
Audio-Sync Video Generation with Multi-Stream Temporal Control

Shuchen Weng^{1†} Haojie Zheng^{1,2†} Zheng Chang³ Si Li³ Boxin Shi^{4,5‡} Xinlong Wang^{1‡}

¹Beijing Academy of Artificial Intelligence

²School of Software and Microelectronics, Peking University

³School of Artificial Intelligence, Beijing University of Posts and Telecommunications

⁴State Key Lab of Multimedia Info. Processing, School of Computer Science, Peking University

⁵Nat'l Eng. Research Ctr. of Visual Tech., School of Computer Science, Peking University

{scweng, wangxinlong}@baai.ac.cn, suimu@stu.pku.edu.cn

{zhengchang98, lisi}@bupt.edu.cn, shiboxin@pku.edu.cn

Abstract

Audio is inherently temporal and closely synchronized with the visual world, making it a naturally aligned and expressive control signal for controllable video generation (*e.g.*, movies). Beyond control, directly translating audio into video is essential for understanding and visualizing rich audio narratives (*e.g.*, Podcasts or historical recordings). However, existing approaches fall short in generating high-quality videos with precise audio-visual synchronization, especially across diverse and complex audio types. In this work, we introduce MTV, a versatile framework for audio-sync video generation. MTV explicitly separates audios into speech, effects, and music tracks, enabling disentangled control over lip motion, event timing, and visual mood, respectively—resulting in fine-grained and semantically aligned video generation. To support the framework, we additionally present DEMIX, a dataset comprising high-quality cinematic videos and demixed audio tracks. DEMIX is structured into five overlapped subsets, enabling scalable multi-stage training for diverse generation scenarios. Extensive experiments demonstrate that MTV achieves state-of-the-art performance across six standard metrics spanning video quality, text-video consistency, and audio-video alignment. Project page: <https://hjzheng.net/projects/MTV/>.

1 Introduction

Audio is a fundamental medium in daily life, crucial for both information delivery (*e.g.*, communication, notifications, and education) and immersive experiences (*e.g.*, enhancing the impact of film visuals). Despite the prevalence of audio-centric platforms (*e.g.*, Podcasts), content presented solely through audio lacks the visual dimension needed to fully convey the richness of events. Since audio is naturally temporal and inherently synchronized with the visual world, researchers [1–3] have devoted considerable attention to translating audios into corresponding videos to enhance audience understanding of rich audio narratives (*e.g.*, historical recordings).

Despite great progress, existing methods face practical limitations in generating high-fidelity cinematic videos with precise synchronization (*e.g.*, pouring water into the transparent cup), primarily due to: (i) **Under-specified audio-visual mapping**. Current approaches handle a wide spectrum of audios and map them to various target scenes (*e.g.*, landscapes [4], dancing [5], music performances [6]). This broad representation scope potentially leads to ambiguous mappings lacking specificity between audio and visual features. (ii) **Inaccurate temporal alignment**. Existing methods primarily focus on

[†] Equal contributions. [‡] Corresponding authors.

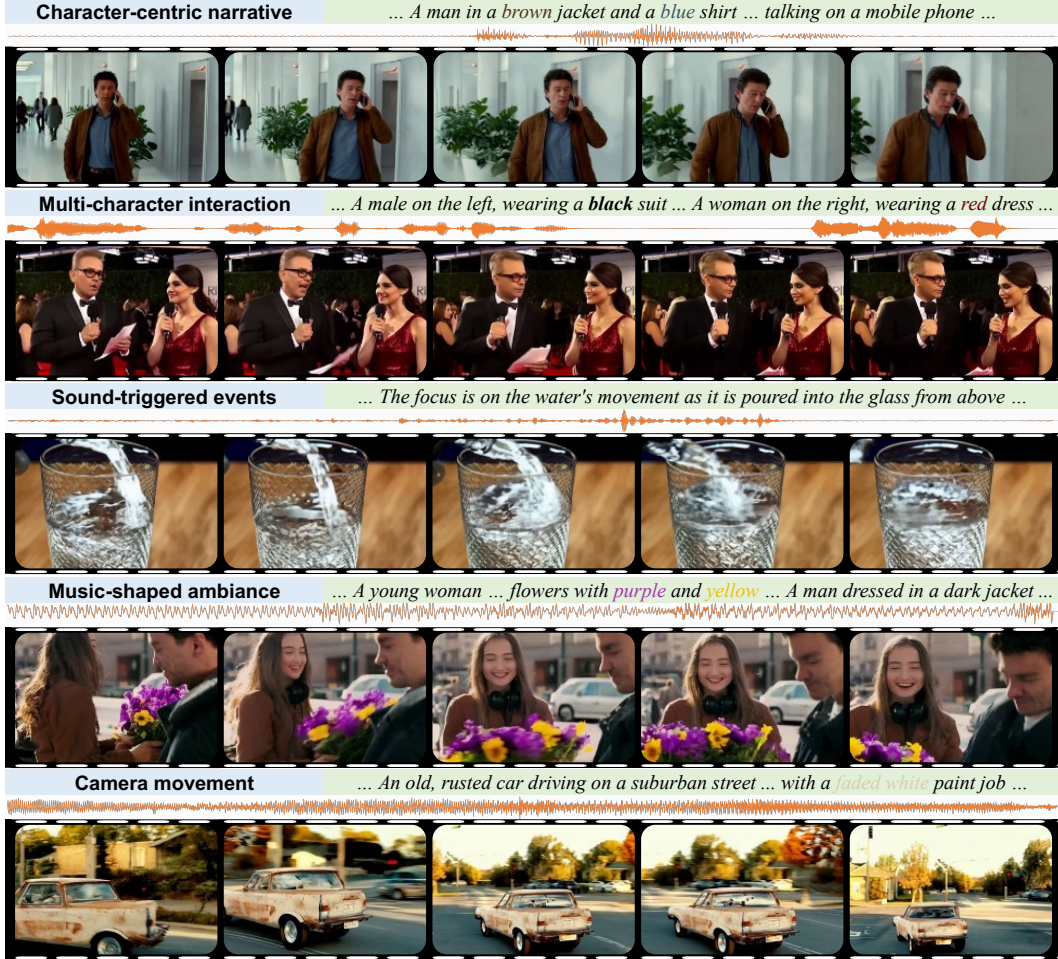


Figure 1: MTV demonstrates versatile audio-sync video generation capabilities following user-provided text descriptions specifying scenes and subjects. Capabilities shown include producing videos centered on targeted characters (1st and 2nd rows) while triggering events with sound effects (3rd row), generating visual mood with accompanying music (4th row), and adaptively handling camera movement (5th row). We present these generated videos in the supplementary materials.

building scene-level semantic consistency (*e.g.*, translating engine sound to a car-centered video), struggling with accurate timing correspondence between individual audio events and their visual features (*e.g.*, speech [7], motion [8], and visual mood [9]).

