

Lightning Fast Caching-based Parallel Denoising Prediction for Accelerating Talking Head Generation

Jianzhi Long¹, Wenhao Sun², Rong-Cheng Tu^{2*}, Dacheng Tao^{2*}

¹The University of Sydney

²Nanyang Technological University

jlon7198@uni.sydney.edu.au, {wenhao006, rongcheng.tu, dacheng.tao}@ntu.edu.sg

Abstract

Diffusion-based talking head models generate high-quality, photorealistic videos but suffer from slow inference, limiting practical applications. Existing acceleration methods for general diffusion models fail to exploit the temporal and spatial redundancies unique to talking head generation. In this paper, we propose a task-specific framework addressing these inefficiencies through two key innovations. First, we introduce **Lightning-fast Caching-based Parallel denoising prediction (LightningCP)**, caching static features to bypass most model layers in inference time. We also enable parallel prediction using cached features and estimated noisy latents as inputs, efficiently bypassing sequential sampling. Second, we propose **Decoupled Foreground Attention (DFA)** to further accelerate attention computations, exploiting the spatial decoupling in talking head videos to restrict attention to dynamic foreground regions. Additionally, we remove reference features in certain layers to bring extra speedup. Extensive experiments demonstrate that our framework significantly improves inference speed while preserving video quality.


Introduction

Talking head generation has received significant attention due to its wide applications in virtual avatars, digital content creation, and real-time communication (Prajwal et al. 2020; Zhang et al. 2023; Wang et al. 2022; Tian et al. 2024). Recent advancements in generative diffusion models (Rombach et al. 2022; Tu et al. 2024) have greatly enhanced the realism and quality of talking head videos (Tian et al. 2024; Xu et al. 2024a; Chen et al. 2025; Zheng et al. 2024), surpassing earlier GAN-based techniques (Wang et al. 2021; Zhang et al. 2023; Wang et al. 2022). However, the computational cost associated with diffusion models—primarily due to iterative denoising steps and large parameter sizes—results in slow inference and hinders real-time deployment (Sun et al. 2025, 2024a). To accelerate diffusion models, previous methods employ model pruning (Fang, Ma, and Wang 2023) and distillation (Salimans and Ho 2022; Meng et al. 2023) techniques, but these often require expensive training.

More recently, DeepCache (Ma, Fang, and Wang 2024) proposes a training-free method that caches and reuses intermediate features to minimize redundant computations,

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

*Corresponding authors.




(a) L_2 Distance.

(b) Cosine similarity matrix.

Figure 1: Analysis of the feature $f_{U_{31}}$ across timesteps in the Hallo model: (a) The L_2 Distance of the feature $f_{U_{31}}$ between consecutive timesteps, (b) The cosine similarity matrix of the feature $f_{U_{31}}$ between all timesteps.

achieving notable acceleration. FasterDiffusion (Li et al. 2023) reuses encoder features to parallelize decoding across multiple timesteps. However, these methods do not consider the unique redundancies in talking head generation, limiting their effectiveness in this task. To this end, we observed several critical redundancies present uniquely in talking head generation and propose a task-specific acceleration framework for maximal efficiency improvement. Fig. 1 illustrates high temporal redundancy in the input feature $f_{U_{31}}$ to the final upsampling layer U_{32} , as measured by both L_2 distance and cosine similarity, particularly at intermediate sampling stages. To exploit this, we propose a caching mechanism termed **Lightning-fast Caching-based Parallel denoising prediction (LightningCP)** that significantly improves inference speed without sacrificing quality. LightningCP caches and reuses $f_{U_{31}}$ across consecutive timesteps to bypass the entire encoder and most of the decoder, reducing model inference to merely a single upsampling layer. We further propose to parallelize the denoising predictions of these consecutive timesteps, significantly improving inference throughput.

Moreover, talking head videos consist of foreground portrait and background environment components, which are associated with unique attention patterns. Fig. 2a and 2b show that foreground tokens are highly correlated with each other while exhibiting low correlation with background tokens; In addition, Fig. 2c, 2d, and 2e demonstrate that the background components of the attention output features remain stable across many timesteps. Based on these insights,



(a) FG noisy latent feature. (b) FG reference feature. (c) BG reference attention. (d) BG audio attention. (e) BG temporal attention.

Figure 2: Average attention scores for foreground noisy latent tokens in the reference attention module (U_{32}) of Hallo, showing attention correlations to (a) FG and BG noisy latent features and (b) FG and BG reference features. The L_2 distance of the BG attention output features in the upsampling layer U_{32} between consecutive timesteps: (c) reference attention, (d) audio attention, and (e) temporal attention. **FG**: foreground. **BG**: background. **std**: standard deviation.

we propose **Decoupled Foreground Attention (DFA)** to perform attention exclusively with foreground tokens while reusing cached background output features, quadratically reducing the attention complexity with respect to foreground token ratio while maintaining generation quality.

In summary, we make the following contributions:

- We introduce a training-free acceleration framework for talking head generation, maximizing inference speed while maintaining or even improving generation quality.
- We propose Lightning-fast Caching-based Parallel denoising prediction (LightningCP), a caching mechanism that reduces UNet inference to a parallelizable upsampling layer by reusing temporally stable decoder features.
- We present Decoupled Foreground Attention (DFA) that reduces the computation complexity of attentions quadratically by leveraging the localization of foreground attention and the stability of background features.
- Extensive experiments demonstrate that our method outperforms existing caching-based acceleration methods in terms of both efficiency and quality.

