

Forecast then Calibrate: Feature Caching as ODE for Efficient Diffusion Transformers

Shikang Zheng^{1,2}, Liang Feng¹, Xinyu Wang^{1,3}, Qinming Zhou^{1,3}, Peiliang Cai¹, Chang Zou¹, Jiacheng Liu¹, Yuqi Lin¹, Junjie Chen¹, Yue Ma⁴, Linfeng Zhang^{1,†}

¹Shanghai Jiao Tong University, China

²South China University of Technology, China

³Tsinghua University, China

⁴The Hong Kong University of Science and Technology, China

Abstract

Diffusion Transformers (DiTs) have demonstrated exceptional performance in high-fidelity image and video generation. To reduce their substantial computational costs, feature caching techniques have been proposed to accelerate inference by reusing hidden representations from previous timesteps. However, current methods often struggle to maintain generation quality at high acceleration ratios, where prediction errors increase sharply due to the inherent instability of long-step forecasting. In this work, we adopt an ordinary differential equation (ODE) perspective on the hidden-feature sequence, modeling layer representations along the trajectory as a feature-ODE. We attribute the degradation of existing caching strategies to their inability to robustly integrate historical features under large skipping intervals. To address this, we propose FoCa (Forecast-then-Calibrate), which treats feature caching as a feature-ODE solving problem. Extensive experiments on image, video generation, and super-resolution tasks demonstrate the effectiveness of FoCa, especially under aggressive acceleration. Without additional training, FoCa achieves near-lossless speedups of 5.50 \times on FLUX, 6.45 \times on HunyuanVideo, 3.17 \times on Inf-DiT, and maintains high quality with a 4.53 \times speedup on DiT.

Introduction

Recent progress in diffusion-based generative models has substantially advanced the capabilities of visual synthesis, achieving unprecedented quality in both image (Ho, Jain, and Abbeel 2020) and video generation (Blattmann et al. 2023). In particular, Diffusion Transformers (DiTs) (Peebles and Xie 2023) have set new performance standards by leveraging transformer-based architectures, which offer superior scalability and modeling capacity. Despite these advantages, DiTs remain computationally intensive due to their iterative sampling process, where each output requires a sequence of denoising steps that scale linearly with both the number of steps and model complexity. This computational overhead presents significant barriers to real-time or resource-constrained deployment.

To alleviate this bottleneck, two main acceleration strategies have been proposed: reducing the number of sampling steps through algorithmic enhancements (Ma, Fang,

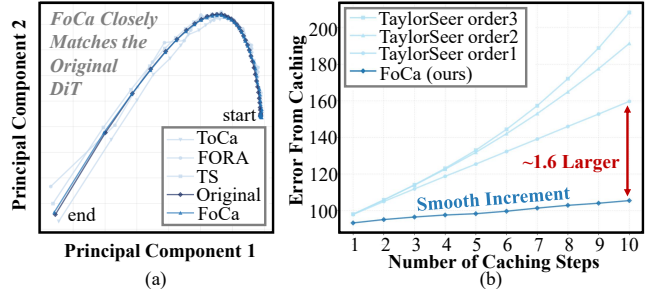


Figure 1: (a) PCA visualization of feature trajectory: FoCa closely follows the original model’s path. (b) Prediction error compared with the original model: FoCa maintains significantly lower error, whereas TaylorSeer’s (TS) higher-order forecasts incur rapidly growing errors.

and Wang 2023), and minimizing per-step computational cost via architectural modifications (Yuan et al. 2024; Zhao et al. 2025). Among these, training-free *feature caching* has emerged as a promising and general-purpose solution. It builds on the observation that intermediate hidden representations exhibit temporal consistency. This temporal coherence allows the model to bypass redundant computations by reusing or predicting features at selected timesteps, effectively reducing latency without retraining the network.

Initial implementations of feature caching were developed in the context of U-Net architectures, where cached activations are reused across nearby timesteps using skip connections (Wimbauer et al. 2024; Ma, Fang, and Wang 2023). As DiTs become the dominant architecture for high-fidelity generation, recent works have extended caching mechanisms to this setting, leading to methods such as FORA, ToCa, and TaylorSeer (TS) (Selvaraju et al. 2024; Zou et al. 2025; Liu et al. 2025). These techniques leverage temporal token similarity to inform feature reuse or extrapolation, but still suffer from error accumulation and instability under long skip intervals.

Limited Utilization of Historical Information. Existing methods fail to fully exploit historical hidden states for stable prediction. Reuse-only approaches simply replicate the most recent feature without incorporating multi-step history or adapting to the evolving feature dynamics



Figure 2: Visualization of the image generated by different methods on prompt: “Zoomed out view of a giraffe and a zebra in the middle of a field covered with colorful flowers.” As the computation ratio increases, the image quality of TaylorSeer and TeaCache degrades, whereas FoCa maintains better performance.

over time—particularly the shifting eigenvalue spectrum of the linearized drift. Forecasting-based methods such as TaylorSeer attempt to transform historical information into local approximations by estimating time derivatives via finite differences and applying truncated Taylor expansions. However, this strategy introduces significant numerical errors, as the finite-difference estimation is highly sensitive to noise and becomes unstable in higher-order terms. Consequently, neither reuse nor current forecasting methods adequately leverage historical information for robust and accurate prediction under large skipping intervals.

