

FilmWeaver: Weaving Consistent Multi-Shot Videos with Cache-Guided Autoregressive Diffusion

Xiangyang Luo^{1,2*}, Qingyu Li^{2†}, Xiaokun Liu², Wenyu Qin², Miao Yang², Meng Wang², Pengfei Wan², Di Zhang², Kun Gai², Shao-Lun Huang^{1†}

¹Tsinghua Shenzhen International Graduate School, Tsinghua University

²Kling Team, Kuaishou Technology

Abstract

Current video generation models perform well at single-shot synthesis but struggle with multi-shot videos, facing critical challenges in maintaining character and background consistency across shots and flexibly generating videos of arbitrary length and shot count. To address these limitations, we introduce **FilmWeaver**, a novel framework designed to generate consistent, multi-shot videos of arbitrary length. First, it employs an autoregressive diffusion paradigm to achieve arbitrary-length video generation. To address the challenge of consistency, our key insight is to decouple the problem into inter-shot consistency and intra-shot coherence. We achieve this through a dual-level cache mechanism: a shot memory caches keyframes from preceding shots to maintain character and scene identity, while a temporal memory retains a history of frames from the current shot to ensure smooth, continuous motion. The proposed framework allows for flexible, multi-round user interaction to create multi-shot videos. Furthermore, due to this decoupled design, our method demonstrates high versatility by supporting downstream tasks such as multi-concept injection and video extension. To facilitate the training of our consistency-aware method, we also developed a comprehensive pipeline to construct a high-quality multi-shot video dataset. Extensive experimental results demonstrate that our method surpasses existing approaches on metrics for both consistency and aesthetic quality, opening up new possibilities for creating more consistent, controllable, and narrative-driven video content. Project Page: <https://filmweaver.github.io>

Introduction

With the continuous advancement of video diffusion models (Blattmann et al. 2023; Kong et al. 2024; Wan et al. 2025), video generation systems have demonstrated remarkable capabilities and found applications across various domains (Xue et al. 2024). However, these models are primarily designed for single video generation. Compared to conventional single-shot video generation, multi-shot videos offer significantly higher practical value in filmmaking, sto-

rytelling, and other creative fields, as they enable the construction of more complex narratives. Nevertheless, this task presents greater challenges, primarily in two aspects: i) maintaining consistency of the same subjects or backgrounds across different shots, and ii) managing shot duration and the number of shots. These requirements cannot be directly addressed by traditional video generation models.

The most straightforward approach to multi-shot generation involves using multiple prompts to generate each shot independently and then concatenating them. However, the coarse-grained textual representations struggle to ensure consistency between shots. To address the consistency problem in multi-shot generation, early methods typically employed complex pipelines involving multiple models (Lin et al. 2023; Long et al. 2024; Yuan et al. 2024; Xiao et al. 2025), decomposing multi-shot video generation into keyframe generation and image-to-video synthesis. These approaches achieved multi-shot consistency through object insertion techniques or attention reference mechanisms (Zhou et al. 2024) during keyframe generation. However, such methods rely heavily on complex pipeline designs and, due to the independent generation of individual segments without considering global temporal coherence, often result in visual discontinuities and abrupt scene transitions.

Recent work has proposed simultaneous multi-shot generation methods that partition the video generated by a video generation model into multiple segments (Kara et al. 2025; Qi et al. 2025; Guo et al. 2025), with each segment corresponding to a shot. These methods are trained on corresponding multi-shot datasets and achieve improved inter-shot consistency. However, since multiple shots share a single sequence, the duration of individual shots is severely constrained. TTT (Dalal et al. 2025) introduces RNN-like mechanisms in the intermediate layers of DiT (Peebles and Xie 2023) to maintain shot consistency, but these lack long-term memory, have fixed shot durations, and incur substantial training costs. LCT (Guo et al. 2025) employs additional positional encoding to achieve multi-shot generation and shot extension, but it requires two-stage training and only supports the MM-DiT (Yang et al. 2024) architecture.

To address these limitations, we introduce a novel cache-guided autoregressive framework for multi-shot video generation. Our approach is founded on a key insight: the ex-

*This work was conducted during the author’s internship at Kling Team, Kuaishou Technology.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

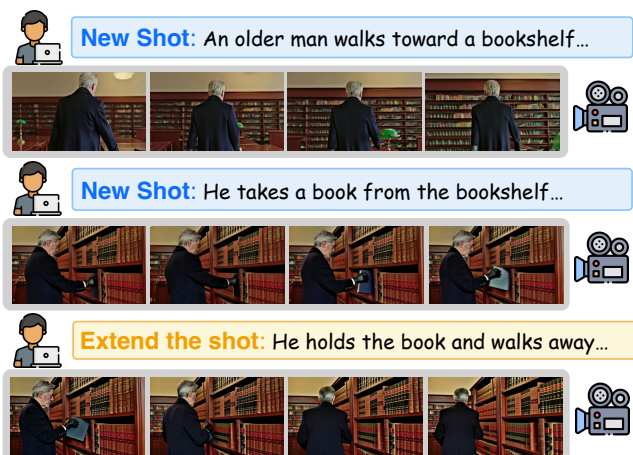


Figure 1: Our framework supports interactive creation of multi-shot sequences. Users can generate new shots or extend existing ones to build a coherent narrative.

explicit decoupling of inter-shot consistency and intra-shot coherence. To this end, we design a dual-level cache system composed of a Shot Cache and a Temporal Cache. The Shot Cache is responsible for maintaining inter-shot consistency; it stores identifying visual features of subjects and backgrounds from preceding shots. When generating a new shot, it allows the model to retrieve relevant context based on the prompt, ensuring character and style continuity across narrative breaks. In parallel, the Temporal Cache manages intra-shot coherence by retaining a memory of the immediately preceding frames within the current shot, guaranteeing fluid motion and preventing visual flickering.