In this paper, we propose the **MTV** framework, enabling **M**ulti-stream **T**emporal control for audio-sync **V**ideo generation to overcome aforementioned issues, with versatile capabilities across scenarios illustrated in Fig. 1. Instead of attempting a direct mapping from composite audios, we explicitly separate audios into distinct controlling tracks (*i.e.*, speech, effects, and music), inspired by CDX’23¹. To provide sufficient high-quality video clips with demixed audio tracks, we contribute a large-scale DEMIX dataset with tailored data processing, including 392K video clips with 1.2K hours. These tracks enable the model to precisely control lip motion, event timing, and visual mood, resolving the ambiguous mapping. To further incorporate rich visual semantics beyond direct audio cues, we leverage features (*e.g.*, subject gesture, scene appearance, camera movement) initially derived from a pretrained text-to-video model [10], and subsequently finetuned using video clips from the DEMIX dataset. To enable the progressive extension of learned high-level video semantic features stage-by-stage, this dataset is structured into five overlapped subsets. A multi-stage training strategy is introduced to learn concrete and localized controls (*e.g.*, lip motion) towards more abstract and global influences (*e.g.*, visual mood), leading to clear audio-visual relationships.

¹<https://www.aicrowd.com/challenges/sound-demixing-challenge-2023>

To achieve accurate temporal alignment, we propose the Multi-Stream Temporal ControlNet (MST-ControlNet) within the MTV framework. The interval stream is designed for specific feature synchronization, which extracts features from the speech and effects tracks. It employs interval interaction blocks to understand each track individually and construct their interplay, maintaining the coherence with inferred semantic features. After that, interval feature injection module inserts features of each track into corresponding time intervals to drive lip motion and event timing. Since visual mood typically covers the entire video clip, the holistic stream is designed for overall aesthetic presentation, which extracts features from the music track using the holistic context encoder. These features then serve as style embeddings, applied uniformly to all frames through global style injection, controlling the visual mood.

We summarize our contributions as follows:

- We present MTV, a versatile audio-sync video generation framework by demixing audio inputs, achieving precise audio-visual mapping and accurate temporal alignment.
- We introduce an audio-sync video generation dataset structured into five overlapped subsets, presenting the multi-stage training strategy for learning audio-visual relationships.
- We propose the multi-stream temporal ControlNet to distinctively process demixed audio tracks and precisely control lip motion, event timing, and visual mood, respectively.

2 Related Works

2.1 Video Diffusion Model

The field of video generation has made significant progress with the adoption of diffusion models. Early approaches [11–13] extend the dynamic modeling capabilities of pretrained text-to-image diffusion models [14] by incorporating temporal layers (*e.g.*, 3D convolutions [15] and temporal attention [16]). However, these methods face inherent challenges in capturing long-range spatial-temporal dependencies due to the convolutional architectures of their backbone (*e.g.*, UNet [17]). To overcome this limitation, Sora report [18] presents the potential of the diffusion transformer (DiT) [19] architecture, prompting a shift towards integrating 3D VAE [20] for spatial-temporal compression and scaling up to train the entire DiT-based model. Further improvement has been achieved by recent foundation models through adaptive layernorm modules [10], progressive scaling [21, 22], and post-training techniques [23]. These advancements in text-to-video models provide a strong foundation and powerful generative priors that could potentially be leveraged for related cross-modal tasks, such as high-quality audio-sync video generation.

2.2 Audio-driven Image Animation

Audio-driven image animation aims to generate dynamic visuals from a static image, synchronized with user-provided audios. Several previous works animate general objects or scenes while maintaining audio-visual consistency. Sound2Sight [24] and CCVS [25] leverage the context of preceding frames to achieve audio-driven subsequent frames generation. TPOS [26] uses audios with variable temporal semantics and amplitude to guide the denoising process. ASVA [27] incorporates a temporal audio control module for effective audio synchronization. Other works concentrate on audio-driven human animation. Talking head [7, 28–30] focus on animating human face images to produce lip motion that synchronize with the speech. Recent works extend animation beyond the head to include half-body movements [31] and introduce pose control for full-body animation [32]. Another specific application is music-to-dance [33, 34], which generates human dance according to the beat of the music. Despite the audio-visual synchronization of these methods, their reliance on static images restricts models’ capability to generate dynamic scenes required for cinematic videos.

2.3 Audio-sync Video Generation

Audio-sync video generation does not require additional images for reference, offering the potential for free scene creation. Early works are designed based on VQGAN [35] and StyleGAN [36], achieving audio control through multi-modal autoregressive transformers [2] and style code alignment [4, 37]. Recently, following the success of diffusion models demonstrating effectiveness in general video generation, researchers have turned their attention. Highlighting the benefit of multi-modal

Table 1: Comparison of DEMIX dataset and previous datasets.

