

Input-Aware Sparse Attention for Real-Time Co-Speech Video Generation

BEIJIA LU, Carnegie Mellon University, USA
ZIYI CHEN, PAII Inc., USA
JING XIAO, PAII Inc., USA
JUN-YAN ZHU, Carnegie Mellon University, USA

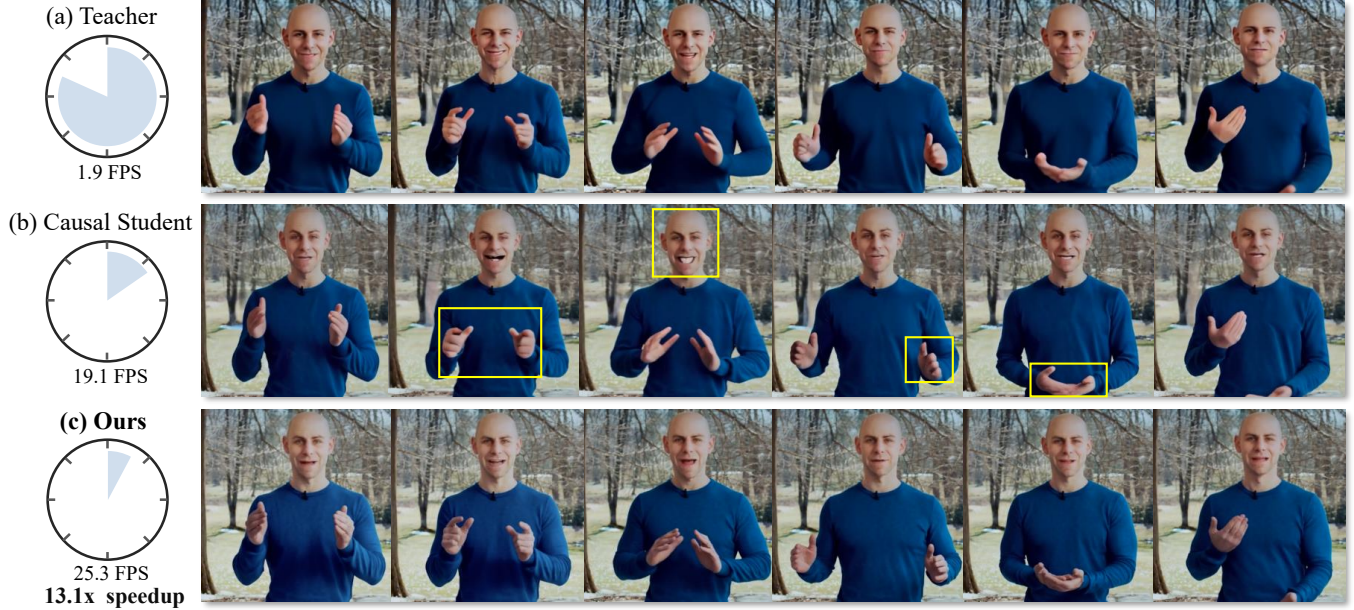


Fig. 1. We introduce a conditional video distillation method for real-time co-speech video generation that leverages human pose conditioning for input-aware sparse attention and distillation loss. Our student model achieves 25.3 FPS, a $13.1\times$ speedup over its teacher model [Zhang et al. 2025], while preserving visual quality. Our method significantly improves motion coherence and lip synchronization over a leading few-step causal student model [Yin et al. 2025], while reducing common visual degradation in the speaker’s face and hands (see yellow box).

Diffusion models can synthesize realistic co-speech video from audio for various applications, such as video creation and virtual agents. However, existing diffusion-based methods are slow due to numerous denoising steps and costly attention mechanisms, preventing real-time deployment. In this work, we distill a many-step diffusion video model into a few-step student model. Unfortunately, directly applying recent diffusion distillation methods degrades video quality and falls short of real-time performance. To address these issues, our new video distillation method leverages input human pose conditioning for both attention and loss functions. We first propose using accurate correspondence between input human pose keypoints to guide attention to relevant regions, such as the speaker’s face, hands, and upper body. This input-aware sparse attention reduces redundant computations

and strengthens temporal correspondences of body parts, improving inference efficiency and motion coherence. To further enhance visual quality, we introduce an input-aware distillation loss that improves lip synchronization and hand motion realism. By integrating our input-aware sparse attention and distillation loss, our method achieves real-time performance with improved visual quality compared to recent audio-driven and input-driven methods. We also conduct extensive experiments showing the effectiveness of our algorithmic design choices.

CCS Concepts: • **Computing methodologies** → **Animation; Computer vision.**

Additional Key Words and Phrases: Co-Speech Video Synthesis, Model acceleration

ACM Reference Format:

Beijia Lu, Ziyi Chen, Jing Xiao, and Jun-Yan Zhu. 2025. Input-Aware Sparse Attention for Real-Time Co-Speech Video Generation. In *SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*, December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3757377.3763831>

1 Introduction

Co-speech video generation, which synthesizes human videos from audio, has enabled diverse applications such as creating realistic

Authors’ Contact Information: Beijia Lu, Carnegie Mellon University, Pittsburgh, USA, beijialu@andrew.cmu.edu; Ziyi Chen, PAII Inc., Palo Alto, USA, ziyi.pai@gmail.com; Jing Xiao, PAII Inc., Palo Alto, USA, xiaojing661@pingan.com.cn; Jun-Yan Zhu, Carnegie Mellon University, Pittsburgh, USA, junyanz@andrew.cmu.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License. *SA Conference Papers '25, Hong Kong, Hong Kong*
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2137-3/25/12
<https://doi.org/10.1145/3757377.3763831>

virtual agents, producing educational video content, and enhancing telepresence. Diffusion models [Ho et al. 2020], with attention mechanisms, excel at this task, producing highly realistic and coherent videos. However, diffusion models are extremely slow for video generation, due to the numerous denoising steps and the quadratic costs associated with full attention over many frames and tokens, making real-time deployment difficult.

One common strategy to reduce computational costs is network distillation [Hinton et al. 2015], where a slow, pre-trained “teacher” video model is distilled into a faster “student” model. Unfortunately, as Figure 1b shows, directly applying recent text-to-video distillation methods, such as CausVid [Yin et al. 2025], to the co-speech generation task proves insufficient. Such methods tend to degrade the video quality, particularly in crucial areas like the speaker’s face and hand regions, which are vital for natural co-speech gestures. Moreover, the speedup achieved still falls short of the demands for real-time performance.