Our framework integrates this dual-level cache system using in-context injection, a mechanism that avoids architectural modifications and thus ensures broad compatibility with existing pre-trained text-to-video models. As shown in Figure 1, this approach provides intuitive interactive control over the narrative flow. This core flexibility is the key that unlocks not only the generation of videos with an arbitrary number of shots and durations but also a range of challenging downstream tasks, such as multi-concept injection (Huang et al. 2025; Luo et al. 2025b) and interactive video extension.

Training this model requires constructing high-quality consistent multi-shot video datasets, which are currently scarce. We propose a data collection pipeline for building high-quality multi-shot datasets, incorporating shot-to-scene shot collection and scene-to-shot multi-level annotation refinement to ensure high-quality consistent caption construction. Given the lack of multi-shot video evaluation methods, we propose evaluation metrics and several baselines based on identity consistency, background consistency, and video quality. In summary, our contributions are as follows:

- We introduce a novel cache-guided autoregressive framework for multi-shot video generation. This architecture uniquely integrates a dual-level cache into the autoregressive flow, enabling the synthesis of coherent videos

of arbitrary length by explicitly managing intra and inter shot consistency.

- The proposed framework exhibits remarkable flexibility, supporting a range of challenging downstream tasks including multi-concept character injection, interactive video extension, showcasing its broad applicability.
- We propose a high-quality consistent shot collection pipeline that includes shot segmentation and clustering, and through multi-level agent annotation, we collect a high-quality multi-shot dataset.
- Extensive experiments demonstrate that our framework surpasses existing methods in video quality and character/background consistency.

Related Works

Single-shot Video Generation

Video Diffusion Video diffusion models (Blattmann et al. 2023; Kong et al. 2024; Wan et al. 2025; Yang et al. 2024) have evolved from their successful applications in image synthesis, with the core challenge being modeling temporal consistency. Early methods (Blattmann et al. 2023; Guo et al. 2023) leverage image generation model priors by inserting additional temporal attention modules to learn motion dynamics, but the video quality is limited. Some approaches (Voleti et al. 2024; Wan et al. 2024) partition the image into a grid of patches to achieve consistent multi-view or video generation. Current video generation models (Kong et al. 2024; Wan et al. 2025) typically employ a 3D-DiT architecture that flatten video height, width, and temporal dimensions into a single dimension for attention computation, further improving generated video quality.

Long Video Generation Extending video generation beyond short clips remains a significant challenge. Early works employ various strategies, such as leveraging latent diffusion models (Rombach et al. 2022), generating variable-length videos from text sequences, and implementing coarse-to-fine architectures (Jin et al. 2024). Other efforts focus on achieving temporal consistency through distributed generation or extending pre-trained models via noise rescheduling techniques (Cai et al. 2025). Recently, research focus has shifted toward long-context video modeling to better utilize historical information (Chen et al. 2024a, 2025). FAR (Gu, Mao, and Shou 2025) proposed a framework conditioned on both long-term and short-term context windows, while FramePack (Zhang and Agrawala 2025) introduces hierarchical compression of context frames. These methods highlight the critical role of context as a form of memory in achieving scene-consistent long video generation. Unlike these approaches, in multi-shot scenarios, we propose a two-level cache system to simultaneously achieve inter-shot consistency and intra-shot coherence.

Multi-shot Generation

Multi-shot generation, particularly in narrative contexts, faces challenges in maintaining consistency. Early works in story visualization center on image generation. To this end,

various techniques are developed; for instance, AutoStudio (Cheng et al. 2024) employ IP-Adapter (Ye et al. 2023) for character identity injection, while StoryDiffusion (Zhou et al. 2024) and Consistory (Tewel et al. 2024) utilizes attention concatenation mechanisms (Luo et al. 2025a) for visual consistency. Early multi-shot video generation methods adopt a two-stage paradigm of keyframe generation followed by I2V video generation. Mora (Yuan et al. 2024) and MovieAgent (Wu, Zhu, and Shou 2025) adopt multi-agent frameworks, while VideodirectorGPT (Lin et al. 2023) introduces layout-to-video generation and VideoStudio (Long et al. 2024) distinguish foreground and background with action information. However, these complex, multi-model paradigms generate segments independently, which often leads to visual discontinuities.

Recent methods enhance inter-shot consistency through simultaneous multi-shot generation by decomposing the generated sequence into multiple shots (Qi et al. 2025; Kara et al. 2025). However, splitting the original video length results in very short duration per shot, reducing practical applicability. TTT (Dalal et al. 2025) introduces RNNs at intermediate DiT layers but lacks long-range dependencies and incurs high training costs. LCT (Guo et al. 2025) and EchoShot (Wang et al. 2025) employ complex positional encoding to distinguish shots. In contrast, we adopt simple positional encoding with a two-level cache system to simultaneously ensure inter-shot consistency and intra-shot coherence through autoregressive single-step training.

