modify a



Figure 9: Ablation study results on test-time training across different numbers of training steps.

stone pillar in the video to a tree, making it consistent with the prompt description. However, excessive training can lead to artifacts and unnatural effects in the video.

5 Conclusion

We present a new method to achieve both motion and appearance customization in a single network using LoRAs. To handle the problem of LoRA merges, we first give a detailed analysis to show the most critical layers for video customization in terms of motion and appearance. Then, we design a novel test-time training process that utilizes the pre-trained single LoRA model as guided to further improve the text-video alignment. Based on the proposed methods, our method achieves state-of-the-art appearance and motion customization performance. We believe our findings about the video diffusion model will also benefit the training and the designing of the text-to-video diffusion model.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62172067 and Grant 62376046, the Natural Science Foundation of Chongqing for Distinguished Young Scholars under Grant CSTB2022NSCQ-JQX0001, in part by the Natural Science Foundation of Chongqing under Grant CSTB2023NSCQ-MSX0341, the Science and Technology Research Program of Chongqing Municipal Education Commission (Grant No. KJQN202200635)

References

- Chen, H.; Xia, M.; He, Y.; Zhang, Y.; Cun, X.; Yang, S.; Xing, J.; Liu, Y.; Chen, Q.; Wang, X.; et al. 2023. Videocrafter1: Open diffusion models for high-quality video generation. *arXiv preprint arXiv:2310.19512*.
- Chen, H.; Zhang, Y.; Cun, X.; Xia, M.; Wang, X.; Weng, C.; and Shan, Y. 2024a. Videocrafter2: Overcoming data limitations for high-quality video diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7310–7320.
- Chen, X.; Liang, C.; Huang, D.; Real, E.; Wang, K.; Pham, H.; Dong, X.; Luong, T.; Hsieh, C.-J.; Lu, Y.; et al. 2024b. Symbolic discovery of optimization algorithms. *Advances in neural information processing systems*, 36.
- Frenkel, Y.; Vinker, Y.; Shamir, A.; and Cohen-Or, D. 2025. Implicit style-content separation using b-lora. In *European Conference on Computer Vision*, 181–198. Springer.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*.
- Gandelsman, Y.; Sun, Y.; Chen, X.; and Efros, A. 2022. Test-Time Training with Masked Autoencoders. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 29374–29385. Curran Associates, Inc.
- Guo, Y.; Yang, C.; Rao, A.; Liang, Z.; Wang, Y.; Qiao, Y.; Agrawala, M.; Lin, D.; and Dai, B. 2023. Animatediff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv preprint arXiv:2307.04725*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- 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*.
- Jiang, Y.; Wu, T.; Yang, S.; Si, C.; Lin, D.; Qiao, Y.; Loy, C. C.; and Liu, Z. 2024. Videobooth: Diffusion-based video generation with image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6689–6700.
- Kumari, N.; Zhang, B.; Zhang, R.; Shechtman, E.; and Zhu, J.-Y. 2023. Multi-concept customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1931–1941.
- Liu, W.; Shen, X.; Li, H.; Bi, X.; Liu, B.; Pun, C.-M.; and Cun, X. 2024a. Depth-aware Test-Time Training for Zero-shot Video Object Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19218–19227.
- Liu, Y.; Cun, X.; Liu, X.; Wang, X.; Zhang, Y.; Chen, H.; Liu, Y.; Zeng, T.; Chan, R.; and Shan, Y. 2024b. Evalcrafter: Benchmarking and evaluating large video generation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22139–22149.
- Pont-Tuset, J.; Perazzi, F.; Caelles, S.; Arbeláez, P.; Sorkine-Hornung, A.; and Van Gool, L. 2017. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*.
- Ren, Y.; Zhou, Y.; Yang, J.; Shi, J.; Liu, D.; Liu, F.; Kwon, M.; and Shrivastava, A. 2024. Customize-a-video: One-shot motion customization of text-to-video diffusion models. *arXiv preprint arXiv:2402.14780*.
- 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.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22500–22510.
- Schneider, S.; Rusak, E.; Eck, L.; Bringmann, O.; Brendel, W.; and Bethge, M. 2020. Improving robustness against common corruptions by covariate shift adaptation. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 11539–11551. Curran Associates, Inc.
- Siarohin, A.; Lathuilière, S.; Tulyakov, S.; Ricci, E.; and Sebe, N. 2019. First Order Motion Model for Image Animation. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Siarohin, A.; Woodford, O.; Ren, J.; Chai, M.; and Tulyakov, S. 2021. Motion Representations for Articulated Animation. In *CVPR*.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations*.
- Soomro, K.; and Zamir, A. R. 2015. Action recognition in realistic sports videos. In *Computer vision in sports*, 181–208. Springer.
- Soomro, K.; Zamir, A. R.; and Shah, M. 2012. UCF101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Sterling, S. 2023. Zeroscope. <https://huggingface.co/cerspense/zeroscope.v2.576w/>. Accessed: 2024-08-10.
- Sun, Y.; Wang, X.; Liu, Z.; Miller, J.; Efros, A.; and Hardt, M. 2020. Test-Time Training with Self-Supervision for Generalization under Distribution Shifts. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 9229–9248. PMLR.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, D.; Shelhamer, E.; Liu, S.; Olshausen, B.; and Darrell, T. 2020. Tent: Fully test-time adaptation by entropy minimization. *arXiv preprint arXiv:2006.10726*.

Wang, L.; Shen, G.; Liang, Y.; Tao, X.; Wan, P.; Zhang, D.; Li, Y.; and Chen, Y. 2024. Motion Inversion for Video Customization. *arXiv preprint arXiv:2403.20193*.

Wang, R.; Sun, Y.; Gandelsman, Y.; Chen, X.; Efros, A. A.; and Wang, X. 2023. Test-time training on video streams. *arXiv preprint arXiv:2307.05014*.

Wei, Y.; Zhang, S.; Qing, Z.; Yuan, H.; Liu, Z.; Liu, Y.; Zhang, Y.; Zhou, J.; and Shan, H. 2024. Dreamvideo: Composing your dream videos with customized subject and motion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6537–6549.

Yadav, P.; Tam, D.; Choshen, L.; Raffel, C. A.; and Bansal, M. 2024. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36.

Yu, L.; Yu, B.; Yu, H.; Huang, F.; and Li, Y. 2024. Language Models are Super Mario: Absorbing Abilities from Homologous Models as a Free Lunch. In *International Conference on Machine Learning*. PMLR.

Yuan, G.; Cun, X.; Zhang, Y.; Li, M.; Qi, C.; Wang, X.; Shan, Y.; and Zheng, H. 2023. Inserting Anybody in Diffusion Models via Celeb Basis. *arXiv preprint arXiv:2306.00926*.

Zhang, W.; Cun, X.; Wang, X.; Zhang, Y.; Shen, X.; Guo, Y.; Shan, Y.; and Wang, F. 2023. Sadtalker: Learning realistic 3d motion coefficients for stylized audio-driven single image talking face animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8652–8661.

Zhao, R.; Gu, Y.; Wu, J. Z.; Zhang, D. J.; Liu, J.-W.; Wu, W.; Keppo, J.; and Shou, M. Z. 2025. Motiondirector: Motion customization of text-to-video diffusion models. In *European Conference on Computer Vision*, 273–290. Springer.