Related Works

Diffusion-based Talking Head Generation

Diffusion-based talking head generation can be classified into two-stage and end-to-end methods. Two-stage methods (Ma et al. 2023; Liu et al. 2024; Li et al. 2025) predict a sequence of intermediate motion representations (e.g. 3DMM, implicit keypoints) from audio, and then generate talking videos from the predicted motion sequence using pretrained GAN-based generators. Although these methods can achieve real-time inference in some cases (Li et al. 2025), their generation quality is limited by the granularity of the motion representation and the rendering quality of the generator. In contrast, end-to-end methods (Tian et al. 2024; Xu et al. 2024a; Chen et al. 2025; Zheng et al. 2024) leverage the powerful generative priors of latent diffusion model (Rombach et al. 2022) and extend it with 3D convolutions, audio attentions, and temporal modules to generate talking head videos. While these methods can achieve higher generation quality, their slow inference speed hinders practical deployment. To bridge this gap, we propose a training-free approach for end-to-end models that yields substantial speedups while preserving quality, taking a significant step forward in real-world application.

Diffusion Model Acceleration

Diffusion models can be accelerated via efficient samplers (Song, Meng, and Ermon 2020; Lu et al. 2022), model pruning (Fang, Ma, and Wang 2023), and distillation (Salimans and Ho 2022; Meng et al. 2023). However, efficient samplers do not exploit feature redundancy within diffusion models, while pruning and distillation incur expensive model training.

To address these limitations, training-free methods have gained popularity in recent years. ToMe (Bolya and Hoffman 2023) accelerates self-attention by removing less informative tokens, but cannot accelerate talking head generation significantly due to the more complex pipeline. DeepCache (Ma, Fang, and Wang 2024) caches high-level features and efficiently updates only low-level features, but requires inefficient iterative denoising. FasterDiffusion (Li et al. 2023) enables parallel decoding for multiple timesteps by reusing encoder features, but involves costly full decoder inference. In contrast, our cache reduces inference to only a single decoder layer and enables parallel denoising of multiple non-key timesteps with a novel input latents estimation technique.

Foreground-Background Decoupling

The decoupling of foreground (FG) and background (BG) has been used in various downstream tasks of diffusion models. FG-BG segmentation masks are used to preserve subject identity in consistent generation (Wang et al. 2025) and multi-prompt generation (Cai et al. 2025), as well as to preserve background in image and video editing (Zhu et al. 2025; Shuai et al. 2024; Sun et al. 2024b). In talking head generation, our method leverages the static nature of talking head background and employs masks to accelerate model inference for the first time.

Preliminaries

Latent Diffusion Models

Latent diffusion models (LDMs) (Rombach et al. 2022) have become the foundation for high-fidelity image and video generation. The diffusion process consists of two main stages: a forward process and a reverse process, both performed in the latent space. The forward process gradually adds Gaussian noise to a clean latent sample z_0 , resulting in a sequence of increasingly noisy latents $\{z_t\}_{t=1}^T$:

$$z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad \epsilon \sim \mathcal{N}(0, I)$$

Here, z_t denotes the noisy latent at timestep t , and ϵ is standard Gaussian noise. Let $\{\beta_t\}_{t=1}^T$ be a predefined noise schedule with $\beta_t \in (0, 1)$, the signal scaling coefficient α_t is defined as $\alpha_t = 1 - \beta_t$, and the cumulative product $\bar{\alpha}_t$ is expressed as $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$.

The reverse process, parameterized by a neural network, aims to reconstruct the original latent z_0 from z_t by iteratively removing noise. The DDIM sampling strategy is widely adopted for efficient and deterministic inference, with the reverse update at each timestep as:

$$z_{t-1} = \lambda_t z_t + \tau_t \cdot \epsilon_\theta(z_t, t, c), \quad \lambda_t = \sqrt{\frac{\alpha_{t-1}}{\alpha_t}}, \quad (1)$$

$$\tau_t = \sqrt{\alpha_{t-1}} \left(\sqrt{\frac{1}{\alpha_{t-1}} - 1} - \sqrt{\frac{1}{\alpha_t} - 1} \right)$$

where $\epsilon_\theta(z_t, t, c)$ is the predicted noise at timestep t , conditioned on z_t , the timestep t , and the condition vector c that is typically a text prompt embedding.

Diffusion-based Talking Head Generation Model

Diffusion-based talking head generation models combine a backbone based on Stable Diffusion (Rombach et al. 2022) with additional modules to perform audio-driven animation. The backbone’s denoising UNet iteratively predicts and removes noise from the noisy latent guided by the reference feature and the audio feature. The reference feature is extracted by Reference Net (Xu et al. 2024b) from the input image to ensure identity consistency throughout denoising. To incorporate the reference feature, a reference attention module extends the self-attention of the denoising Unet by concatenating the noisy latent with the reference feature as key and value inputs. An audio attention module injects motion information into the denoising process by performing cross-attention between the noisy latent and audio feature, producing lip movements and head motions synced with the driving audio. A temporal layer performs self-attention on the frame dimension, capturing subtle transitions and enhancing frame-to-frame coherence. Finally, the denoised latent features are decoded by the VAE decoder into the talking head video frames.

At each timestep t of the denoising process, the latent representation $z_t \in \mathcal{R}^{b,c,f,h,w}$ of a video clip is progressively refined, where (b, c, f, h, w) represent batch size, channels, number of frames, image height, and image width. This iterative update can be formally expressed as follows:

$$z_{t-1} = \lambda_t z_t + \tau_t \cdot \epsilon_\theta(z_t, t, c), \quad c = \{R, A, c_{\text{others}}\} \quad (2)$$

Specifically, the condition feature $c = \{R, A, c_{\text{others}}\}$ includes reference feature R , audio features A , and other customized condition features c_{others} which may include face identity embedding (Xu et al. 2024a) or emotion features (Zheng et al. 2024). In the denoising process, the spatial attention modules concatenate the reference feature R directly with noisy latent. Audio features A spans the f frames within the video clip. Finally, the noisy latent is reshaped

from $b \times c \times f \times h \times w$ to $(b \times h \times w) \times f \times c$ as the input to the temporal layer. A key distinction between talking head generation and general video generation is that the former produces videos of varying lengths determined by the input audio clip duration. To accommodate this, the driving audio clip is segmented into fixed-length intervals, each corresponding to individual video clips of length f , which are then combined to form the complete video output.

