

ProCache: Constraint-Aware Feature Caching with Selective Computation for Diffusion Transformer Acceleration

Fanpu Cao^{1*}, Yaofu Chen^{1*†}, Zeng You¹, Wei Luo^{2,3}

¹South China University of Technology, China

²South China Agricultural University, China

³Pazhou Laboratory, China

chenyaofu@scut.edu.cn

Abstract

Diffusion Transformers (DiTs) have achieved state-of-the-art performance in generative modeling, yet their high computational cost hinders real-time deployment. While feature caching offers a promising training-free acceleration solution by exploiting temporal redundancy, existing methods suffer from two key limitations: (1) uniform caching intervals fail to align with the non-uniform temporal dynamics of DiT, and (2) naive feature reuse with excessively large caching intervals can lead to severe error accumulation. In this work, we analyze the evolution of DiT features during denoising and reveal that both feature changes and error propagation are highly time- and depth-varying. Motivated by this, we propose ProCache, a training-free dynamic feature caching framework that addresses these issues via two core components: (i) a constraint-aware caching pattern search module that generates non-uniform activation schedules through offline constrained sampling, tailored to the model’s temporal characteristics; and (ii) a selective computation module that selectively compute within deep blocks and high-importance tokens for cached segments to mitigate error accumulation with minimal overhead. Extensive experiments on PixArt- α and DiT demonstrate that ProCache achieves up to $1.96\times$ and $2.90\times$ acceleration with negligible quality degradation, significantly outperforming prior caching-based methods.

Code — <https://github.com/macovaseas/ProCache>

1 Introduction

Diffusion models have achieved remarkable success in various visual generation tasks, including image and video synthesis (Rombach et al. 2022; Brooks et al. 2024). Traditional diffusion models predominantly employed CNN-based U-Net architectures (Rombach et al. 2022), while recent transformer-based diffusion models (DiT) (Peebles and Xie 2023; Brooks et al. 2024) have established state-of-the-art performance by scaling up model parameters and training data. However, such superior performance comes at the cost of significant computational overhead, which severely limits the deployment of DiT in real-world applications.

*These authors contributed equally.

†Corresponding author.

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

To address this, plenty of efforts have been made to accelerate DiT, including layer/token pruning (Fang, Ma, and Wang 2023; Fang et al. 2024; Zhang et al. 2025), quantization (Shang et al. 2023; So et al. 2023; Liu and Zhang 2024), and knowledge distillation (Li et al. 2023; Salimans and Ho 2022). However, these approaches typically require additional post-training cost and struggle to achieve satisfactory performance under high acceleration ratios. In contrast, feature caching methods (Ma, Fang, and Wang 2023; Wimbauer et al. 2024; Selvaraju et al. 2024; Liu et al. 2025) offer a training-free alternative by leveraging temporal redundancy in denoising steps. They cache computed features for reuse in subsequent steps, significantly improving inference speed without any training requirement. Their plug-and-play nature has attracted substantial interest.

However, these caching-based methods have their natural limitations. First, uniform activation intervals (*e.g.*, full computation every N steps) fail to capture the dynamic characteristics in denoising process. Although such rigid schedules mitigate error accumulation through periodic updates, they overlook the time-varying nature of feature evolution across steps. Second, the exponential decay of feature similarity with increasing timestep gaps causes significant error accumulation when reusing stale features. Existing approaches like Δ -DiT (Chen et al. 2024) and FORA (Selvaraju et al. 2024) directly reuse features without refinement, leading to quality degradation at high acceleration ratios.

To better understand these limitations, we conduct a quantitative analysis of feature error across DiT blocks throughout the denoising trajectory. Our investigation reveals two key observations. First, as illustrated in Figure 1, although feature errors grow progressively over diffusion steps, their accumulation varies significantly across network depths. In particular, deeper blocks exhibit substantially higher error magnitudes. Second, model outputs exhibit distinct temporal dynamics across denoising steps, with gradual stabilization in the early and middle stages followed by rapid changes (aligning with findings in (Liu et al. 2024; So et al. 2023; Liu et al. 2025) and also evident in our empirical studies of Figure 3). This suggests that a uniform caching schedule fails to align with the evolving nature of the denoising process and may lead to inaccurate caching reuse in fast-changing phases. This suggests that a uniform caching strategy, as adopted in prior work, is fundamentally suboptimal.