To address these limitations, we propose a conditional video distillation method for fast co-speech generation with the following key insight. Unlike general text-to-video distillation, co-speech generation benefits from readily available input human pose conditioning. We leverage this pose information through an input-aware sparse attention mechanism, which significantly reduces computation while preserving video quality. Pose keypoints allow us to identify key regions (e.g., the speaker’s face, hands, upper body). This enables our sparse attention to selectively focus on tokens within these regions and their corresponding areas in similar frames, rather than performing dense, costly full attention across the entire video frames. Beyond guiding attention, pose information also pinpoints critical regions, such as the face and hands. This also enables us to design distillation losses that prioritize perceptual quality and accuracy in these areas, crucial for convincing co-speech video. Building on these insights, our method incorporates input-aware sparse attention and distillation loss into an efficient video distillation method. The sparse attention mechanism focuses computation on pose-defined relevant regions, which reduces redundancy and improves motion coherence. The input-aware distillation loss further enhances visual quality, particularly lip synchronization and hand motion realism.

We have conducted extensive experiments on both the publicly available TalkShow dataset [Yi et al. 2023] and a more comprehensive, newly curated dataset. On both datasets, our method demonstrates significant improvements. It achieves real-time performance, with a $13.1\times$ speedup compared to the teacher model, while concurrently improving upon the teacher model’s visual quality by 10% in lip synchronization and 2% in motion coherence. Our approach consistently outperforms recent audio-driven and pose-driven methods in terms of generation quality, lip synchronization, and hand motion realism, while being significantly faster. We further include a thorough ablation study to quantitatively and qualitatively analyze the effectiveness of both sparse attention and distillation loss. To our knowledge, our work presents the first real-time diffusion-based co-speech avatar. Our code and models are available at <https://beijia11.github.io/IASA>.

2 Related Work

Co-Speech Video Generation. Co-speech video generation aims to synthesize human videos driven by audio inputs, ensuring temporal synchronization between speech and human motions, such as facial expressions and hand gestures. One common approach divides this task into two subtasks: generating human motions from speech, and subsequently creating videos from these synthesized motions. Some works solely focus on one of these subtasks [Chen et al. 2024; Liu et al. 2024b; Tian et al. 2024; Yi et al. 2023; Zhang et al. 2025], while others address the task as a whole [Corona et al. 2024; Ginosar et al. 2019; Hogue et al. 2024; Liu et al. 2025; Mahapatra et al. 2024; Meng et al. 2025].

Regarding audio-to-motion generation, Speech2Gesture [Ginosar et al. 2019] uses Generative Adversarial Networks (GANs) [Goodfellow et al. 2014] to generate human 2D keypoints. TalkShow [Yi et al. 2023] and EMAGE [Liu et al. 2024b] use VQ-VAEs [Van Den Oord et al. 2017] for generating human 3D meshes, while DiffShog [Chen et al. 2024] replaces VQ-VAE with diffusion models [Ho et al. 2020; Song et al. 2021]. Recent methods [He et al. 2024; Hogue et al. 2024; Liu et al. 2022] focus on generating unsupervised motion representations [Siarohin et al. 2019a,b] to further improve the output fidelity. In our work, we reuse the EMAGE module [Liu et al. 2024b] and focus on accelerating the motion-to-video part, as audio-to-motion accounts for a small portion of the overall computational costs.

For the motion-to-video task, conditional GANs [Isola et al. 2017; Wang et al. 2018] have been widely used to inject various motion representations into videos: Speech2Gesture [Ginosar et al. 2019] uses 2D keypoints, Mahapatra et al. [2024] use 3D textured human meshes, and methods like ANGIE [Liu et al. 2022] and S2G-MDDiffusion [He et al. 2024] use implicit motion representations [Wang et al. 2021]. However, GAN-based methods often struggled with generalization, tending to overfit to specific identities seen during training. To improve generalization across unseen speakers and motions, recent works have increasingly adopted diffusion models. MM-Diffusion [Ruan et al. 2023] uses a joint diffusion architecture to model multimodal correlations. Other approaches [Guan et al. 2024; Meng et al. 2025; Tian et al. 2024; Zhang et al. 2025] have introduced explicit intermediate pose representations as conditioning signals within the diffusion framework. While diffusion-based methods outperform GANs in quality and generalization, their iterative denoising process and costly attention mechanisms hinder real-time deployment. Our work aims to significantly accelerate diffusion-based co-speech avatars.

Diffusion Model Distillation. Diffusion models [Ho et al. 2020; Sohl-Dickstein et al. 2015; Song et al. 2021] require a large number of denoising steps to achieve high-quality generation, limiting their applicability in low-latency applications. To accelerate inference, a range of distillation techniques have been proposed to reduce the number of steps while preserving visual quality. Trajectory-preserving methods aim to approximate the denoising trajectory of a teacher model [Luhman and Luhman 2021; Salimans and Ho 2022]. For example, Luhman and Luhman [2021] train a single-step student to match the teacher’s output at the final step, while Consistency Distillation [Kim et al. 2024] trains the student to map any points on an ODE trajectory to its origin. Rectified flow [Liu et al. 2024a] trains

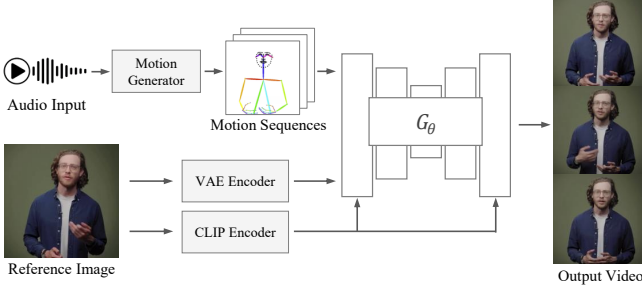


Fig. 2. **Our two-stage co-speech video generation pipeline.** In Stage 1, an audio input and a reference image are fed into an audio-to-motion generator [Liu et al. 2024b] to produce motion sequences represented by dense pose keypoints. In Stage 2, these motion sequences are fed into our efficient student video generation network G_θ . The network is conditioned on features from the reference image, which are separately encoded by the VAE encoder [Rombach et al. 2022] and CLIP encoder [Radford et al. 2021], to synthesize the final video. In our work, we focus on accelerating the video generation, significantly speeding up the teacher model [Zhang et al. 2025].

a student model on the linear interpolation path of noise-image pairs obtained from the teacher.

In contrast to trajectory-based methods, Distribution Matching Distillation (DMD) [Yin et al. 2024a,b] and adversarial distillation [Kang et al. 2024; Parmar et al. 2024; Sauer et al. 2024a,b] supervise the student at the distribution level by minimizing an approximate reverse KL divergence [Franceschi et al. 2023; Luo et al. 2023] or through adversarial learning [Goodfellow et al. 2014; Isola et al. 2017]. This enables architectural flexibility between teacher and student, and provides stronger global supervision. Our method builds upon DMD to distill a bidirectional teacher into an efficient student. However, DMD loss alone is insufficient for our task, which motivates us to introduce new modules that leverage input pose sequences.