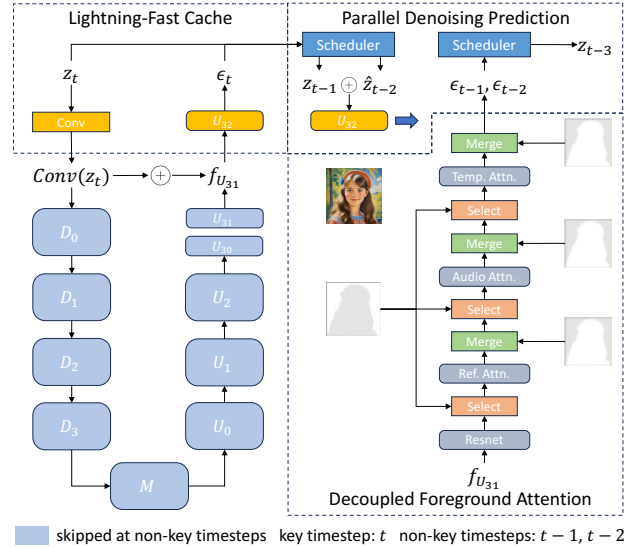


Figure 3: The pipeline of the accelerated talking head model. At key timestep t , we perform full model inference and cache feature $f_{U_{31}}$. At non-key timesteps $t-1$ and $t-2$, we reuse cached $f_{U_{31}}$ and bypass the encoder (D_0, D_1, D_2, D_3), midblock M , and all of the decoder (U_0, U_1, U_2, U_3) except its last layer U_{32} . Moreover, denoising prediction at $t-1$ and $t-2$ can be executed in parallel and further accelerated through decoupled foreground attention.

Method

In this section, we detail the proposed acceleration methods for diffusion-based talking head generation. Our approach introduces Lightning-fast Caching-based Parallel denoising prediction (LCP) and Decoupled Foreground Attention (DFA), significantly reducing inference times while maintaining or even improving video quality. In addition, auxiliary techniques, including input latent estimation and reference feature removal, are proposed to further enhance the visual quality of specific diffusion-based talking head models.

Lightning-fast Caching-based Parallel Denoising Prediction

As shown in Fig. 3, diffusion-based talking head generation models typically adopt a UNet architecture, with four encoder (downsampling) blocks D_0, D_1, D_2, D_3 , a mid-block M , and four decoder (upsampling) blocks U_0, U_1, U_2, U_3 . In each denoising step, the input latent is passed sequentially through the encoder, mid-block, and decoder for denoising

prediction. Since this full forward pass is repeated for every diffusion timestep, the process incurs substantial computation cost, especially with large latent diffusion models.

We observe that for most denoising timesteps in talking head generation, all downsampling operations in the encoder and the majority of the upsampling layers in the decoder are redundant. By caching and reusing high-level decoder features across timesteps, it is feasible to bypass the encoder entirely and all but the final upsampling layer of the decoder. As illustrated in Fig. 1, the input feature $f_{U_{31}}$ to the final layer U_{32} of the last upsampling block exhibits substantial temporal similarity across consecutive denoising steps. Specifically, Fig. 1a depicts the feature L_2 distance between adjacent timesteps, while Fig. 1b presents the feature cosine similarity matrix computed across all timesteps. Both metrics confirm that $f_{U_{31}}$ maintains high temporal similarity, especially in the middle timesteps. This insight motivates us to cache $f_{U_{31}}$ at selected key timesteps and reuse it in subsequent non-key timesteps, significantly reducing computation cost without compromising the fidelity of the generated output.


Lightning-fast Cache Motivated by this observation, we propose a lightweight caching strategy termed Lightning-fast Cache, which enables highly efficient inference by partially bypassing the U-Net architecture. Fig. 3 shows the pipeline of our proposed method. At each key timestep t , we execute a full forward pass of the UNet and cache the intermediate feature $f_{U_{31}}^t$. Then, for the following non-key timesteps $t-1$ and $t-2$, we skip all encoder and early decoder computations by reusing the cached $f_{U_{31}}^t$.

At non-key timesteps like $t-1$, the denoising network ϵ_θ simplifies to a subnetwork $\hat{\epsilon}_\theta$, computing only the final upsampling operation U_{32} using the cached $f_{U_{31}}^t$ and a lightweight convolution of the current latent z_t :

$$\epsilon_{t-1} = \hat{\epsilon}_\theta(f_{U_{31}}^t, z_{t-1}, t-1, c) \quad (3)$$

$$= U_{32}(f_{U_{31}}^t, \text{conv}(z_{t-1}), t-1, c) \quad (4)$$

where $\text{conv}(z_t)$ denotes a shallow convolution applied to the latent z_t to generate the skip-connection input for U_{32} , and c denotes additional conditioning (e.g. audio feature, reference feature).



(a) Input latents.

(b) Predicted noise.

Figure 4: L_2 distance between consecutive timesteps for: (a) input latents, (b) predicted noise. t_{thresh} is the threshold timestep after which input latents estimation is applied.

Parallel Denoising Prediction With Lightning-fast Cache, we further accelerate denoising prediction of consecutive non-key timesteps in parallel on multiple GPUs, because inputs

of these non-key timesteps share the same cached feature and differ only in their timestep-specific variables.

Given a key timestep t , we cache the intermediate feature $f_{U_{31}}^t$ and reuse it across a block of consecutive non-key timesteps $\{t-1, t-2, \dots, t-N\}$. Then, the denoising equation of non-key timestep $t-i$ becomes:

$$z_{t-1-i} = \lambda_{t-i} z_{t-i} + \tau_{t-i} \cdot \hat{\epsilon}_\theta(f_{U_{31}}^t, z_{t-i}, t-i, c), \quad (5)$$

s.t. $1 \leq i \leq N-1$.

However, there is a key challenge in parallelization: when $i > 1$, the input latents z_{t-i} for each non-key timestep are not available in advance. A straightforward workaround is to approximate all unknown input latents using the latent z_{t-1} computed at the key timestep.

