

Motion-Adaptive Temporal Attention for Lightweight Video Generation with Stable Diffusion

Rui Hong; George Mason University; Fairfax, VA, USA
Shuxue Quan; Independent Researcher

Abstract

We present a motion-adaptive temporal attention mechanism for parameter-efficient video generation built upon frozen Stable Diffusion models. Rather than treating all video content uniformly, our method dynamically adjusts temporal attention receptive fields based on estimated motion content: high-motion sequences attend locally across frames to preserve rapidly changing details, while low-motion sequences attend globally to enforce scene consistency. We inject lightweight temporal attention modules into all UNet transformer blocks via a cascaded strategy—global attention in down-sampling and middle blocks for semantic stabilization, motion-adaptive attention in up-sampling blocks for fine-grained refinement. Combined with temporally correlated noise initialization and motion-aware gating, the system adds only 25.8M trainable parameters (2.9% of the base UNet) while achieving competitive results on WebVid validation when trained on 100K videos. We demonstrate that the standard denoising objective alone provides sufficient implicit temporal regularization, outperforming approaches that add explicit temporal consistency losses. Our ablation studies reveal a clear trade-off between noise correlation and motion amplitude, providing a practical inference-time control for diverse generation behaviors.

Introduction

Text-to-video generation has emerged as a frontier challenge in generative modeling. While text-to-image diffusion models such as Stable Diffusion [1] have achieved remarkable quality, extending them to temporally coherent video remains difficult. The core challenge is twofold: maintaining visual consistency across frames while producing meaningful motion that aligns with the text prompt.

Existing approaches to video diffusion fall broadly into two categories. Full video diffusion models [2, 3, 4] train end-to-end on large-scale video datasets, achieving strong results but at enormous computational cost. Parameter-efficient methods such as AnimateDiff [6] instead inject trainable temporal layers into frozen image diffusion models, learning motion patterns from video data while preserving the spatial generation quality of the base model. However, these methods treat all video content uniformly, applying identical temporal attention regardless of motion characteristics. This one-size-fits-all strategy is suboptimal: a slow panning landscape and a fast-moving animal require fundamentally different temporal modeling.

We propose a **Motion-Adaptive Temporal UNet** that addresses this limitation through three key innovations. First, we introduce **motion-adaptive attention bias**, where the temporal attention mechanism dynamically adjusts its receptive field based on estimated motion content—high-motion sequences attend to

nearby frames to preserve detail, while low-motion sequences attend globally to enforce consistency. Second, we employ **motion-aware gating** that modulates the strength of temporal contributions at each network layer according to motion intensity. Third, we adopt a **cascaded injection strategy** in which down-sampling and middle blocks apply global temporal stabilization, while up-sampling blocks apply motion-adaptive local refinement.

Our temporal modules are injected into a frozen Stable Diffusion UNet, adding only 25.8M trainable parameters (2.9% of the base model). Combined with temporally correlated noise initialization, the approach produces temporally coherent videos while adapting to diverse motion patterns. Training is conducted on a 100K subset of WebVid [11], demonstrating strong parameter efficiency compared to methods such as AnimateDiff, whose motion modules require 417M parameters. Code is available at <https://github.com/hongrui16/motion-adaptive-video>.

Related Work

Latent Diffusion Models

Latent Diffusion Models (LDMs) [1] perform the diffusion process in a compressed latent space learned by a variational autoencoder (VAE), greatly reducing computational cost compared to pixel-space diffusion. Stable Diffusion, an open-source LDM, uses a UNet architecture with spatial self-attention and text-conditioned cross-attention to generate high-quality images from text prompts. Our method builds upon this frozen image generation backbone.

Video Diffusion Models

Early video diffusion models [2, 3] extended image diffusion to video by jointly modeling spatial and temporal dimensions, typically requiring end-to-end training on large video corpora. Video LDM [4] demonstrated that temporal layers could be added to pretrained image LDMs, training only the temporal components while keeping spatial layers frozen. This paradigm was further developed by AnimateDiff [6], which proposed domain-agnostic motion modules that can be combined with various personalized text-to-image models. More recent work such as Stable Video Diffusion [5] and ModelScopeT2V [7] has continued to advance temporal modeling within the latent diffusion framework.