Video Generation Acceleration. Accelerating video diffusion models [Blattmann et al. 2023; Ho et al. 2022; Peebles and Xie 2023] is critical due to their high computational cost. Common strategies [Lin et al. 2025; Lin and Yang 2024; Zhai et al. 2024] include reducing denoising steps using model distillation [Kang et al. 2024; Luhman and Luhman 2021; Yin et al. 2024b] or accelerating attention computation. This includes causal attention, inherent in autoregressive models [Alonso et al. 2024; Chen et al. 2025; Jin et al. 2024], where the current frame only attends to the past frames. This also enables streaming video generation and long video synthesis [Yin et al. 2025]. Other methods use attention sparsity [Child et al. 2019], for instance, by dynamically identifying spatiotemporal patterns [Xi et al. 2025]. However, existing methods do not fully exploit domain-specific cues, such as those available in co-speech video generation. In contrast, our method leverages input human pose information to design sparse attention and distillation loss. Our method outperforms generic video distillation methods by a large margin.

3 Method

We introduce a real-time, audio-to-video generation model that creates a co-speech video from an audio clip and a speaker’s reference

image. Our key idea is to distill a slow teacher video model into a fast student model using input-aware sparse attention and distillation. Below, we first describe our co-speech video generation pipeline in Section 3.1. Then, we introduce a sparse attention mechanism conditioned on the input human poses in Section 3.2, which dramatically reduces computation costs. Finally, Section 3.3 describes our distillation objective, which further increases the quality of crucial regions such as faces and hands.

3.1 Co-Speech Gesture Video Generation

As shown in Figure 2, we adopt a two-stage pipeline for co-speech gesture video generation, which consists of generating motion sequences from audio and synthesizing the corresponding gesture videos. For the first stage, we build on the EMAGE framework [Liu et al. 2024b], which converts speech audio and a reference image into temporally aligned upper body motion sequences, including facial expressions, hand gestures, and body movements. These motion sequences are represented as dense pose keypoints to capture structural motion information. In the second stage, we leverage input-conditioned video diffusion models [Zhang et al. 2025] as our teacher model, which takes a reference image and a motion sequence to synthesize temporally coherent and identity-consistent videos. This pipeline forms the basis for our proposed acceleration method, as described in the following sections.

3.2 Input-Aware Sparse Attention

We operate on a video sequence consisting of T frames, where each frame is represented by N tokens. Let $t_q, t_k \in \{1, \dots, T\}$ be the indices for a query and a key frame, and $i, j \in \{1, \dots, N\}$ be the respective token indices within these frames. To eliminate irrelevant interactions and accelerate attention computation, we employ an attention mask $\mathbf{M}(t_q, i, t_k, j) \in \{0, -\infty\}$, where $\mathbf{M}(t_q, i, t_k, j) = 0$ indicates that token i in frame t_q is permitted to attend to token j in frame t_k . Conversely, $\mathbf{M}(t_q, i, t_k, j) = -\infty$ denotes that attention between the token pair is suppressed. We next introduce global and local attention masking mechanisms designed to specify which frames and which regions the model should attend to, respectively.

Input-Aware Global Attention Masking. To ensure our temporal attention mechanism focuses on the relevant past information and adheres to causality, we introduce an input-aware global attention mask $\mathbf{M}_{\text{global}}$. This mask guides the attention computation by identifying a pertinent subset of historical frames for each current frame t_q .

For each frame $t \in \{1, \dots, T\}$, we represent its pose with B upper-body keypoints, denoted as $\mathbf{P}_t \in \mathbb{R}^{B \times 2}$. To quantify pose similarity between the current frame t_q and a historical frame $t_k < t_q$, we consider a global transformation matrix τ applied to \mathbf{P}_{t_k} to compensate for the subject’s global movement. The similarity $S(t_q, t_k)$ is then computed as the minimum alignment error, defined by:

$$S(t_q, t_k) = \min_{\tau} \|\mathbf{P}_{t_q} - \tau(\mathbf{P}_{t_k})\|_2, \quad (1)$$

where τ is a rigid transformation matrix. For each current frame t_q , we select the top- K most similar historical frames to construct the attention mask accordingly. Let $\mathcal{S}_{t_q} = \{S(t_q, t_k) \mid t_k < t_q\}$ be the

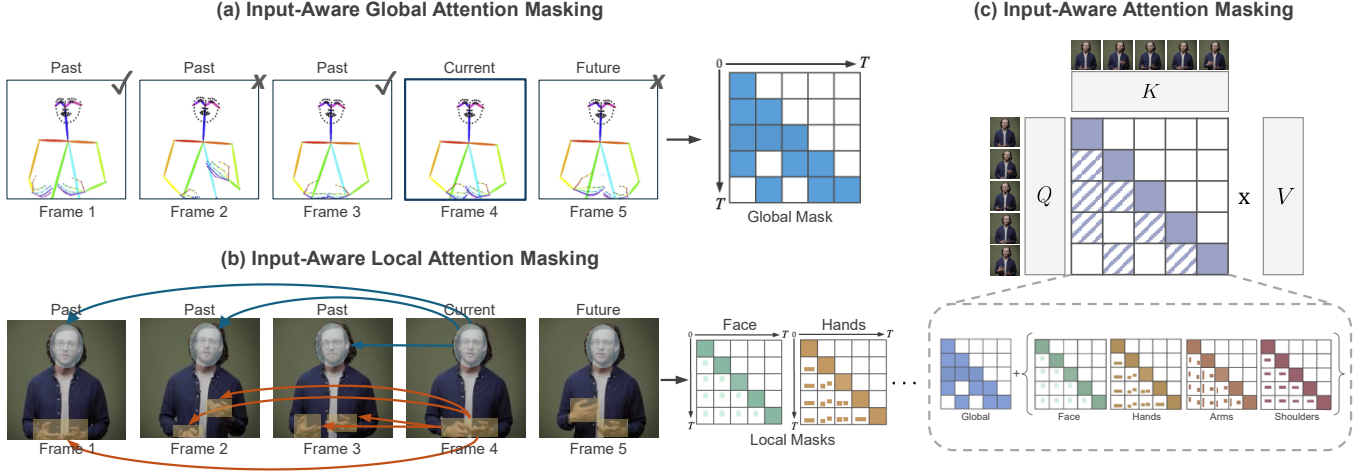


Fig. 3. **Input-Aware Sparse Attention.** Our attention mechanism selectively focuses on tokens within salient body regions and their corresponding areas in temporally relevant frames. (a) We first apply global masking, which restricts attention to the K most similar past frames based on pose similarity. (b) Then local masking limits inter-frame attention to matched regions (e.g., face, hands) to enhance temporal coherence. (c) Our input-aware attention masking integrates both global and local masks to form an efficient and structured sparse attention pattern.

set of similarity scores with previous frames, and $S_{t_q}^{(K)}$ be the K -th smallest value in S_{t_q} .