Method

Our proposed framework, FilmWeaver, introduces an autoregressive video generation paradigm capable of producing multi-shot videos of arbitrary length and shot count. The core of our method is a novel dual-level cache mechanism designed to maintain both intra-shot coherence and inter-shot consistency. This is complemented by a strategic two-stage training process and a meticulous data curation pipeline to enable robust learning.

Autoregressive Generation with Dual-Level Cache

As illustrated in Figure 2, our framework’s central innovation is a dual-level cache that provides contextual information to the diffusion model. This framework comprises two key components: a Temporal Cache, which serves as a short-term memory to ensure the fluid coherence of generated content within a single shot, and a Shot Cache, which provides long-term memory to maintain the consistency of core concepts, such as characters and backgrounds across different shots. Both caches are injected into the model via in-context learning. This modifies the conditioning of the denoiser, which now takes the text prompt \mathbf{c}_{text} , the Temporal Cache C_{temp} , and the Shot Cache C_{shot} as input. The model is thus trained with the following objective:

$$\mathcal{L} = \mathbb{E}_{\mathbf{v}_0, \mathbf{c}_{\text{text}}, \epsilon, t} \left[\|\epsilon - \epsilon_{\theta}(\mathbf{v}_t, t, \mathbf{c}_{\text{text}}, C_{\text{temp}}, C_{\text{shot}})\|^2 \right]. \quad (1)$$

This approach conditions the generation process on relevant past information without altering the model architecture.

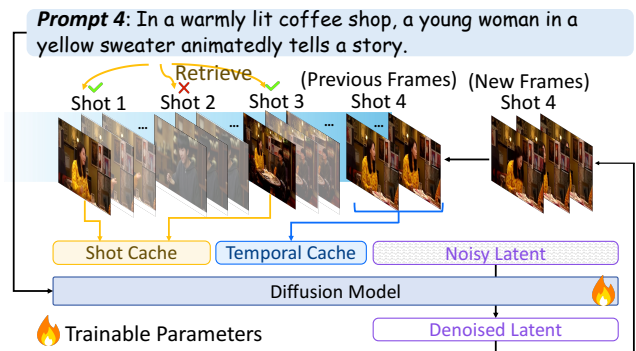


Figure 2: The framework of FilmWeaver. New video frames are generated autoregressively and consistency is enforced via a dual-level cache mechanism: a Shot Cache for long-term concept memory, populated through prompt-based keyframes retrieval from past shots, and a Temporal Cache for intra-shot coherence.

Temporal Cache for Intra-Shot Coherence. To ensure seamless continuity within a single shot, we employ a Temporal Cache. This cache functions as a sliding window, storing conditioning frames from the immediate past of the current generation chunk. As new frames are generated, they are added to the cache, while the oldest frames are discarded. Given the high temporal redundancy in videos, storing all past frames is computationally prohibitive. Therefore, inspired by recent works (Gu, Mao, and Shou 2025; Zhang and Agrawala 2025), we implement a differential compression strategy. Frames closer to the current generation window are preserved with higher fidelity, while frames further in the past are progressively compressed. This approach efficiently retains essential motion and content information while minimizing computational overhead.

Shot Cache for Inter-Shot Consistency. To ensure inter-shot consistency, we introduce the Shot Cache. When generating a new shot, this cache is built by selecting the top-K keyframes from previous shots that are most semantically relevant to the new text prompt. Relevance is determined by computing the cosine similarity between the CLIP embeddings (Radford et al. 2021) of the prompt and each candidate keyframe. The resulting cache provides a concise yet highly relevant visual summary of the narrative history, guiding the model to maintain consistency. The retrieval process is formulated as:

$$C_{\text{shot}} = \arg \text{top-k} (\text{sim}(\phi_T(\mathbf{c}_{\text{text}}), \phi_I(kf))), \quad (2)$$

where C_{shot} is the resulting Shot Cache, \mathcal{KF} is the set of all keyframes (kf) from previous shots, and \mathbf{c}_{text} is the text prompt for the new shot. Furthermore, ϕ_T and ϕ_I are the pretrained CLIP text and image encoders, respectively, $\text{sim}(\cdot, \cdot)$ denotes the cosine similarity function, and the $\arg \text{top-k}$ operator selects the K keyframes with the highest similarity scores.

Inference Stages and Modes As illustrated in Figure 3, our multi-shot generation process unfolds across four se-

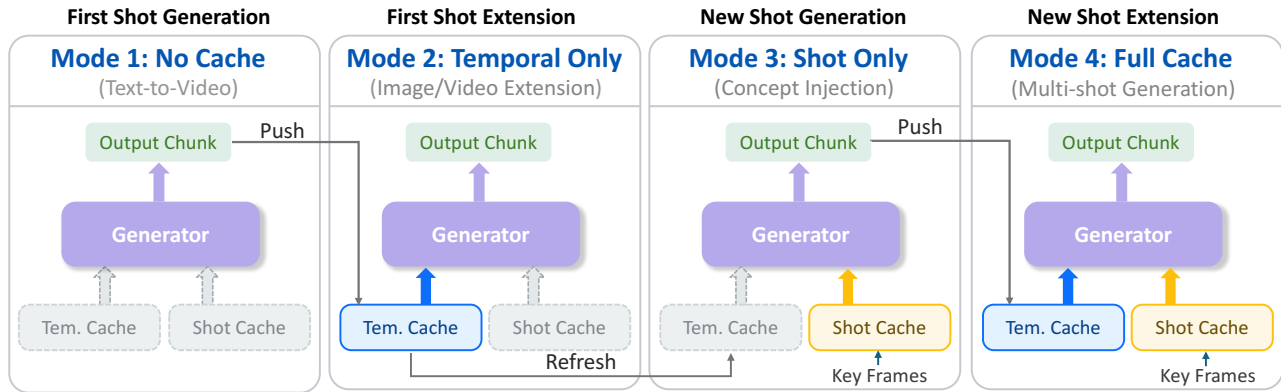


Figure 3: The multi-shot generation process involves four distinct scenarios, each corresponding to a different inference mode: (1) No Cache initializes generation and populates the caches; (2) Temporal Only extends the current shot using temporal context; (3) Shot Only transitions to a new shot by injecting key frames from previous shots; and (4) Full Cache combines both contexts.

quential stages. Each stage is governed by a distinct inference mode, determined by the state of our dual-level cache, enabling versatile and consistent video creation.