Motion-Aware Generation

Several works have explored conditioning video generation on motion signals. MotionCtrl [8] conditions on camera trajectories, while DragAnything [9] allows point-based motion control. MCDiff [10] uses stroke-based motion conditioning for controllable video synthesis. However, these methods require explicit

motion inputs at inference time. Our approach instead estimates motion characteristics from the content itself and adapts the temporal modeling accordingly, requiring no additional user input beyond the text prompt.

Method Overview

Figure 1 illustrates the overall pipeline. Our method builds upon a pretrained Stable Diffusion model, which operates in the latent space of a frozen VAE. During training, input video frames are reshaped to $\mathbf{x} \in \mathbb{R}^{(B \cdot T) \times C \times H \times W}$ (batch size B , $T=8$ frames, $C=3$ RGB channels, $H=W=256$) and encoded by the VAE encoder \mathcal{E} into latent representations $\mathbf{z}_0 \in \mathbb{R}^{(B \cdot T) \times c \times h \times w}$, where $h=w=32$, $c=4$. Noise is added via the forward diffusion process, and a UNet denoiser ε_θ learns to predict the added noise. Text conditioning is provided through a frozen CLIP text encoder via cross-attention. At inference, the denoising process iteratively removes noise from random latents, and the VAE decoder \mathcal{D} reconstructs video frames from the denoised latents.

Our key modification is the injection of lightweight, trainable temporal attention modules into the UNet blocks. Within each UNet transformer block, the original frozen spatial processing (spatial self-attention, text cross-attention, and feed-forward layers) is preserved as a black box. After the spatial transformer produces per-frame features, our temporal attention block operates across the frame dimension, enabling inter-frame communication at every denoising step. This is wrapped in a `SpatialTemporalWrapper` that is transparent to the rest of the UNet.

Temporal Attention Block

Each injected temporal attention block operates on the intermediate UNet feature tensor after spatial processing. Let $\mathbf{z} \in \mathbb{R}^{(B \cdot T) \times c \times h \times w}$ denote these features, which are reshaped to $\tilde{\mathbf{z}} \in \mathbb{R}^{(B \cdot h \cdot w) \times T \times c}$ so that attention is computed across the temporal dimension T independently at each spatial position. Figure 2 illustrates the two variants.

Global Temporal Attention

The global block applies standard multi-head temporal self-attention uniformly across all frames. Learnable temporal positional encodings $\mathbf{P} \in \mathbb{R}^{1 \times T \times c}$ are added to the input before attention, giving $\hat{\mathbf{z}} = \tilde{\mathbf{z}} + \mathbf{P}$. The block then computes:

$$\mathbf{z}^{(1)} = \tilde{\mathbf{z}} + \text{Attn}(\text{LN}(\hat{\mathbf{z}})) \quad (1)$$

$$\mathbf{z}^{(2)} = \mathbf{z}^{(1)} + \text{FFN}(\text{LN}(\mathbf{z}^{(1)})) \quad (2)$$

where $\text{LN}(\cdot)$ is layer normalization, $\text{Attn}(\cdot)$ is multi-head self-attention along the temporal dimension T , and $\text{FFN}(\cdot)$ is a two-layer feed-forward network. The output $\mathbf{z}^{(2)}$ is reshaped back to $(B \cdot T) \times c \times h \times w$.

Adaptive Temporal Attention

The adaptive block extends the global formulation with two motion-conditioned components, given motion score $m \in [0, 1]$ estimated from the video content.