Unfortunately, reusing the same input latent z_{t-1} across multiple timesteps $t-2, \dots, t-N$ can lead to significant error accumulation, ultimately degrading video quality. As illustrated in Fig. 4a, the feature similarity between consecutive input latents decreases substantially over time when using samplers like DDIM (Song, Meng, and Ermon 2020). This compounding discrepancy highlights the risk of not updating the input latents during parallel denoising prediction.

To address this issue, we propose **Input Latents Estimation**, which estimates input latents for consecutive non-key timesteps in parallel prediction by leveraging the temporal feature similarity of the predicted noise. As illustrated in Fig. 4b, the predicted noise remain stable across the timesteps where the input latents show increasing discrepancies. Specifically, we use the diffusion scheduler to estimate the input latents $\hat{z}_{t-2}, \dots, \hat{z}_{t-N}$ for non-key timesteps $\{t-2, \dots, t-N\}$ using the predicted noise $\epsilon_\theta(z_t, t, c)$ at key timestep t through Eq. 6a, and then use them to compute the actual latents $\{z_{t-2}, \dots, z_{t-N}\}$ via Eq. 6b. For timestep $t-2$, \hat{z}_{t-2} is estimated using the latent z_{t-1} computed at key timestep t ; For subsequent timesteps after $t-2$, we iteratively estimate input latents \hat{z}_{t-i} using the previously estimated latent \hat{z}_{t-i+1} as the reference.

$$\hat{z}_{t-i} = \begin{cases} \lambda_{t-1} z_{t-1} + \tau_{t-1} \cdot \epsilon_\theta(z_t, t, c) & \text{for } i = 2 \\ \lambda_{t-i+1} \hat{z}_{t-i+1} + \tau_{t-i+1} \cdot \epsilon_\theta(z_t, t, c) & \text{for } i > 2 \end{cases} \quad (6a)$$

$$z_{t-1-i} = \lambda_{t-i} z_{t-i} + \tau_{t-i} \cdot \hat{\epsilon}_\theta(f_{U_{31}}^t, \hat{z}_{t-i}, t-i, c), \quad (6b)$$

s.t. $1 \leq i \leq N-1$.

This two-step refinement allows us to recover higher-quality output latents by using more accurate model inputs. In practice, we apply Input Latents Estimation after a designated threshold timestep t_{thresh} . In the early denoising phase ($t < t_{\text{thresh}}$), the input latent z_t exhibits high similarity across steps, so parallel prediction without correction suffices. Beyond this threshold, the quality benefit of estimating \hat{z}_{t-i} becomes significant. Our experiments confirm that Input Latents Estimation greatly improves video fidelity, particularly for high-speed samplers such as DDIM (Song, Meng, and Ermon 2020), while incurring negligible extra computation.

Decoupled Foreground Attention

To further accelerate inference, we propose Decoupled Foreground Attention (DFA) to streamline the reference attention, self-attention in audio module, and temporal attentions in U_{32} . The computation cost of these attentions scales quadratically with the sequence length of the video tokens, which is as large as $64 \times 64 = 4096$ in U_{32} . To reduce this cost, DFA restricts the attention computation to the foreground components of the input features. This reduction of spatial tokens decrease both query and key length, effectively reducing the attention complexity from $\mathcal{O}(L^2)$ to $\mathcal{O}(L_f^2)$, where L_f is the number of foreground tokens.

Our design is based on two empirical observations:

- **Foreground-Localized Attention:** Foreground attention maps are primarily concentrated within the foreground region, indicating spatial decoupling between dynamic facial motion and static background content. Fig. 2a and 2b visualizes this phenomenon in the reference attention module of the upsampling layer U_{32} . Across all the foreground noisy latent tokens, we calculate the average attention weight sum per token with regard to tokens of 4 distinct groups: foreground noisy latent, background noisy latent, foreground reference feature, and background reference feature. The resulting analysis shows strong localization of attention within foreground tokens and weak correlation between foreground and background tokens.
- **Temporal Redundancy in Background:** Background attention features remain stable across consecutive timesteps, making them suitable for caching and reuse. Fig. 2 visualizes the L_2 distance for the background component of the attention output features in the upsampling layer U_{32} . The analysis on reference (Fig. 2c), audio (Fig. 2d), and temporal attentions (Fig. 2e) collectively reveal high feature similarity, particularly in the middle denoising timesteps, supporting the strategy of caching background features to reduce computational redundancy.

Foreground Masking for Token Reduction. We use a 64×64 face segmentation mask to separate the foreground tokens. The mask is downsampled from the reference image’s face segmentation mask, which is obtained using an off-the-shelf face parsing model (Yakhyokhuja 2024). Let $M \in \{0, 1\}^{H \times W}$ denote the binary foreground segmentation mask, where an element value of 1 indicates that the corresponding spatial location belongs to the foreground region. Given flattened attention features $Q, K, V \in \mathbb{R}^{L \times d}$, where $L = H \cdot W$, we select the foreground token subset as $Q_{\mathcal{F}}, K_{\mathcal{F}}, V_{\mathcal{F}} \in \mathbb{R}^{L_f \times d}$, where $L_f = \|M\|_1 = \sum_{ij} M_{ij}$.

The attention output is then computed over the reduced foreground region:

$$A_{\mathcal{F}} = \text{softmax} \left(\frac{Q_{\mathcal{F}} K_{\mathcal{F}}^{\top}}{\sqrt{d}} \right) V_{\mathcal{F}} \in \mathbb{R}^{L_f \times d}. \quad (7)$$

This leads to a quadratic reduction in attention complexity—from $\mathcal{O}(L^2)$ to $\mathcal{O}(L_f^2)$ —proportional to the foreground token ratio. In reference attention, we apply the same foreground mask to both the noise and reference features. Similarly, in the audio and temporal attentions, the computation is restricted to foreground tokens.