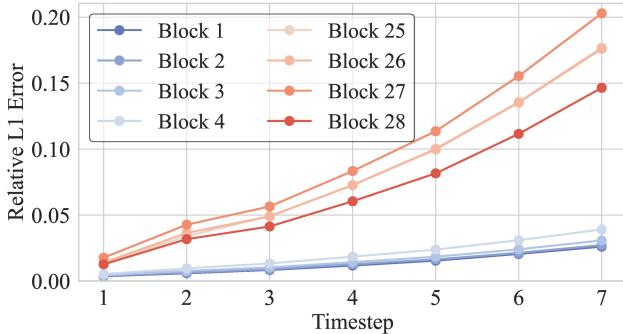


Figure 1: Evolution of relative L1 error across diffusion steps in DiT blocks, which is computed from 10 samples on PixArt- α . Errors grow progressively, with deeper blocks (e.g., Block 25–28) showing significantly higher magnitudes than shallower ones (e.g., Block 1–4), highlighting the non-uniform error accumulation across network depths.

Inspired by these, we propose **ProCache**, a training-free dynamic caching framework for efficient diffusion transformer inference. Our ProCache consists of two key components: 1) Constraint-aware caching pattern search mechanism replaces uniform-interval schedules with offline-optimized, non-uniform activation patterns. By analyzing the evolution of transformer activations across timesteps, we identify phases of high semantic sensitivity and design a constrained sampling procedure to generate caching patterns that maximize reuse in stable stages while enforcing frequent computation during critical one. 2) Selective computation strategy that mitigates error accumulation in long reuse intervals through lightweight, structured updates. Instead of full computation, we selectively refresh only a small set of transformer blocks and tokens. This partial update focuses computational effort where it matters most, effectively maintaining semantic consistency with minimal overhead.

Contributions: 1) We analyze that DiT features evolve non-uniformly during inference, *i.e.*, stable in early stages but highly dynamic in later timesteps, with error propagation concentrated in deep layers, revealing the critical mismatch of uniform caching. 2) We propose ProCache, a training-free acceleration framework that solves this via constraint-aware caching pattern search and selective computation, achieving adaptive speedup without quality loss. 3) Experiments show ProCache delivers up to $1.96\times$ and $2.90\times$ acceleration on PixArt- α and DiT, respectively, with negligible degradation—outperforming prior methods by a large margin.

2 Related Work

Diffusion models (DMs) have recently adopted Transformer based architectures, outperforming U-Net based counterparts across multiple domains (Chen et al. 2023; Ma et al. 2024b; Zheng et al. 2024), which introduces substantial computational overhead during inference, rendering real-time deployment challenging. Efforts to mitigate computational demands have explored model compression techniques such as pruning (Fang, Ma, and Wang 2023; Bolya

and Hoffman 2023; Fang et al. 2024; Zhang et al. 2025), quantization (Shang et al. 2023; So et al. 2023; Liu and Zhang 2024), and knowledge distillation (Salimans and Ho 2022; Li et al. 2023). These approaches, however, require resource-intensive retraining or fine-tuning to maintain generation quality. Consequently, recent research has shifted toward training-free acceleration methods, including *feature caching* and *sampling timestep reduction* methods.

Feature Caching methods accelerate diffusion models by reusing intermediate features across sampling steps, exploiting temporal redundancies in denoising computations. Early approaches like DeepCache (Ma, Fang, and Wang 2023) and Faster Diffusion (Li et al. 2024) targeted U-Net architectures, reusing low-resolution features or skipping encoder computations. However, they rely on U-Net’s hierarchical structure limits applicability to transformer-based diffusion models. Recent techniques address this gap through specialized reuse strategies: FORA (Selvaraju et al. 2024) and Δ -DiT (Chen et al. 2024) leverage attention feature reuse and MLP representation sharing, while PAB (Zhao et al. 2025) optimizes attention head computations within DiT blocks. L2C (Ma et al. 2024a) introduces dynamic computation routing through trainable layer-wise skip decisions, achieving higher acceleration ratios at the cost of substantial training overhead. Despite these advances, these methods apply identical caching strategies across all tokens, ignoring fine-grained feature dynamics. To address this, token-level optimization methods like ToCa (Zou et al. 2025) selectively caching tokens based on importance scores.