Motion-adaptive attention bias. A distance-based bias matrix $\mathbf{b}_m \in \mathbb{R}^{T \times T}$ is injected into the attention logits, where

$i, j \in \{1, \dots, T\}$ index the query and key frames:

$$\mathbf{b}_m(i, j) = -\lambda(m) \cdot |i - j|, \quad \lambda(m) = 0.05 + 0.45m \quad (3)$$

$$\mathbf{z}^{(1)} = \tilde{\mathbf{z}} + \text{Attn}(\text{LN}(\hat{\mathbf{z}}), \mathbf{b}_m) \quad (4)$$

where $\lambda(m)$ is the motion-dependent decay coefficient. When $m \rightarrow 1$ (high motion), λ is large and attention decays steeply with frame distance, focusing each position on nearby frames. When $m \rightarrow 0$ (low motion), λ is small and attention spans all frames globally.

Motion-aware gating. After the feed-forward step (identical to Eq. 2), a learned scalar gate $g(m) \in (0, 1)$ modulates the strength of the temporal residual:

$$\mathbf{z}_{\text{out}} = \tilde{\mathbf{z}} + g(m) \cdot (\mathbf{z}^{(2)} - \tilde{\mathbf{z}}), \quad g(m) = \sigma(\phi(m)) \quad (5)$$

where $\sigma(\cdot)$ is the sigmoid function and $\phi(m)$ is a two-layer MLP with scalar output. The gate parameters are initialized such that $g(m) \approx 0$, so the block begins as a near-identity mapping and gradually learns its temporal contribution.

Zero Initialization

All output projection layers (W_O and the final FFN layer) are zero-initialized. Combined with the residual connections, this ensures that at initialization both block variants are exact identity mappings, following the principle established in ControlNet [16], and do not disturb the pretrained UNet weights at the start of training.

Cascaded Injection Strategy

We inject temporal attention into all transformer blocks of the UNet, but with a cascaded strategy that assigns different modes based on network depth. Down-sampling and middle blocks use **global mode**: standard temporal self-attention (no distance bias) stabilizes the semantic content and background consistency across frames at lower spatial resolutions. Up-sampling blocks use **adaptive mode**: motion-adaptive attention bias protects fine-grained features from temporal blurring, particularly in regions with significant motion. This design is motivated by the observation that low-resolution features encode scene-level semantics that should be consistent across all frames, while high-resolution features encode local details that are sensitive to inter-frame motion.

Temporally Correlated Noise

Standard diffusion training and inference use independently sampled Gaussian noise for each frame. However, when the UNet processes each frame largely independently (before temporal attention takes effect), independent noise leads to completely divergent scene content across frames.

We introduce temporally correlated noise via an autoregressive process:

$$\varepsilon_f = \rho \cdot \varepsilon_{f-1} + \sqrt{1 - \rho^2} \cdot \xi_f, \quad \xi_f \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (6)$$

where $f \in \{1, \dots, T\}$ is the frame index, $\rho \in [0, 1]$ is the correlation coefficient, and $\varepsilon_0 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Each ε_f remains marginally standard Gaussian, but adjacent frames share correlated noise