Method	Year	Modality		Scene			Audio component				Specifications	
		Text	Audio	People	Objects	Cinematic	Speech	Effects	Music	Demix	Clips	Hours
UCF-101 [38]	2012	–	✓	✓	–	–	–	✓	✓	–	13K	27
HIMV-200K [39]	2017	–	✓	✓	✓	✓	–	–	✓	–	200K	–
AudioSet [40]	2017	–	✓	✓	✓	–	✓	✓	✓	–	2.1M	5.8K
VoxCeleb2 [41]	2018	–	✓	✓	–	–	✓	–	–	–	150K	2.4K
VGGSound [42]	2020	–	✓	✓	✓	–	✓	✓	✓	–	200K	550
WebVid-10M [43]	2021	✓	–	✓	✓	–	–	–	–	–	10.7M	52K
Landscape [4]	2022	–	✓	–	✓	–	–	✓	–	–	9K	26
InternVid [44]	2024	✓	✓	✓	✓	–	✓	✓	✓	–	7.1M	760K
Ours (DEMIX)	2025	✓	✓	✓	✓	✓	✓	✓	✓	✓	392K	1.2K

conditions, TA2V [6] demonstrates that conditioning on both text descriptions and audio inputs significantly enhances the quality of generated videos. To achieve audio-visual alignment at both global and temporal levels, TempoTokens [1] designs a lightweight adapter for text-to-video generation model. Introducing a unified diffusion architecture, MM-Diffusion [5] enables both joint audio-video generation and zero-shot audio-sync video generation. Leveraging diffusion-based latent aligners for open-domain audio-visual generation, Xing *et al.*[3] achieve the audio-sync video editing and open-domain content creation. Although great progress has been made, audio-sync video generation still faces under-specific audio-visual mapping and inaccurate temporal alignment. Therefore, achieving cinematic quality remains challenging.

3 Dataset

We introduce the DEMIX dataset, tailored for training demixed audio-sync video generation models.

Data source. The training data is sourced from three aspects: (i) 65 hours of talking head videos from CelebV-HQ [45]; (ii) 4,923 hours of cinematic videos from MovieBench [46] (69h), Condensed Movies [47] (1,270h), and Short-Films 20K [48] (3,584h); and (iii) 8,903 hours film-related videos from YouTube. All collected videos include their accompanying audio tracks.

Video filtering. Following previous video generation models [10, 12, 49], we use PySceneDetect [50] to segment video into single-shot clips. Audiobox-aesthetics [51] is further used to assess the quality of accompanying audio, removing clips with low scores. For the left video clips, we annotate each one with text descriptions using LLaVA-Video [52].

Demixing filtering. To improve audio demixing reliability, we employ a dual-demixing comparison strategy, comparing demixing outputs from MVSEP [53] (speech, effects, music) and Spleeter [54] (speech, others). After that, we calculate the L1 distance between the speech tracks. Next, the ‘others’ track from Spleeter is conditionally compared: to the effects track from MVSEP if music is silent (below -45dB), and to the music track if effects are silent. Clips are discarded only if high L1 distances are found on any of the comparable pairs.

Voice-over filtering. To build clear audio-visual relationships for cinematic videos, we first detect whether people are present in the videos using YOLO [55]. Next, we perform speaker diarization for the accompanying audio using Scribe [56] to identify active speaker segments and count the number of speakers. After that, we detect the active speaker from videos for each frame using TalkNet [57]. As a result, we can discard clips where speech occurs in the audio but the video analysis detects neither a visible person nor an active speaker in the corresponding frames.

Subset division. To facilitate multi-stage training for versatile audio-sync video generation models, the filtered DEMIX data is structured into five overlapped subsets. The basic face subset comprises all talking head videos. The remaining cinematic and film-related videos are then categorized to form the other subsets: assignment to single character or multiple characters depends on the annotated human count, while assignment to sound event or visual mood occurs if the respective effects or music track is non-silent.

Data statistics. After data collection and filtering, our DEMIX dataset includes 18K basic face, 54K single character, 39K multiple characters, 166K sound event, and 195K visual mood data, tailored for




Figure 2: The pipeline of our MTV framework. (a-c) MTV is built on a pretrained text-to-video model [10] that provides strong generative priors for synthesizing diverse visual scenarios. (d) Explicitly separated audio tracks (*i.e.*, speech, effects, music) are fed into our proposed multi-stream temporal ControlNet to ensure synchronization for lip motion, event timing, and visual mood. (e) The MTV framework is trained on our contributed DEMIX dataset with five overlapped subsets and tailored text structures, enabling a multi-stage training strategy for audio-sync video generation.

cinematic videos, totaling non-overlapped 392K clips with 1.2K hours, accompanied by demixed audio tracks². For comprehensive evaluation, we hold out 1K video clips from the dataset to form the testing set. We provide an additional comparison with existing audio-related datasets [4, 38–44] in Tab. 1, highlighting that ours is tailored for versatile audio-sync video generation using demixed audio tracks, while robustly covering scenarios with people, objects, and cinematic visuals.

4 Method

This section begins with an overview of our MTV framework for audio-sync video generation (Sec. 4.1). Next, we detail the Multi-stream Temporal ControlNet (MST-ControlNet), including the interval stream for specific feature synchronization, and the holistic stream for overall aesthetic presentation (Sec. 4.2). Finally, we present the multi-stage training strategy for effectively learning audio-visual relationships (Sec. 4.3).

4.1 Overview

MTV generates audio-sync videos based on user-provided text descriptions y (specifying the scenes and subjects) and demixed audio tracks $a = \{a^s, a^e, a^m\}$ (representing speech, effects, and music) to respectively drive the lip motion, event timing, and visual mood. The pipeline is illustrated in Fig. 2.

Video compression. As presented in Fig. 2 (a), MTV is equipped with a pretrained spatio-temporal variational autoencoder (VAE) encoder $\mathcal{E}$ to map video clips x into latent code $z_0 = \mathcal{E}(x)$. After that, its corresponding VAE decoder $\mathcal{D}$ is used to reconstruct video clips from the latent code $x = \mathcal{D}(z_0)$.

Denoising network. As presented in Fig. 2 (b), we concatenate the text embeddings f^y and noised latent code z_t before feeding them into the network to ensure the video-text correspondence. The expert Adaptive LayerNorm (AdaLN) [10] then independently processes text and video features within this unified sequence. Next, 3D full-attention is used to interact semantics of text embeddings with corresponding video features. After being extracted by MST-ControlNet, audio cues are integrated via the interval feature injection and holistic style injection mechanisms. Finally, a feed-forward network (FFN) is used to refine the resulting video features.

Denoising process. As presented in Fig. 2 (c), MTV finally generates audio-sync videos by iteratively denoising latent codes. During training, at each time step $t \in \{0, \dots, T\}$, Gaussian noise $\epsilon_t \sim$

²Dataset samples are visualized in the supplementary materials.