Sampling Timesteps Reduction methods constitute another prominent approach for accelerating diffusion models, operating by minimizing the number of iterative denoising steps required during generation. While early approaches like DDIM (Song, Meng, and Ermon 2022) established deterministic non-Markovian reverse processes for significant step reduction, subsequent advances have refined this paradigm: The DPM-Solver series (Lu et al. 2022, 2023) introduced high-order adaptive solvers for optimized differential equation solving, while Rectified Flow (Lipman et al. 2023) employs recurrent flow fields to straighten probability paths. Specialized extensions further enhance applicability, YONOS (Noroozi et al. 2024) enables single-step super-resolution via scale distillation, DC-Solver (Zhao et al. 2024) addresses predictor-corrector misalignment through dynamic compensation, and InvSR (Yue, Liao, and Loy 2025) facilitates arbitrary-step sampling via noise predictor networks. Notably, these techniques are orthogonal to the feature caching methods that our work advances.

3 Proposed Method

3.1 Problem Statement and Overview

Let \mathcal{G} denote a pretrained Diffusion Transformer (DiT) (Peebles and Xie 2023) model with L layers, *i.e.*, $\mathcal{G} = g_1 \circ g_2 \circ \dots \circ g_L$. Each layer l consists of three components: self-attention (SA), cross-attention (CA), and a multilayer perceptron (MLP), *i.e.*, $g_l = \mathcal{F}_l^{\text{SA}} \circ \mathcal{F}_l^{\text{CA}} \circ \mathcal{F}_l^{\text{MLP}}$. All these three operations can be unified under a common paradigm: $\mathcal{F}(\mathbf{x}) = \mathbf{x} + \text{AdaLN} \circ f(\mathbf{x})$, where $f(\mathbf{x})$ represents either

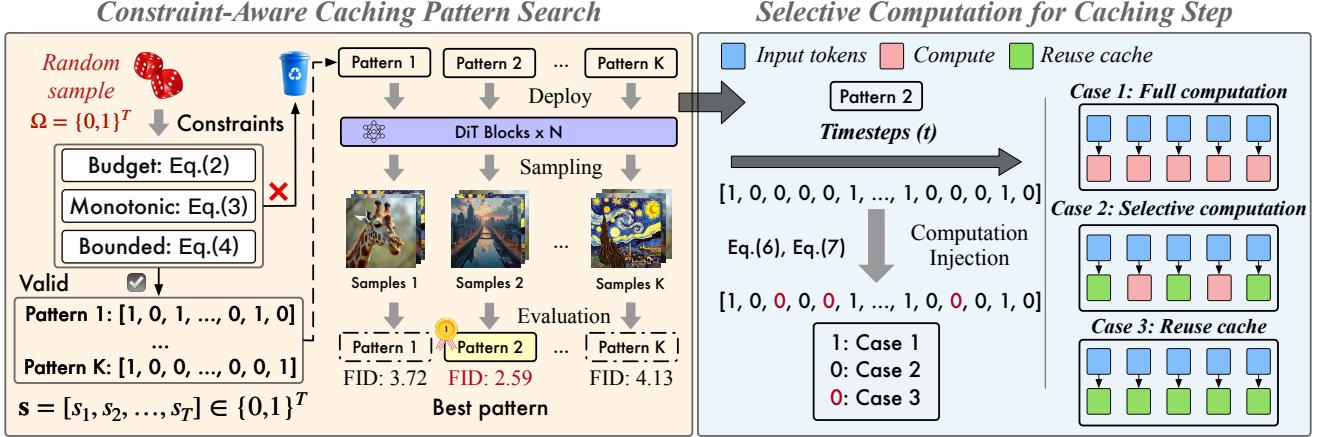


Figure 2: Overall pipeline of ProCache. 1) ProCache explores valid caching patterns under three principled constraints and selects the optimal strategy via lightweight offline sampling based on quality metrics (e.g., FID). 2) It then inserts partial recomputation within contiguous cache steps, selectively updating high-importance tokens in deeper layers.

attention or MLP transformations, and AdaLN (Peebles and Xie 2023) ensures stable feature normalization across varying noise conditions. During inference, DiT generates outputs progressively through multiple timesteps from 1 to T . However, this process can be time-consuming due to the sequential computation over many timesteps.