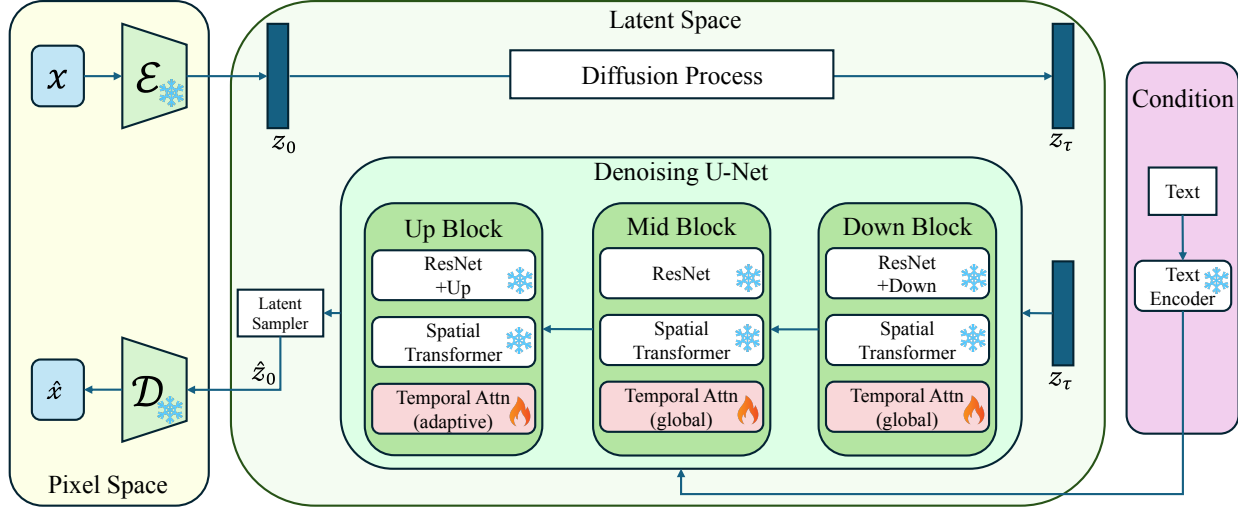


Figure 1. Architecture overview. The overall pipeline follows the Latent Diffusion Model framework: video frames are encoded into latent space by a frozen VAE encoder (\mathcal{E}), processed by the denoising UNet with injected temporal attention modules (shown in red), conditioned on text via CLIP, and the VAE decoder \mathcal{D} reconstructs the frames back to pixel space. All spatial components (ResNet, Spatial Transformer) are frozen (snowflake); only temporal attention blocks are trained (flame). Down and mid blocks use global temporal attention; up blocks use motion-adaptive attention.

structure. This provides a strong prior for scene consistency without constraining the learned temporal dynamics.

Correlated noise is applied during both training and inference. The correlation coefficient ρ controls a trade-off between scene consistency and motion amplitude, with $\rho = 0.6$ used during training and varied at inference for different generation behaviors (see ablation in Section).

Training Objective

The primary training objective is the standard denoising loss:

$$\mathcal{L}_{\text{denoise}} = \mathbb{E}_{\mathbf{z}_0, \varepsilon, t} \left[\|\varepsilon - \varepsilon_{\theta}(\mathbf{z}_t, t, \mathbf{c})\|^2 \right] \quad (7)$$

where \mathbf{z}_0 are the clean latent encodings of video frames, t is the diffusion timestep, $\varepsilon = [\varepsilon_1, \dots, \varepsilon_T]$ is the temporally correlated noise from Eq. 6, and $\mathbf{z}_t = \sqrt{\bar{\alpha}_t} \mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_t} \varepsilon$ is the noisy latent at timestep t . \mathbf{c} is the text conditioning. Only the temporal attention parameters are updated; the base UNet, VAE, and text encoder remain frozen.

We additionally explored a temporal consistency loss that penalizes inter-frame differences in the predicted clean latents, weighted by both motion intensity and noise level:

$$\mathcal{L}_{\text{temp}} = \mathbb{E}_{t, f} \left[w_m \cdot w_t \cdot \|\hat{\mathbf{z}}_0^{(f+1)} - \hat{\mathbf{z}}_0^{(f)}\|^2 \right] \quad (8)$$

where $f \in \{1, \dots, T-1\}$ is the frame index, $\hat{\mathbf{z}}_0^{(f)} = (\mathbf{z}_t^{(f)} - \sqrt{1 - \bar{\alpha}_t} \varepsilon_f^{\theta}) / \sqrt{\bar{\alpha}_t}$ is the clean latent of frame f estimated via the posterior mean, $w_m = 0.1 + 0.9 \cdot (1 - m)$ assigns higher weight to low-motion videos, and $w_t = \bar{\alpha}_t$ down-weights the loss at high noise levels. The total loss is:

$$\mathcal{L} = \mathcal{L}_{\text{denoise}} + \lambda_{\text{temp}} \cdot \mathcal{L}_{\text{temp}} \quad (9)$$

As shown in our ablation (Section), we find that the denoising loss alone ($\lambda_{\text{temp}} = 0$) produces better results than including the temporal consistency term, as the temporal attention modules learn sufficient implicit temporal regularization from the video training data.