- **First Shot Generation (No Cache).** The process begins with both caches empty ($C_{\text{temp}} = \emptyset, C_{\text{shot}} = \emptyset$). The model operates as a standard text-to-video (T2V) generator, creating the initial video chunk from the text prompt and populating the caches for subsequent stages.
- **First Shot Extension (Temporal Only).** To generate subsequent chunks within the same shot, the model leverages the populated Temporal Cache ($C_{\text{temp}} \neq \emptyset, C_{\text{shot}} = \emptyset$). This ensures high temporal coherence and fluid motion, ideal for applications like video extension or image-to-video (I2V) generation.
- **New Shot Generation (Shot Only).** To transition to a new shot, the Temporal Cache is cleared while the Shot Cache is populated with keyframes from previous shots ($C_{\text{temp}} = \emptyset, C_{\text{shot}} \neq \emptyset$). This mode generates a new scene that maintains long-term consistency of key elements like characters and backgrounds. We can also artificially set reference frames to achieve multiple concept injection.
- **New Shot Extension (Full Cache).** When extending the new shot, both caches are active ($C_{\text{temp}} \neq \emptyset, C_{\text{shot}} \neq \emptyset$). The generator leverages the Temporal Cache for immediate coherence and the Shot Cache for long-term consistency, seamlessly blending both contexts.

Our training process is designed to accommodate all four scenarios, ensuring the model operates robustly across the entire multi-shot generation sequence.

Training Strategy

To simplify the learning task and promote stable convergence, we adopt a two-stage training strategy, coupled with data augmentation techniques to reduce the model’s over-reliance on the cache.

Progressive Training Curriculum. Our strategy unfolds in two progressive stages. The first stage focuses exclusively

on teaching the model to generate long, coherent single-shot videos with only text input. During this phase, the Shot Cache is disabled (i.e., its inputs are zeroed out), and the model is trained only with the Temporal Cache. This foundational step allows the model to master intra-shot dynamics without the added complexity of cross-shot consistency. Building on this foundation, the second stage activates the Shot Cache and fine-tunes the model on a mixed curriculum that includes all four cache-based scenarios. This progressive approach, where the model is pre-trained on the simpler long-video task, allows it to tackle the more complex challenge of multi-shot consistency with greater efficiency and faster convergence.

Data Augmentation. We observe that during training, the model develops an over-reliance on the provided visual context, leading to a “copy-paste” phenomenon. This behavior significantly reduces motion dynamism and diminishes the model’s adherence to the textual prompt. To mitigate this, we employ several targeted regularization techniques. First, we apply negative sampling to the Shot Cache by randomly introducing irrelevant keyframes during training. This compels the model to discriminate between useful and distracting context, guided by the prompt. Furthermore, inspired by prior strategies, we introduce noise to both caches to discourage exact replication. However, we find that excessive noise in the Temporal Cache degrades video coherence. We therefore adopt an asymmetric noising strategy: a substantial noise level, corresponding to a random diffusion timesteps between 100 to 400, is applied to the Shot Cache, while a much milder noise level (0–100 timesteps) is applied to the Temporal Cache. This dual approach effectively curbs over-reliance on visual context, thereby enhancing the ability of prompt following and generative quality.

Multi-Shot Data Curation

A major obstacle in multi-shot video generation is the lack of high-quality, consistently annotated datasets. To address this, we developed a comprehensive data curation pipeline as illustrated in Figure 5. For the initial Shot Splitting stage,



Figure 4: Examples of FilmWeaver’s downstream applications. Our framework enables multi-concept Injection (Top), generating dynamic videos from reference images while faithfully preserving subject identity. It also excels at complex video extension tasks (Bottom), where the narrative is steered mid-sequence by modifying a text prompt. The seamless transition from a “snowboard” to a “surfboard” scene, while maintaining the athlete’s appearance.

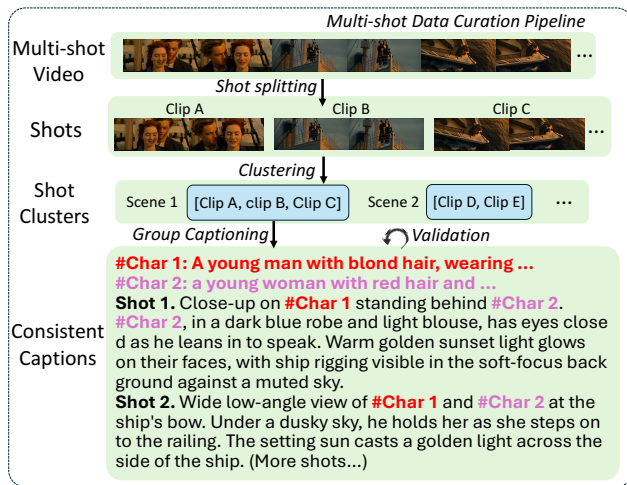


Figure 5: The pipeline of Multi-shot data curation, which first segments videos into shots and clusters them into coherent scenes. We then introduce a Group Captioning strategy that jointly describes all shots within a scene, enforcing consistent attributes for characters and objects. This process, finalized with a validation step, yields a high-quality dataset of video-text pairs with strong temporal coherence.