To accelerate inference, we can exploit temporal redundancy by reusing features across timesteps (Selvaraju et al. 2024). Specifically, instead of recomputing $\mathcal{F}(x_t^l)$ at every timestep t , we reuse previously computed features by approximating $\mathcal{F}(x_{t-k}^l) := \mathcal{F}(x_t^l)$, where x_t^l represents the features at timestep t and layer l and k represents the offset in timesteps. However, existing approaches typically adopt a uniform interval for feature reuse across all timesteps, failing to account for the non-uniform temporal dynamics inherent in DiT. Furthermore, these methods propagate cached features across multiple denoising steps without adaptive refinement. As a result, such strategies may lead to suboptimal approximations, limiting acceleration gains due to unnecessary recomputation or inaccurate feature reuse.

Method Overview. In this paper, we propose *ProCache*, a training-free caching framework for efficient DiT inference. As illustrated in Figure 2, our ProCache consists of two core components: 1) We propose a constraint-aware caching pattern search method that aims to replace uniform-interval caching pattern with a model-tailored one (c.f. Section 3.2). We devise an offline constrained sampling pipeline that searches for promising caching patterns under three principled constraints. The resulting pattern maximizes feature reuse in stable denoising steps while ensuring frequent computation during high-sensitivity ones. 2) We devise a selective computation framework to mitigate error accumulation in long reuse intervals (c.f. Section 3.3). Instead of fully computing features, we selectively compute only a small subset of transformer blocks and tokens at chosen timesteps. This effectively refreshes critical semantic content while maintaining low computational overhead. The overall pipeline of ProCache is shown in Algorithm 1.

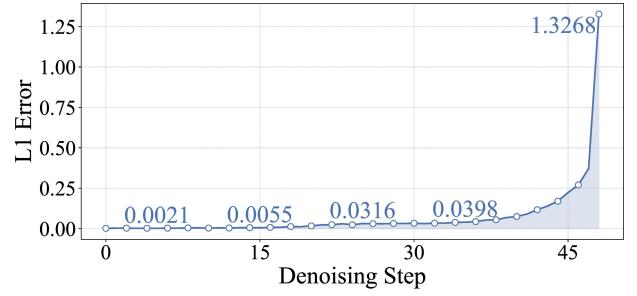


Figure 3: Output L1 error between the current step and the previous step in DiT-XL/2 across the denoising process.

3.2 Constraint-Aware Caching Pattern Search

The aforementioned limitations of uniform-interval caching raise a crucial question: *are all timesteps equally conducive to feature reuse?* To investigate this, we measure the temporal discrepancy in DiT-XL-2 outputs via $\|\mathcal{G}_t(\mathbf{x}_t) - \mathcal{G}_{t-1}(\mathbf{x}_{t-1})\|$ across the denoising process. As shown in Figure 3, feature divergence remains small in early and mid stages but grows sharply in later steps, following an approximately exponential trend. This reveals that uniform caching misaligns with DiT’s temporal dynamics, causing either excessive computation in stable phases or error accumulation in fast-evolving stages. Instead, effective caching should adapt to the non-uniform evolution of features.

Motivated by this, we propose a constraint-aware caching pattern search framework that finds a customized, non-uniform caching schedule for each model that reduces feature reuse in critical phases and maximizes it where stable. To this end, we first represent each caching pattern as a binary sequence $\mathbf{s} = [s_1, s_2, \dots, s_T] \in \{0, 1\}^T$, where $s_t = 1$ denotes computation at step t , and $s_t = 0$ denotes feature reuse from the cache. Let $\mathcal{T}(\mathbf{s}) = \{t \in [T] \mid s_t = 1\}$ denote the set of timesteps where computation is performed, and let $M = |\mathcal{T}(\mathbf{s})| = \sum_{t=1}^T s_t$ be the total number of acti-

Algorithm 1: DiT Inference Pipeline of l -th Block with our ProCache. Note: Before the inference, we have identified an effective caching pattern \mathbf{s}^* via the procedure in Sec. 3.2.