Motion Estimation

During training, per-video motion scores are computed from the input frames as the mean absolute pixel difference between consecutive frames, normalized to $[0, 1]$ using dataset-level statistics (5th and 95th percentiles for robust normalization). At inference, the motion level is estimated automatically from prompt semantics—keywords such as “running” or “dancing” map to high motion, while “still” or “landscape” map to low—or can be set manually by the user.

Experiments

Setup

Architecture

Temporal attention blocks are injected into all 16 transformer blocks of the UNet: 6 in the down-sampling path, 1 in the middle block, and 9 in the up-sampling path. Each temporal block uses 4 attention heads with dimension 64 per head and a feed-forward inner dimension of 512. This adds 25.8M trainable parameters (2.9% of the 860M base UNet). We evaluate with two frozen base models: SD 1.5 (sd-legacy/stable-diffusion-v1-5) and SD 2.1-base (stabilityai/stable-diffusion-2-1-base), the 512×512 variant distinct from the 768×768 stable-diffusion-2-1 model. Our main comparison uses SD 2.1-base.

Training Data

We train on a 100K-video subset of WebVid-10M [11] for all experiments. We sample 8 frames per video at a frame interval

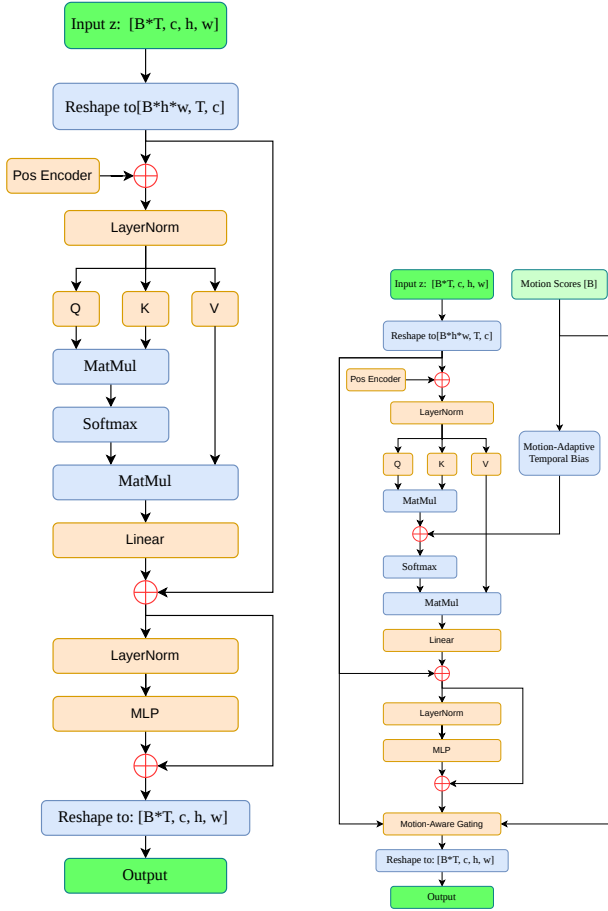


Figure 2. Temporal attention block variants. Left: Global mode (down/mid blocks) — uniform temporal self-attention, no motion conditioning. Right: Adaptive mode (up blocks) — motion-adaptive attention bias b_m and motion-aware gate $g(m)$ are both conditioned on the per-video motion score m .

of 2 and resize to 256×256 . Motion scores are pre-computed for the entire training set using robust percentile normalization.

Training Details

Training uses AdamW with learning rate 1×10^{-4} , cosine schedule with 500 warm-up steps, and gradient norm clipping at 1.0. We train for 30 epochs with batch size 55 per GPU across A100 GPUs using mixed-precision (FP16) and gradient checkpointing. The noise correlation coefficient is $\rho = 0.6$ during training.

Inference

We use DPM-Solver++ [15] with Karras noise schedule for 50 denoising steps, classifier-free guidance scale of 7.5, and resolution 256×256 with 8 frames per video. The noise correlation coefficient is adjustable at inference time without retraining.