Merging with Cached Background. Once the updated foreground attention output $A_{\mathcal{F}}$ is obtained, we reconstruct the complete attention feature map by merging it with the cached background output $A_{\mathcal{B}}^{(t^*)} \in \mathbb{R}^{L_b \times d}$ from the most recent key timestep t^* , where $L_b = L - L_f$ denotes the number of background tokens. The final combined feature is given by:

$$A = \text{Merge}(A_{\mathcal{F}}, A_{\mathcal{B}}^{(t^*)}) \in \mathbb{R}^{L \times d}, \quad (8)$$

By narrowing the expensive attentions to the foreground and reusing temporally stable background features, DFA achieves further acceleration with negligible impact on generation quality.

Reference Feature Removal Reference attention incurs the highest computation cost in the entire pipeline because the concatenation of reference feature makes the key and value twice as long. Our empirical study shows that reference features in certain layers do not affect video quality, an effect that varies among different talking head models. In our experiments, we apply reference feature removal on all of the tested talking head models to achieve additional acceleration, while maintaining or even improving the video quality.

Experiments

Experimental Settings

Models, Datasets, and Evaluation Metrics We evaluate our method on four diffusion-based talking head generation models: Hallo (Xu et al. 2024a), MEMO (Zheng et al. 2024), EchoMimic, and Accelerated EchoMimic (EchoMimic Acc.) (Chen et al. 2025), which accelerates inference by reducing the number of diffusion steps and removing classifier-free guidance (Ho and Salimans 2022). For testing, we use 100 videos from the HDTF dataset (Zhang et al. 2021) and 100 videos from the MEAD dataset (Wang et al. 2020).

To assess video quality, we use standard evaluation metrics for talking head generation: Fréchet Video Distance (FVD) (Unterthiner et al. 2018), Fréchet Inception Distance (FID) (Heusel et al. 2017), expression-FID (E-FID) (Tian et al. 2024), and Lip Sync (Prajwal et al. 2020). Specifically, E-FID quantifies the expressiveness of facial motion by computing FID between 3DMM (Deng et al. 2019) expression coefficients extracted from generated and ground-truth videos. To evaluate generation speed, we report latency as the UNet inference time required to generate a fixed-length video clip. Since the final video length varies with input audio duration, latency and FLOPs are measured per clip for fair comparison. Speedup is computed as the latency of the base model divided by the latency of the accelerated model.

Baselines We compare our method with two training-free diffusion acceleration baselines: DeepCache (Ma, Fang, and Wang 2024) and FasterDiffusion (Li et al. 2023). For fair comparison, we use identical key timesteps across all acceleration methods for each model, as timestep selection substantially affects efficiency and output quality. All experiments are run using A800 GPUs. We use a single GPU to test latency of the base model and DeepCache, while using multi-GPUs to test latency of FasterDiffusion and our method.

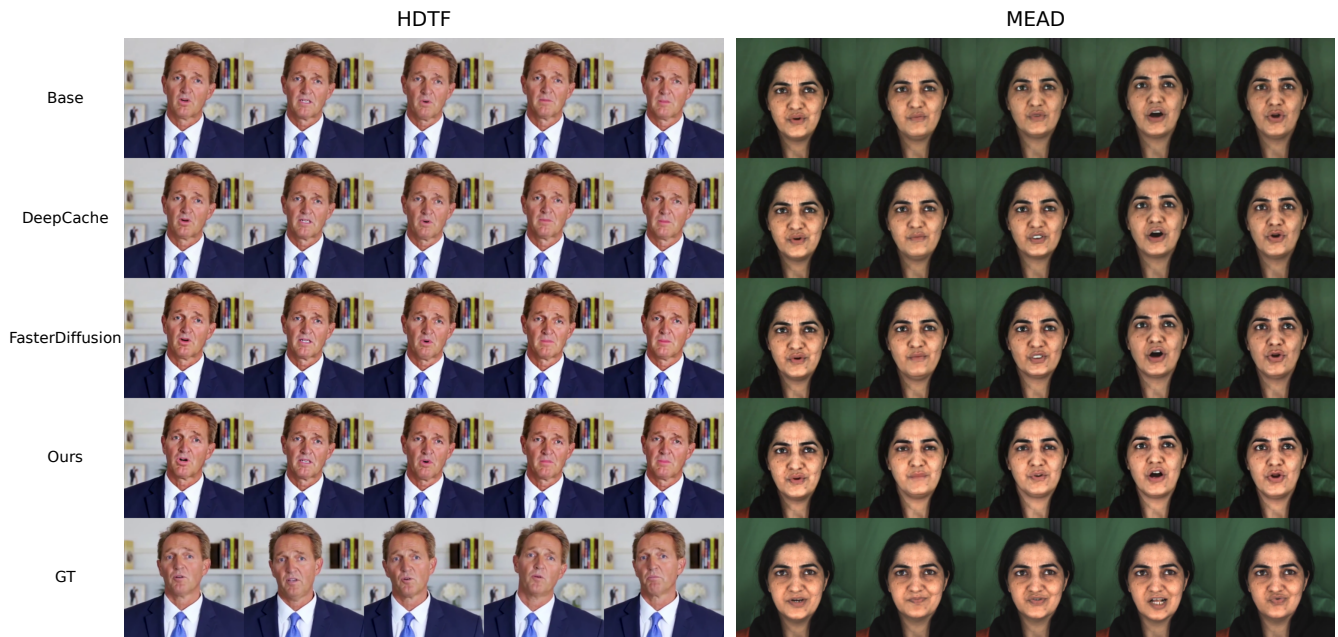


Figure 5: Qualitative result on HDTF and MEAD datasets using the Hallo model. From top to bottom: results from the base model, DeepCache, FasterDiffusion, our proposed method, and Ground Truth (GT).

