
NaRCan: Natural Refined Canonical Image with Integration of Diffusion Prior for Video Editing

Ting-Hsuan Chen Jiewen Chan Hau-Shiang Shiu
Shih-Han Yen Chang-Han Yeh Yu-Lun Liu
National Yang Ming Chiao Tung University

Abstract

We propose a video editing framework, NaRCan, which integrates a hybrid deformation field and diffusion prior to generate high-quality natural canonical images to represent the input video. Our approach utilizes homography to model global motion and employs multi-layer perceptrons (MLPs) to capture local residual deformations, enhancing the model’s ability to handle complex video dynamics. By introducing a diffusion prior from the early stages of training, our model ensures that the generated images retain a high-quality natural appearance, making the produced canonical images suitable for various downstream tasks in video editing, a capability not achieved by current canonical-based methods. Furthermore, we incorporate low-rank adaptation (LoRA) fine-tuning and introduce a noise and diffusion prior update scheduling technique that accelerates the training process by 14 times. Extensive experimental results show that our method outperforms existing approaches in various video editing tasks and produces coherent and high-quality edited video sequences. Code and video results are available at koi953215.github.io/NaRCan_page.

1 Introduction

Video editing has always been a fascinating research area. For example, style transfer transforms the original video into a completely new style, enriching the viewing experience. Other tasks include dynamic segmentation and handwriting, which all demonstrate the broad application value of video editing across various fields. Currently, diffusion model technology is becoming increasingly mature and is known for its powerful generative capabilities and frequent use in video editing. However, in video-to-video tasks, maintaining temporal consistency presents a significant challenge, particularly when applying image-based diffusion models to video editing tasks. When these models, originally designed for image generation, are applied frame-by-frame to videos, they often produce temporally inconsistent results due to their frame-independent processing nature. This limitation has motivated numerous research efforts to enhance temporal coherence through various techniques, such as optical flow guidance, latent space alignment, and cross-attention mechanisms, aiming to produce high-quality video sequences while preserving temporal consistency across frames.

Yet, even with solutions addressing temporal consistency, diffusion-based methods [16, 4, 19, 8] still encounter significant challenges in precise localized editing tasks. This is where canonical-based methods demonstrate their unique advantages. By consolidating video content into a single representative image, canonical-based methods [3, 1, 5, 45] enable intuitive and precise spatial control over editing operations while inherently maintaining temporal consistency. This unified representation allows direct application of image-based editing techniques while ensuring that modifications propagate coherently throughout the video sequence, making these methods particularly effective for a wide range of video editing tasks.

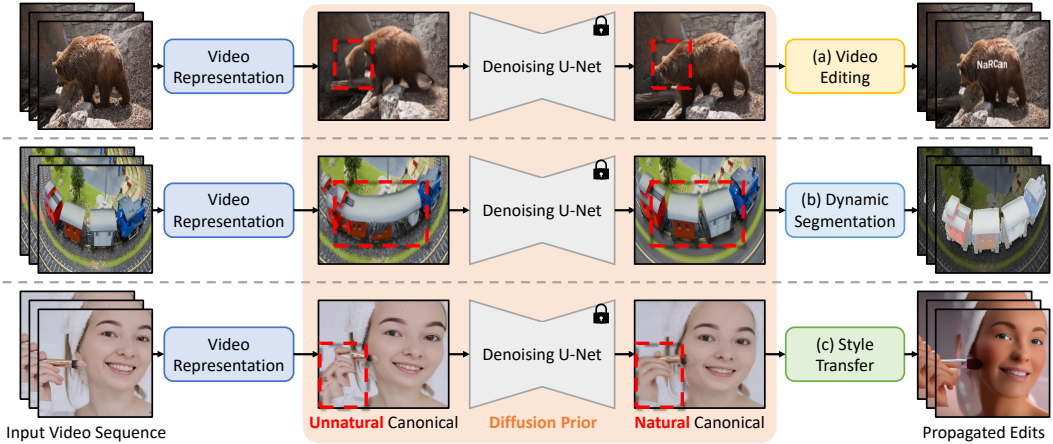


Figure 1: **Video representation with diffusion prior.** Given an RGB video, we can represent the video using a canonical image. However, the canonical image and reconstruction training process focuses only on reconstruction quality and could produce an unnatural canonical image. This could cause problems with downstream tasks such as prompt-based video editing. In the bottom example, if the hand is distorted in the canonical image, the image editor, such as ControlNet [75], may not recognize it and could introduce an irrelevant object instead. In this paper, we propose introducing the *diffusion prior* from a LoRA [18] fine-tuned diffusion model to the training pipeline and constraining the canonical image to be natural. Our method facilitates several downstream tasks, such as (a) video editing, (b) dynamic segmentation, and (c) video style transfer.

While previous canonical-based approaches have established the fundamental utility of canonical images in video editing, our observations suggest that enhancing the naturalness of these representations can significantly improve their effectiveness in various editing scenarios. We observe that existing methods primarily focus on reconstruction quality without explicit constraints to ensure natural canonical image generation. To address this opportunity for enhancement, we propose NaRCan, a novel hybrid deformation field network architecture that incorporates diffusion priors (Figure 1) into the training pipeline (Figure 2). This integration enables our method to generate high-fidelity natural canonical images across diverse scenarios while maintaining temporal consistency. Through comprehensive experimental evaluation, we demonstrate that NaRCan achieves superior performance compared to state-of-the-art video editing methods in both qualitative and quantitative metrics.

Our main contributions are three-fold:

- We have designed a novel deformation field structure to represent object variations throughout an entire scene. Compared to other canonical-based models, our model demonstrates superior expressive capability and faster convergence speed.
- We have effectively integrated the diffusion prior into our pipeline, enabling our method to generate high-quality natural canonical images in any scenario. Additionally, we designed a dynamic scheduling method that significantly accelerates the entire training process.
- We thoroughly evaluate our method to show the state-of-the-art video editing performance.

2 Related Work

Implicit neural representation. Implicit neural representation [43] using coordinate-based MLP is an outstanding way to represent a video, capable of obtaining a canonical image to represent the entire video [45]. Recent methods employ hash grid encoding [42] or positional encoding [40] combined with MLP. The approach in [5] is more effective at handling spatial information, but the resulting canonical image exhibits severe distortion and warping with in-the-wild videos [11]. Therefore, we propose a hybrid deformation field method composed of homography [9] and residual deformation MLP. This model design fits the deformation information in videos better than existing methods.

Consistent video editing. There are generally three approaches to video editing: (1) propagation-based, (2) layered representation-based, and (3) canonical-based. The first approach, propagation-

based, focuses on propagating information across frames [35, 20, 22, 23, 57, 63, 67]. This method can easily produce inaccurate results due to occlusions and error propagation. The second approach, layered representation-based, separates a video into foreground and background, obtaining canonical images for both [36, 25, 73, 34, 37, 38]. We can edit the entire video by editing these canonical images and then synthesizing the video afterward. However, this method heavily relies on masks. If the mask-RCNN captures incorrect targets in the preprocessing stage, the fitted foreground can be incorrect, especially in scenes with large camera movement.