we employ an expert model (Chen et al. 2024b) to segment the source videos into individual shots. These shots are then grouped into coherent multi-shot sequences through Scene Clustering, a process that utilizes a sliding window to evaluate the CLIP similarity between clips. Following clustering, scenes undergo a rigorous filtering protocol to ensure data quality, removing clips shorter than one second and scenes featuring more than three distinct individuals. Our annotation pipeline begins with a Group Captioning strategy. We provide the entire sequence of shots from a scene to Gemini 2.5 Pro, which generates descriptions for all shots simultaneously. This joint-captioning approach can maintain consistent descriptions of the same character across different shots, while also accurately reflecting cinematic changes like transitions from a wide shot to a close-up. Subsequently,

each generated annotation undergoes a Validation step. In this stage, each individual shot and its corresponding caption are fed back into the model to verify its accuracy and refine any ambiguous or generic phrasing.

Extended Applications

The architecture of FilmWeaver enables a versatile range of downstream applications. As illustrated in Figure 4, our framework’s can easily handle multi-concept injection and video extension. First, for multi-concept injection (Top), we manually populate the Shot Cache with provided concept images, which denotes the Mode 3 in Figure 3. FilmWeaver can then generate a coherent, dynamic scene that seamlessly integrates these multiple subjects while faithfully preserving their individual identities. Furthermore, the framework excels at Dynamic Video Extension (Bottom). This is achieved by modifying the text prompt mid-sequence while preserving the Temporal Cache to maintain continuity. This process corresponds to Mode 2 (Temporal Only) in Figure 3, allowing the narrative to be steered on the fly. For instance, it can fluidly transition a video from the context of “an athlete snowboarding on snowy slopes” to “surfing on ocean waves” in direct response to the new prompt. The model’s ability to leverage its distinct cache modes for such sophisticated tasks highlights its potential for advanced video editing and creative narrative exploration.

Experiments

Implementation Details

For our implementation, we adapt the HunyuanVideo (Kong et al. 2024). The batch size is set to 16. The first stage is trained for 10K steps, and the second stage is trained for 10K steps.

Comparison with Existing Methods

To evaluate FilmWeaver’s performance, we compare it against two main categories of existing multi-shot video generation methods. The first category utilized the full pipeline to generate multi-shot videos, as represented by VideoStudio (Long et al. 2024). The second category denoted key-frames based methods comprises two-stage