$\mathcal{N}(0, 1)$ is added to the clean latent code z_0 to produce a noised latent code $z_t = \sqrt{\bar{\alpha}_t}z_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon_t$. A diffusion transformer ϵ_θ is trained to predict the noise ϵ_t , given the noised latent code z_t , demixed audio tracks a , denoising time step t , and text descriptions y . The diffusion transformer is trained by minimizing the loss:

$$\mathcal{L}_{\text{dm}} = \mathbb{E}_{t, z_0, \epsilon_t \sim \mathcal{N}(0, 1)} [\|\epsilon_t - \epsilon_\theta(z_t, a, t, y)\|^2]. \quad (1)$$

For inference, we iteratively denoise a randomly sampled noise $z_T \sim \mathcal{N}(0, 1)$ to obtain the latent code z'_0 to generate video clips with the VAE decoder $x' = \mathcal{D}(z'_0)$.

4.2 Multi-stream Temporal ControlNet

After explicitly separating audios into speech, effects, and music tracks, we propose the MST-ControlNet to achieve accurate temporal alignment by respectively controlling lip motion, event timing, and visual mood. As presented in Fig. 2 (d), the architecture consists of an audio encoding module followed by two specialized streams.

Audio encoding. Given demixed audio tracks $a = \{a^s, a^e, a^m\}$, we initially extract their corresponding features $\{f^s, f^e, f^m\}$ from the demixed tracks using wav2vec [58]. After that, speech and effect features are fed into the interval stream for specific feature synchronization. Instead, music features are fed into the holistic stream for overall aesthetic presentation.

Interval stream. We design the interval stream to interval-wise control the lip motion and event timing. Specifically, we separately process speech features f^s and effect features f^e with a stack of linear layers and concatenate them before feeding them into N interval interaction blocks. Within each block, these features are processed independently (via AdaLN, Gate, and FFN) to refine per-track understanding. To model their interplay at each time interval i , the corresponding speech features f_i^s and effects features f_i^e are jointly processed by a self-attention $[\tilde{f}_i^s, \tilde{f}_i^e] = \text{SelfAttn}([f_i^s, f_i^e])$. This interaction also maintains the coherence with inferred semantic features. Finally, interacted speech features $\tilde{f}^s$ and effects features $\tilde{f}^e$ are integrated into their corresponding time intervals via the interval feature injection mechanism:

$$h_i^s = \text{CrossAttn}(h_i, \tilde{f}_i^s), \quad h_i^e = \text{CrossAttn}(h_i, \tilde{f}_i^e), \quad (2)$$

where h_i represents the video latent code at i -th interval. $\text{CrossAttn}(\cdot, \cdot)$ means a cross-attention, where the latent code serves as the query and the audio features as the key and value. Let M be the number of intervals, the resulting latent code is then updated as $h' = \{h_i^s + h_i^e\}_{i=1}^M$.

Holistic stream. The holistic stream is designed to control the visual mood for the entire video clip. Specifically, we process the music features f^m through a holistic context encoder, comprising three linear layers and a 1D convolutional layer to extract features representing the visual mood. Since the environmental ambiance typically covers the entire video clip, an average pooling is applied to merge all the intervals and transform them into holistic music features $\tilde{f}^m$. Next, these features are regarded as style embeddings. By independently transforming these features into scale factor $\gamma^m = \text{Linear}(\tilde{f}^m)$ and shift factor $\beta^m = \text{Linear}(\tilde{f}^m)$, we modulate the video latent code h' uniformly across all intervals via the holistic style injection:

$$h^m = h' \odot (\gamma^m + 1) + \beta^m, \quad (3)$$

where h^m is the modulated latent code, fed into the denoising network to refine video features.

4.3 Multi-stage training strategy

As the dataset is structured as five overlapped subsets, we introduce the multi-stage training strategy to progressively scale up the model stage-by-stage.

Text structure. As presented in Fig. 2 (e), we create a template to structure text descriptions, enabling our MTV framework to be compatible with these distinct training subsets. Specifically, this template begins with a sentence indicating the number of participants (e.g., “Two person conversation”), based on Scribe [56] speaker counts. It then consists of subsequent entries for each individual, starting with a unique identifier (e.g., *Person1*, *Person2*) followed by their respective appearance description. Following these individual entries, an explicit identifier for the currently active speaker is specified. Finally, a sentence provides an overall description of the scene. Notably, when there is no active speaker in the video, only the overall description will be provided.

Table 2: Quantitative experiment results of comparison and ablation. $\uparrow$ ($\downarrow$) means higher (lower) is better. Throughout the paper, best performances are highlighted in **bold**.

Method	FVD $\downarrow$	Temp-C (%) $\uparrow$	Text-C (%) $\uparrow$	Audio-C (%) $\uparrow$	Sync-C $\uparrow$	Sync-D $\downarrow$
Comparison with state-of-the-art methods						
MM-Diffusion [5]	879.77	94.15	15.61	5.43	1.53	11.21
TempoTokens [1]	795.88	93.13	24.68	6.71	1.45	10.48
Xing <i>et al.</i> [3]	805.23	93.30	24.51	7.30	1.55	10.50
Ours (MTV)	626.06	95.40	26.55	26.22	3.17	9.43
Ablation study						
W/o SE	667.81	95.30	26.49	24.68	2.46	9.55
W/o SI	626.46	94.84	25.50	19.64	2.53	9.76
W/o TB	698.36	95.14	26.37	24.50	2.31	9.78

Training schedule. We train the model from concrete and localized controls towards more abstract and global influences. Initially, we train the model to learn lip motion using the basic face subset. It then learns human pose, scene appearance, and camera movement on the single character subset. To handle scenarios with multiple speakers, we subsequently train the model on the multiple characters subset. Following this, our training focus shifts to event timing and extending subject understanding from humans to objects using the sound event subset. Finally, we train the model on the environmental ambiance subset to improve its representation of visual mood.