The third approach, canonical-based, typically uses MLP to obtain the deformation information of each pixel to form a canonical image [45]. Transferring the video to canonical space maintains temporal consistency while editing and supports various downstream tasks, such as super-resolution and segmentation. CoDeF [45] uses this approach, but canonical images can deform severely in videos with significant camera or object movement. CoDeF suggests that using a group model could resolve these issues. However, using group CoDeF requires masks for training data obtained from SAM-track [12]. Incorrect masks for training data can result in an unnatural canonical image of the video. Even with correct masks for the foreground object, the canonical image might still be corrupted, rendering these images ineffective for video editing.

Video processing via generative models. Some works utilize GAN inversion [65, 78, 60, 74, 48, 26, 69] to edit images or videos. Today, numerous generative models exist for editing images. Some methods, such as GLIDE [44], DALL-E [55, 54], stable diffusion [56], and Imagen [58], are trained on millions of images, resulting in incredible generative abilities. Other methods like SDEdit [39], ControlNet [75], and LDM [2] use conditions to achieve better editing results. Instruction-based video editing methods like InstructPix2Pix [3] and another work [10] often yield sub-optimal results for different editing operations. Techniques like LoRA [18] can assist in fine-tuning to find better weights for editing. Additionally, many zero-shot diffusion-based methods [72, 16, 76, 4, 8, 19, 27] do not require model training but still need constraints to maintain temporal consistency. To address temporal consistency concerns, most methods like Tune-A-Video [70], Text2Video-Zero [27], FateZero [52], and Vid2Vid-Zero [66] incorporate cross-attention mechanisms. Some works propose training for video editing, such as Imagen Video [17] and Make-A-Video [59], but these require large datasets and significant computational resources. Unlike these methods, MeDM [13] uses a flow-coding algorithm to solve this problem. Canonical-based design does not require another mechanism to maintain temporal consistency, however. Once the canonical image is edited, the changes can be propagated to every frame using a deformation field.

Lifting the naturalness of canonical image by diffusion models. The diffusion prior has been applied in various domains. Reconfusion [71] is a few-shot novel view synthesis work that introduces a diffusion model before optimizing sampled novel views. Dreamfusion [50], a text-to-3D work, introduces score distillation sampling (SDS) loss, referencing a 2D diffusion model to optimize 3D outputs. This approach inspires us to utilize the diffusion model to improve performance. Other methods like [6, 7, 21, 41, 33, 61, 64, 68, 79] also employ diffusion prior for text-to-3D tasks.

Several diffusion models focus on text-to-image generation [53, 56, 58, 62]. Additionally, some diffusion models can refine corrupted images to make them appear more natural. We propose adding a diffusion model to our pipeline (Figure 2) to enhance the naturalness of our canonical images. Our goal is to improve the restoration of canonical images using the diffusion model. While we create canonical images using a hybrid deformation field, this method might not always deliver optimal performance, especially in scenarios with dramatic motion changes or severe non-rigid transformations. The capabilities of the hybrid deformation field are still limited in such cases. Therefore, we aim to introduce a diffusion model to make our canonical images more natural, enhancing the effectiveness of video editing. This research has significant practical implications as it can improve the quality and realism of video editing, benefiting various industries such as film production, advertising, and virtual reality.

3 Method

In this section, we first introduce our hybrid deformation modeling by combining homography and deformation MLP in Section 3.1. Subsequently, we elaborate on how we integrate the diffusion prior from a LoRA fine-tuned latent diffusion model to ensure the naturalness of our canonical image representation in Section 3.2. Finally, we provide an additional way to improve the quality of our

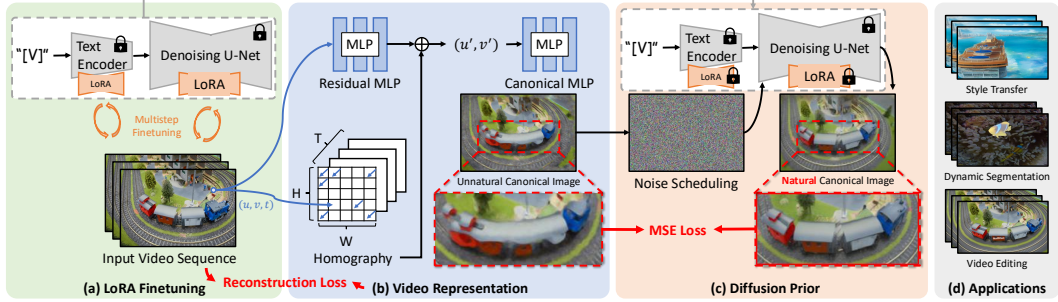


Figure 2: **Our proposed framework.** Given an input video sequence, our method aims to represent the video with a *natural* canonical image, which is a crucial representation for versatile downstream applications. (a) First, we fine-tune the LoRA weights of a pre-trained latent diffusion model on the input frames. (b) Second, we represent the video using a canonical MLP and a deformation field, which consists of homography estimation and residual deformation MLP for non-rigid residual deformations. By relying entirely on the reconstruction loss, the canonical MLP often fails to represent a natural canonical image, causing problems for downstream applications. *E.g.*, image-to-image translation methods such as ControlNet [75] may not be able to recognize that there is a train in the canonical image. (c) Therefore, we leverage the fine-tuned latent diffusion model to regularize and correct the unnatural canonical image into a natural one. Specifically, we sophisticatedly design a noise scheduling corresponding to the frame reconstruction process. (d) The natural and artifacts-free canonical image can then be facilitated to various downstream tasks such as video style transfer, dynamic segmentation, and editing, such as adding handwritten characters of “NaRCan”.

video representation by separating multiple canonical images and describing the necessary changes for downstream tasks in Section 3.3. Figure 2 shows our proposed framework.

3.1 Hybrid Deformation Field for Deformation Modeling

Traditional methods often rely on direct predictions of Δu and Δv , which means the displacement of pixel points u, v at time t , by using an MLP $g(\cdot, \cdot, \cdot)$ and query the RGB color by another canonical image MLP $f(\cdot, \cdot)$:

$$\Delta u, \Delta v = g(u, v, t), \quad [R, G, B] = f(u + \Delta u, v + \Delta v), \quad (1)$$

supplemented with a TVFlow regularization term to prevent overfitting by constraining the magnitude of spatial deformations. While this regularization effectively prevents extreme deformations, it merely imposes magnitude constraints without providing meaningful guidance for modeling complex spatial transformations. To address this limitation, we propose a hybrid deformation field architecture composed of a trainable homography matrix $H(u, v, t)$ and residual deformation MLP:

$$u', v' = H(u, v, t) + g(u, v, t), \quad [R, G, B] = f(u', v'). \quad (2)$$

Unlike the simple constraint-based approach of TVFlow, our method leverages homography to provide global displacement information as structural guidance for the subsequent residual deformation MLP. This hierarchical design enables the residual deformation MLP to learn and express the deformation field more accurately and effectively by building upon the initial geometric transformation provided by the homography matrix.

3.2 Diffusion Prior

The canonical image encompasses all information within the entire video and can be reconstructed for each original video frame using the deformation field outlined in Section 3.1. Thus, editing solely the canonical image yields a temporally coherent edited video. However, editing tasks such as drawing or writing on objects or editing based on ControlNet [75] require a natural image as input to produce meaningful edited images. Existing canonical base methods do not incorporate mechanisms to ensure that the generated image is natural; instead, they rely solely on the model’s ability to learn a natural image. However, when encountering scenarios with camera movement or significant changes in video content, these existing techniques cannot adapt to such drastic variations. The model may generate a canonical image that is nearly impossible to edit (rendering it devoid of any subsequent value).