Instability Under Large Skipping Intervals. As the interval between executed timesteps increases, existing caching methods struggle to maintain accurate hidden state estimation. In reuse-based approaches, simply copying outdated features leads to growing mismatches as temporal distance increases. Similarly, forecasting methods such as Taylor series extrapolation suffer from accumulated approximation errors when extrapolating over longer spans. These inaccuracies compound with larger skip lengths, resulting in significant deviations from the correct denoising trajectory. Empirically, this leads to rapid degradation in generation quality, as shown in Figure 2, limiting the practical effectiveness of current caching strategies.

To solve these problems, we reformulate feature caching as solving a feature-ODE, providing a principled approach to modeling hidden-state evolution. This perspective enables the application of multi-step solvers that leverage historical hidden states directly, without relying on high-order derivative approximations. Specifically, we introduce **FoCa** (Forecast-then-Calibrate) framework, which combines a Backward Differentiation Formula based predictor with a lightweight Heun corrector. The forecast component aggregates previous features for stable forecasting, while the Heun calibration incorporates the previous full activation as a reference, damping overshoot on larger skip intervals.

FoCa provides robust and efficient prediction across di-

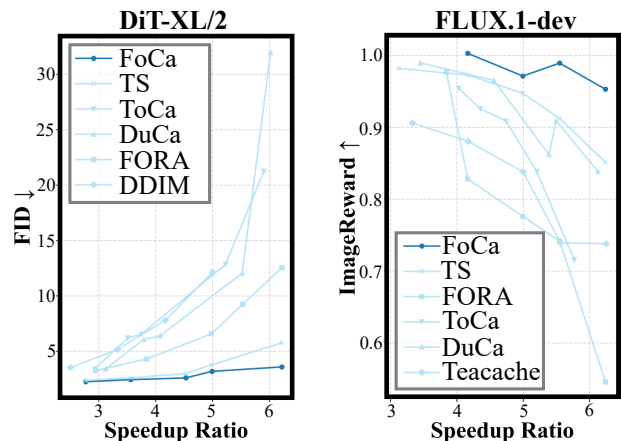


Figure 3: Comparison between previous caching methods and FoCa. FoCa shows significantly better performance.

verse tasks and architectures. Without additional training or architectural changes, it achieves near-lossless acceleration of **5.50** \times on FLUX for text-to-image generation, **6.45** \times on HunyuanVideo for text-to-video generation, and **3.17** \times on Inf-DiT for super-resolution. FoCa also maintains high image quality with **4.53** \times acceleration on DiT-XL/2. In summary, our main contributions are:

- **ODE Perspective on Feature Caching.** We regard the time-indexed feature tensors as trajectories of a feature-ODE. Framing caching in this perspective lets us apply multistep ODE solvers for prediction, eliminating the need for noisy high-order derivatives while stably propagating hidden states across skipped timesteps.
- **FoCa Framework.** We propose a training-free predictor–corrector framework combining a backward differentiation formula forecasting with Heun calibration, which remains stable even under large acceleration ratios.

- **Outstanding Performance.** We evaluate CoT across diverse architectures and tasks on multiple resolutions, including Drawbench on FLUX, Vbench on Hunyuan-Video. In all settings, HyCa delivers state-of-the-art performance.

Related Works

Diffusion models (Sohl-Dickstein et al. 2015; Ho, Jain, and Abbeel 2020) have become a dominant framework for high-fidelity image and video generation. While early implementations primarily adopted U-Net-based architectures (Ronneberger, Fischer, and Brox 2015), the introduction of Diffusion Transformers (DiTs) has enabled greater scalability and expressiveness. DiTs have since powered numerous state-of-the-art generative systems (Chen et al. 2023, 2024; hp-caitech 2024). Despite progress, the sequential nature of diffusion sampling remains a major obstacle for real-time and large-scale applications. To address this, recent research has focused on two main acceleration strategies: reducing the number of sampling steps and optimizing the efficiency of the denoising network.

Sampling Step Reduction. Several approaches aim to shorten the sampling trajectory while preserving output quality. DDIM (Song, Meng, and Ermon 2022) introduces a deterministic sampling formulation that enables fewer steps. The DPM-Solver family (Lu et al. 2022; Zheng et al. 2023) applies high-order numerical solvers to more accurately approximate the reverse diffusion process. Rectified Flow (Liu, Gong, and Liu 2022) constructs direct generative flows via optimal transport, and Consistency Models (Song et al. 2023) enforce self-consistency constraints to enable single-step or few-step generation. While effective, these methods generally require redesigning the sampling algorithm or retraining the model, limiting their applicability to pretrained diffusion systems.

Denoising Network Acceleration. An alternative line of work focuses on reducing per-step computational cost through model-level optimizations, broadly categorized into model compression and feature caching strategies.