We define input-aware global attention mask $M_{\text{global}}(t_q, t_k)$ as:

$$M_{\text{global}}(t_q, i, t_k, j) = \begin{cases} 0, & \text{if } S(t_q, t_k) \leq S_{t_q}^{(K)} \\ -\infty, & \text{otherwise.} \end{cases} \quad (2)$$

When applied to the attention logits prior to the softmax operation, this mask ensures that only the top- K most similar historical frames contribute to the attention computation for frame t_q . We call this the global attention mask because its value is constant for all token pairs (i, j) between any two frames t_q and t_k , as it depends only on these frame indices.

Input-Aware Local Attention Masking. To further enforce local consistency and focus temporal attention on relevant body parts in human subjects, we introduce an input-aware mask M_{local} . The mask partitions tokens into coherent local regions, as defined by keypoint locations estimated via the rigid moving least squares transformation [Schaefer et al. 2006]. This formulation supports dense attention within each frame while constraining inter-frame attention to correspondences across homologous local regions. Let $\mathcal{R} = \{\text{faces, hands, arms, bodies, shoulders}\}$ be the set of local regions. Each region $r \in \mathcal{R}$ corresponds to a fixed subset of token indices $\mathcal{I}_r \subseteq \{1, \dots, N\}$. We then define the local attention mask $M_{\text{local}}(t_q, i, t_k, j)$ as:

$$M_{\text{local}}(t_q, i, t_k, j) = \begin{cases} 0, & \text{if } t_q = t_k \\ 0, & \text{if } t_q \neq t_k \text{ and } \exists r \in \mathcal{R} \text{ s.t. } i \in \mathcal{I}_r \text{ and } j \in \mathcal{I}_r \\ -\infty, & \text{otherwise.} \end{cases} \quad (3)$$

This formulation ensures that inter-frame attention focuses on locally corresponding areas (e.g., face regions attending to other face regions), while intra-frame attention remains dense to capture local context.

Input-Aware Sparse Attention Mechanism. We adopt M as our final mask for attention, where $M = M_{\text{global}} + M_{\text{local}}$. This design enables the attention mechanism to adapt dynamically to the input signals, focusing on critical regions (e.g., faces and hands) while minimizing redundant computation in static areas, such as the background.

To ensure efficient execution on modern hardware, attention is computed over blocks rather than individual 1×1 tokens. Specifically, we apply a pooling operation with a block size of 128×128 [Guo et al. 2024] to construct the sparse attention mask by Flashinfer [Ye et al. 2025], which substantially reduces the number of attention queries while preserving fine-grained structure.

3.3 Input-Aware Model Distillation

Now we distill the original multi-step teacher video diffusion model with full attention into an efficient student model with input-aware sparse attention. This enables real-time video generation while maintaining perceptual quality.

DMD Loss. Following variational score distillation [Wang et al. 2023] and distribution matching distillation (DMD) [Yin et al. 2024a,b], the student model G_θ is trained to match the latent distributions of the teacher in fewer timesteps. Specifically, the student minimizes the reverse Kullback-Leibler (KL) divergence between the student’s data distribution $p_{\text{gen}}(\mathbf{x}_t)$ and the teacher’s data distribution $p_{\text{data}}(\mathbf{x}_t)$ at randomly sampled timesteps t :

$$\begin{aligned}\nabla_{\theta}\mathcal{L}_{\text{DMD}} &\triangleq \mathbb{E}_t \left[\nabla_{\theta} \text{KL}(p_{\text{gen},t} \| p_{\text{data},t}) \right] \\ &\approx -\mathbb{E}_t \left[(s_{\text{data}}(\mathbf{x}_t, t) - s_{\text{gen}}(\mathbf{x}_t, t)) \frac{\partial G_{\theta}}{\partial \theta} \right].\end{aligned}\quad (4)$$

where $s_{\text{data}}(\mathbf{x}_t, t)$ and $s_{\text{gen}}(\mathbf{x}_t, t)$ denote the score functions of the teacher and student models at timestep t for the input \mathbf{x}_t , respectively. This formulation encourages the student to approximate the teacher’s denoising behavior with fewer steps, achieving a balance between efficiency and fidelity.

Input-Aware Distillation Loss. We compute an input-aware loss over all regions by applying masks \mathbf{m}_r to the full ground-truth image \mathbf{x} and the generated image $\hat{\mathbf{x}}$. These masks are specific to each image, constructed from its pose information to delineate the same region type r . The unified loss is formulated as follows:

$$\mathcal{L}_{\text{region}} = \sum_{r \in \mathcal{R}} \lambda_r \mathcal{L}_r(\mathbf{m}_r \odot \mathbf{x}, \mathbf{m}_r \odot \hat{\mathbf{x}}), \quad (5)$$

where \odot denotes the Hadamard product (element-wise multiplication), λ_r balances the importance of each region, and $\mathcal{L}_r(\cdot, \cdot)$ denotes the reconstruction loss for each region. Specifically, for the face region, we extract feature embeddings using a pre-trained ArcFace network [Deng et al. 2019a] and compute L2 distance between the two, while for other regions, LPIPS metric [Zhang et al. 2018] is used. This approach leverages specialized metrics for distinct regions. Face-specific and gesture-specific losses have been used in prior works in video generation [Richardson et al. 2021; Zhang et al. 2025]. Here, we adopt them for the video distillation task.

Final Objective. Combining the distribution matching distillation loss and the input-aware distillation loss, the overall training objective for the student model G_{θ} is

$$G_{\theta}^* = \arg \min_{G_{\theta}} \mathcal{L}_{\text{DMD}}(G_{\theta}) + \lambda_{\text{region}} \mathcal{L}_{\text{region}}(G_{\theta}), \quad (6)$$

where λ_{region} is a scalar hyperparameter controlling the trade-off between matching the teacher’s latent distributions and preserving semantic regional consistency.

4 Experiments

4.1 Experimental Setups

Datasets. We evaluate our method on two datasets: the public TalkShow dataset [Yi et al. 2023], which contains four speakers with varied backgrounds and irregular camera movements, and our curated YouTube Talking Video dataset. Our YouTube dataset features 45 speakers with diverse appearances in half-body and headshot views, set against clean backgrounds with static cameras.

To extract human poses, we use DWPose [Yang et al. 2023] and apply a strict filtering pipeline to ensure high-quality pose information. We retain only frames where (1) facial keypoint confidence exceeds 0.9, (2) upper-body keypoint confidence exceeds 0.8, and (3) upper-body keypoint visibility is at least 90%. This filtering process yields reliable upper-body and frontal-face cues for training.

Metrics. We evaluate our method using a comprehensive set of metrics, including image and video fidelity (FID [Heusel et al. 2017], FVD [Unterthiner et al. 2019], E-FID [Deng et al. 2019b]), image quality (PSNR, SSIM [Wang et al. 2004]), audio-lip synchronization (Sync-C, Sync-D) [Prajwal et al. 2020], hand motion quality and diversity (HKC, HKD) [Lin et al. 2024], and inference efficiency (FPS). Specifically, Sync-C and Sync-D measure the alignment between speech and lip movements, while E-FID evaluates the alignment of emotional expressions between the generated and real videos. HKC and HKD evaluate the average detection confidence and variance of the hand keypoints in the generated sequence. The reported FPS is measured on an NVIDIA H100 GPU, excluding I/O latency, to accurately reflect the model’s performance.