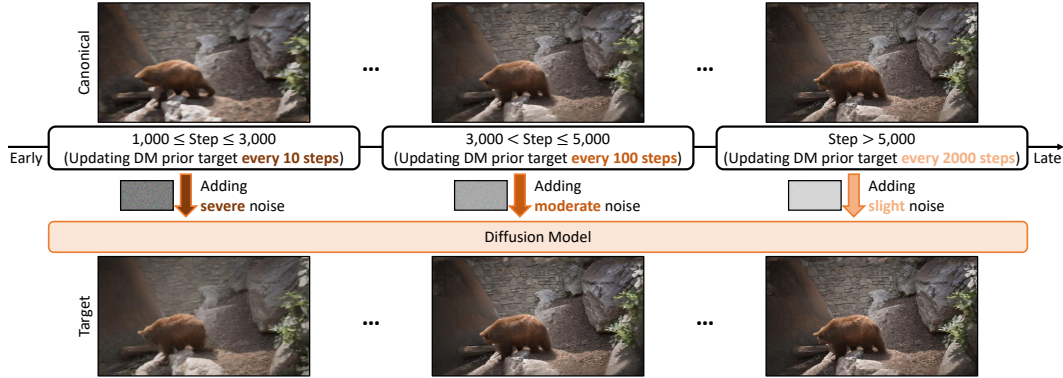


Figure 3: **Noise and diffusion prior update scheduling.** Initially, our model fits object outlines before the fields converge and without the diffusion prior, resulting in unnatural elements in the canonical image due to complex non-rigid objects. Upon introducing the diffusion prior with increased noise and update frequency, the model learns to generate natural, high-quality images, leading to convergence. Thus, the strength of noise and the update frequency will also decrease. Moreover, it’s worth mentioning that update scheduling cuts training time from 4.8 hours to 20 minutes.

To address these challenges, we introduce diffusion priors, which successfully mitigate this issue. Our method can generate high-quality canonical images through diffusion priors, providing valuable inputs for various video editing tasks.

LoRA fine-tuning. To enhance the current diffusion model’s ability to represent all video content better, we introduced a special token specific to this scene. We then fine-tuned the LoRA weight of the pre-trained diffusion model. This ensures that the diffusion model generates high-quality natural canonical images tailored to the testing sequence rather than randomly generating natural images that do not belong to the scene.

Noise and diffusion prior update scheduling. While the hybrid deformation field technique mentioned in Section 3.1 already produces better canonical images than other existing methods, it still needs to improve. As Figure 9(b) depicts, canonical images generated solely relying on homography and residual deformation MLP still exhibit various degrees of distortion and unnatural characteristics. Therefore, integrating diffusion priors becomes imperative. We extract the canonical region currently observed by the model and calculate diffusion loss with the target image generated by the diffusion model to ensure the generation of natural canonical images. However, generating a target image at each step would significantly prolong the training phase. Hence, we propose a hierarchical update scheduling to accelerate the process, which is shown in Figure 3. In the initial stages of training, when the deformation field has not yet converged, more substantial noise is introduced to allow the diffusion model to dominate the scene of the canonical image. Simultaneously, the frequency of generating target images needs to be denser. As training progresses and the deformation field becomes more stable, the noise intensity and frequency of generating target images decrease accordingly. This hierarchical scheduling approach ensures that the final canonical image approaches the quality of per-step updates while speeding up the training process by 14 times (4.8 hours to 20 minutes.) We opt for diffusion loss over SDS because using SDS, as mentioned in the Reconfusion [71] paper, is more prone to generate artifacts.

3.3 Separated NaRCan

When encountering overly complex scenes, relying solely on a single natural canonical image representing the entire scenario is impractical and unrealistic. Hence, we need to segment the original video into multiple segments $\{S_1, \dots, S_k\}$ and train dedicated residual deformation MLPs $\{R_1, \dots, R_k\}$ for each segment to obtain k natural canonical images $\{C_1, \dots, C_k\}$. It is worth noting that S_i and S_{i+1} have an overlap, referred to as the overlap window. Frames within this region are obtained using linear interpolation shown in Figure 4. This method ensures that excellent temporal consistency is maintained when switching from canonical image C_i to C_{i+1} . Table 1 demonstrates that our temporal consistency surpasses all existing methods even after segmentation. Additionally,




Figure 4: **Linear interpolation.** After using the grid trick [18] to obtain the highly consistent canonical images C_k and C_{k+1} , we interpolate all frames within the overlap window. As time progresses, the weight for reconstructing each frame gradually shifts from referencing C_k to solely referencing C_{k+1} . We achieve editing results with remarkable temporal consistency through this linear interpolation approach. Please refer to our supplementary material for more video results.

we adopt different processing approaches for various downstream tasks to ensure that Separated NaRCan can adeptly adapt to these tasks.

Style transfer. With multiple canonical images obtained, we utilize the grid trick [14, 24] to ensure sufficient consistency in style and content across the k canonical images. Specifically, we concat the canonical images into a larger image (with 2×2 canonical images) and use ControlNet [75] to perform text-guided image editing.

Video editing. When addressing video editing tasks such as handwriting, we leverage the pre-trained optical flow models to compute the flow between the k canonical images and use this flow to warp the editing content from C_1 to the C_k image.

4 Experiments

4.1 Experimental Setup

We conduct experiments to underscore the robustness and versatility of our proposed method. Our representation is robust with a variety of deformations, encompassing rigid and non-rigid objects, as well as complex scenarios such as smog and waves. We commence the introduction of the diffusion model at the 1000th iteration. From iteration 1000 to iteration 3000, the noise intensity is set at 0.4, and the target image generation frequency is every 10 iterations. Subsequently, spanning from iteration 3001 to iteration 5000, the noise intensity is adjusted to 0.3, with the target image generation frequency occurring every 100 iterations. Beyond the 5000th iteration mark, the noise intensity decreases to 0.2, and the target image is generated every 2000 iterations. The total iteration is 12000 iterations, and Figure 3 is shown to visualize the process of noise scheduling. On a single NVIDIA RTX4090 GPU, the average training duration is approximately 20 minutes when utilizing 100 video frames. When evaluating the temporal consistency in Table 1, we compare our separated NaRCan with other compared methods by setting $k = 3$, *i.e.*, we represent the sequence using three canonical images. By adjusting the training parameters accordingly, the optimization duration can be varied from 20 minutes to an hour.

4.2 Evaluation

Video editing. We run our method CoDeF [45], Hashing-nvd [5], CCEdit [15], and MeDM [13] on the DAVIS [49] and BalanceCC Benchmark that CCEdit proposed for evaluating the results of video editing. In Figure 5, we show the comparison of visual results. The BalanceCC Benchmark provides a unified text prompt for each scene, and we ensure that all video clips from this benchmark are restricted to a maximum of 100 frames. We utilize ControlNet [75] to edit the video using the unified text prompt provided by the BalanceCC Benchmark for a fair comparison. We executed MeDM and CCEdit within their provided environment settings and pre-trained models. Moreover,

Table 1: **Quantitative results on the BalanceCC [15] dataset.** There are 100 videos in BalanceCC. To ensure a representative distribution similar to BalanceCC, we randomly select 50 videos from BalanceCC and calculate warping and interpolation errors, which are the metrics for temporal consistency. Our method outperforms these baseline methods in terms of temporal consistency.