Model Compression. Compression-based techniques accelerate inference by simplifying network structures. These include structured pruning (Fang, Ma, and Wang 2023; Zhu et al. 2024), quantization (Kim et al. 2025; Li et al. 2023), and token-level reductions such as merging or pruning (Bolya and Hoffman 2023; Cheng, Chen, and Jia 2025; Zhang et al. 2025). While these methods can significantly reduce runtime, they typically require retraining or fine-tuning to maintain generation quality, posing challenges for deployment on large-scale pretrained models.

Feature Caching. Training-free caching methods have gained attention for their simplicity and compatibility with existing networks. Early approaches such as DeepCache and FasterDiffusion sped up U-Nets by reusing features across nearby timesteps. Follow-up work extended caching to transformers: FORA and Δ -DiT store block outputs,

while ToCa, DuCa and RAS reuse tokens or regions adaptively. TaylorSeer replaced direct reuse with Taylor extrapolation, yet it fails to exploit past feature information effectively and can become unstable at larger timesteps because of its higher-order derivatives. In this work, we propose a more robust forecasting-then-calibrate framework that remains accurate and stable over longer skip intervals, thus delivering higher speed-ups without retraining or quality loss.

Method

Preliminary

Diffusion Models. Diffusion models operate by simulating a Markovian forward and reverse stochastic process that transitions data between clean images and Gaussian noise. The forward process progressively adds Gaussian noise to the data through a time-dependent variance schedule:

$$x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon_t, \quad (1)$$

where $\epsilon_t \sim \mathcal{N}(0, I)$ represents isotropic Gaussian noise, and α_t is a monotonically decreasing sequence controlling the noise level. After T timesteps, the sample becomes indistinguishable from pure noise. The reverse process learns to denoise the sample using a neural network $\epsilon_\theta(x_t, t)$ to predict the noise:

$$x_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) + \sigma_t \epsilon, \quad (2)$$

where $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$, and σ_t controls the noise level during denoising. This reverse process can also be described by a reverse-time stochastic differential equation or its corresponding ordinary differential equation.

Feature Caching in Diffusion Models. Feature caching accelerates diffusion inference by avoiding redundant computations across timesteps. Two main strategies exist: *reuse* and *forecast*. The **reuse-based** method directly substitutes features from a previous timestep:

$$\mathcal{F}(x_{t-k}^l) := \mathcal{F}(x_t^l), \quad \forall k \in [1, N-1], \quad (3)$$

achieving $(N-1) \times$ acceleration by skipping intermediate steps. However, errors quickly accumulate as the temporal gap increases, degrading output quality. The **forecast-based** method, exemplified by TaylorSeer, predicts future features via Taylor expansion:

$$\mathcal{F}_{\text{pred},m}(x_{t-k}^l) = \mathcal{F}(x_t^l) + \sum_{i=1}^m \frac{\Delta^i \mathcal{F}(x_t^l)}{i! \cdot N^i} (-k)^i \quad (4)$$

where $\Delta^i \mathcal{F}$ is the i -th order finite difference. While more accurate than direct reuse, this method is sensitive to noise under long intervals or complex dynamics, due to the limitations of local smoothness assumptions and the instability of high-order derivatives.

Overview of FoCa

To reduce the high inference cost of diffusion models, we propose **FoCa**, a predictor-corrector framework designed to




Figure 4: The computation of FoCa for x_{t-1} , x_{t-2} , x_{t-3} , respectively. For each skipping step, FoCa predicts a future hidden state using the two most recent timestep, then applies a Heun corrector that blends the prediction with the most recent fully-computed feature so that enables accurate and stable prediction under large skip intervals.

accelerate diffusion inference by leveraging historical features for stable and accurate forecasting as shown in Fig. 4. Unlike prior methods that rely on local extrapolation and suffer from high-order instability, FoCa integrates multi-step prediction with a lightweight correction mechanism to support longer skip intervals without degrading generation quality. The core workflow consists of three stages:

- **Multi-step Forecasting.** FoCa uses a second-order backward differentiation scheme to predict future hidden features based on the two most recent timesteps. This multi-step strategy captures temporal evolution more accurately than single-step reuse methods and high-order local derivative approximation methods.
- **Heun-style Calibration.** To mitigate potential oscillation or over-extrapolation from the forecasting step, FoCa introduces a correction mechanism that blends the prediction with the latest full-computation feature. This dampens overshoot and stabilizes the inference process under large skipping intervals.
- **Prediction Execution.** Finally, use the corrected feature for the subsequent denoising step unless the next timestep is scheduled for full computation.

By combining high-accuracy forecasting with stability-aware correction, FoCa offers a principled and retraining-free solution for accelerating diffusion transformers across image, video, and super-resolution tasks.

An ODE Perspective on Hidden-Feature Evolution. During reverse-time denoising in diffusion models, the hidden features evolve smoothly along the sampling trajectory. Let $x_t^l = G_l(x_t)$ be the layer- l activation and $\mathcal{F}(x_t^l)$ its downstream feature. Since the network blocks are differentiable and x_t evolves continuously, the composite map $t \mapsto \mathcal{F}(x_t^l)$ is differentiable; by the chain rule and the probability-flow ODE for x_t ,

$$\frac{d}{dt} \mathcal{F}(x_t^l) = g_\theta(\mathcal{F}(x_t^l), t), \quad (5)$$

where g_θ implicitly aggregates Jacobians of preceding layers. Although g_θ is intractable, the feature samples $\{\mathcal{F}(x_{t_k}^l)\}$ on a reverse-time grid permit classical linear multistep integration using only cached values. This insight motivates the use of backward differentiation formula to predict $\mathcal{F}(x_{t_{k+1}}^l)$ from the two most recently stored tensors, thereby skipping full forward passes while retaining accuracy.