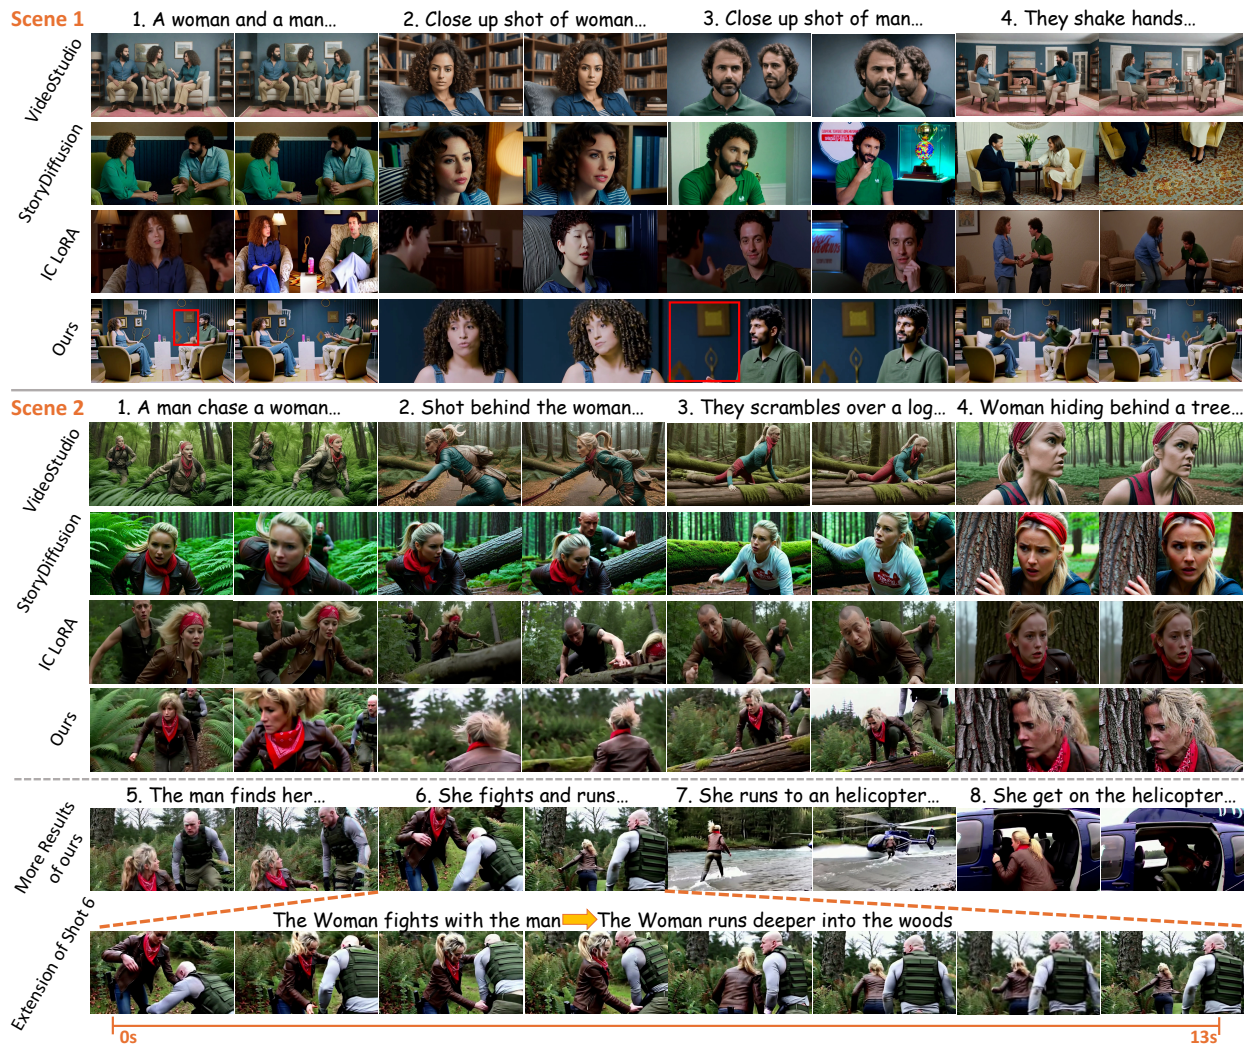


Figure 6: Qualitative comparison of multi-shot video generation. We compare two scenarios where existing methods tend to confuse character features and fail to maintain background consistency, while our method handles these challenges effectively. Additionally, we showcase more shots generated by our framework in Scene 2, forming a continuous narrative. For the sixth shot in Scene 2, we demonstrate our video extension capability by creating a coherent long video using two prompts.

pipelines that first generate keyframes and then animate them using an I2V model. For these pipelines, we use StoryDiffusion (Zhou et al. 2024) and IC-LoRA (Huang et al. 2024) for consistent keyframe generation, followed by the powerful Hunyuan I2V model for the animation step. Since IC-LoRA generates a limited number of images per run, we invoke it multiple times to produce the required keyframes. Given that VideoStudio and IC-LoRA require structured inputs, we use Gemini 2.5 Pro (Comanici et al. 2025) to reformat the prompts accordingly.

Qualitative Comparison. We present a comprehensive qualitative comparison in Figure 6, evaluating our method against leading approaches across two distinct and challenging narrative scenarios. In Scene 1, a conversation scene with alternating wide and close-up shots, competing meth-

ods exhibit severe consistency failures. They struggle with multi-person identity preservation, leading to feature entanglement where attributes like clothing and facial structures are blended between characters and their backgrounds are inconsistent. In contrast, our method successfully preserves the distinct appearance of each individual and maintains a stable background across all shots. As highlighted by the red box in the fourth line, details such as the wall art behind the man in Shot 3 remain perfectly consistent with the Shot 1. Similarly, our method robustly preserves character identities during the dynamic action in Scene 2, overcoming the appearance inconsistencies common to competitor models.

Crucially, our framework’s capabilities extend far beyond short, four-shot sequences. We demonstrate this by generating a continuous and coherent 8-shot narrative for Scene 2, showcasing its strength in long-form storytelling. Moreover,

we explicitly highlight our video extension capability within “Extension of Shot 6,” where we use a second prompt to seamlessly continue the action within a single shot, creating a longer and more dynamic clip. These results underscore our model’s superior ability to generate consistent, controllable, and lengthy narrative-driven video content. For additional examples, please refer to the appendix.

Method	Visual Quality		Consistency(%)		Text Alignment	
	Aes.↑	Incep.↑	Char.↑	All.↑	Char. ↑	All. ↑
VideoStudio	32.02	6.81	73.34	62.40	20.88	31.52
StoryDiffusion	35.61	<u>8.30</u>	70.03	67.15	20.21	30.86
IC-LoRA	31.78	6.95	72.47	71.19	<u>22.16</u>	28.74
Ours	<u>33.69</u>	8.57	74.61	75.12	23.07	<u>31.23</u>

Table 1: Quantitative comparison with existing methods.

Quantitative Comparison. These visual observations are substantiated by our quantitative evaluation, summarized in Table 1. For quantitative assessment, we evaluate performance across three key dimensions: Visual Quality, Consistency, and Text Alignment. For Visual Quality, we adopt the evaluation protocol from MovieBench (Wu et al. 2025), which contains Aesthetics Score (Aes.) and Inception Score (Incep.). For Consistency and Text Alignment, we report metrics for both character-specific (Char.) and overall (All.) aspects.

Quantitative results in Table 1 validate our method’s superiority. We achieve state-of-the-art results in both Consistency metrics and character-level Text Alignment. Crucially, our method also attains the highest Inception Score, demonstrating top-tier visual quality, and remains highly competitive across all other metrics. We posit that the visual quality of our method can be further enhanced through improved data curation and optimized training strategies.