Training details. We initialize our spatial-temporal VAE and DiT backbone with pretrained weights from CogVideoX [10] and train our model to generate audio-sync videos at a 480×720 resolution. For each stage, we train our model for 40K steps on 24 NVIDIA A800 GPUs using the Adam-based optimizer [59] with a learning rate of 1×10^{-5} , where MST-ControlNet and attention layers of the backbone are trainable. For inference, our model requires 280s to generate a 49-frame audio-sync video on a NVIDIA A100 GPU.

5 Experiments

5.1 Comparison with state-of-the-art methods

As audio-sync video generation is an emerging task, the relevant comparison methods are still developing. We compare our method with three recent state-of-the-art approaches in our DEMIX dataset. For TempoTokens [1] and Xing *et al.* [3], we evaluate them using both text descriptions and corresponding audios as their original configuration. Since MM-Diffusion [5] can only support audio inputs and its training focuses on specific landscape and dancing, we finetune it to ensure a fair comparison. 50 videos are randomly selected from the testing set for evaluation.

Quantitative comparisons. As presented in Tab. 2, we quantitatively evaluate performance across three main aspects: (i) Visual quality is assessed using Frechét Video Distance (FVD) [60]. (ii) Temporal consistency (Temp-C) is measured by calculating similarity between consecutive frames using CLIP [61]. (iii) We examine text-video alignment via Text Consistency (Text-C) [62], audio-video alignment using Audio Consistency (Audio-C) [63], and specifically lip motion synchronization with Sync-C and Sync-D [64]. As a result, our framework outperforms state-of-the-art methods across all six quantitative metrics. These metric details are provided in the supplementary materials.

Qualitative comparisons. As presented in Fig. 3, qualitative comparisons with state-of-the-art methods [1, 3, 5] highlight the advantages of our framework. For instance, even after finetuning MM-Diffusion [5] for over 320K steps using the official code on 8 NVIDIA A100 GPUs, it still struggles with generating cinematic videos. TempoTokens [1] struggles to generate cinematic videos for complex text-specified scenarios, resulting in unrealistic human expressions (Fig. 3 left). Xing *et al.* [3] find it difficult to effectively achieve audio synchronization for specific event timing, leading to incorrect rendering of human gestures for guitar performance (Fig. 3 right). In contrast, our MTV framework faithfully generates audio-sync videos with cinematic quality.



Figure 3: Visual comparison results with state-of-the-art methods for audio-sync video generation.



Figure 4: Ablation study results of different MST-ControlNet variants.

5.2 Ablation Study

To evaluate the effectiveness of key components within MST-ControlNet, we conduct ablation studies against three baseline configurations, as shown in Fig. 4 and Tab. 2.

W/o SE (Separate Extraction). We extract all features from demixed audio tracks using interval interaction blocks. This prevents music features from shaping the overall aesthetic presentation, leading to reduced visual mood (Fig. 4 left, degraded FVD and Temp-C).

W/o SI (Separate Injection). We extract features from demixed audio tracks by their respective encoders. These features are then concatenated and injected into the denoising network via a shared cross-attention. This reduces conditional consistency (Fig. 4 left, decreased Text-C and Audio-C).

W/o TB (Training Backbone). We freeze all weights of DiT backbone and only train our proposed MST-ControlNet to preserve more generative priors. This impairs the specific feature synchronization, especially the lip motion synchronization (Fig. 4 right, reduced Sync-C and Sync-D).

Table 3: User study results. Ours (MTV) clearly produces a higher score than state-of-the-art methods.

Subjective criteria	MM-Diffusion [5]	TempoTokens [1]	Xing <i>et al.</i> [3]	Ours (MTV)
Semantic consistency	0.96%	13.60%	11.28%	74.16%
Motion fluency	0.64%	8.96%	12.56%	77.84%
Overall preference	0.72%	12.00%	12.40%	74.88%

Table 4: Quantitative experiment results with alternative pre-trained components.

Method	FVD ↓	Temp-C (%) ↑	Text-C (%) ↑	Audio-C (%) ↑	Sync-C ↑	Sync-D ↓
CogVideoX+Wav2Vec	626.06	95.40	26.55	26.22	3.17	9.43
CogVideoX+Beats	598.53	95.91	26.25	25.28	3.02	9.52
Wan14B+Wav2Vec	353.61	96.36	27.23	26.49	3.08	9.56

5.3 User Study

To better evaluate our method from a human perception perspective, we conduct three subjective user study experiments in Tab. 3. We present videos generated by our method and all baselines to participants and ask them to choose the best one based on the following criteria: (i) **Semantic consistency**. How well the video content aligns with the text description. (ii) **Motion fluency**. The realism and temporal coherence of the motion. (iii) **Overall preference**. How good the holistic quality of the video is. For each study, we randomly select 50 text descriptions from the test set, and the evaluations are conducted by 25 volunteers. The table below shows the percentage of times each method is chosen as the winner. Our method is consistently favored by human observers and has achieved the highest scores across all three subjective criteria.

5.4 Analysis of Pre-trained Components

We evaluate the robustness of our proposed method by integrating it with alternative pre-trained components. Specifically, we test replacing the audio encoder (Wav2Vec/BEATs) and the video backbone (CogVideoX/Wan14B) in Tab. 4.

BEATs. Since Wav2Vec [58] is a common setting for speech encoding (*e.g.*, Hallo3 [7]), this baseline only replaces it with BEATs [65] for both the effects and music tracks. As shown in Tab. 4, this baseline achieves comparable (or slightly better) video-related metrics (*i.e.*, FVD and Temp-C) but shows a slight degradation on audio-related metrics (*i.e.*, Audio-C, Sync-C, and Sync-D), suggesting that our current choice of Wav2Vec [58] is a robust and effective one for this task.

Wan14B. Since Wan14B [21] shares a similar DiT-based structure with CogVideoX [10], we can integrate our proposed MST-ControlNet into it without architectural changes. Specifically, our interval feature injection and holistic style injection modules are added after each text cross-attention layer. The quantitative results below show this baseline achieves better performance on video- and text-related metrics (*i.e.*, FVD, Temp-C, and Text-C) due to the stronger capabilities of the Wan14B [21], while achieving comparable performance on all audio-related metrics (*i.e.*, Audio-C, Sync-C, and Sync-D).