BDF2 Forecasting. The second-order backward differentiation formula (BDF2) is a linear multistep method. It approximates the derivative at the next step, $k+1$, using the latest two points ($\mathcal{F}(x_{k-1}^l)$, $\mathcal{F}(x_k^l)$):

$$\mathcal{F}^{(1)}(x_{k+1}^l) = \frac{3\mathcal{F}(x_{k+1}^l) - 4\mathcal{F}(x_k^l) + \mathcal{F}(x_{k-1}^l)}{2h_k}. \quad (6)$$

This formula is implicit because the right-hand side depends on the unknown $\mathcal{F}(x_{k+1}^l)$ through the function $\mathcal{F}^{(1)}$. To create a practical, explicit forecasting scheme, we approximate the required future derivative $\mathcal{F}^{(1)}(x_{k+1}^l)$ by extrapolating from the most recent step’s derivative, $\mathcal{F}^{(1)}(x_k^l)$. This yields our predictor:

$$\hat{\mathcal{F}}(x_{k+1}^l) = \frac{4}{3}\mathcal{F}(x_k^l) - \frac{1}{3}\mathcal{F}(x_{k-1}^l) + \frac{2h_k}{3}\mathcal{F}^{(1)}(x_k^l). \quad (7)$$

Heun Calibration. We enhance our BDF2 forecast by applying a Heun-style corrector step. The Heun method, also known as the explicit trapezoidal rule, improves the prediction by averaging the slopes at the beginning and the predicted end of the interval. First, we compute the BDF2 forecast $\hat{\mathcal{F}}(x_{k+1}^l)$ using Eq. 7. Then, we apply the Heun corrector to obtain the final, more accurate state $\mathcal{F}_c(x_{k+1}^l)$:

$$\mathcal{F}_c(x_{k+1}^l) = \mathcal{F}(x_k^l) + \frac{h_k}{2}[\mathcal{F}^{(1)}(x_{k-N}^l) + \mathcal{F}^{(1)}(x_{k+1}^l)]. \quad (8)$$

Where (x_{k-N}^l) is the latest fully computed step.

Proposition 1 (FoCa ensures stable prediction under large interval, proof in Appendix A.1). *Given a smooth feature evolution $\mathcal{F}(x_t^l)$ in diffusion models, FoCa achieves a prediction error bound that is independent of the number of skipped steps k :*

$$\|\hat{\mathcal{F}}(x_k^l) - \mathcal{F}^*(x_k^l)\| \leq \left(\frac{1-\rho^k}{1-\rho}\right) \tau_{\max} \leq \frac{\tau_{\max}}{1-\rho} \quad (9)$$

In contrast, the error of TaylorSeer and reuse methods increases with k , leading to degraded accuracy in long-range prediction. This highlights FoCa’s robustness under aggressive caching and large acceleration ratios.

Practical Impact. This predictor-corrector framework combines the strengths of both methods. The BDF2 predictor leverages historical information for initial prediction.

Method	Efficient Attention	Acceleration				ImageReward \uparrow	CLIP \uparrow Score
		Latency(s) \downarrow	Speed \uparrow	FLOPs(T) \downarrow	Speed \uparrow		
[dev]: 50 steps	✓	25.82	1.00 \times	3719.50	1.00 \times	0.9898	32.404
60% steps	✓	16.70	1.55 \times	2231.70	1.67 \times	0.9663	32.312
Δ -DiT ($\mathcal{N} = 2$)	✓	17.80	1.45 \times	2480.01	1.50 \times	0.9444	32.273
Δ -DiT ($\mathcal{N} = 3$)	✓	13.02	1.98 \times	1686.76	2.21 \times	0.8721	32.102
34% steps	✓	9.07	2.85 \times	1264.63	3.13 \times	0.9453	32.114
Chipmunk	✓	12.72	2.02 \times	1505.87	2.47 \times	0.9936	32.776
FORA ($\mathcal{N} = 3$)	✓	10.16	2.54 \times	1320.07	2.82 \times	0.9776	32.266
ToCa ($\mathcal{N} = 6$)	✗	13.16	1.96 \times	924.30	4.02 \times	0.9802	32.083
DuCa ($\mathcal{N} = 5$)	✓	8.18	3.15 \times	978.76	3.80 \times	0.9955	32.241
TaylorSeer ($\mathcal{N} = 4, O = 2$)	✓	9.24	2.80 \times	1042.27	3.57 \times	0.9857	32.413
FoCa ($\mathcal{N} = 5$)	✓	7.46	3.46\times	893.54	4.16\times	1.0029	32.948
22% steps	✓	6.04	4.28 \times	818.29	4.55 \times	0.8183	31.772
FORA ($\mathcal{N} = 4$)	✓	8.12	3.14 \times	967.91	3.84 \times	0.9730	32.142
ToCa ($\mathcal{N} = 8$)	✗	11.36	2.27 \times	784.54	4.74 \times	0.9451	31.993
DuCa ($\mathcal{N} = 7$)	✓	6.74	3.83 \times	760.14	4.89 \times	0.9757	32.066
TeaCache ($l = 0.8$)	✓	7.21	3.58 \times	892.35	4.17 \times	0.8683	31.704
TaylorSeer ($\mathcal{N} = 5, O = 2$)	✓	7.46	3.46 \times	893.54	4.16 \times	0.9768	32.467
FoCa ($\mathcal{N} = 7$)	✓	6.36	4.05\times	670.44	5.54\times	0.9891	32.920
FORA ($\mathcal{N} = 7$)	✓	7.71	3.34 \times	670.14	5.55 \times	0.7418	31.519
ToCa ($\mathcal{N} = 12$)	✗	10.34	2.50 \times	644.70	5.77 \times	0.7155	31.808
DuCa ($\mathcal{N} = 10$)	✓	9.22	2.80 \times	606.91	6.13 \times	0.8382	31.759
TeaCache ($l = 1.2$)	✓	6.12	4.22 \times	669.27	5.56 \times	0.7394	31.704
TaylorSeer ($\mathcal{N} = 7, O = 2$)	✓	6.38	4.05 \times	670.44	5.54 \times	0.9128	32.128
FoCa ($\mathcal{N} = 8$)	✓	5.88	4.39\times	596.07	6.24\times	0.9502	32.706