Require: The DiT block \mathcal{F}_l , searched caching pattern \mathbf{s}^* , selective layer set $\mathcal{U}^{\text{cmpt}}$, selective updated timestep set $\mathcal{T}^{\text{cmpt}}$, token selection ratio p , #denoising steps T

- 1: **for** $t = 1$ **to** T **do**
- 2: **if** $\mathbf{s}_t^* = 1$ **then**
- 3: // Case 1: Computation without cache
- 4: $\mathcal{F}_l(\mathbf{x}_t) = \mathbf{x}_t + f_l(\mathbf{x}_t)$; store cache $\mathbf{c} = f_l(\mathbf{x}_t)$
- 5: // AdaLN is omitted for simplicity, same below
- 6: **else**
- 7: **if** $l \in \mathcal{U}^{\text{cmpt}}$ and $t \in \mathcal{T}^{\text{cmpt}}$ **then**
- 8: // Case 2: Selective computation
- 9: Calculate $\mathcal{I}^{\text{cmpt}}$ via Eqn.(9) with p
- 10: $\mathcal{F}_l(\mathbf{x}_t) = \mathbf{x}_t + f_l(\mathbf{x}_t; \mathbf{c}; \mathcal{I}^{\text{cmpt}})$; where \mathbf{c} is from cache // We only compute on important tokens, while the remaining tokens use cache
- 11: Store cache $\mathbf{c} = f_l(\mathbf{x}_t; \mathbf{c}; \mathcal{I}^{\text{cmpt}})$
- 12: **else**
- 13: // Case 3: No computation, use cache
- 14: $\mathcal{F}_l(\mathbf{x}_t) = \mathbf{x}_t + \mathbf{c}$; where \mathbf{c} is from cache
- 15: **end if**
- 16: **end if**
- 17: **end for**
- 18: **return** $\mathcal{F}_l(\mathbf{x}_t)$ // as the input for next block computation

vated steps. We define the ordered activation timestamps as $t_1 < t_2 < \dots < t_M$, obtained by sorting $\mathcal{T}(\mathbf{s})$ in ascending order. We define the i -th reuse interval v_i as:

$$v_i = t_{i+1} - t_i - 1, \quad \text{for } i = 1, 2, \dots, M-1, \quad (1)$$

which denotes the number of consecutive timesteps with cached features between two computed steps.

When sampling candidate sequences from the search space $\Omega = \{0, 1\}^T$, we hope the candidates satisfied the following three constraints, designed to balance computational efficiency, generation stability, and alignment with DiT's temporal dynamics:

1. **Budget constraint:** To ensure significant acceleration, the total number of activation timesteps must remain below a predefined budget $B \ll T$:

$$M = \sum_{t=1}^T s_t \leq B. \quad (2)$$

2. **Monotonic constraint:** As shown in Figure 3, feature changes in DiT are slow in early stages but accelerate in later stages. To preserve fidelity where it matters most, reuse intervals should be non-increasing over time:

$$v_{i+1} \leq v_i, \quad \forall i = 1, \dots, M-2. \quad (3)$$

3. **Bounded constraint:** Excessively long reuse intervals can lead to error accumulation, while overly short ones undermine efficiency. Therefore, each interval is constrained within a feasible range:

$$v^{\min} \leq v_i \leq v^{\max}, \quad \forall i = 1, \dots, M-1, \quad (4)$$

where $v^{\min} \geq 0$ and $v^{\max} \geq v^{\min}$ are hyper-parameters.

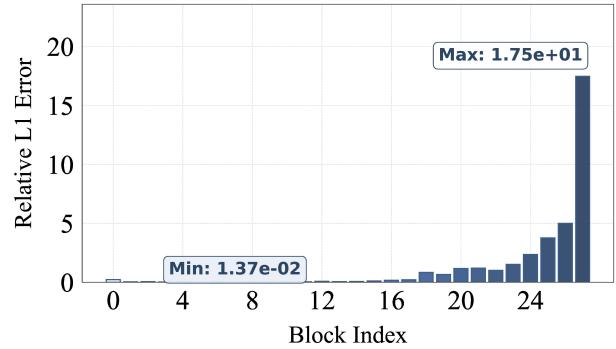


Figure 4: Relative L1 error of features across different blocks in DiT-XL/2, measured at step 20.

Formally, we define the constrained search space as the set of all patterns that satisfy all three above constraints:

$$\mathcal{C} = \{\mathbf{s} \in \{0, 1\}^T \mid \mathbf{s} \text{ satisfies Eqns.(2), (3), and (4)}\}. \quad (5)$$

By restricting the search to this constrained space, we efficiently eliminate invalid patterns while preserving effective variations that align with the temporal dynamics of DiT.