4.2 Implementation Details

Similar to the teacher model, we train our model across speakers with diverse visual and motion characteristics. The training procedure consists of two stages. First, we fine-tune the teacher model by incorporating our input-aware sparse attention, utilizing 4 NVIDIA H100 GPUs for approximately 6 hours. Second, we distill the student model from the adapted teacher, which takes about 10 hours on 8 NVIDIA H100 GPUs.

4.3 Comparisons

We compare our method with recent open-source methods in both audio-driven and pose-driven video generation settings. For audio-driven baselines MMDiffusion [Ruan et al. 2023] and S2G-MDD [He et al. 2024], which are trained per speaker, we perform our evaluation on the same TalkShow [Yi et al. 2023] dataset to ensure a fair comparison. For pose-driven baselines, we compare against AnimateAnyone [Hu et al. 2023], EchoMimicV2 [Meng et al. 2025], and MimicMotion [Zhang et al. 2025] on our YouTube Talking dataset. To evaluate the complete co-speech pipeline, we use our audio-to-motion module to generate pose sequences as input to these methods. EchoMimicV2 [Meng et al. 2025] accepts audio as input, its video generation strictly relies on externally provided pose sequences rather than audio-driven motion prediction. Therefore, we categorize it as a pose-driven method.

Quantitative Results. As detailed in Table 1, our method achieves approximately 3× faster inference compared to existing audio-driven and pose-driven baselines. In addition to speed, our approach produces higher-quality and more realistic results. Compared with audio-driven methods, our model not only maintains high generation quality but also substantially improves Sync-C and HKC. More specifically, the lip synchronization confidence significantly improves from 4.36 to 7.26. Additionally, compared to S2G-MDD, our method improves HKC from 0.956 to 0.968 on the test set.

Compared to pose-driven methods, our approach outperforms others in both lip synchronization and overall motion quality. Remarkably, our student model, when compared with its teacher model MimicMotion [Zhang et al. 2025], not only achieves an impressive 13.1× inference acceleration without sacrificing generation quality, but also further enhances motion and synchronization quality. Specifically, our model substantially improves HKC from 0.928 to




Fig. 4. **Qualitative comparison of audio-driven methods.** We show body animation results conditioned on the same input audio. Our method not only achieves accurate lip synchronization but also generates clearer hand gestures.




Fig. 5. **Qualitative comparison of pose-driven methods.** All methods are conditioned on the same motion sequence and reference image. Our approach produces more realistic faces, hands, and body movements.



Fig. 6. **Additional Qualitative Results.** This figure presents a selection of further examples generated by our method. Given only a single static reference image and an input audio clip, our model effectively synthesizes highly realistic and expressive video outputs. These results visually demonstrate its capability to produce natural facial expressions, fluid body movements, and accurate lip synchronization in real time.

Table 1. **Quantitative comparison with state-of-the-art methods.** We evaluate audio-driven methods on the TalkShow dataset [Yi et al. 2023] and pose-driven methods on the YouTube Talking Video dataset. Metrics include image and video fidelity, audio-lip synchronization and hand motion accuracy. **Bold** indicates the best result within each group of audio- or pose-driven methods.

Method	FPS \uparrow	FID \downarrow	FVD \downarrow	SSIM \uparrow	PSNR \uparrow	HKC \uparrow	HKD \uparrow	Sync-C \uparrow	Sync-D \downarrow	E-FID \downarrow
Audio-Driven										
MMDiffusion [Ruan et al. 2023]	2.82	81.25	985.06	0.549	17.62	0.946	23.24	3.44	8.02	4.71
S2G-MDD [He et al. 2024]	6.89	68.11	883.48	0.552	19.00	0.956	23.44	4.36	7.59	3.39
Ours	25.31	66.81	823.58	0.596	19.40	0.968	24.26	7.26	6.98	3.16
Pose-Driven										
AnimateAnyone [Hu et al. 2023]	2.07	61.23	759.83	0.752	20.94	0.921	24.63	2.00	9.38	5.08
EchoMimicv2 [Meng et al. 2025]	8.92	57.98	611.65	0.800	22.02	0.939	24.83	7.19	7.02	2.97
MimicMotion [Zhang et al. 2025]	1.93	56.32	628.03	0.823	23.89	0.928	24.82	4.56	7.34	3.38
Ours	25.31	57.01	610.34	0.829	23.42	0.948	24.83	7.28	6.99	3.01



Fig. 7. **Qualitative ablation on model components.** While the combination of global attention and causal attention improves efficiency, it introduces degradation in the hand and mouth regions. Our input-aware distillation not only restores visual fidelity and temporal consistency, but also achieves real-time performance with natural and synchronized motion.

0.948 and Sync-C from 4.56 to 7.28, demonstrating enhanced hand motion confidence and lip synchronization.

Qualitative Results. As shown in Figure 4, our method demonstrates clear improvements over existing audio-driven methods in lip-audio synchronization, expressive hand gestures, and overall visual quality. Specifically, our generated videos exhibit better lip-audio synchronization, where the lip movements align more naturally with the speech content. Additionally, the overall visual quality is significantly higher, producing sharper and more realistic

Table 2. **Quantitative ablation on model components.** We analyze the effect of progressively incorporating architectural components. Introducing input-aware attention and distillation substantially accelerate inference but degrade quality. In contrast, our input-aware distillation achieves improved motion realism accuracy while maintaining real-time performance.

Method	FPS \uparrow	SSIM \uparrow	PSNR \uparrow	HKC \uparrow	Sync-C \uparrow
Teacher	1.93	0.823	23.89	0.928	4.56
+finetuned	1.93	0.835	24.01	0.945	7.15
+global attention	3.28	0.834	23.89	0.943	7.14
+local attention	4.42	0.831	23.89	0.943	7.14
+ \mathcal{L}_{DMD}	25.31	0.826	22.99	0.941	7.08
+ \mathcal{L}_{Region} (Ours)	25.31	0.829	23.87	0.948	7.28

appearances compared to the blurry or less realistic results from the baselines. Figure 5 shows that our method generates more natural lip and hand animations than pose-driven baselines. Existing pose-driven methods often produce stiff or unnatural movements in these critical regions. In contrast, our model maintains high fidelity and realism, with lifelike facial and hand animations, while achieving faster inference.

4.4 Ablation Study

We conduct ablation studies on the YouTube Talking Video dataset.