5.5 Application

As presented in Fig. 5, our model support four typical scenarios: (i) By integrating text-to-video generative priors and learned audio-visual synchronized capabilities, our model can create vivid virtual characters. (ii) Given user-provided images and taking them as arbitrary keyframes, our model can drive the image according to the given audios. (iii) Although our model generates video segments of 49 frames, it can achieve long video generation by using the generated frame to initialize the next segment. (iv) Following training-free approaches [66], our model can generate scene transitions guided by providing time-varying text descriptions.

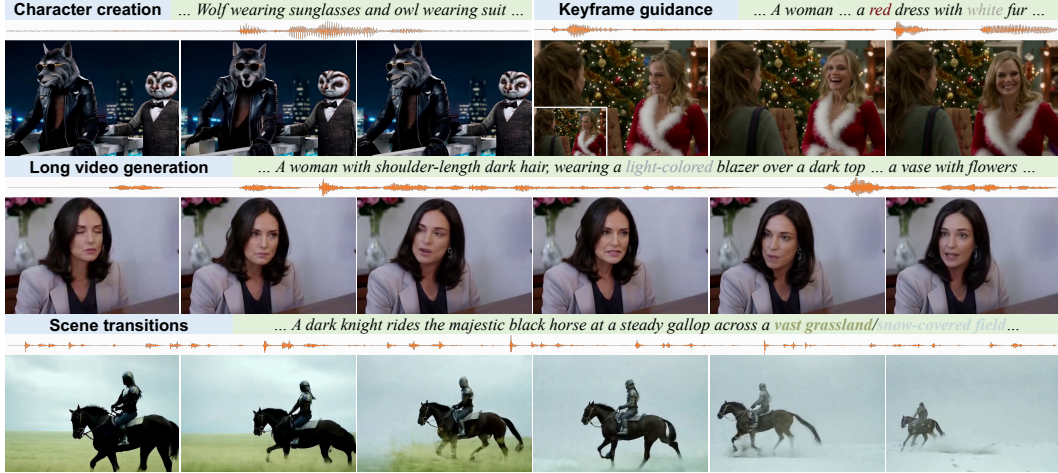


Figure 5: Examples of versatile application scenarios for our proposed MTV framework.



Figure 6: Examples of controllability study for text descriptions and demixed audios.

5.6 Controllability

As shown in Fig. 6, leveraging control from both text descriptions and the three demixed audio tracks (*i.e.*, speech, effects, music), our model can offer controllability across following four key aspects: (i) Modifying the text descriptions while keeping all audio tracks fixed allows the visual scene appearance to be edited without affecting the audio synchronization. (ii) Given a demixed speech track, the model enables precise control over the synchronized lip motion of the generated character. (iii) Similarly, with a demixed effects track, the model accurately synchronizes event timing with the sound effects. (iv) By changing the demixed music track, the model creates different visual moods for the generated video.

6 Conclusion

In this work, we presented MTV, a versatile framework for audio-sync video generation. MTV leverages generative priors from pretrained text-to-video models [10] and is trained on our contributed DEMIX dataset that provides sufficient cinematic videos with demixed audio tracks. Equipped with our proposed MST-ControlNet, MTV is able to independently control lip motion, event timing, and visual mood. Combined with a multi-stage training strategy for effective learning of complex audio-visual relationships, MTV achieves state-of-the-art performance across six evaluation metrics.

Limitation. Although our approach demonstrates the potential of using demixed audio tracks for precise video control, it is fundamentally limited by the scope of categories provided by upstream audio demixing techniques [53, 54]. We believe the capabilities of audio-sync video generation methods will further progress with advancements in audio demixing methods.

Acknowledgement. This work is supported by National Natural Science Foundation of China (Grant No. 62136001). We thank all the insightful reviewers for the helpful suggestions, and the colleagues at Beijing Academy of Artificial Intelligence for their support throughout this project.

References

- [1] G. Yariv, I. Gat, S. Benaïm, L. Wolf, I. Schwartz, and Y. Adi, “Diverse and aligned audio-to-video generation via text-to-video model adaptation,” in *AAAI*, 2024.
- [2] S. Ge, T. Hayes, H. Yang, X. Yin, G. Pang, D. Jacobs, J.-B. Huang, and D. Parikh, “Long video generation with time-agnostic vqgan and time-sensitive transformer,” in *ECCV*, 2022.
- [3] Y. Xing, Y. He, Z. Tian, X. Wang, and Q. Chen, “Seeing and hearing: Open-domain visual-audio generation with diffusion latent aligners,” in *CVPR*, 2024.
- [4] S. H. Lee, G. Oh, W. Byeon, J. Bae, C. Kim, W. J. Ryoo, S. H. Yoon, J. Kim, and S. Kim, “Sound-guided semantic video generation,” in *ECCV*, 2022.
- [5] L. Ruan, Y. Ma, H. Yang, H. He, B. Liu, J. Fu, N. J. Yuan, Q. Jin, and B. Guo, “MM-Diffusion: Learning multi-modal diffusion models for joint audio and video generation,” in *CVPR*, 2023.
- [6] M. Zhao, W. Wang, T. Chen, R. Zhang, and R. Li, “TA2V: Text-audio guided video generation,” *TMM*, 2024.
- [7] J. Cui, H. Li, Y. Zhan, H. Shang, K. Cheng, Y. Ma, S. Mu, H. Zhou, J. Wang, and S. Zhu, “Hallo3: Highly dynamic and realistic portrait image animation with video diffusion transformer,” in *CVPR*, 2025.
- [8] S. Qian, Z. Tu, Y. Zhi, W. Liu, and S. Gao, “Speech drives templates: Co-speech gesture synthesis with learned templates,” in *ICCV*, 2021.
- [9] B. M.-K. Ng, S. R. Sudhoff, H. Li, J. Kamphuis, T. Nadolsky, Y. Chen, K. Y.-J. Yun, and Y.-H. Lu, “Visualize music using generative arts,” in *CAI*, 2024.
- [10] Z. Yang, J. Teng, W. Zheng, M. Ding, S. Huang, J. Xu, Y. Yang, W. Hong, X. Zhang, G. Feng, *et al.*, “CogVidex: Text-to-video diffusion models with an expert transformer,” in *ICLR*, 2025.
- [11] P. Esser, J. Chiu, P. Atighehchian, J. Granskog, and A. Germanidis, “Structure and content-guided video synthesis with diffusion models,” in *ICCV*, 2023.
- [12] A. Blattmann, T. Dockhorn, S. Kulal, D. Mendelevitch, M. Kilian, D. Lorenz, Y. Levi, Z. English, V. Voleti, A. Letts, *et al.*, “Stable video diffusion: Scaling latent video diffusion models to large datasets,” *arXiv preprint arXiv:2311.15127*, 2023.
- [13] Y. He, T. Yang, Y. Zhang, Y. Shan, and Q. Chen, “Latent video diffusion models for high-fidelity long video generation,” *arXiv preprint arXiv:2211.13221*, 2022.
- [14] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *CVPR*, 2022.
- [15] T. Soo Kim and A. Reiter, “Interpretable 3D human action analysis with temporal convolutional networks,” in *CVPR workshops*, 2017.
- [16] G. Bertasius, H. Wang, and L. Torresani, “Is space-time attention all you need for video understanding?,” in *ICML*, 2021.
- [17] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *MICCAI*, 2015.
- [18] Y. Liu, K. Zhang, Y. Li, Z. Yan, C. Gao, R. Chen, Z. Yuan, Y. Huang, H. Sun, J. Gao, *et al.*, “Sora: A review on background, technology, limitations, and opportunities of large vision models,” *arXiv preprint arXiv:2402.17177*, 2024.
- [19] W. Peebles and S. Xie, “Scalable diffusion models with transformers,” in *ICCV*, 2023.
- [20] L. Yu, J. Lezama, N. B. Gundavarapu, L. Versari, K. Sohn, D. Minnen, Y. Cheng, V. Birodgar, A. Gupta, X. Gu, A. G. Hauptmann, B. Gong, M.-H. Yang, I. Essa, D. A. Ross, and L. Jiang, “Language model beats diffusion – tokenizer is key to visual generation,” in *ICLR*, 2024.