Method	Visual Quality		Consistency(%)		Text Alignment	
	Aes.↑	Incep.↑	Char.↑	All.↑	Char. ↑	All. ↑
w/o A	30.04	7.77	72.36	75.92	21.88	28.12
w/o S	33.92	8.63	68.11	65.44	22.41	31.79
w/o T	31.61	8.36	70.79	70.57	20.21	30.70
Ours	33.69	8.57	74.61	75.12	23.07	31.23

Table 2: Quantitative results of ablation study.

Ablation Studies

We conduct ablation studies to validate our core components: the Shot Cache (S), the Temporal Cache (T), and the noise augmentation (A) strategy. As shown in Figure 7, removing the Shot Cache (w/o S) leads to a severe loss of long-term visual consistency, failing to preserve character appearance and scene style. Removing the Temporal Cache (w/o T) results in temporally incoherent and disjointed motion. Furthermore, Figure 8 demonstrates that without noise augmentation (w/o A), the model’s over-reliance on visual context



Figure 7: Qualitative ablation study of our dual-level cache. Without the shot cache (w/o S), the model fails to maintain visual style and the clothes of character. Without the temporal cache (w/o T), the generated sequence lacks coherence, resulting in disjointed motion. Our full method successfully preserves both appearance and motion continuity.



Figure 8: Qualitative ablation study on noise augmentation. Without noise augmentation, the model over-relies on past frames, hindering the ability of prompt following, which is crucial in video extension. Applying noise reduces this dependency and improves the ability of prompt following.

from the Temporal Cache hinders its ability to adapt to dynamic prompts, thus reducing prompt controllability during video extension. Table 2 quantitatively corroborates these findings. The variants without the caches show a significant drop in consistency metrics, while the absence of noise augmentation leads to a notable decrease in text alignment scores. These results confirm that each component is essential: the Shot Cache for inter-shot consistency, the Temporal Cache for intra-shot coherence, and noise augmentation for robust prompt adherence.

Conclusion

We present FilmWeaver, a novel cache-guided autoregressive framework addressing consistency and duration in multi-shot video generation. Our key contribution is a dual-level cache mechanism combining a retrieval-based Shot Cache for long-term consistency with a Temporal Cache for intra-shot motion coherence. Supported by noise augmentation for better prompt adherence, this architecture enables generating coherent videos of arbitrary length. Experiments on a curated dataset show FilmWeaver significantly outperforms existing methods in visual consistency and aesthetic quality. Its flexibility allows direct application to tasks like video extension. FilmWeaver represents a substantial advance in creating complex, controllable narrative videos, opening new possibilities for automated storytelling.

References

- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; et al. 2023. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv preprint arXiv:2311.15127*.
- Cai, M.; Cun, X.; Li, X.; Liu, W.; Zhang, Z.; Zhang, Y.; Shan, Y.; and Yue, X. 2025. Ditctrl: Exploring attention control in multi-modal diffusion transformer for tuning-free multi-prompt longer video generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 7763–7772.
- Chen, B.; Martí Monsó, D.; Du, Y.; Simchowitz, M.; Tedrake, R.; and Sitzmann, V. 2024a. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural Information Processing Systems*, 37: 24081–24125.
- Chen, G.; Lin, D.; Yang, J.; Lin, C.; Zhu, J.; Fan, M.; Zhang, H.; Chen, S.; Chen, Z.; Ma, C.; et al. 2025. Skyreels-v2: Infinite-length film generative model. *arXiv preprint arXiv:2504.13074*.
- Chen, T.-S.; Siarohin, A.; Menapace, W.; Deyneka, E.; Chao, H.-w.; Jeon, B. E.; Fang, Y.; Lee, H.-Y.; Ren, J.; Yang, M.-H.; et al. 2024b. Panda-70m: Captioning 70m videos with multiple cross-modality teachers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13320–13331.
- Cheng, J.; Lu, X.; Li, H.; Zai, K. L.; Yin, B.; Cheng, Y.; Yan, Y.; and Liang, X. 2024. Autostudio: Crafting consistent subjects in multi-turn interactive image generation. *arXiv preprint arXiv:2406.01388*.
- Comanici, G.; Bieber, E.; Schaekermann, M.; Pasupat, I.; Sachdeva, N.; Dhillon, I.; Blistein, M.; Ram, O.; Zhang, D.; Rosen, E.; et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Dalal, K.; Kocejka, D.; Xu, J.; Zhao, Y.; Han, S.; Cheung, K. C.; Kautz, J.; Choi, Y.; Sun, Y.; and Wang, X. 2025. One-minute video generation with test-time training. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 17702–17711.
- Gu, Y.; Mao, W.; and Shou, M. Z. 2025. Long-context autoregressive video modeling with next-frame prediction. *arXiv preprint arXiv:2503.19325*.
- 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*.
- Guo, Y.; Yang, C.; Yang, Z.; Ma, Z.; Lin, Z.; Yang, Z.; Lin, D.; and Jiang, L. 2025. Long context tuning for video generation. *arXiv preprint arXiv:2503.10589*.
- Huang, L.; Wang, W.; Wu, Z.-F.; Shi, Y.; Dou, H.; Liang, C.; Feng, Y.; Liu, Y.; and Zhou, J. 2024. In-context lora for diffusion transformers. *arXiv preprint arXiv:2410.23775*.
- Huang, Y.; Yuan, Z.; Liu, Q.; Wang, Q.; Wang, X.; Zhang, R.; Wan, P.; Zhang, D.; and Gai, K. 2025. Conceptmaster: Multi-concept video customization on diffusion transformer models without test-time tuning. *arXiv preprint arXiv:2501.04698*.
- Jin, Y.; Sun, Z.; Li, N.; Xu, K.; Jiang, H.; Zhuang, N.; Huang, Q.; Song, Y.; Mu, Y.; and Lin, Z. 2024. Pyramidal flow matching for efficient video generative modeling. *arXiv preprint arXiv:2410.05954*.
- Kara, O.; Singh, K. K.; Liu, F.; Ceylan, D.; Rehg, J. M.; and Hinz, T. 2025. ShotAdapter: Text-to-Multi-Shot Video Generation with Diffusion Models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 28405–28415.
- Kong, W.; Tian, Q.; Zhang, Z.; Min, R.; Dai, Z.; Zhou, J.; Xiong, J.; Li, X.; Wu, B.; Zhang, J.; et al. 2024. Hunyuan-video: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*.
- Lin, H.; Zala, A.; Cho, J.; and Bansal, M. 2023. Videodirector: Consistent multi-scene video generation via llm-guided planning. *arXiv preprint arXiv:2309.15091*.
- Long, F.; Qiu, Z.; Yao, T.; and Mei, T. 2024. VideoStudio: Generating Consistent-Content and Multi-Scene Videos. In *European Conference on Computer Vision*, 468–485. Springer.
- Luo, X.; Cheng, J.; Xie, Y.; Zhang, X.; Feng, T.; Liu, Z.; Ma, F.; and Yu, F. 2025a. Object Isolated Attention for Consistent Story Visualization. *arXiv preprint arXiv:2503.23353*.
- Luo, X.; Zhu, Y.; Liu, Y.; Lin, L.; Wan, C.; Cai, Z.; Li, Y.; and Huang, S.-L. 2025b. CanonSwap: High-Fidelity and Consistent Video Face Swapping via Canonical Space Modulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10064–10074.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Qi, T.; Yuan, J.; Feng, W.; Fang, S.; Liu, J.; Zhou, S.; He, Q.; Xie, H.; and Zhang, Y. 2025. Mask²DiT: Dual Mask-based Diffusion Transformer for Multi-Scene Long Video Generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 18837–18846.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- 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.
- Tewel, Y.; Kaduri, O.; Gal, R.; Kasten, Y.; Wolf, L.; Chechik, G.; and Atzmon, Y. 2024. Training-free consistent text-to-image generation. *ACM Transactions on Graphics (TOG)*, 43(4): 1–18.