Performance Analysis

Table 1 presents results on HDTF (Zhang et al. 2021) and MEAD (Wang et al. 2020) datasets, comparing our method with base models and acceleration baselines (Ma, Fang, and Wang 2024; Li et al. 2023) across multiple diffusion-based talking head models (Xu et al. 2024a; Zheng et al. 2024; Chen et al. 2025). Our method achieves the most inference acceleration, with the highest speedup on Hallo (3.15 \times) (Xu et al. 2024a) and EchoMimic (3.10 \times) (Chen et al. 2025). While MEMO (Zheng et al. 2024) and EchoMimic Acc. (Chen et al. 2025) have fewer diffusion steps—limiting the number of non-key timesteps—our method is still the fastest on every model. Furthermore, our approach achieves the lowest FLOPs in all settings, demonstrating superior computational efficiency via targeted optimizations: caching temporally stable decoder features and restricting high-cost attention operations to the dynamic foreground.

More importantly, this speedup is achieved without compromising visual quality. As shown in Fig. 5, our method maintains the visual fidelity, identity consistency, and facial movement expressiveness of the base model. In addition, it facilitates the generation of more synchronized lip movements compared to the base model and the acceleration baselines. On the HDTF dataset, our method consistently yields the lowest FVD across all models and the highest Sync scores in most cases, while maintaining FID and E-FID comparable to the best-performing baselines. On MEAD, our approach achieves the lowest FID and the highest Sync scores in most cases. These results highlight the effectiveness of the proposed mechanisms: re-using cached decoder feature in Lightning-fast Cache preserves the feature’s temporal coherency across timesteps, and Decoupled Foreground Atten-

tion preserves motion expressiveness. While other baselines may show competitive performance in isolated settings (e.g., DeepCache on Hallo or FasterDiffusion on EchoMimic Acc.), they often perform poorly on other models or datasets. In contrast, our method demonstrates consistent quality across all settings. Additionally, FasterDiffusion requires 3 extra key timesteps on EchoMimic to avoid severe video quality degradation, significantly diminishing its speedup. In comparison, our method retains both high video quality and high speed under the original settings, further validating its robustness and generalization ability.

Ablation Study

To assess the contribution of each component of our method, we conduct ablations on HDTF (Zhang et al. 2021) and MEAD (Wang et al. 2020) using three variants: (1) Lightning Fast Caching-based Parallel denoising prediction (LCP), (2) LCP + Decoupled Foreground Attention (DFA), and (3) full model with LCP + DFA + reference feature removal (rm), as shown in Table 2.

Adding LCP to the base model significantly accelerates inference speed while maintaining decent video quality. DFA reduces FLOPs and introduces additional speedup across all scenarios, notably achieving a 8.9% FLOPs reduction on Hallo and pushes speedup above 3. Additionally, it retains or even improves video quality, indicating that removing background tokens allow the model to preserve essential interactions within the face region. Reference feature removal further enhances lip sync, suggesting that eliminating redundant static appearance features in certain layers enables better modeling of fine-grained lip motion.

We also evaluate the effect of input latent estimation on a

Method	Variants	HDTF				MEAD				FLOPS ↓ ($\times 10^{12}/\text{clip}$)	Latency ↓ (s/clip)	Speedup ↑
		FVD ↓	FID ↓	E-FID ↓	Sync ↑	FVD ↓	FID ↓	E-FID ↓	Sync ↑			
Hallo	Base	135.32	4.69	0.157	7.38	153.42	6.17	0.297	6.03	2158.08	23.692	-
	DeepCache	129.31	4.62	0.157	7.52	137.56	6.26	0.276	6.15	1041.28	12.884	1.84
	FasterDiffusion	140.56	5.63	0.200	7.47	168.72	7.33	0.285	6.14	1561.64	<u>9.032</u>	<u>2.62</u>
	Ours	121.80	4.89	0.178	7.66	158.62	6.64	0.277	6.43	671.40	7.528	3.15
MEMO	Base	82.07	3.40	0.159	8.10	82.90	5.63	0.161	6.06	919.86	14.934	-
	DeepCache	82.52	3.42	0.162	8.17	85.28	5.66	0.170	6.09	491.14	9.212	1.62
	FasterDiffusion	98.12	4.17	0.232	8.23	93.49	5.89	0.221	6.24	712.54	6.740	<u>2.22</u>
	Ours	80.23	3.46	0.163	8.17	79.18	5.54	0.171	6.12	406.08	6.416	2.33
EchoMimic	Base	115.78	4.43	0.116	5.75	166.82	5.43	0.335	4.89	965.19	10.491	-
	DeepCache	165.42	4.57	0.093	5.86	234.59	6.65	0.384	4.83	437.97	5.460	1.92
	FasterDiffusion [†]	126.27	4.69	0.109	5.90	175.23	5.67	0.467	4.79	820.05	4.788	2.19
	Ours	116.94	4.55	0.103	6.08	195.57	5.45	0.391	4.93	286.77	3.387	3.10
EchoMimic (Acc.)	Base	145.35	5.06	0.196	6.27	213.72	6.27	0.361	5.33	90.27	1.0458	-
	DeepCache	146.56	5.10	0.203	6.40	214.52	6.15	0.367	5.49	67.49	0.8268	1.26
	FasterDiffusion	143.41	5.92	0.226	6.52	157.62	6.65	0.331	5.54	68.93	0.7680	1.36
	Ours	136.14	5.18	0.201	6.58	192.90	5.82	0.359	5.61	60.52	0.7236	1.45

Table 1: Quantitative results on HDTF and MEAD. [†] We include 3 extra key timesteps for FasterDiffusion on EchoMimic to avoid severe video quality degradation.