Ablation Study on Model Components. We conduct an ablation study to analyze the contribution of each component, with results in Table 2 and Figure 7. Leveraging the teacher model as the baseline, we observe strong motion and lip synchronization performance after being fine-tuned on the co-speech dataset. However, the inference speed remains a bottleneck at 1.93 FPS. While input-aware global attention and input-aware local attention maintain generation quality without degradation, the speed improvements are limited. Direct application of DMD distillation achieves a significant speedup but results in quality degradation, particularly noticeable artifacts on faces and hands. By incorporating our input-aware distillation, the




Fig. 8. **Ablation on Efficiency.** We compare the inference time with the incorporation of input-aware attention and distillation. Our method significantly reduces computation while preserving quality, enabling real-time performance.

model achieves real-time performance at 25.31 FPS while maintaining generation quality comparable to the finetuned Teacher model.

Ablation Study on Efficiency. We present a detailed breakdown of the inference time across different architectural variants by decomposing the total runtime into four components: Attention, Linear, Norm, and Others in Figure 8. The baseline Teacher model requires 103.6 seconds to process an 8-second video. By introducing Global Attention, we reduce the time to 60.9 seconds, mainly due to reductions in attention and linear computations. Local Attention further decreases the time to 45.2 seconds. Finally, applying distillation reduces the runtime to 7.9 seconds, a 13.1 \times speedup compared to the Teacher model, primarily attributed to the substantial reduction in attention costs. This demonstrates the effectiveness of our sparse attention strategy in enabling real-time co-speech video generation.

4.5 Results

Additional Qualitative Results. In Figure 6, we present additional qualitative examples generated by our method to illustrate its visual quality and expressiveness. Conditioned on a single static reference image and an input audio clip, our model synthesizes temporally coherent video sequences with high appearance fidelity and natural motion. The results demonstrate lifelike facial expressions, smooth upper body movements, and precise synchronization of the audio lip, even under variations in speaker identity, pose, and expression. These examples highlight the robustness of our method in capturing fine-grained motion dynamics while preserving speaker identity across frames. In addition, they complement the quantitative evaluations by providing intuitive evidence that our method can generate visually convincing talking videos in real time, underscoring its potential for applications in human-computer interaction and digital content creation.

Long Video Generation. We adopt a progressive strategy for long video generation that explicitly enforces temporal smoothness across extended sequences. At each denoising step, the video is partitioned into consecutive segments, and each segment is generated independently by conditioning the trained model on a fixed reference image together with its corresponding sub-sequence of poses. To guarantee consistency between adjacent segments, the overlapped frames are progressively fused according to their temporal indices.

Table 3. **User Preference Study.** Results of a user study evaluating the perceptual quality of video generation. Participants were asked to compare our method against baseline approaches along three aspects: Lip Synchronization, Motion Realism, and Video Quality. The table reports preference rates (%) across all participants for each method.

Method	User Preference		
	Lip Sync	Motion Realism	Video Quality
MMDiffusion	2.3%	11.1%	2.2%
S2G-MDD	13.2%	15.6%	6.7%
Ours	84.5%	73.3%	91.1%
AnimateAnyone	8.9%	16.7%	18.9%
EchoMimicV2	31.1%	24.4%	24.4%
MimicMotion	12.2%	26.7%	27.8%
Ours	47.8%	32.2%	28.9%
Ground Truth	54.5%	46.7%	56.6%
Ours	45.5%	53.3%	43.4%

This progressive design not only mitigates the accumulation of temporal artifacts but also preserves global coherence throughout the sequence, enabling our method to stably generate videos of several minutes in length without noticeable degradation in visual quality or motion dynamics.

User Preference Study. We conduct a user preference study on Amazon Mechanical Turk (AMT) with 30 participants to evaluate the perceptual quality of our generated videos. Specifically, each participant is asked to perform three separate pairwise comparisons over 72 videos: (1) our method vs. audio-driven generation, (2) our method vs. pose-driven generation, and (3) our method vs. ground truth (GT). In each comparison, participants are presented with videos generated from identical reference images and audio inputs, displayed in randomized order without revealing the underlying method. For each trial, users are asked to indicate their preference in terms of lip synchronization, motion realism, and overall video quality. As reported in Table 3, our method is favored by users in comparison to baselines, and its performance is comparable to the ground truth.

5 Discussion and Limitations

Limitations. Our current method still has several limitations. First, we cannot handle videos with dynamic backgrounds or complex scene motion, as our input-aware sparse attention assumes a relatively static camera and background. This limits the applicability of our approach to studio-like or controlled environments. Second, while our method improves lip synchronization and hand motion realism, it still struggles with subtle gestures, such as finger articulation, as shown in Figure 9. Third, our current system is trained and evaluated primarily on monolingual English-speaking datasets. Extending to multilingual or code-switching scenarios may require additional modeling of phoneme-to-gesture dynamics across languages. Lastly, although our model runs in real-time, it still requires a relatively powerful GPU for inference, which may restrict its deployment on edge devices or mobile platforms.

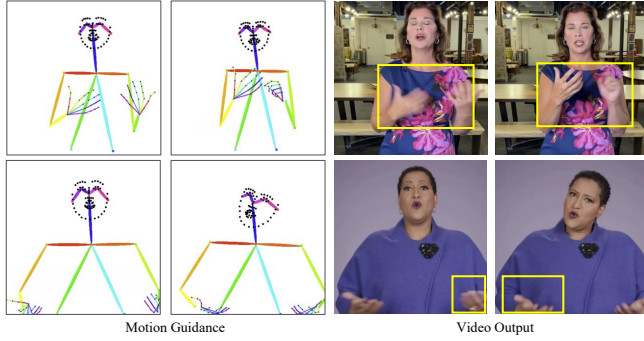


Fig. 9. **Limitations.** Given the same motion input (left), we observe that when fingers undergo fine-grained movements, our model tends to produce blurry or distorted hand regions.

Ethics Considerations. Our research advances generative models for human video synthesis. We recognize that this technology may also be misused for malicious purposes such as spreading misinformation, damaging reputations, or producing deceptive content. Prior studies [Cozzolino et al. 2021; Wang et al. 2020] on deepfake detection have shown that generated videos often contain detectable artifacts, underscoring the importance of video forensics in mitigating these risks. However, current forensic techniques remain limited, and the continued advancement of generative models necessitates parallel progress in detection to ensure reliability.

Conclusion. In this work, we introduced a new conditional video distillation framework specifically designed for fast and high-quality co-speech video generation. Our core approach leverages pose information as a natural and effective conditioning signal, guiding both the attention and the supervision mechanisms of our student model. To achieve this, we proposed an input-aware sparse attention module that intelligently focuses computation on dynamically important regions across frames, alongside an input-aware distillation loss which prioritizes visual fidelity in critical areas like the face and hands. Collectively, these components enabled our distilled model to achieve real-time inference speeds, all while not only preserving but significantly enhancing motion realism and lip synchronization quality compared to the teacher model.

Acknowledgments

We would like to thank Kangle Deng, Muyang Li, Nupur Kumari, Sheng-Yu Wang, Maxwell Jones, Gaurav Parmar for their insightful feedback and input that contributed to the finished work. The project is partly supported by Ping An Research.