Table 1: Quantitative comparison in text-to-image generation on FLUX.

The Heun corrector then refines this prediction by incorporating the most recent fully computed step, effectively damping oscillations and reducing error accumulation. This hybrid approach allows larger acceleration ratios while maintaining high-fidelity generation, outperforming methods that rely on either single-step reuse or less accurate high-order derivative extrapolation schemes.




Figure 5: Visual comparison of 5.5 \times accelerated FLUX.

Experiments

Experiment Settings. We evaluate our method on four representative diffusion-based generative models: the text-to-image model **FLUX.1-dev** (Black Forest Labs 2024), the

super-resolution model **Inf-DiT** (Yang et al. 2024), the text-to-video model **HunyuanVideo** (Sun et al. 2024), and the class-conditional image generation model **DiT-XL/2** (Peebles and Xie 2023). For text-to-image generation, we follow the DrawBench (Saharia et al. 2022) protocol and evaluate photorealism and text alignment using ImageReward (Xu et al. 2023) and CLIP Score (Hessel et al. 2021). For text-to-video generation, we adopt the VBench (Huang et al. 2023) benchmark, which assesses video quality across multiple human-aligned dimensions. For class-conditional image generation, we evaluate on ImageNet (Russakovsky et al. 2015) using FID-50k (Heusel et al. 2017) and Inception Score (IS) as standard quality metrics. For super-resolution, we use DIV8K (Gu et al. 2019) and measures the quality with PSNR and SSIM. Please refer to Appendix for details.

Results on Text-to-Image Generation. As summarized in Table 1, our method achieves a consistently superior speed-quality trade-off on FLUX.1-dev across all acceleration levels. With a moderate cache interval ($\mathcal{N} = 5$), FoCa attains the highest ImageReward score of **1.0029** while achieving a 4.16 \times FLOPs reduction, outperforming all reuse-based baselines, and the prediction-based method TaylorSeer (0.9857 at 3.57 \times). Under higher compression ($\mathcal{N} = 7$), FoCa retains a strong ImageReward of **0.9891** at a 5.54 \times speedup, surpassing TaylorSeer’s performance at $\mathcal{N} = 5, O = 2$ (0.9857), while other baselines degrade notably. At the most aggressive setting ($\mathcal{N} = 8$), FoCa still achieves **0.9502** ImageReward at 6.24 \times acceleration, sig-

Method	Efficient Attention	Acceleration				VBench ↑ Score(%)
		Latency(s) ↓	Speed ↑	FLOPs(T) ↓	Speed ↑	
Original	✓	334.96	1.00×	29773.0	1.00×	80.66
DDIM-22%	✓	87.01	3.85×	6550.1	4.55×	78.74
FORA ($\mathcal{N} = 5$)	✓	83.78	4.00×	5960.4	5.00×	78.83
ToCa ($\mathcal{N} = 5$)	✗	93.80	3.57×	7006.2	4.25×	78.86
DuCa ($\mathcal{N} = 5$)	✓	87.48	3.83×	6483.2	4.62×	78.72
TeaCache ($l = 0.5$)	✓	61.47	5.45×	5359.1	5.56×	78.32
TaylorSeer ($\mathcal{N} = 7, O = 1$)	✓	72.76	4.60×	4611.3	6.45×	79.37
FoCa ($\mathcal{N} = 7$)	✓	72.18	4.64×	4611.3	6.45×	79.68

Table 2: Quantitative comparison in text-to-video generation on VBench.