Effective Caching Patterns Derivation. To find the final caching pattern for a given DiT model, we follow an automated, offline procedure that combines constrained sampling with lightweight evaluation. We first generate K candidate patterns $\mathbf{s}^{(1)}, \mathbf{s}^{(2)}, \dots, \mathbf{s}^{(K)}$ from the constrained search space \mathcal{C} . Then, we evaluate each candidate on a small representative dataset on a quality metric such as FID. We adopt the pattern that achieves the best performance as the final pattern \mathbf{s}^* . This process is fully training-free, and completes less than one hour on a single GPU.

3.3 Selective Computation for Caching Step

Given the searched caching pattern \mathbf{s}^* , we incorporate lightweight computations to mitigate error accumulation over multiple denoising steps. Specifically, we define the set of timesteps for computation based solely on the structure of \mathbf{s}^* . Let $\ell(t)$ denote the starting timestep of the maximal contiguous zero block containing t , defined as:

$$\ell(t) = \max \{\tau \leq t \mid s_\tau = 1\} + 1, \quad (6)$$

with $\ell(t) = 1$ if no such τ exists. We apply selective computation at every second position within each zero block. Formally, the set of timesteps selected for updates is:

$$\mathcal{T}^{\text{cmpt}} = \{t \mid s_t = 0, t - \ell(t) + 1 \text{ is even}\}. \quad (7)$$

At each $t \in \mathcal{T}^{\text{cmpt}}$, instead of reusing all cached features, we compute only a subset of layers and tokens, while keeping the rest cached. Formally, the update rule is:

$$\mathbf{x}_i^{(l)} = \begin{cases} f^{(l)}(\mathbf{x}_i^{(l-1)}), & \text{if } l \in \mathcal{U}^{\text{cmpt}} \text{ and } i \in \mathcal{I}^{\text{cmpt}} \\ \text{Cache}^{(l)}(\mathbf{x}_i), & \text{otherwise} \end{cases} \quad (8)$$

where $\mathbf{x}_i^{(l)}$ denotes token \mathbf{x}_i at layer l , $\mathcal{U}^{\text{cmpt}}$ is the set of layers to compute, $\mathcal{I}^{\text{cmpt}}$ denotes the selected token index,

Method	Latency(s) \downarrow	FLOPs(T) \downarrow	Speed \uparrow	FID \downarrow	sFID \downarrow	Precision \uparrow	Recall \uparrow	Inception Score \uparrow
DDPM-250 steps	49.564	118.68	-	2.31	4.98	0.82	0.58	243.42
DDIM-50 steps	4.549	23.74	1.00 \times	2.43	4.40	0.80	0.59	241.25
DDIM-25 steps	2.263	11.87	2.00 \times	3.18	4.74	0.79	0.58	232.01
DDIM-20 steps	1.862	9.49	2.50 \times	3.81	5.15	0.78	0.58	221.43
DDIM-17 steps	1.563	8.07	2.94 \times	4.58	5.76	0.77	0.56	208.72
Δ-DiT($\mathcal{N} = 3$)	2.572	16.46	1.47 \times	3.75	5.70	0.77	0.54	207.57
FORA ($\mathcal{N} = 3$)	2.191	8.59	2.76 \times	3.88	6.43	0.79	0.54	229.02
ToCa ($\mathcal{N} = 3$)	2.087	10.23	2.32 \times	3.04	5.14	0.79	0.56	230.70
ToCa ($\mathcal{N} = 4$)	2.063	8.73	2.72 \times	3.64	5.55	0.78	0.56	223.25
ProCache (Ours)	1.725	8.18	2.90\times	2.96	4.93	0.80	0.57	232.85

Table 1: Quantitative comparison on class-to-image generation on ImageNet with DiT-XL/2.

$f^{(l)}(\cdot)$ denotes the computation operator at the l -th layer, and Cache $^{(l)}(\cdot)$ retrieves the cached activation. Notably, our selective computation strategy introduces only about 3% additional latency, as we compute on a small set of caching steps, a subset of layers (25%), and a subset of tokens (7–30%). We detail the construction of \mathcal{U}^{cmt} and \mathcal{I}^{cmt} below.