Method	Venue	Short-term $E_{warp} \downarrow$	Long-term $E_{warp} \downarrow$	Interpolation error \downarrow
Hashing-nvd [5]	ICCV 2023	0.0070	0.0495	9.204
CoDeF [45]	CVPR 2024	0.0085	0.0785	8.721
MeDM [13]	AAAI 2024	0.0072	0.0583	9.941
Ours	-	0.0065	0.0484	8.365

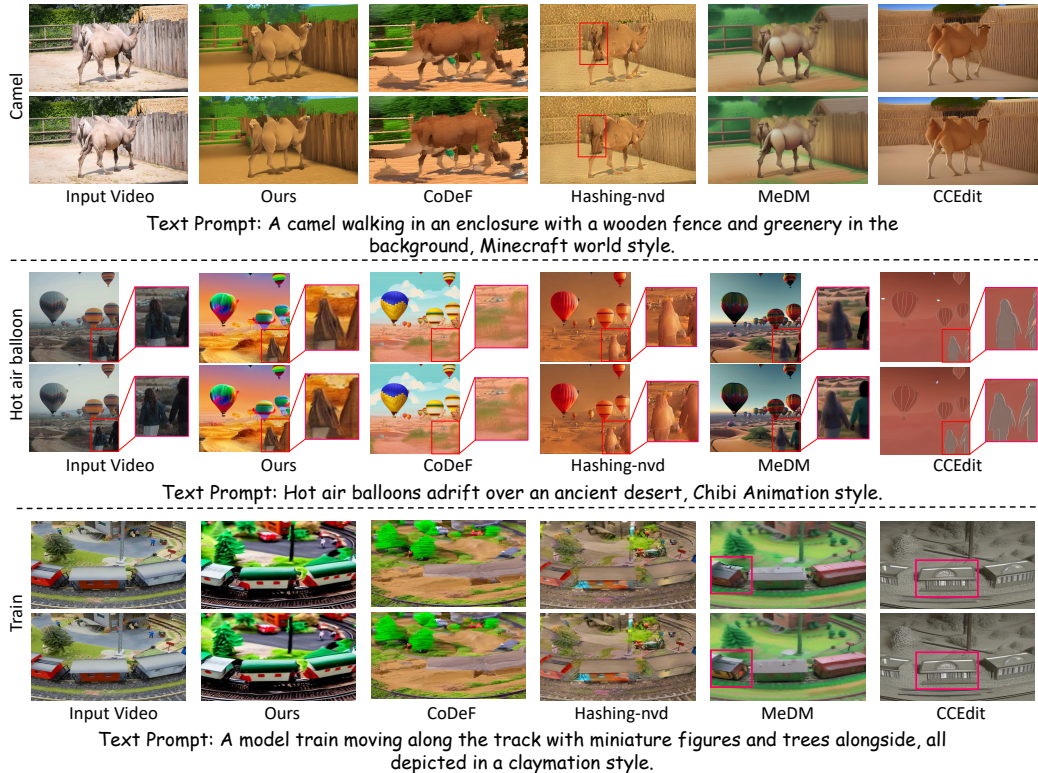


Figure 5: **Qualitative comparisons on text-guided video-to-video translation.** Our method achieves prompt alignment, synthesis quality, and temporal consistency best. Zoom in for the best view, and please refer to the supplementary materials for video comparisons. (a) In the camel scene, MeDM [13] fails to generate clear-textured images to ensure temporal consistency, while CCEdit [15] fails to correctly identify the second camel in the background. (b) CoDeF [45] misses capturing the presence of a person in the bottom right corner, Hashing-nvd [5] exhibit noticeable contours due to masking, and both MeDM and CCEdit suffer from temporal inconsistency issues. For instance, in MeDM, the person transitions from wearing black clothes to blue clothes. (c) MeDM and CCEdit still exhibit temporal inconsistency issues, such as significant color, texture, and structure changes. Other methods almost entirely lose the original train information or appear as unnatural artifacts.

Hashing-nvd is a video decomposition work that outputs two images representing foreground and background. To maintain the consistent style of these two images for video editing, We utilize the grid trick proposed by RAVE [24] to tackle this problem. Finally, there are also some video-editing results on DAVIS [49], and the corresponding text prompts originated from the BalanceCC Benchmark. For the evaluation of temporal consistency, as shown in Table 1, we utilize short-term [29, 77] and long-term [30] warping errors, along with interpolation error [32], as our primary metrics.

Metrics for evaluation. Since our main focus is text-guided video editing, we conduct user studies on edited videos compared with other methods, like CoDeF [45], MeDM [13], and Hashing-nvd [5].




Figure 6: **User Study.** Our method achieves the highest user preference ratios across all three aspects, compared with MeDM [13], CoDeF [45], Hashing-nvd [5].

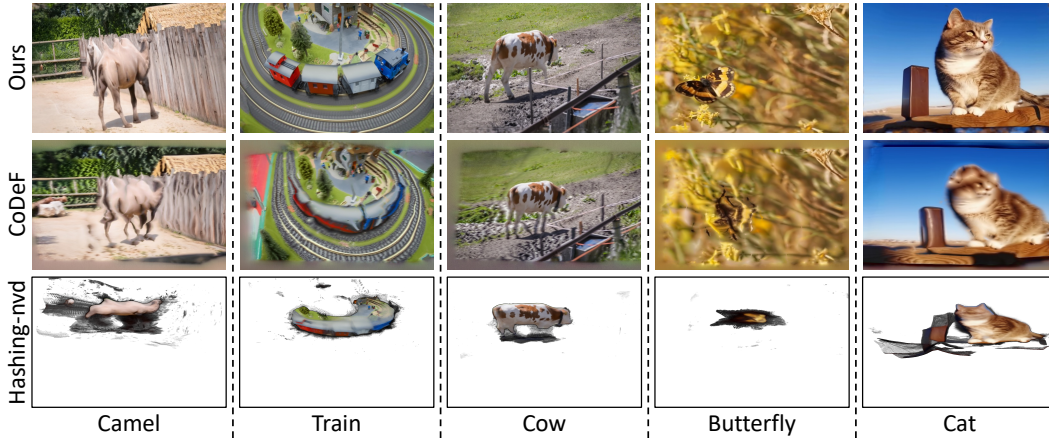


Figure 7: **Qualitative comparisons on the canonical image.** Our method generates more natural canonical images through a fine-tuned diffusion prior compared with CoDeF [45], Hashing-nvd [5]. The capability of the canonical image to represent input frames plays a crucial role in downstream applications. (Hashing-nvd consists of two canonical images. Here, we have selected the canonical image representing the foreground.)

In the user study, 36 participants were shown two edited videos, the original video on each page, and the text prompt used to edit the videos. There are three critical questions for users to answer. (1) *Which video has better temporal consistency?* (2) *Which video aligns better with the text prompt?* (3) *Which video has better overall quality?* Figure 6 summarizes the user study results, demonstrating that our method outperforms in all three dimensions.

Comparison of canonical images. In Figure 7, we run our method, CoDeF [45], and Hashing-nvd [5] on DAVIS [49] and BalanceCC [15]. Note that these are the canonical images of the input video, and we show the foreground atlas for Hashing-nvd. There will be two output atlases for the background and foreground of Hashing-nvd. For comparison, we only show the foreground atlas for Hashing-nvd. As a result, shown in Figure 7, the output canonical images of our method are more natural than others, even in the scenes with the dramatic motion of the objects, such as scenes named "train" and "butterfly." We can clearly see that our canonical images preserve the original object information well in the above two scenarios.