Method	FLOPs(T) ↓	Speed ↑	FID ↓	IS ↑
DDIM-50 steps	23.74	1.00×	2.32	241.25
DDIM-25 steps	11.87	2.00×	3.18	232.01
L2C ($NFE = 30$)	11.55	2.05×	2.61	237.83
Δ -DiT ($\mathcal{N} = 2$)	18.04	1.31×	2.69	225.99
Δ -DiT ($\mathcal{N} = 3$)	16.14	1.47×	3.75	207.57
FoCa ($\mathcal{N} = 2$)	12.35	1.92×	2.17	239.94
DDIM-20 steps	9.49	2.50×	3.81	221.43
FORA ($\mathcal{N} = 3$)	8.58	2.77×	3.55	229.02
ToCa ($\mathcal{N} = 3$)	10.23	2.32×	2.87	235.21
DuCa ($\mathcal{N} = 3$)	9.58	2.48×	2.88	233.37
SmoothCache	8.57	2.77×	4.15	231.71
TaylorSeer ($\mathcal{N} = 3$)	8.56	2.77×	2.35	236.22
FoCa ($\mathcal{N} = 3$)	8.56	2.77×	2.25	237.70
DDIM-12 steps	5.70	4.17×	7.80	184.50
FORA ($\mathcal{N} = 5$)	5.24	4.53×	6.58	193.01
ToCa ($\mathcal{N} = 6$)	6.34	3.75×	6.55	189.53
DuCa ($\mathcal{N} = 5$)	6.27	3.78×	6.06	198.46
TaylorSeer ($\mathcal{N} = 5$)	5.24	4.53×	2.74	231.42
FoCa ($\mathcal{N} = 5$)	5.24	4.53×	2.60	230.76

Table 3: Quantitative comparison in class-to-image generation on ImageNet with DiT-XL/2.

nificantly outperforming other methods. Qualitative comparisons in Figure 5 further illustrate that FoCa preserves better image quality, even under large skipping intervals.

Results on Text-to-Video Generation. Table 2 demonstrates that FoCa delivers the best overall trade-off on HunyuanVideo. With a cache window of $\mathcal{N} = 7$, FoCa holds a **6.45×** speedup ratio, while achieving the highest VBench score of **79.68%**. In particular, compared to the original sampler (80.66%), FoCa attains near-lossless quality with more computational savings 6.45×. Qualitative comparisons on Fig. 7 further confirm that FoCa preserves video quality, spatial relationship, background scenery better.

Results on Class-to-Image Generation. On DiT-XL/2, FoCa establishes a stronger quality–efficiency frontier than other training-free baselines, and it also surpasses the training-based L2C at comparable budgets. With interval ($\mathcal{N} = 2$), FoCa’s FID even improves over the original 50-step DDIM at 1.92× speedup; at $\mathcal{N} = 3$ it achieves FID 2.25 (2.77×), outperforming TaylorSeer FID 2.35. Under higher

Method	Latency(s) ↓	FLOPs(T) ↓	Speed ↑	PSNR ↑
Original	451.3	124004	1.00×	30.85
AB-Cache	197.2	45177	2.74×	29.77
TaylorSeer	199.5	45184	2.74×	29.67
FoCa	161.7	39096	3.17×	31.03

Table 4: Quantitative comparison on Inf-DiT for generative image super-resolution from 512×512 to 2048×2048.



Figure 6: Ablation study of Heun Correction. Adding the Heun step consistently reduces FID and boosts Inception Score across all speedup ratios.

compression ($\mathcal{N} = 5$), FoCa maintains FID 2.60 at 4.53×, whereas most method collapse.

Results on Image-Super-Resolution. We have evaluated our method and the other caching methods for image super-resolution. On Inf-DiT, FoCa yields a **3.17×** speedup with even better PSNR of **31.03 dB**. In contrast, TaylorSeer achieves only 29.66 dB, confirming that noisy finite-difference forecasting degrades high-frequency detail. These results also demonstrate that FoCa’s robustness beyond text-conditional synthesis.

Ablation Study. We evaluate the impact of Heun correction in FoCa on DiT-XL/2. As shown in Figure 6, we compare the BDF2-only with FoCa at varying activation intervals. While BDF2-only forecasting deteriorates rapidly under large skip intervals, FoCa consistently achieves lower FID and higher Inception Score, demonstrating improved stability and accuracy. These results highlight the importance of integrating stability-aware correction into multi-step caching schemes, particularly on aggressive acceleration.



Figure 7: Visualization of different caching methods on HunyuanVideo. FoCa maintains high-quality generation under the higher acceleration ratio, while other methods suffer from issues such as missing details, and mistaken spatial relationships.

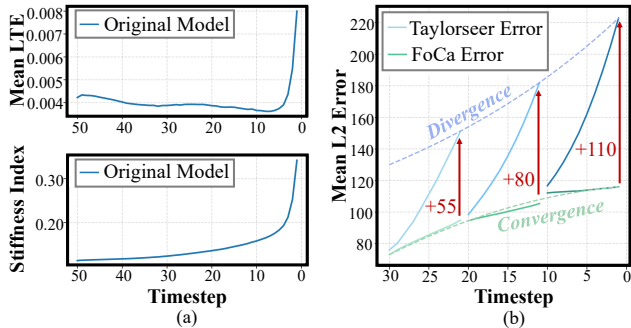


Figure 8: Stiffness Analysis: (a) LTE and Stiffness Index of FLUX across timesteps, revealing a sharp rise in the late stage that indicates a numerically stiff region. LTE denotes local truncation error, which measures numerical integration accuracy. The Stiffness Index measures the numerical stiffness of the model. (b) Comparison of ten-step forecasting error starting from various timesteps.