Evaluation Metrics

We evaluate on 400 uniformly sampled captions from the WebVid validation set using the following metrics. **CLIP-SIM:**

Comparison with AnimateDiff on WebVid validation (400 samples, $\rho=0.6$, $\lambda_{temp}=0$, 100K training videos). † AnimateDiff re-trained on the same 100K WebVid subset with $T=8$, 256×256 , matching inference settings. ‡ Base SD 1.5 without temporal modules; high MM reflects incoherent frame-independent generation.

Method	CLIP \uparrow	FC	MM \uparrow	FID \downarrow	FVD \downarrow
AnimateDiff \dagger	0.308	0.969	0.062	81.63	79.61
SD 1.5 \ddagger	0.267	0.777	0.326	180.48	211.02
Ours (SD 1.5)	0.302	0.980	0.051	89.96	85.99
Ours (SD 2.1)	0.303	0.983	0.040	92.09	92.49

average CLIP [12] cosine similarity between each frame and its text prompt, measuring text-video alignment. **Frame Consistency (FC):** average CLIP cosine similarity between consecutive frame embeddings, measuring temporal coherence; a very low FC indicates incoherence, while an excessively high FC may reflect static output rather than good generation quality. **Motion Magnitude (MM):** mean absolute pixel difference between consecutive frames in pixel space, measuring motion amplitude. **FVD:** Fréchet Video Distance [13] measuring overall video quality against real WebVid videos. **FID:** Fréchet Inception Distance [14] computed per-frame, measuring single-frame image quality.

Comparison with AnimateDiff

Table 1 compares our method against AnimateDiff [6] under equal training conditions. To ensure a fair comparison, we retrain AnimateDiff’s motion modules on the same 100K-video WebVid subset with identical settings ($T=8$ frames, 256×256 , same evaluation protocol). We also include a baseline using SD 1.5 with correlated noise but without temporal modules to isolate the contribution of our temporal attention. Removing the temporal modules causes FID to rise from 89.96 to 180.48 and FVD from 85.99 to 211.02, while frame consistency drops sharply from 0.980 to 0.777, confirming that temporal attention is the primary driver of inter-frame coherence. AnimateDiff achieves better FID and FVD under these equal-data conditions, which we attribute to its larger motion module capacity (417M parameters). Importantly, our temporal modules add only 25.8M trainable parameters— $16 \times$ fewer—making our approach substantially more parameter-efficient.

Ablation Studies

Unless otherwise noted, ablation experiments use SD 2.1 trained on the 100K subset, evaluated on 400 WebVid validation captions. The base model comparison (Table 4) additionally includes SD 1.5 under identical settings.

Effect of Temporal Consistency Loss

Table 2 compares training with the combined loss ($\lambda_{temp} = 2.0$) versus denoising loss only ($\lambda_{temp} = 0$). Removing the explicit temporal loss produces videos with larger motion amplitude (0.040 vs. 0.025). The temporal attention mechanism provides sufficient implicit regularization through the denoising objective

Effect of temporal consistency loss (λ_{temp}). SD 2.1, 100K, $\rho=0.6$.

λ_{temp}	CLIP \uparrow	FC	MM \uparrow	FID \downarrow	FVD \downarrow
2.0	0.305	0.990	0.025	95.82	97.21
0	0.303	0.983	0.040	92.09	92.49

Effect of noise correlation ρ at inference. SD 2.1, 100K, $\lambda_{\text{temp}}=0$.

ρ	CLIP \uparrow	FC	MM \uparrow	FID \downarrow	FVD \downarrow
0.4	0.278	0.929	0.158	105.85	108.22
0.6	0.303	0.983	0.040	92.09	92.49
0.85	0.298	0.999	0.008	130.87	140.11

applied to temporally coherent video data, making the explicit consistency term unnecessary and even counterproductive for motion diversity.

Effect of Noise Correlation

Table 3 reports the effect of the noise correlation coefficient ρ at inference time, using a model trained with $\rho=0.6$. Higher correlation suppresses motion amplitude: at $\rho=0.85$, motion magnitude collapses to 0.008 (effectively static video). Lower correlation allows more diverse inter-frame content. This parameter provides a user-controllable trade-off at inference time without retraining.