References

Eloi Alonso, Adam Jelley, Vincent Micheli, Anssi Kanervisto, Amos Storkey, Tim Pearce, and Francois Fleuret. 2024. Diffusion for World Modeling: Visual Details Matter in Atari. In *Conference on Neural Information Processing Systems (NeurIPS)*.
 Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, Varun Jampani, and Robin Rombach. 2023. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets. *arXiv preprint arXiv:2311.15127* (2023).
 Boyuan Chen, Diego Marti Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. 2025. Diffusion forcing: Next-token prediction meets full-sequence

diffusion. *Conference on Neural Information Processing Systems (NeurIPS)* 37 (2025), 24081–24125.
 Junming Chen, Yunfei Liu, Jianan Wang, Ailing Zeng, Yu Li, and Qifeng Chen. 2024. Diffshg: A diffusion-based approach for real-time speech-driven holistic 3d expression and gesture generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 7352–7361.
 Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509* (2019).
 Enric Corona, Andrei Zanfir, Eduard Gabriel Bazavan, Nikos Kolotouros, Thimo Alldieck, and Cristian Sminchisescu. 2024. Vlogger: Multimodal diffusion for embodied avatar synthesis. *arXiv preprint arXiv:2403.08764* (2024).
 Davide Cozzolino, Andreas Rössler, Justus Thies, Matthias Nießner, and Luisa Verdoliva. 2021. Id-Reveal: Identity-aware deepfake video detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 15108–15117.
 Jiansheng Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. 2019a. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4690–4699.
 Yu Deng, Jialong Yang, Sicheng Xu, Dong Chen, Yunde Jia, and Xin Tong. 2019b. Accurate 3D Face Reconstruction with Weakly-Supervised Learning: From Single Image to Image Set. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*.
 Jean-Yves Franceschi, Mike Gartrell, Ludovic Dos Santos, Thibaut Issenhardt, Emmanuel de Bezenac, Mickaël Chen, and Alain Rakotomamonjy. 2023. Unifying GANs and Score-Based Diffusion as Generative Particle Models. In *Conference on Neural Information Processing Systems (NeurIPS)*.
 Shiry Ginosar, Amir Bar, Gefen Kohavi, Caroline Chan, Andrew Owens, and Jitendra Malik. 2019. Learning individual styles of conversational gesture. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3497–3506.
 Ian J Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Conference on Neural Information Processing Systems (NeurIPS)* 27 (2014).
 Jiazhi Guan, Quanwei Yang, Kaisiyuan Wang, Hang Zhou, Shengyi He, Zhiliang Xu, Haocheng Feng, Errui Ding, Jingdong Wang, Hongtao Xie, et al. 2024. TALK-Act: Enhance Textual-Awareness for 2D Speaking Avatar Reenactment with Diffusion Model. In *ACM SIGGRAPH Asia*.
 Junxian Guo, Haotian Tang, Shang Yang, Zhekai Zhang, Zhijian Liu, and Song Han. 2024. Block Sparse Attention. <https://github.com/mit-han-lab/Block-Sparse-Attention>.
 Xu He, Qiaochu Huang, Zhensong Zhang, Zhiwei Lin, Zhiyong Wu, Sicheng Yang, Mingli Li, Zhiyi Chen, Songcen Xu, and Xiaofei Wu. 2024. Co-Speech Gesture Video Generation via Motion-Decoupled Diffusion Model. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2017. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Conference on Neural Information Processing Systems (NeurIPS)*.
 Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
 Jonathan Ho, Ajay Jain, and Pieter Abbeel. 2020. Denoising Diffusion Probabilistic Models. In *Conference on Neural Information Processing Systems (NeurIPS)*.
 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. 2022. Video diffusion models. *arXiv:2204.03458* (2022).
 Steven Hogue, Chenxu Zhang, Hamza Daruger, Yapeng Tian, and Xiaohu Guo. 2024. DiffTED: One-shot Audio-driven TED Talk Video Generation with Diffusion-based Co-Speech Gestures. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*.
 Li Hu, Xin Gao, Peng Zhang, Ke Sun, Bang Zhang, and Liefeng Bo. 2023. Animate Anyone: Consistent and Controllable Image-to-Video Synthesis for Character Animation. *arXiv preprint arXiv:2311.17117* (2023).
 Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
 Yang Jin, Zhicheng Sun, Ningyuan Li, Kun Xu, Hao Jiang, Nan Zhuang, Quzhe Huang, Yang Song, Yadong Mu, and Zhouchen Lin. 2024. Pyramidal Flow Matching for Efficient Video Generative Modeling. *arXiv preprint arXiv:2410.05954* (2024).
 Minguk Kang, Richard Zhang, Connelly Barnes, Sylvain Paris, Suha Kwak, Jaesik Park, Eli Shechtman, Jun-Yan Zhu, and Taesung Park. 2024. Distilling diffusion models into conditional gans. In *European Conference on Computer Vision*. Springer, 428–447.
 Dongjun Kim, Chieh-Hsin Lai, Wei-Hsiang Liao, Naoki Murata, Yuhta Takida, Toshimitsu Uesaka, Yutong He, Yuki Mitsuhashi, and Stefano Ermon. 2024. Consistency Trajectory Models: Learning Probability Flow ODE Trajectory of Diffusion. In *International Conference on Learning Representations (ICLR)*.
 Gaojie Lin, Jianwen Jiang, Chao Liang, Tianyun Zhong, Jiaqi Yang, and Yanbo Zheng. 2024. Cyberhost: Taming audio-driven avatar diffusion model with region codebook attention. *arXiv preprint arXiv:2409.01876* (2024).
 Shanchuan Lin, Xin Xia, Yuxi Ren, Ceyuan Yang, Xuefeng Xiao, and Lu Jiang. 2025. Diffusion adversarial post-training for one-step video generation. *arXiv preprint*