Analysis on the stiffness of feature trajectories. In numerical analysis, a differential equation is considered stiff when certain components evolve at vastly different timescales, causing instability for standard integration methods unless extremely small step sizes are used. In this work, we model the evolution of hidden features in diffusion models as an ODE system and solve it using a BDF2-Heun predictor-corrector scheme. This method is particularly well-suited for stiff ODEs, making it a natural fit for handling the challeng-

ing dynamics of diffusion sampling. To explore the applicability of our method, we investigate the stiffness of the hidden feature trajectories. Our stiffness diagnosis in Fig. 8(a) reveals a late-stage region (from step 10 to 0) where LTE and Stiffness Index both rise sharply. This indicates high sensitivity to step-size perturbations and numerical instability, which we define as the stiff zone in the evolution of hidden states. To further analyze, Fig. 8(b) shows forecasting error under controlled prediction horizons. Specifically, we execute full computation up to timestep 30, 20, and 10, and forecast the next 10 steps using either TaylorSeer or FoCa. As prediction begins deeper into the stiff zone, TaylorSeer exhibits rapidly diverging errors due to over-extrapolation, while FoCa maintains bounded error growth through stable multi-step aggregation and Heun-style correction. This results in convergent error curves and consistent generation quality even under aggressive skipping.

Conclusion

We presented **FoCa**, a training-free acceleration framework that reinterprets hidden feature state prediction as solving an ODE and couples BDF2 predictor with a Heun calibrator. This predictor-corrector design utilizes historical features effectively and expands the stability region into the stiff zone where existing caching methods collapse. Performing *near-lossless* speedups in diverse settings: **5.50** \times on FLUX, **6.45** \times on HunyuanVideo, **3.17** \times on Inf-DiT and state-of-the-art trade-offs on ImageNet with DiT-XL/2 (up to 4.53 \times with superior FID). FoCa provides a novel acceleration perspective from the view of ODE in the feature space.

Acknowledgments

This paper was partially sponsored by Transsion Holdings.