Effect of Base Model

Table 4 compares Stable Diffusion 1.5 and 2.1 as base models under otherwise identical settings (100K training, $\rho=0.6$, $\lambda_{\text{temp}}=0$). SD 1.5 achieves slightly better FID and FVD, likely because the WebVid reference statistics are closer to SD 1.5’s training distribution, and produces slightly larger motion amplitude. Both models benefit substantially from temporal attention modules.

Qualitative Results

Figure 3 shows the first and last frame (frames 1 and 8) of generated sequences from our method and the retrained AnimateDiff baseline for two representative WebVid validation prompts. Displaying the temporal endpoints highlights scene consistency: a method with poor temporal coherence would show visible subject or background drift between frame 1 and frame 8. Our method maintains consistent subject identity and scene appearance across the full sequence, with natural motion appropriate to the content.

Discussion

Implicit vs. explicit temporal regularization. A key finding is that the denoising loss alone suffices for temporal learning, and explicit temporal smoothness losses are counterproductive. We hypothesize that the temporal attention modules learn to produce consistent noise predictions because inconsistent predictions are penalized by the denoising loss when trained on video frames that are inherently temporally coherent. The explicit temporal loss over-constrains the system by penalizing all inter-frame variation,

Effect of base model. 100K, $\rho=0.6$, $\lambda_{\text{temp}}=0$.

Base Model	CLIP \uparrow	FC	MM \uparrow	FID \downarrow	FVD \downarrow
SD 1.5	0.302	0.980	0.051	89.96	85.99
SD 2.1	0.303	0.983	0.040	92.09	92.49

suppressing natural motion diversity.

Role of correlated noise. Temporally correlated noise serves as a crucial initialization prior. Without it, the UNet—which processes each frame through frozen spatial attention independently—produces completely divergent scene content across frames, creating an impossible task for the lightweight temporal attention to correct. Correlated noise ensures that frames start from similar noise patterns, providing a foundation that the temporal attention can then refine. The correlation coefficient provides a practical inference-time knob that does not require retraining.

Motion amplitude trade-off. Our ablations reveal that ρ directly controls motion amplitude. At $\rho=0.85$, videos are effectively static (MM=0.008); at $\rho=0.4$, motion is rich (MM=0.158) but at some cost to CLIP alignment (0.278). The operating point $\rho=0.6$ strikes a practical balance for diverse video prompts.

Limitations. Our method generates 8 frames at 256×256 , which is below the resolution and duration of state-of-the-art video generators trained on orders of magnitude more data. The FID and FVD scores remain above AnimateDiff’s even under equal training data, reflecting the capacity advantage of AnimateDiff’s larger motion modules. Future work will explore scaling to higher resolutions, longer sequences, and integration with more modern architectures such as DiT-based diffusion models [17].

Conclusion

We presented a motion-adaptive temporal attention mechanism for parameter-efficient video generation with frozen Stable Diffusion models. By injecting lightweight temporal modules (25.8M parameters, 2.9% of the base UNet) with motion-adaptive attention bias, motion-aware gating, and a cascaded global-to-adaptive injection strategy, our method produces temporally coherent videos while adapting to diverse motion characteristics. Combined with temporally correlated noise, the approach achieves competitive results with $16\times$ fewer parameters than AnimateDiff under identical training conditions (100K videos, $T=8$). Our analysis reveals that implicit temporal regularization through the denoising objective alone outperforms explicit temporal consistency losses, and that the noise correlation coefficient provides a convenient inference-time control for balancing consistency and motion amplitude—a trade-off not typically exposed by fixed-architecture methods.

References

- [1] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, High-resolution image synthesis with latent diffusion models, Proc. CVPR, pg. 10684–10695 (2022).
- [2] J. Ho, T. Salimans, A. Gritsenko, W. Chan, M. Norouzi, and D. J. Fleet, Video diffusion models, Proc. NeurIPS, 35, 8633–8646 (2022).