Layer Selection: Focus on Deep Blocks. To determine the set of layers to compute \mathcal{U}^{cmt} , we analyze feature error propagation across transformer blocks during the denoising process. As shown in Figure 4, errors accumulate predominantly in deeper layers, where high-level semantic refinement occurs. While shallow layers exhibit strong temporal stability and minimal deviation over timesteps. Motivated by this, we define \mathcal{U}^{cmt} as the indices of the deepest $D = r * L$ layers: $\mathcal{U}^{\text{cmt}} = \{l \mid (L - D) \leq l \leq L\}$, where r is a computed ratio and L is the total number of blocks.

Token Selection: Prioritizing Important Tokens. To determine the set of tokens to compute \mathcal{I}^{cmt} , we consider their dynamic influence during the generation process. Tokens with large feature changes play a more critical role in shaping the output structure, and thus benefit more from fresh computation. To quantify this, we measure token importance based on the magnitude of its attention module output. Let \mathbf{x}_i denote the i -th token in the sequence, and let \mathbf{v}_i be its corresponding output from the attention module. We define its importance score as: $T(\mathbf{x}_i) = \|\mathbf{v}_i\|_2$. We then select the top- $p\%$ most important tokens to form \mathcal{I}^{cmt} :

$$\mathcal{I}^{\text{cmt}} = \{i \mid \text{rank}(T(\mathbf{x}_i)) \leq p\% \times N\}, \quad (9)$$

where N is the number of tokens. In practice, p is chosen to retain only a small fraction (7–30%), ensuring minimal computational overhead while preserving key information.

Note that for the cross-attention and FFN modules, we apply this token selection strategy to compute only the top- $p\%$ most important tokens. However, in the self-attention module, each token attends to all others, making partial computation prone to error propagation. Thus, we do not employ this token selection strategy for self-attention module.

4 Experiments

4.1 Experiment Setup

Baselines. We employ three SoTA vision generative models on two tasks: PixArt- α (Chen et al. 2023) with 20 DPM-Solver++ (Lu et al. 2023) steps for text-to-image generation; FLUX.1-dev and FLUX.1-schnell (Labs 2024) utilizing 50 and 4 Rectified Flow (Lipman et al. 2023) steps respectively for text-to-image generation; and DiT-XL/2 (Peebles and Xie 2023) with 50 DDIM (Song, Meng, and Ermon 2022) steps for class-conditional image generation.

Implementation Details. For DiT-XL/2, we perform class-conditional generation on the full 1,000-class ImageNet (Deng et al. 2009), synthesizing 50,000 images at 256×256 resolution. PixArt- α text-to-image experiments employed 30,000 randomly selected captions from the COCO-2017 (Lin et al. 2015) validation set, generating corresponding images at 256×256 resolution. For FLUX family, we generated high-resolution 1024×1024 images on PartiPrompts (Yu et al. 2022) containing 1,632 text prompts. We set the sampling budget K to 5 and $B \in \{7, 17\}$ in our caching pattern search. In selective computation, we compute on the deep 25–50% blocks. v^{\min} and v^{\max} are typically set in the range of 2–5. We put more details in the Appendix.

4.2 Quantitative Evaluations

Comparisons with SoTA on ImageNet. In Table 1, our ProCache achieves a $2.90\times$ speedup while achieving the best FID and sFID scores of 2.96 and 4.93, respectively. Compared to the uniform caching strategy like FORA, our ProCache achieves nearly a 30% improvement in FID and sFID metrics while enhancing the speedup ratio by 5%. This result demonstrates that ProCache outperforms other acceleration techniques. Moreover, ProCache maintains competitive Precision, Recall and Inception Score, surpassing existing acceleration methods and reinforcing its advantage.