Downstream video processing. To evaluate the performance of our method’s handwritten video editing, we compare with CoDeF [45] and Hashing-nvd [5], which produce canonical images of the video for users to write the characters on the canonical image to accomplish video editing. We write “NaRCan” to test the performance of these three methods on two scenes, “gold-fish” and “train”, in the BalanceCC Benchmark [15] in Figure 8(a). We extract the same frame of videos for comparison.

For dynamic video segmentation, we segment the mask using the Segment Anything Model (SAM) [28] based on the learned canonical image of each method and propagate it to the sequence.

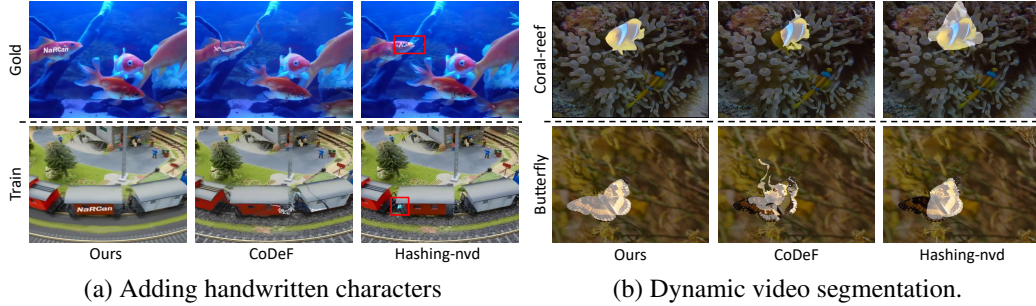


Figure 8: **Qualitative comparisons on (a) adding handwritten characters and (b) dynamic video segmentation.** Our method represents a natural image via diffusion prior, thus can achieve temporally consistent video editing and able to precisely edit desired areas.

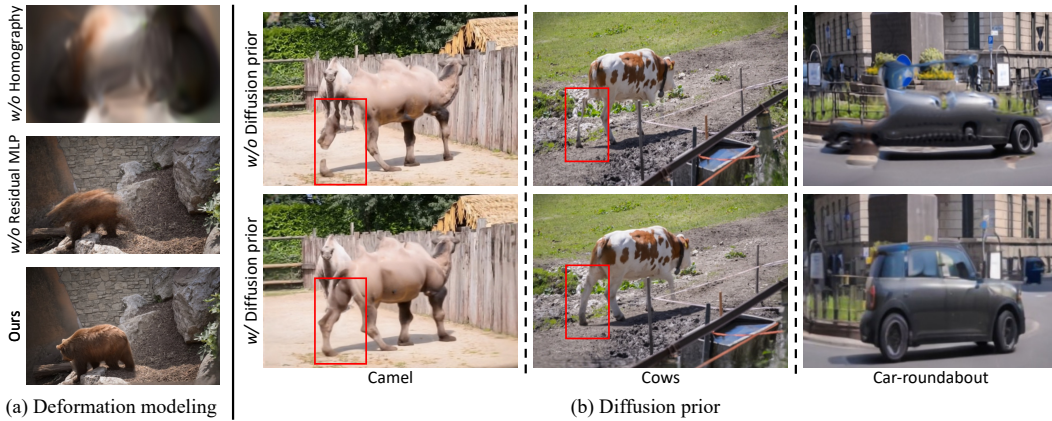


Figure 9: **Ablation studies.** (a) Deformation modeling: (*Top*) We show that canonical images without homography modeling fail to generate a faithful image as the capacity of residual deformation MLP could dominate the training process and still achieve near-perfect frame reconstruction. (*Mid*) On the contrary, without residual deformation MLP, our method cannot model local non-rigid transformation, resulting in blurry foreground objects. (*Bottom*) Combining homography and residual deformation MLP has the best of both worlds and achieves the best canonical image representation. (b) Diffusion prior: (*Top*) Without diffusion prior to regularizing the canonical image, the training process relies only on the frame reconstruction and could sacrifice the faithfulness of the canonical image. (*Bottom*) Our fine-tuned diffusion prior effectively corrects the canonical image to faithfully represent the input frames and results in natural canonical images.

The target of the mask in these two scenes are the clownfish named “coral-reef” and the flying butterfly in the scene called “butterfly.” We use white to mark the mask for better visibility in Figure 8(b).

4.3 Ablation Study

Homography & Residual Deformation MLP. In this section, we conducted ablation experiments focusing on both homography and residual deformation MLP. Figure 9(a) clearly illustrates that using only MLP to fit the deformation field [31, 46, 47, 51] results in unsuitable canonical images for downstream tasks. Without the global information homography provides, the model encounters difficulties in converging the diffusion loss and instead focuses solely on optimizing the reconstruction loss. Conversely, if we rely solely on homography to express the deformation field, homography’s expressive power is limited in capturing detailed variations in non-rigid objects. As a result, only approximate and blurred outcomes can be obtained.

Diffusion prior. Relying on homography and the residual deformation MLP achieves relatively better canonical images than previous methods. However, the lack of assistance from the diffusion prior still prevents the stable generation of high-quality natural canonical images. Figure 9(b) demonstrates the significance of supervising canonical images with the diffusion prior.

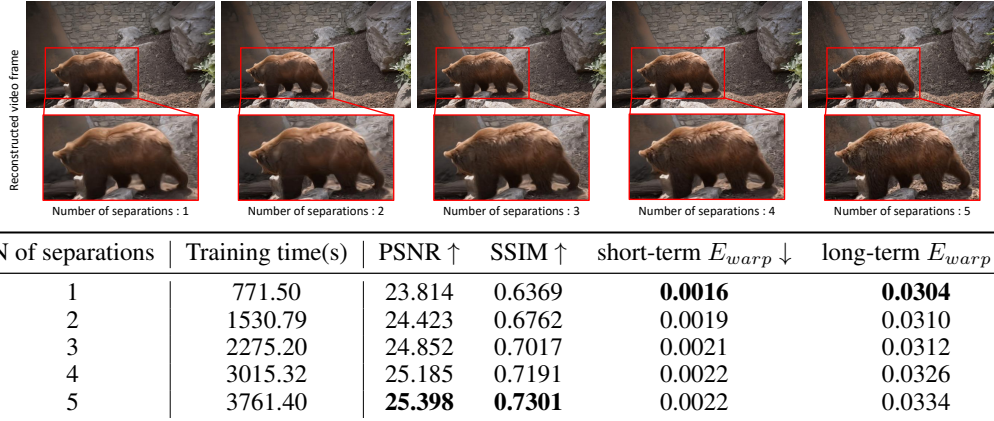


Figure 10: **Trade-off between reconstruction quality and temporal consistency with varying separations.** The visual results demonstrate that increasing the number of separations (from 1 to 5) improves the reconstruction quality of video frames. However, the table reveals a trade-off: as the number of separations increases, temporal consistency decreases. Our empirical findings suggest that using 3 separations typically achieves a balance between reconstruction quality and temporal consistency for most scenarios.



Figure 11: **Failure cases.**

Impact of Separation Numbers. To investigate the optimal configuration of our Separated NaRCan, we conduct ablation studies by varying the number of canonical images from one to five. Our analysis reveals a clear trade-off between reconstruction quality and temporal consistency, as illustrated in Figure 10. While increasing the number of separations improves frame reconstruction fidelity by allowing more detailed scene representation, it introduces additional transition regions that may compromise temporal coherence. Through extensive experimentation, we find that using three canonical images strikes an optimal balance between representational capacity and temporal stability for most video scenarios.