Voleti, V.; Yao, C.-H.; Boss, M.; Letts, A.; Pankratz, D.; Tochilkin, D.; Laforte, C.; Rombach, R.; and Jampani, V. 2024. Sv3d: Novel multi-view synthesis and 3d generation from a single image using latent video diffusion. In *European Conference on Computer Vision*, 439–457. Springer.

Wan, C.; Luo, X.; Luo, H.; Cai, Z.; Song, Y.; Zhao, Y.; Bai, Y.; He, Y.; and Gong, Y. 2024. Grid: Omni Visual Generation. *arXiv preprint arXiv:2412.10718*.

Wan, T.; Wang, A.; Ai, B.; Wen, B.; Mao, C.; Xie, C.-W.; Chen, D.; Yu, F.; Zhao, H.; Yang, J.; et al. 2025. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*.

Wang, J.; Sheng, H.; Cai, S.; Zhang, W.; Yan, C.; Feng, Y.; Deng, B.; and Ye, J. 2025. EchoShot: Multi-Shot Portrait Video Generation. *arXiv preprint arXiv:2506.15838*.

Wu, W.; Liu, M.; Zhu, Z.; Xia, X.; Feng, H.; Wang, W.; Lin, K. Q.; Shen, C.; and Shou, M. Z. 2025. Moviebench: A hierarchical movie level dataset for long video generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 28984–28994.

Wu, W.; Zhu, Z.; and Shou, M. Z. 2025. Automated movie generation via multi-agent cot planning. *arXiv preprint arXiv:2503.07314*.

Xiao, J.; Yang, C.; Zhang, L.; Cai, S.; Zhao, Y.; Guo, Y.; Wetzstein, G.; Agrawala, M.; Yuille, A.; and Jiang, L. 2025. Captain Cinema: Towards Short Movie Generation. *arXiv preprint arXiv:2507.18634*.

Xue, H.; Luo, X.; Hu, Z.; Zhang, X.; Xiang, X.; Dai, Y.; Liu, J.; Zhang, Z.; Li, M.; Yang, J.; et al. 2024. Human motion video generation: A survey. *Authorea Preprints*.

Yang, Z.; Teng, J.; Zheng, W.; Ding, M.; Huang, S.; Xu, J.; Yang, Y.; Hong, W.; Zhang, X.; Feng, G.; et al. 2024. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*.

Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.

Yuan, Z.; Liu, Y.; Cao, Y.; Sun, W.; Jia, H.; Chen, R.; Li, Z.; Lin, B.; Yuan, L.; He, L.; et al. 2024. Mora: Enabling generalist video generation via a multi-agent framework. *arXiv preprint arXiv:2403.13248*.

Zhang, L.; and Agrawala, M. 2025. Packing input frame context in next-frame prediction models for video generation. *arXiv preprint arXiv:2504.12626*.

Zhou, Y.; Zhou, D.; Cheng, M.-M.; Feng, J.; and Hou, Q. 2024. Storydiffusion: Consistent self-attention for long-range image and video generation. *Advances in Neural Information Processing Systems*, 37: 110315–110340.