- [3] U. Singer, A. Polyak, T. Hayes, X. Yin, J. An, S. Zhang, Q. Hu, H. Yang, O. Ashual, O. Gafni, D. Parikh, S. Gupta, and Y. Taigman, Make-A-Video: Text-to-video generation without text-video data, arXiv preprint arXiv:2209.14792 (2022).
- [4] A. Blattmann, R. Rombach, H. Ling, T. Dockhorn, S. W. Kim, S. Fidler, and K. Kreis, Align your latents: High-resolution video synthesis with latent diffusion models, Proc. CVPR, pg. 22563–22575 (2023).
- [5] A. Blattmann, T. Dockhorn, S. Kulal, D. Mendelevitch, M. Kilian, D. Lorenz, Y. Levi, Z. English, V. Voleti, A. Letts, V. Jampani, and R. Rombach, Stable Video Diffusion: Scaling latent video diffusion models to large datasets, arXiv preprint arXiv:2311.15127 (2023).
- [6] Y. Guo, C. Yang, A. Rao, Z. Liang, Y. Wang, Y. Qiao, M. Agrawala, D. Lin, and B. Dai, AnimateDiff: Animate your personalized text-to-image diffusion models without specific tuning, Proc. ICLR (2024).
- [7] J. Wang, H. Yuan, D. Chen, Y. Zhang, X. Wang, and S. Zhang, ModelScope text-to-video technical report, arXiv preprint arXiv:2308.06571 (2023).
- [8] Z. Wang, Z. Yuan, X. Wang, T. Chen, M. Xia, P. Luo, and Y. Shan, MotionCtrl: A unified and flexible motion controller for video generation, Proc. ACM SIGGRAPH (2024).
- [9] W. Wu, Z. Li, Y. Gu, R. Zhao, Y. He, D. J. Zhang, M. Z. Shou, Y. Li, T. Gao, and D. Zhang, DragAnything: Motion control for anything using entity representation, Proc. ECCV (2024).
- [10] T.-S. Chen, C. H. Lin, H.-Y. Tseng, T.-Y. Lin, and M.-H. Yang, Motion-conditioned diffusion model for controllable video synthesis, arXiv preprint arXiv:2304.14404 (2023).
- [11] M. Bain, A. Nagrani, G. Varol, and A. Zisserman, Frozen in time: A joint video and image encoder for end-to-end retrieval, Proc. ICCV, pg. 1728–1738 (2021).
- [12] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, Learning transferable visual models from natural language supervision, Proc. ICML, pg. 8748–8763 (2021).
- [13] T. Unterthiner, S. van Steenkiste, K. Kurach, R. Marinier, M. Michalski, and S. Gelly, Towards accurate generative models of video: A new metric and challenges, arXiv preprint arXiv:1812.01717 (2019).
- [14] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, GANs trained by a two time-scale update rule converge to a local Nash equilibrium, Proc. NeurIPS, 30 (2017).
- [15] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu, DPM-Solver++: Fast solver for guided sampling of diffusion probabilistic models, arXiv preprint arXiv:2211.01095 (2022).
- [16] L. Zhang, A. Rao, and M. Agrawala, Adding conditional control to text-to-image diffusion models, Proc. ICCV, pg. 3836–3847 (2023).
- [17] W. Peebles and S. Xie, Scalable diffusion models with transformers, Proc. ICCV, pg. 4195–4205 (2023).

Author Biography

Rui Hong is a PhD student at George Mason University. His research interests include sign language translation and generation, 3D pose estimation, multimodality AI, and generative AI.

Shuxue Quan is an independent researcher with interests in artificial intelligence, computer vision, camera systems, and AR/VR/XR/MR technologies.

(a) “Time lapse of sky clouds storm”



(b) “Pacific black duck, *anas superciliosa*, relaxing”




Figure 3. Qualitative comparison on two WebVid validation prompts. Each pair shows frame 1 and frame 8 of the 8-frame generated sequence at 256×256 , highlighting temporal consistency across the full generation. AnimateDiff+ retrained on 100K videos, $T=8$.

JOIN US AT THE NEXT EI!

electronic IMAGING

Imaging across applications . . . Where industry and academia meet!



- **SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •**
- **INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •**

www.electronicimaging.org