5 Conclusion

In this paper, we introduce NaRCan, a video editing framework, integrating diffusion priors and LoRA [18] fine-tuning to produce high-quality natural canonical images. This method tackles the challenges of maintaining the natural appearance of the canonical image and reduces training times with new noise-scheduling techniques. The results show NaRCan’s advantage in managing complex video dynamics and its potential for wide use in various multimedia applications.

Limitations. Our method relies on LoRA [18] fine-tuning to enhance the diffusion model’s ability to represent the current scene. However, LoRA fine-tuning is time-consuming (about 30 minutes). Additionally, our training pipeline includes diffusion loss, which increases the training time. In cases of extreme changes in video scenes, diffusion loss sometimes fails to guide the model in generating high-quality natural images (Figure 11). These limitations point out the challenge of balancing model adaptability with computational efficiency and effectiveness in varied conditions.

Acknowledgments and Disclosure of Funding

This research was funded by the National Science and Technology Council, Taiwan, under Grants NSTC 112-2222-E-A49-004-MY2 and 113-2628-E-A49-023-. The authors are grateful to Google, NVIDIA, and MediaTek Inc. for generous donations. Yu-Lun Liu acknowledges the Yushan Young Fellow Program by the MOE in Taiwan.

References

- [1] Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kasten, and Tali Dekel. Text2live: Text-driven layered image and video editing. In *ECCV*, 2022.
- [2] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent diffusion models. In *CVPR*, 2023.
- [3] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023.
- [4] Duygu Ceylan, Chun-Hao P Huang, and Niloy J Mitra. Pix2video: Video editing using image diffusion. In *ICCV*, 2023.
- [5] Cheng-Hung Chan, Cheng-Yang Yuan, Cheng Sun, and Hwann-Tzong Chen. Hashing neural video decomposition with multiplicative residuals in space-time. In *ICCV*, 2023.
- [6] Dave Zhenyu Chen, Haoxuan Li, Hsin-Ying Lee, Sergey Tulyakov, and Matthias Nießner. Scenetex: High-quality texture synthesis for indoor scenes via diffusion priors. *arXiv preprint arXiv:2311.17261*, 2023.
- [7] Rui Chen, Yongwei Chen, Ningxin Jiao, and Kui Jia. Fantasia3d: Disentangling geometry and appearance for high-quality text-to-3d content creation. In *ICCV*, 2023.
- [8] Weifeng Chen, Jie Wu, Pan Xie, Hefeng Wu, Jiashi Li, Xin Xia, Xuefeng Xiao, and Liang Lin. Control-a-video: Controllable text-to-video generation with diffusion models. In *ICLR*, 2024.
- [9] Yue Chen, Xingyu Chen, Xuan Wang, Qi Zhang, Yu Guo, Ying Shan, and Fei Wang. Local-to-global registration for bundle-adjusting neural radiance fields. In *CVPR*, 2023.
- [10] Jiaxin Cheng, Tianjun Xiao, and Tong He. Consistent video-to-video transfer using synthetic dataset. In *ICLR*, 2024.
- [11] Ka Leong Cheng, Qiuyu Wang, Zifan Shi, Kecheng Zheng, Yinghao Xu, Hao Ouyang, Qifeng Chen, and Yujun Shen. Learning naturally aggregated appearance for efficient 3d editing. *arXiv preprint arXiv:2312.06657*, 2023.
- [12] Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. Segment and track anything. *arXiv preprint arXiv:2305.06558*, 2023.
- [13] Ernie Chu, Tzuhsuan Huang, Shuo-Yen Lin, and Jun-Cheng Chen. Medm: Mediating image diffusion models for video-to-video translation with temporal correspondence guidance. In *AAAI*, 2024.
- [14] Jan-Niklas Dihlmann, Andreas Engelhardt, and Hendrik Lensch. Signerf: Scene integrated generation for neural radiance fields. *arXiv preprint arXiv:2401.01647*, 2024.
- [15] Ruoyu Feng, Wenming Weng, Yanhui Wang, Yuhui Yuan, Jianmin Bao, Chong Luo, Zhibo Chen, and Baining Guo. Ccredit: Creative and controllable video editing via diffusion models. In *CVPR*, 2024.
- [16] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Tokenflow: Consistent diffusion features for consistent video editing. In *ICLR*, 2024.

- [17] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- [18] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *ICLR*, 2021.
- [19] Zhihao Hu and Dong Xu. Videocontrolnet: A motion-guided video-to-video translation framework by using diffusion model with controlnet. *arXiv preprint arXiv:2307.14073*, 2023.
- [20] Allan Jabri, Andrew Owens, and Alexei Efros. Space-time correspondence as a contrastive random walk. In *NeurIPS*, 2020.
- [21] Ajay Jain, Ben Mildenhall, Jonathan T Barron, Pieter Abbeel, and Ben Poole. Zero-shot text-guided object generation with dream fields. In *CVPR*, 2022.
- [22] Varun Jampani, Raghudeep Gadde, and Peter V Gehler. Video propagation networks. In *CVPR*, 2017.
- [23] Ondřej Jamriška, Šárka Sochorová, Ondřej Texler, Michal Lukáč, Jakub Fišer, Jingwan Lu, Eli Shechtman, and Daniel Šykora. Stylizing video by example. *ACM TOG*, 2019.
- [24] Ozgur Kara, Bariscan Kurtkaya, Hidir Yesiltepe, James M Rehg, and Pinar Yanardag. Rave: Randomized noise shuffling for fast and consistent video editing with diffusion models. In *CVPR*, 2024.
- [25] Yoni Kasten, Dolev Ofri, Oliver Wang, and Tali Dekel. Layered neural atlases for consistent video editing. *ACM TOG*, 2021.
- [26] Kai Katsumata, Duc Minh Vo, Bei Liu, and Hideki Nakayama. Revisiting latent space of gan inversion for robust real image editing. In *WACV*, 2024.
- [27] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *ICCV*, 2023.
- [28] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *ICCV*, 2023.
- [29] Wei-Sheng Lai, Jia-Bin Huang, Oliver Wang, Eli Shechtman, Ersin Yumer, and Ming-Hsuan Yang. Learning blind video temporal consistency. In *ECCV*, 2018.
- [30] Chenyang Lei, Yazhou Xing, and Qifeng Chen. Blind video temporal consistency via deep video prior. In *NeurIPS*, 2020.
- [31] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *CVPR*, 2022.
- [32] Xirui Li, Chao Ma, Xiaokang Yang, and Ming-Hsuan Yang. Vidtome: Video token merging for zero-shot video editing. In *CVPR*, 2024.
- [33] Chen-Hsuan Lin, Jun Gao, Luming Tang, Towaki Takikawa, Xiaohui Zeng, Xun Huang, Karsten Kreis, Sanja Fidler, Ming-Yu Liu, and Tsung-Yi Lin. Magic3d: High-resolution text-to-3d content creation. In *CVPR*, 2023.
- [34] Sharon Lin, Matthew Fisher, Angela Dai, and Pat Hanrahan. Layerbuilder: Layer decomposition for interactive image and video color editing. *arXiv preprint arXiv:1701.03754*, 2017.
- [35] Yu-Lun Liu, Wei-Sheng Lai, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Hybrid neural fusion for full-frame video stabilization. In *ICCV*, 2021.