References

- Black Forest Labs. 2024. FLUX. <https://github.com/black-forest-labs/flux>.
- Blattmann, A.; Dockhorn, T.; Kulal, S.; Mendelevitch, D.; Kilian, M.; Lorenz, D.; Levi, Y.; English, Z.; Voleti, V.; Letts, A.; Jampani, V.; and Rombach, R. 2023. Stable Video Diffusion: Scaling Latent Video Diffusion Models to Large Datasets. *arXiv preprint arXiv:2311.15127*.
- Bolya, D.; and Hoffman, J. 2023. Token Merging for Fast Stable Diffusion. *arXiv preprint arXiv:2303.17604*.
- Chen, J.; Huang, Z.; Zhang, L.; Wang, J.; Li, J.; Li, X.; Liu, J.; and Chen, T. 2023. PixArt-alpha: Fast Training of Diffusion Transformer for High-Fidelity Text-to-Image Synthesis. *arXiv preprint arXiv:2310.00426*.
- Chen, J.; Huang, Z.; Zhang, L.; Wang, J.; Li, J.; Li, X.; Liu, J.; and Chen, T. 2024. PixArt-Sigma: Scalable and Efficient Diffusion Transformer for High-Resolution Image Synthesis. *arXiv preprint arXiv:2403.09068*.
- Cheng, X.; Chen, Z.; and Jia, Z. 2025. CAT Pruning: Cluster-Aware Token Pruning For Text-to-Image Diffusion Models. *arXiv preprint arXiv:2502.00433*.
- Fang, G.; Ma, X.; and Wang, X. 2023. Structural Pruning for Diffusion Models. *arXiv preprint arXiv:2305.10924*.
- Gu, S.; Lugmayr, A.; Danelljan, M.; Fritsche, M.; Lamour, J.; and Timofte, R. 2019. DIV8K: DIVerse 8K Resolution Image Dataset. In *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, 512–520.
- Hessel, J.; Holtzman, A.; Forbes, M.; Bras, R. L.; and Choi, Y. 2021. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. *arXiv preprint arXiv:2104.08718*.
- 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. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising Diffusion Probabilistic Models. *arXiv preprint arXiv:2006.11239*.
- hpcaitech. 2024. Open-Sora: An Open-Source Project for Efficient Video Production. <https://github.com/hpcaitech/Open-Sora>.
- Huang, Z.; He, Y.; Yu, J.; Zhang, F.; Si, C.; Jiang, Y.; Zhang, Y.; Wu, T.; Jin, Q.; Chanpaisit, N.; Wang, Y.; Chen, X.; Wang, L.; Lin, D.; Qiao, Y.; and Liu, Z. 2023. VBench: Comprehensive Benchmark Suite for Video Generative Models. *arXiv preprint arXiv:2311.17982*.
- Kim, S.; Lee, H.; Cho, W.; Park, M.; and Ro, W. W. 2025. Ditto: Accelerating Diffusion Model via Temporal Value Similarity. *arXiv preprint arXiv:2501.11211*.
- Li, X.; Liu, Y.; Lian, L.; Yang, H.; Dong, Z.; Kang, D.; Zhang, S.; and Keutzer, K. 2023. Q-Diffusion: Quantizing Diffusion Models. *arXiv preprint arXiv:2302.04304*.
- Liu, J.; Zou, C.; Lyu, Y.; Chen, J.; and Zhang, L. 2025. From Reusing to Forecasting: Accelerating Diffusion Models with TaylorSeers. *arXiv preprint arXiv:2403.06923*.
- Liu, X.; Gong, C.; and Liu, Q. 2022. Flow Straight and Fast: Learning to Generate and Transfer Data with Rectified Flow. *arXiv preprint arXiv:2209.03003*.
- 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. *arXiv preprint arXiv:2206.00927*.
- Ma, X.; Fang, G.; and Wang, X. 2023. DeepCache: Accelerating Diffusion Models for Free. *arXiv preprint arXiv:2312.00858*.
- Peebles, W.; and Xie, S. 2023. Scalable Diffusion Models with Transformers. *arXiv preprint arXiv:2212.09748*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv preprint arXiv:1505.04597*.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M. S.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115: 211–252.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, S. K. S.; Ayan, B. K.; Mahdavi, S. S.; Lopes, R. G.; Salimans, T.; Ho, J.; Fleet, D. J.; and Norouzi, M. 2022. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. *arXiv preprint arXiv:2205.11487*.
- Selvaraju, P.; Ding, T.; Chen, T.; Zharkov, I.; and Liang, L. 2024. FORA: Fast-Forward Caching in Diffusion Transformer Acceleration. *arXiv preprint arXiv:2407.01425*.
- Sohl-Dickstein, J.; Weiss, E. A.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. *arXiv preprint arXiv:1503.03585*.
- Song, J.; Meng, C.; and Ermon, S. 2022. Denoising Diffusion Implicit Models. *arXiv preprint arXiv:2010.02502*.
- Song, Y.; Dhariwal, P.; Chen, M.; and Sutskever, I. 2023. Consistency Models. *arXiv preprint arXiv:2303.01469*.
- Sun, X.; et al. 2024. Hunyuan-Large: An Open-Source MoE Model with 52 Billion Activated Parameters by Tencent. *arXiv preprint arXiv:2411.02265*.
- Wimbauer, F.; Wu, B.; Schoenfeld, E.; Dai, X.; Hou, J.; He, Z.; Sanakoyeu, A.; Zhang, P.; Tsai, S.; Kohler, J.; Rupprecht, C.; Cremers, D.; Vajda, P.; and Wang, J. 2024. Cache Me if You Can: Accelerating Diffusion Models through Block Caching. *arXiv preprint arXiv:2312.03209*.
- Xu, J.; Liu, X.; Wu, Y.; Tong, Y.; Li, Q.; Ding, M.; Tang, J.; and Dong, Y. 2023. ImageReward: Learning and Evaluating Human Preferences for Text-to-Image Generation. *arXiv preprint arXiv:2304.05977*.
- Yang, Z.; Jiang, H.; Hong, W.; Teng, J.; Zheng, W.; Dong, Y.; Ding, M.; and Tang, J. 2024. Inf-DiT: Upsampling Any-Resolution Image with Memory-Efficient Diffusion Transformer. *arXiv preprint arXiv:2405.04312*.

Yuan, Z.; Zhang, H.; Lu, P.; Ning, X.; Zhang, L.; Zhao, T.; Yan, S.; Dai, G.; and Wang, Y. 2024. DiTFastAttn: Attention Compression for Diffusion Transformer Models. *arXiv preprint arXiv:2406.08552*.

Zhang, E.; Tang, J.; Ning, X.; and Zhang, L. 2025. Training-Free and Hardware-Friendly Acceleration for Diffusion Models via Similarity-based Token Pruning. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Zhao, X.; Jin, X.; Wang, K.; and You, Y. 2025. Real-Time Video Generation with Pyramid Attention Broadcast. *arXiv preprint arXiv:2408.12588*.

Zheng, K.; Lu, C.; Chen, J.; and Zhu, J. 2023. DPM-Solver-v3: Improved Diffusion ODE Solver with Empirical Model Statistics. *arXiv preprint arXiv:2310.13268*.

Zhu, H.; Tang, D.; Liu, J.; Lu, M.; Zheng, J.; Peng, J.; Li, D.; Wang, Y.; Jiang, F.; Tian, L.; Tiwari, S.; Sirasao, A.; Yong, J.-H.; Wang, B.; and Barsoum, E. 2024. DiP-GO: A Diffusion Pruner via Few-step Gradient Optimization. *arXiv preprint arXiv:2410.16942*.

Zou, C.; Liu, X.; Liu, T.; Huang, S.; and Zhang, L. 2025. Accelerating Diffusion Transformers with Token-wise Feature Caching. *arXiv preprint arXiv:2410.05317*.