- [21] A. Wang, B. Ai, B. Wen, C. Mao, C.-W. Xie, D. Chen, F. Yu, H. Zhao, J. Yang, J. Zeng, *et al.*, “Wan: Open and advanced large-scale video generative models,” *arXiv preprint arXiv:2503.20314*, 2025.
- [22] W. Kong, Q. Tian, Z. Zhang, R. Min, Z. Dai, J. Zhou, J. Xiong, X. Li, B. Wu, J. Zhang, *et al.*, “Hunyuanvideo: A systematic framework for large video generative models,” *arXiv preprint arXiv:2412.03603*, 2024.
- [23] G. Ma, H. Huang, K. Yan, L. Chen, N. Duan, S. Yin, C. Wan, R. Ming, X. Song, X. Chen, *et al.*, “Step-Video-T2V technical report: The practice, challenges, and future of video foundation model,” *arXiv preprint arXiv:2502.10248*, 2025.
- [24] M. Chatterjee and A. Cherian, “Sound2Sight: Generating visual dynamics from sound and context,” in *ECCV*, 2020.
- [25] G. Le Moing, J. Ponce, and C. Schmid, “CCVS: Context-aware controllable video synthesis,” in *NeurIPS*, 2021.
- [26] Y. Jeong, W. Ryoo, S. Lee, D. Seo, W. Byeon, S. Kim, and J. Kim, “The power of sound (TPoS): Audio reactive video generation with stable diffusion,” in *ICCV*, 2023.
- [27] L. Zhang, S. Mo, Y. Zhang, and P. Morgado, “Audio-synchronized visual animation,” in *ECCV*, 2024.
- [28] J. Jiang, C. Liang, J. Yang, G. Lin, T. Zhong, and Y. Zheng, “Loopy: Taming audio-driven portrait avatar with long-term motion dependency,” in *ICLR*, 2025.
- [29] H. Wei, Z. Yang, and Z. Wang, “AniPortrait: Audio-driven synthesis of photorealistic portrait animation,” *arXiv preprint arXiv:2403.17694*, 2024.
- [30] W. Zhang, X. Cun, X. Wang, Y. Zhang, X. Shen, Y. Guo, Y. Shan, and F. Wang, “SadTalker: Learning realistic 3d motion coefficients for stylized audio-driven single image talking face animation,” in *CVPR*, 2023.
- [31] G. Lin, J. Jiang, C. Liang, T. Zhong, J. Yang, and Y. Zheng, “CyberHost: A one-stage diffusion framework for audio-driven talking body generation,” in *ICLR*, 2025.
- [32] G. Lin, J. Jiang, J. Yang, Z. Zheng, and C. Liang, “OmniHuman-1: Rethinking the scaling-up of one-stage conditioned human animation models,” *arXiv preprint arXiv:2502.01061*, 2025.
- [33] W. Xuanchen, W. Heng, L. Dongnan, and W. Cai, “Dance any beat: Blending beats with visuals in dance video generation,” in *WACV*, 2025.
- [34] Z. Chen, H. Xu, G. Song, Y. Xie, C. Zhang, X. Chen, C. Wang, D. Chang, and L. Luo, “X-dancer: Expressive music to human dance video generation,” *arXiv preprint arXiv:2502.17414*, 2025.
- [35] P. Esser, R. Rombach, and B. Ommer, “Taming transformers for high-resolution image synthesis,” in *CVPR*, 2021.
- [36] T. Karras, S. Laine, and T. Aila, “A style-based generator architecture for generative adversarial networks,” in *CVPR*, 2019.
- [37] D. Jeong, S. Doh, and T. Kwon, “Träumerai: Dreaming music with stylegan,” *arXiv preprint arXiv:2102.04680*, 2021.
- [38] K. Soomro, A. R. Zamir, and M. Shah, “Ucf101: A dataset of 101 human actions classes from videos in the wild,” *arXiv preprint arXiv:1212.0402*, 2012.
- [39] S. Hong, W. Im, and H. S. Yang, “Content-based video-music retrieval using soft intra-modal structure constraint,” *arXiv preprint arXiv:1704.06761*, 2017.
- [40] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “AudioSet: An ontology and human-labeled dataset for audio events,” in *ICASSP*, 2017.