- [36] Yu-Lun Liu, Chen Gao, Andreas Meuleman, Hung-Yu Tseng, Ayush Saraf, Changil Kim, Yung-Yu Chuang, Johannes Kopf, and Jia-Bin Huang. Robust dynamic radiance fields. In *CVPR*, 2023.
- [37] Erika Lu, Forrester Cole, Tali Dekel, Weidi Xie, Andrew Zisserman, David Salesin, William T Freeman, and Michael Rubinstein. Layered neural rendering for retiming people in video. *SIGGRAPH Asia*, 2020.
- [38] Erika Lu, Forrester Cole, Tali Dekel, Andrew Zisserman, William T Freeman, and Michael Rubinstein. Omnimate: Associating objects and their effects in video. In *CVPR*, 2021.
- [39] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *ICLR*, 2022.
- [40] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [41] Nasir Mohammad Khalid, Tianhao Xie, Eugene Belilovsky, and Tiberiu Popa. Clip-mesh: Generating textured meshes from text using pretrained image-text models. In *SIGGRAPH Asia*, 2022.
- [42] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 2022.
- [43] Seonghyeon Nam, Marcus A Brubaker, and Michael S Brown. Neural image representations for multi-image fusion and layer separation. In *ECCV*, 2022.
- [44] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *ICML*, 2022.
- [45] Hao Ouyang, Qiuyu Wang, Yuxi Xiao, Qingyan Bai, Juntao Zhang, Kecheng Zheng, Xiaowei Zhou, Qifeng Chen, and Yujun Shen. Codef: Content deformation fields for temporally consistent video processing. In *CVPR*, 2024.
- [46] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, 2021.
- [47] Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M. Seitz. Hypernerf: A higher-dimensional representation for topologically varying neural radiance fields. *ACM TOG*, 2021.
- [48] Gaurav Parmar, Yijun Li, Jingwan Lu, Richard Zhang, Jun-Yan Zhu, and Krishna Kumar Singh. Spatially-adaptive multilayer selection for gan inversion and editing. In *CVPR*, 2022.
- [49] 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.
- [50] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2023.
- [51] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, 2021.
- [52] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. In *ICCV*, 2023.
- [53] 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 *ICML*, 2021.

- [54] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021.
- [55] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [56] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.
- [57] Manuel Ruder, Alexey Dosovitskiy, and Thomas Brox. Artistic style transfer for videos. In *GCPR*, 2016.
- [58] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *NeurIPS*, 2022.
- [59] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. In *ICLR*, 2023.
- [60] Haorui Song, Yong Du, Tianyi Xiang, Junyu Dong, Jing Qin, and Shengfeng He. Editing out-of-domain gan inversion via differential activations. In *ECCV*, 2022.
- [61] Junshu Tang, Tengfei Wang, Bo Zhang, Ting Zhang, Ran Yi, Lizhuang Ma, and Dong Chen. Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior. In *ICCV*, 2023.
- [62] Luming Tang, Nataniel Ruiz, Qinghao Chu, Yuanzhen Li, Aleksander Holynski, David E Jacobs, Bharath Hariharan, Yael Pritch, Neal Wadhwa, Kfir Aberman, et al. Realfill: Reference-driven generation for authentic image completion. In *CVPR*, 2024.
- [63] Ondřej Texler, David Futschik, Michal Kučera, Ondřej Jamriška, Šárka Sochorová, Mencei Chai, Sergey Tulyakov, and Daniel Šykora. Interactive video stylization using few-shot patch-based training. *ACM TOG*, 2020.
- [64] Haochen Wang, Xiaodan Du, Jiahao Li, Raymond A Yeh, and Greg Shakhnarovich. Score jacobian chaining: Lifting pretrained 2d diffusion models for 3d generation. In *CVPR*, 2023.
- [65] Tengfei Wang, Yong Zhang, Yanbo Fan, Jue Wang, and Qifeng Chen. High-fidelity gan inversion for image attribute editing. In *CVPR*, 2022.
- [66] Wen Wang, Yan Jiang, Kangyang Xie, Zide Liu, Hao Chen, Yue Cao, Xinlong Wang, and Chunhua Shen. Zero-shot video editing using off-the-shelf image diffusion models. *arXiv preprint arXiv:2303.17599*, 2023.
- [67] Xiaolong Wang, Allan Jabri, and Alexei A Efros. Learning correspondence from the cycle-consistency of time. In *CVPR*, 2019.
- [68] Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. In *NeurIPS*, 2023.
- [69] Zhixiang Wang, Yu-Lun Liu, Jia-Bin Huang, Shin’ichi Satoh, Sizhuo Ma, Gurunandan Krishnan, and Jian Wang. Disco: Portrait distortion correction with perspective-aware 3d gans. *IJCV*, 2024.
- [70] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, 2023.
- [71] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P Srinivasan, Dor Verbin, Jonathan T Barron, Ben Poole, et al. Reconfusion: 3d reconstruction with diffusion priors. In *CVPR*, 2024.
- [72] Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. Rerender a video: Zero-shot text-guided video-to-video translation. In *SIGGRAPH Asia*, 2023.

- [73] Vickie Ye, Zhengqi Li, Richard Tucker, Angjoo Kanazawa, and Noah Snavely. Deformable sprites for unsupervised video decomposition. In *CVPR*, 2022.
- [74] Ahmet Burak Yildirim, Hamza Pehlivan, Bahri Batuhan Bilecen, and Aysegul Dundar. Diverse inpainting and editing with gan inversion. In *ICCV*, 2023.
- [75] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023.
- [76] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. In *ICLR*, 2024.
- [77] Shangchen Zhou, Peiqing Yang, Jianyi Wang, Yihang Luo, and Chen Change Loy. Upscale-a-video: Temporal-consistent diffusion model for real-world video super-resolution. In *CVPR*, 2024.
- [78] Jiapeng Zhu, Yujun Shen, Deli Zhao, and Bolei Zhou. In-domain gan inversion for real image editing. In *ECCV*, 2020.
- [79] Joseph Zhu and Peiye Zhuang. Hifa: High-fidelity text-to-3d with advanced diffusion guidance. *arXiv preprint arXiv:2305.18766*, 2023.

A Appendix

A.1 Single NaRCan compared with Separated NaRCan

Dissecting canonical-based methods. In this section, we will delve deeper into two different canonical-based methods: CoDeF and Hashing-nvd. As shown clearly in Figure 12, CoDeF generates poor-quality canonical images due to the lack of supervision and constraints from the diffusion prior. In contrast, our method can consistently produce high-quality canonical images regardless of the video length, effectively preparing for subsequent downstream editing tasks. The drawbacks of Hashing-nvd are also apparent. This method relies heavily on the Mask-RCNN technique, often resulting in inaccurate or incorrect foreground and background segmentation. Consequently, the final canonical images generated are difficult to edit or inapplicable to techniques such as ControlNet.

Parameter settings of Separated NaRCan. Subsequently, as shown in Figure 13, From our experiments, we found that when using Separated NaRCan, it is crucial to limit the number of segmentations in the video. This is because the editing information on the edited canonical image is propagated through warping using flows. If the flow is inaccurate or contains errors, excessive warping can lead to severe cumulative errors, significantly damaging the edited content.