Method	Variants	HDTF				MEAD				FLOPS ↓ ($\times 10^{12}/\text{clip}$)	Latency ↓ (s/clip)	Speedup ↑
		FVD ↓	FID ↓	E-FID ↓	Sync ↑	FVD ↓	FID ↓	E-FID ↓	Sync ↑			
Hallo	LCP	124.62	4.99	0.189	7.53	151.8	6.54	0.268	6.17	745.72	8.052	2.94
	LCP + DFA	128.31	5.10	0.198	7.57	154.48	6.79	0.274	6.19	679.56	7.688	3.08
	LCP + DFA + rm	121.80	4.89	0.178	7.66	158.62	6.64	0.277	6.43	671.44	7.528	3.15
MEMO	LCP	83.32	3.45	0.169	8.21	80.26	5.55	0.168	6.17	411.02	6.640	2.25
	LCP + DFA	80.49	3.43	0.162	8.16	79.57	5.50	0.168	6.11	406.82	6.500	2.30
	LCP + DFA + rm	80.23	3.46	0.163	8.17	79.18	5.54	0.171	6.12	406.08	6.416	2.33
EchoMimic	LCP	136.87	4.50	0.098	5.96	218.06	5.86	0.384	4.84	319.26	3.447	3.04
	LCP + DFA	116.33	4.64	0.099	5.98	180.79	5.56	0.388	4.94	291.21	3.405	3.08
	LCP + DFA + rm	116.94	4.55	0.103	6.08	195.57	5.45	0.391	4.93	286.77	3.387	3.10
EchoMimic (Acc.)	LCP	144.09	5.15	0.207	6.50	201.47	5.99	0.363	5.50	62.69	0.7368	1.42
	LCP + DFA	138.24	5.20	0.206	6.47	179.33	5.94	0.359	5.53	61.33	0.7326	1.43
	LCP + DFA + rm	136.14	5.18	0.201	6.58	192.90	5.82	0.359	5.61	60.52	0.7236	1.45

Table 2: Ablation study results on HDTF and MEAD. **LCP**: Lightning-fast Caching-based Parallel denoising prediction. **DFA**: Decoupled Foreground Attention. **rm**: reference feature removal.

Method	Variants	HDTF			
		FVD ↓	FID ↓	E-FID ↓	Sync ↑
Hallo	LCP w/o estimate	208.46	7.14	0.355	7.64
	LCP w/ estimate	200.33	7.04	0.352	7.63
EchoMimic	LCP w/o estimate	460.21	6.66	0.226	5.98
	LCP w/ estimate	225.16	6.22	0.233	6.05

Table 3: Ablation study results of input latents estimation on HDTF. **w/ estimate**: with input latents estimation. **w/o estimate**: without input latents estimation.

subset of HDTF. Table 3 shows that without this estimation, severe degradation in FVD, FID, and E-FID occurs, confirming that the drastic decrease in the similarity of input latent features across timesteps leads to significant error accumulation. Therefore, it is necessary to alleviate this error by estimating the input latents using the temporally stable pre-

dicted noise, which significantly improves generation quality with negligible computational overhead.

Conclusion

In this paper, we introduce Lightning-fast Caching-based Parallel denoising prediction (LightningCP), a training-free acceleration framework for diffusion-based talking head generation. By caching and reusing high-level decoder features, our approach enables efficient parallel inference across multiple denoising timesteps. Additionally, Decoupled Foreground Attention (DFA) further enhances inference efficiency by restricting attention to the foreground region. We also employ auxiliary techniques to improve generation quality via input latents estimation and introduce extra speedup via reference feature removal. Extensive experiment validates the superior efficiency and quality of our method, which provides a practical and plug-and-play solution for accelerating diffusion-based talking head generation.

Ethical Statement

This research is intended to support the development of ethical virtual humans and creative applications in entertainment and communication. However, as with other generative talking head technologies, there is a risk that the proposed acceleration framework could be misused for malicious purposes such as fraud, impersonation, or defamation (Liu et al. 2025). Moreover, the increasing prevalence of AI-generated media may contaminate public datasets and negatively impact the training of downstream AI models (Tu et al. 2025c,a), potentially leading to degraded or unsafe system behavior.

To mitigate these risks, we recommend that any generated talking head videos be clearly watermarked or labeled as AI generated. We are also committed to assist the ongoing community efforts to develop robust evaluations and safety protocols (Tu et al. 2025b), for promoting safe and trustworthy deployment of generative technologies.

Acknowledgements

This research is supported by the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore through Alibaba-NTU Global e-Sustainability CorpLab (ANGEL).