- [41] J. S. Chung, A. Nagrani, and A. Zisserman, “VoxCeleb2: Deep speaker recognition,” *arXiv preprint arXiv:1806.05622*, 2018.
- [42] H. Chen, W. Xie, A. Vedaldi, and A. Zisserman, “VggSound: A large-scale audio-visual dataset,” in *ICASSP*, 2020.
- [43] M. Bain, A. Nagrani, G. Varol, and A. Zisserman, “Frozen in time: A joint video and image encoder for end-to-end retrieval,” in *ICCV*, 2021.
- [44] Y. Wang, Y. He, Y. Li, K. Li, J. Yu, X. Ma, X. Li, G. Chen, X. Chen, Y. Wang, *et al.*, “InternVid: A large-scale video-text dataset for multimodal understanding and generation,” in *ICLR*, 2024.
- [45] H. Zhu, W. Wu, W. Zhu, L. Jiang, S. Tang, L. Zhang, Z. Liu, and C. C. Loy, “CelebV-HQ: A large-scale video facial attributes dataset,” in *ECCV*, 2022.
- [46] W. Wu, M. Liu, Z. Zhu, X. Xia, H. Feng, W. Wang, K. Q. Lin, C. Shen, and M. Z. Shou, “MovieBench: A hierarchical movie level dataset for long video generation,” *arXiv preprint arXiv:2411.15262*, 2024.
- [47] M. Bain, A. Nagrani, A. Brown, and A. Zisserman, “Condensed movies: Story based retrieval with contextual embeddings,” in *ACCV*, 2020.
- [48] R. Ghermi, X. Wang, V. Kalogeiton, and I. Laptev, “Short film dataset (SFD): A benchmark for story-level video understanding,” *arXiv preprint arXiv:2406.10221*, 2024.
- [49] H. Chen, Y. Zhang, X. Cun, M. Xia, X. Wang, C. Weng, and Y. Shan, “VideoCrafter2: Overcoming data limitations for high-quality video diffusion models,” in *CVPR*, 2024.
- [50] B. Castellano, “Video cut detection and analysis tool.” <https://github.com/Breakthrough/PySceneDetect>.
- [51] A. Tjandra, Y.-C. Wu, B. Guo, J. Hoffman, B. Ellis, A. Vyas, B. Shi, S. Chen, M. Le, N. Zacharov, *et al.*, “Meta audiobox aesthetics: Unified automatic quality assessment for speech, music, and sound,” *arXiv preprint arXiv:2502.05139*, 2025.
- [52] Y. Zhang, J. Wu, W. Li, B. Li, Z. Ma, Z. Liu, and C. Li, “Video instruction tuning with synthetic data,” *arXiv preprint arXiv:2410.02713*, 2024.
- [53] R. Solovyev, “Cinematic sound demixing.” <https://github.com/ZFTurbo/MVSEP-CDX23-Cinematic-Sound-Demixing>.
- [54] R. Hennequin, A. Khelif, F. Voituret, and M. Moussallam, “Spleeter: a fast and efficient music source separation tool with pre-trained models,” *Journal of Open Source Software*, 2020.
- [55] G. Jocher, J. Qiu, and A. Chaurasia, “Ultralytics YOLO.” <https://github.com/ultralytics/ultralytics>.
- [56] Elevenlabs, “Meet scribe.” <https://elevenlabs.io/blog/meet-scribe>.
- [57] S. Beliaev and B. Ginsburg, “TalkNet 2: Non-autoregressive depth-wise separable convolutional model for speech synthesis with explicit pitch and duration prediction,” *arXiv preprint arXiv:2104.08189*, 2021.
- [58] A. Baevski, H. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: a framework for self-supervised learning of speech representations,” in *NeurIPS*, 2020.
- [59] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [60] T. Unterthiner, S. Van Steenkiste, K. Kurach, R. Marinier, M. Michalski, and S. Gelly, “Towards accurate generative models of video: A new metric & challenges,” *arXiv preprint arXiv:1812.01717*, 2018.
- [61] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, “Learning transferable visual models from natural language supervision,” in *ICML*, 2021.

- [62] J. Wang, C. Wang, K. Huang, J. Huang, and L. Jin, “VideoCLIP-XL: Advancing long description understanding for video clip models,” *arXiv preprint arXiv:2410.00741*, 2024.
- [63] R. Girdhar, A. El-Nouby, Z. Liu, M. Singh, K. V. Alwala, A. Joulin, and I. Misra, “ImageBind: One embedding space to bind them all,” in *CVPR*, 2023.
- [64] J. S. Chung and A. Zisserman, “Out of time: automated lip sync in the wild,” in *ACCV Workshops*, 2017.
- [65] S. Chen, Y. Wu, C. Wang, S. Liu, D. Tompkins, Z. Chen, W. Che, X. Yu, and F. Wei, “BEATs: Audio pre-training with acoustic tokenizers,” in *ICML*, 2023.
- [66] M. Cai, X. Cun, X. Li, W. Liu, Z. Zhang, Y. Zhang, Y. Shan, and X. Yue, “DitCtrl: Exploring attention control in multi-modal diffusion transformer for tuning-free multi-prompt longer video generation,” *arXiv preprint arXiv:2412.18597*, 2024.

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 introduction part reflect the main contributions and scope.

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: The limitations are discussed in Sec. 6.

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: [\[NA\]](#)

Justification: The paper does not include theoretical results.

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: The model framework and experiment settings are described in detail.

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: [No]

Justification: Both the dataset and the codes will be released upon acceptance.

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, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: The experiments are presented in detail.

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: Error bars are not reported because it would be too computationally expensive.

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: These can be found in Sec. 4.3.

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: The authors are fully aware of 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: [NA]

Justification: This paper has no societal impact.

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: The paper poses no such risks.

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: They are cited clearly.

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: Codes and the dataset will be released upon acceptance.

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: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing nor research with human subjects.

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: [\[NA\]](#)

Justification: The paper does not involve crowdsourcing nor research with human subjects.

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.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

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

Justification: We clearly describe the usage of LLMs to annotate our dataset.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.