Figure 12: **Canonical analysis.** (a) Compared to existing canonical-based methods, our approach robustly generates high-quality natural canonical images regardless of the video length. (b) Our method accurately preserves the correct foreground and background information from the original scenes, avoiding severe distortions or warping and preventing generating content inconsistent with the scene. For example, in “kite-surfing”, Hashing-nvd erroneously generates an additional person.

A.2 Video Comparisons

We provide an interactive HTML interface to browse video results for comparisons. Specifically, we provide video comparisons on three different tasks: (1) ControlNet style transfer, (2) dynamic segmentation, and (3) adding handwritten characters. We compare our proposed method, NaRCan, with state-of-the-art methods: Hashing-nvd [5], CoDeF [45], and MeDM [13]. We also visualize the optimized canonical images if available for reference.

A.3 User Study

To conduct our user study, we employ GitHub Pages in conjunction with a Google Form to facilitate user evaluations of video quality. Each evaluation session comprises 15 scenes, each of which contains 3 questions. For these evaluations, we randomly select 15 scenes from a pool of 100 scenes in BalanceCC Benchmark [15]. Each evaluation page (Figure 14) presents a video edited using our method alongside a randomly selected video sourced from compared methods: MeDM [13], CoDeF [45], or Hashing-nvd [5].




Figure 13: **Separated NaRCan and per-frame warping.** Since we have the flexibility to separate into multiple canonical images, we conduct an experiment to determine how many canonical images are optimal. We test using Separated NaRCan with segmentation equal to 3 and segmentation equal to N , where N equals the number of frames in the video. The results show that the editing information is damaged and has significant displacement due to inaccurate optical flow.

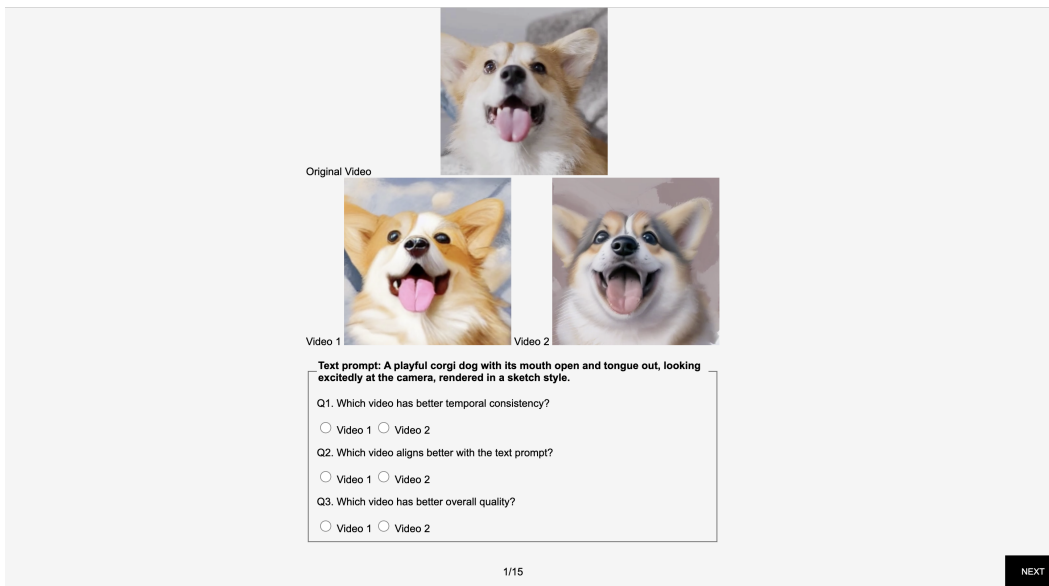


Figure 14: **User Study Website.** The video above is the original video. The two videos below are the ones being compared. We ask participants to determine which of the given videos best matches the description provided in the questions.

A.4 Grid Trick

We adopt the technique named “grid trick” from RAVE [24]. In that paper, a method called grid trick is introduced, where multiple images are merged together and then fed to ControlNet [75] for editing to obtain style transfer images with sufficient consistency in content and color tones. Therefore, when using Separated NaRCan, we only need to apply the grid trick to our k canonical images to achieve this desired consistency among the canonical images easily [24].

Finally, coupled with the linear interpolation method mentioned in our paper, we can ensure that the video maintains sufficient temporal consistency, thereby successfully creating a high-quality style transfer video.

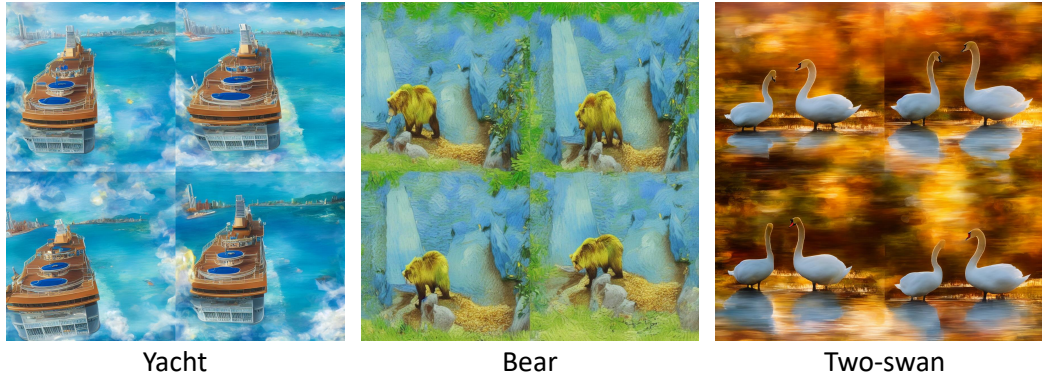


Figure 15: **Canonical images after using the grid trick.** Using Separated NaRCAN, we will obtain multiple canonical images, and through the grid trick, we can generate a high-quality and consistent style transfer canonical image. Therefore, Separated NaRCAN still demonstrates excellent performance in this task.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and the introduction briefly introduce the motivation and the main contribution of our work. Our work utilizes the diffusion model to lift the performance of video editing, which is also mentioned in the abstract and the Section 1 introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discussed our limitations in 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.

- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: First, we conducted a user study on our method and other baseline methods, including CoDeF [45], MeDM [13], Hashing-nvd [5], CCEdit [15]. We designed three questions focused on temporal consistency, text alignment, and quality and tested them on approximately 36 participants. Secondly, we evaluated these methods, including ours, using three metrics to demonstrate that our performance is better.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: We explain our experimental setting in Sec 4.1, and we will provide the metric and the main code. So, we think that will not be a concern. And on the other hand, we will provide the link to our user study and the statistic we show is based on the people we tested, not created by us.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.

- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We provide it in the zip file

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Yes, we would show our hyperparameters in Sec 4.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: No, we didn't show it.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We show the machine we used in Sec 4.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We have improved the temporal consistency and quality in processing various video tasks, such as style transfer, dynamic segmentation, and other techniques typically applied to images.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: We don't release any data or models that have a high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.

- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: The paper explicitly mentions and respects the licenses and terms of use for all external assets. The creators of the code, data, and models are properly credited in the references section, and the terms of use for these assets are adhered to as specified by their respective licenses.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: The new assets introduced in the paper are thoroughly documented in [Sec/ 3](#).

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[Yes\]](#)

Justification: We conduct user study in [Figure 6](#) and also provided screenshot in supplementary material.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [No]

Justification: Our paper doesn't bring potential risks incurred by study participants.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.