References

- Bolya, D.; and Hoffman, J. 2023. Token merging for fast stable diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4599–4603.
- Cai, M.; Cun, X.; Li, X.; Liu, W.; Zhang, Z.; Zhang, Y.; Shan, Y.; and Yue, X. 2025. Ditctrl: Exploring attention control in multi-modal diffusion transformer for tuning-free multi-prompt longer video generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 7763–7772.
- Chen, Z.; Cao, J.; Chen, Z.; Li, Y.; and Ma, C. 2025. Echomimic: Lifelike audio-driven portrait animations through editable landmark conditions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 2403–2410.
- Deng, Y.; Yang, J.; Xu, S.; Chen, D.; Jia, Y.; and Tong, X. 2019. Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 0–0.
- Fang, G.; Ma, X.; and Wang, X. 2023. Structural Pruning for Diffusion Models. In *Advances in Neural Information Processing Systems*. ArXiv:2305.10924.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Li, S.; Hu, T.; Khan, F. S.; Li, L.; Yang, S.; Wang, Y.; Cheng, M.-M.; and Yang, J. 2023. Faster diffusion: Rethinking the role of unet encoder in diffusion models. *CoRR*.
- Li, T.; Zheng, R.; Yang, M.; Chen, J.; and Yang, M. 2025. Ditto: Motion-space diffusion for controllable realtime talking head synthesis. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 9704–9713.
- Liu, J.; Liang, S.; Zhao, S.; Tu, R.; Zhou, W.; Liu, A.; Tao, D.; and Lam, S. K. 2025. T2V-OptJail: Discrete Prompt Optimization for Text-to-Video Jailbreak Attacks. *arXiv preprint arXiv:2505.06679*.
- Liu, T.; Chen, F.; Fan, S.; Du, C.; Chen, Q.; Chen, X.; and Yu, K. 2024. Anitalker: animate vivid and diverse talking faces through identity-decoupled facial motion encoding. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 6696–6705.
- Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2022. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural Information Processing Systems*, 35: 5775–5787.
- Ma, X.; Fang, G.; and Wang, X. 2024. Deepcache: Accelerating diffusion models for free. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15762–15772.
- Ma, Y.; Zhang, S.; Wang, J.; Wang, X.; Zhang, Y.; and Deng, Z. 2023. Dreamtalk: When expressive talking head generation meets diffusion probabilistic models. *arXiv preprint arXiv:2312.09767*, 2(3).
- Meng, C.; Rombach, R.; Gao, R.; Kingma, D.; Ermon, S.; Ho, J.; and Salimans, T. 2023. On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14297–14306.
- Prajwal, K.; Mukhopadhyay, R.; Namboodiri, V. P.; and Jawahar, C. 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.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Salimans, T.; and Ho, J. 2022. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*.
- Shuai, X.; Ding, H.; Ma, X.; Tu, R.; Jiang, Y.-G.; and Tao, D. 2024. A survey of multimodal-guided image editing with text-to-image diffusion models. *arXiv preprint arXiv:2406.14555*.
- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Sun, W.; Tu, R.-C.; Ding, Y.; Jin, Z.; Liao, J.; Liu, S.; and Tao, D. 2025. VORTA: Efficient Video Diffusion via Routing Sparse Attention. *arXiv preprint arXiv:2505.18809*.
- Sun, W.; Tu, R.-C.; Liao, J.; Jin, Z.; and Tao, D. 2024a. Asymmr: Video diffusion transformers acceleration with asymmetric reduction and restoration. *arXiv preprint arXiv:2412.11706*.
- Sun, W.; Tu, R.-C.; Liao, J.; and Tao, D. 2024b. Diffusion model-based video editing: A survey. *arXiv preprint arXiv:2407.07111*.

- Tian, L.; Wang, Q.; Zhang, B.; and Bo, L. 2024. Emo: Emote portrait alive generating expressive portrait videos with audio2video diffusion model under weak conditions. In *European Conference on Computer Vision*, 244–260. Springer.
- Tu, R.-C.; Ji, Y.; Jiang, J.; Kong, W.; Cai, C.; Zhao, W.; Wang, H.; Yang, Y.; and Liu, W. 2025a. Global and local semantic completion learning for vision-language pre-training. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Tu, R.-C.; Ma, Z.-A.; Lan, T.; Zhao, Y.; Huang, H.-Y.; and Mao, X.-L. 2025b. Automatic evaluation for text-to-image generation: Task-decomposed framework, distilled training, and meta-evaluation benchmark. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 22340–22361.
- Tu, R.-C.; Mao, X.-L.; Liu, J.-Y.; Ma, Z.-A.; Lan, T.; and Huang, H. 2025c. Prospective layout-guided multi-modal online hashing. *IEEE Transactions on Image Processing*.
- Tu, R.-C.; Sun, W.; Jin, Z.; Liao, J.; Huang, J.; and Tao, D. 2024. Spagent: Adaptive task decomposition and model selection for general video generation and editing. *arXiv preprint arXiv:2411.18983*.
- Unterthiner, T.; Van Steenkiste, S.; Kurach, K.; Marinier, R.; Michalski, M.; and Gelly, S. 2018. Towards accurate generative models of video: A new metric & challenges. *arXiv preprint arXiv:1812.01717*.
- Wang, K.; Wu, Q.; Song, L.; Yang, Z.; Wu, W.; Qian, C.; He, R.; Qiao, Y.; and Loy, C. C. 2020. Mead: A large-scale audio-visual dataset for emotional talking-face generation. In *European Conference on Computer Vision*, 700–717. Springer.
- Wang, M.; Ding, H.; Peng, J.; Zhao, Y.; Chen, Y.; and Wei, Y. 2025. Characonsist: Fine-grained consistent character generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16058–16067.
- Wang, S.; Li, L.; Ding, Y.; Fan, C.; and Yu, X. 2021. Audio2head: Audio-driven one-shot talking-head generation with natural head motion. *arXiv*.
- Wang, S.; Li, L.; Ding, Y.; and Yu, X. 2022. One-shot talking face generation from single-speaker audio-visual correlation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2531–2539.
- Xu, M.; Li, H.; Su, Q.; Shang, H.; Zhang, L.; Liu, C.; Wang, J.; Yao, Y.; and Zhu, S. 2024a. Hallo: Hierarchical audio-driven visual synthesis for portrait image animation. *arXiv preprint arXiv:2406.08801*.
- Xu, Z.; Zhang, J.; Liew, J. H.; Yan, H.; Liu, J.-W.; Zhang, C.; Feng, J.; and Shou, M. Z. 2024b. Magicanimate: Temporally consistent human image animation using diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1481–1490.
- Yakhyokhuja, V. 2024. face-parsing. <https://github.com/yakhyo/face-parsing>. GitHub repository.
- Zhang, W.; Cun, X.; Wang, X.; Zhang, Y.; Shen, X.; Guo, Y.; Shan, Y.; and Wang, F. 2023. SadTalker: Learning Realistic 3D Motion Coefficients for Stylized Audio-Driven Single Image Talking Face Animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8652–8661.
- Zhang, Z.; Li, L.; Ding, Y.; and Fan, C. 2021. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3661–3670.
- Zheng, L.; Zhang, Y.; Guo, H.; Pan, J.; Tan, Z.; Lu, J.; Tang, C.; An, B.; and Yan, S. 2024. MEMO: Memory-Guided Diffusion for Expressive Talking Video Generation. *arXiv preprint arXiv:2412.04448*.
- Zhu, T.; Zhang, S.; Shao, J.; and Tang, Y. 2025. Kv-edit: Training-free image editing for precise background preservation. *arXiv preprint arXiv:2502.17363*.