- arXiv:2501.08316 (2025).
- Shanchuan Lin and Xiao Yang. 2024. AnimateDiff-Lightning: Cross-Model Diffusion Distillation. arXiv:2403.12706
- Haiyang Liu, Xingchao Yang, Tomoya Akiyama, Yuantian Huang, Qiaoge Li, Shigeru Kuriyama, and Takafumi Taketomi. 2025. TANGO: Co-Speech Gesture Video Reenactment with Hierarchical Audio Motion Embedding and Diffusion Interpolation. In *International Conference on Learning Representations (ICLR)*.
- Haiyang Liu, Zihao Zhu, Giorgio Becherini, Yichen Peng, Mingyang Su, You Zhou, Xuefei Zhe, Naoya Iwamoto, Bo Zheng, and Michael J Black. 2024b. Emage: Towards unified holistic co-speech gesture generation via masked audio gesture modeling. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xian Liu, Qianyi Wu, Hang Zhou, Yuanqi Du, Wayne Wu, Dahua Lin, and Ziwei Liu. 2022. Audio-Driven Co-Speech Gesture Video Generation. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Xingchao Liu, Xiwen Zhang, Jianzhu Ma, Jian Peng, and Qiang Liu. 2024a. InstaFlow: One Step is Enough for High-Quality Diffusion-based Text-to-Image Generation. In *International Conference on Learning Representations (ICLR)*.
- Eric Luhman and Troy Luhman. 2021. Knowledge Distillation in Iterative Generative Models for Improved Sampling Speed. arXiv preprint arXiv:2101.02388 (2021).
- Weijian Luo, Tianyang Hu, Shifeng Zhang, Jiacheng Sun, Zhenguo Li, and Zhihua Zhang. 2023. Diff-Instruct: A Universal Approach for Transferring Knowledge from Pre-Trained Diffusion Models. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Aniruddha Mahapatra, Richa Mishra, Renda Li, Ziyi Chen, Boyang Ding, Shoulei Wang, Jun-Yan Zhu, Peng Chang, Mei Han, and Jing Xiao. 2024. Co-Speech Gesture Video Generation with 3D Human Meshes. In *European Conference on Computer Vision (ECCV)*.
- Rang Meng, Xingyu Zhang, Yuming Li, and Chenguang Ma. 2025. EchoMimicV2: Towards Striking, Simplified, and Semi-Body Human Animation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Gaurav Parmar, Taesung Park, Srinivasa Narasimhan, and Jun-Yan Zhu. 2024. One-step image translation with text-to-image models. arXiv preprint arXiv:2403.12036 (2024).
- William Peebles and Saining Xie. 2023. Scalable Diffusion Models with Transformers. In *IEEE International Conference on Computer Vision (ICCV)*. 4172–4182. doi:10.1109/ICCV51070.2023.00387
- KR Prajwal, Rudrabha Mukhopadhyay, Vinay P Nambodiri, and CV Jawahar. 2020. A lip sync expert is all you need for speech to lip generation in the wild. In *Proceedings of the 28th ACM international conference on multimedia*. 484–492.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*. 8748–8763.
- Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. 2021. Encoding in Style: a StyleGAN Encoder for Image-to-Image Translation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 10684–10695.
- Ludan Ruan, Yiyang Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin Jin, and Baining Guo. 2023. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 10219–10228.
- Tim Salimans and Jonathan Ho. 2022. Progressive distillation for fast sampling of diffusion models. arXiv preprint arXiv:2202.00512 (2022).
- Axel Sauer, Frederic Boesel, Tim Dockhorn, Andreas Blattmann, Patrick Esser, and Robin Rombach. 2024a. Fast high-resolution image synthesis with latent adversarial diffusion distillation. In *ACM SIGGRAPH Asia*. 1–11.
- Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. 2024b. Adversarial diffusion distillation. In *European Conference on Computer Vision (ECCV)*. Springer, 87–103.
- Scott Schaefer, Travis McPhail, and Joe Warren. 2006. Image deformation using moving least squares. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 533–540.
- Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. 2019a. Animating arbitrary objects via deep motion transfer. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2377–2386.
- Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. 2019b. First order motion model for image animation. *Conference on Neural Information Processing Systems (NeurIPS)* 32 (2019).
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Learning Representations (ICLR)*. pmlr, 2256–2265.
- Jiaming Song, Chenlin Meng, and Stefano Ermon. 2021. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations (ICLR)*.
- Linrui Tian, Siqi Hu, Qi Wang, Bang Zhang, and Liefeng Bo. 2024. EMO2: End-Effector Guided Audio-Driven Avatar Video Generation. arXiv preprint arXiv:2501.10687 (2024).
- Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. 2019. FVD: A new metric for video generation. In *International Conference on Learning Representations (ICLR) Workshop*.
- Aaron Van Den Oord, Oriol Vinyals, et al. 2017. Neural discrete representation learning. *Conference on Neural Information Processing Systems (NeurIPS)* 30 (2017).
- Sheng-Yu Wang, Oliver Wang, Richard Zhang, Andrew Owens, and Alexei A Efros. 2020. CNN-generated images are surprisingly easy to spot...for now. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018. Video-to-Video Synthesis. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Ting-Chun Wang, Arun Mallya, and Ming-Yu Liu. 2021. One-shot free-view neural talking-head synthesis for video conferencing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 10039–10049.
- Zhou Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. 2004. Image Quality Assessment: From Error Visibility to Structural Similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. doi:10.1109/TIP.2003.819861
- Zhengyi Wang, Cheng Lu, Yikai Wang, Fan Bao, Chongxuan Li, Hang Su, and Jun Zhu. 2023. Prolificdreamer: High-fidelity and diverse text-to-3d generation with variational score distillation. *Conference on Neural Information Processing Systems (NeurIPS)* 36 (2023), 8406–8441.
- Haocheng Xi, Shuo Yang, Yilong Zhao, Chenfeng Xu, Muiyang Li, Xiuyu Li, Yujun Lin, Han Cai, Jintao Zhang, Dacheng Li, et al. 2025. Sparse VideoGen: Accelerating Video Diffusion Transformers with Spatial-Temporal Sparsity. In *International Conference on Machine Learning (ICML)*.
- Zhendong Yang, Ailing Zeng, Chun Yuan, and Yu Li. 2023. Effective whole-body pose estimation with two-stages distillation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 4210–4220.
- Zihao Ye, Lequn Chen, Ruihang Lai, Wuwei Lin, Yineng Zhang, Stephanie Wang, Tianqi Chen, Baris Kasicki, Vinod Grover, Arvind Krishnamurthy, and Luis Ceze. 2025. FlashInfer: Efficient and Customizable Attention Engine for LLM Inference Serving. In *Proceedings of the 2025 Conference on Machine Learning and Systems (MLSys)*.
- Hongwei Yi, Hualin Liang, Yifei Liu, Qiong Cao, Yandong Wen, Timo Bolkart, Dacheng Tao, and Michael J Black. 2023. Generating holistic 3d human motion from speech. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 469–480.
- Tianwei Yin, Michael Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and William T Freeman. 2024a. Improved Distribution Matching Distillation for Fast Image Synthesis. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Tianwei Yin, Michael Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. 2024b. One-Step Diffusion with Distribution Matching Distillation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and Xun Huang. 2025. From Slow Bidirectional to Fast Autoregressive Video Diffusion Models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yuanhao Zhai, Kevin Lin, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Chung-Ching Lin, David Doermann, Junsong Yuan, and Lijuan Wang. 2024. Motion Consistency Model: Accelerating Video Diffusion with Disentangled Motion-Appearance Distillation. arXiv preprint arXiv:2406.06890 (2024).
- Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yuang Zhang, Jiaxi Gu, Li-Wen Wang, Han Wang, Junqi Cheng, Yuefeng Zhu, and Fangyuan Zou. 2025. MimicMotion: High-Quality Human Motion Video Generation with Confidence-aware Pose Guidance. In *International Conference on Machine Learning (ICML)*.