Comparisons with SoTA on MS-COCO30k. To evaluate ProCache’s performance in text-to-image generation tasks, we conducted comprehensive experiments on the COCO30K dataset. As shown in Table 3, we compare our ProCache with three SoTA training-free acceleration methods across various configuration settings. For comparative

Models	Method	Latency(s) \downarrow	FLOPs(T) \downarrow	Speed \uparrow	Image Reward \uparrow
FLUX.1-dev (Labs 2024)	Euler-50 steps	33.85	3719.50	1.00 \times	1.202
	68% steps	23.02	2529.26	1.47 \times	1.200
	FORA (Selvaraju et al. 2024)	20.82	2483.32	1.51 \times	1.196
	ToCa ($\mathcal{N} = 2$)	19.88	2458.06	1.51 \times	1.202
FLUX.1-schnell (Labs 2024)	ProCache (Ours)	18.73	2415.25	1.54\times	1.207
	LCM-4 steps	2.882	277.88	1.00 \times	1.133
	75% steps	2.162	208.41	1.33 \times	1.132
	FORA¹ (Selvaraju et al. 2024)	2.365	225.60	1.23 \times	1.129
	FORA² (Selvaraju et al. 2024)	2.365	225.60	1.23 \times	1.124
	FORA³ (Selvaraju et al. 2024)	2.365	225.60	1.23 \times	1.123
ToCa ($\mathcal{N} = 2$)	ProCache (Ours) ($\mathcal{N} = 2$)	1.890	181.30	1.53 \times	1.134
	ProCache (Ours) ($\mathcal{N} = 2$)	1.817	177.26	1.56\times	1.138

Table 2: Quantitative comparison in text-to-image generation for FLUX on Image Reward. FORA¹, FORA², and FORA³ represent different step-skipping configurations where the model bypasses step 2, 3, and 4 respectively during generation.



Figure 5: Image generation samples at 1024×1024 resolutions under $1.56 \times$ speed-up ratios.

Method	Latency(s) \downarrow	FLOPs(T) \downarrow	Speed \uparrow	COCO2017	
				FID \downarrow	CLIP \uparrow
PixArt-α	2.142	11.18	1.00 \times	28.12	16.29
50% steps	1.044	5.59	2.00 \times	37.62	15.81
Δ-DiT	1.724	7.68	1.54 \times	28.91	16.41
FORA¹	1.575	5.66	1.98 \times	29.63	16.40
FORA²	1.301	6.05	2.79 \times	29.84	16.42
ToCa	1.473	6.33	1.77 \times	28.02	16.43
ProCache	1.215	5.70	1.96\times	27.66	16.45

Table 3: Quantitative comparison in text-to-image generation for PixArt- α on MS-COCO2017. FORA¹ and FORA² denote configurations with different activation intervals, corresponding to $\mathcal{N} = 2$ and $\mathcal{N} = 3$, respectively.

analysis, we also included the 10-step DPM-Solver++ sampling approach. The quantitative results demonstrate that under an acceleration ratio of approximately 2.0, our ProCache achieves the best trade-off between generation quality and inference speed, outperforming existing methods by securing the lowest FID-30k score and the highest CLIP score on the MS-COCO2017 generation task.

Comparisons with SoTA on PartiPrompts. We further evaluate our ProCache on state-of-the-art text-to-image dif-



Figure 6: Qualitative comparisons of 512×512 image generations between the original sampler (top, $1.00 \times$ speed) and ProCache (bottom, $2.90 \times$ speed-up).

fusion models featuring higher resolutions and fewer denoising steps, demonstrating its effectiveness under advanced generation conditions. Since this model can generate images in just four steps, our ProCache does not use the pattern search strategy and instead adopts a fixed interval of 2 steps for generation. In Table 2, our ProCache achieves superior acceleration trade-off compared to existing methods, maintaining comparable Image Reward scores to other baselines at an approximately around $1.55 \times$ speedup ratio on both the 50-step FLUX.1-dev and 4-step FLUX.1-schnell models. In contrast to FORA that exhibit quality degradation even at $1.5 \times$ acceleration, ProCache preserves perceptual quality metrics while doubling inference efficiency.

Method	Speed↑	FID↓	sFID↓
Default ($B = 13$)	$2.93\times$	4.75	8.43
+ Searched Pattern	$2.93\times$	3.15	5.12
+ Selective Computation	$2.90\times$	3.28	5.95
ProCache (Ours)	$2.90\times$	2.94	4.93

(a) Ablations on the proposed components.

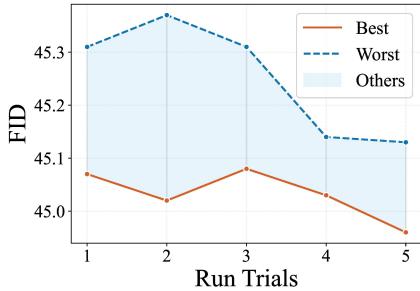
K	FID↓	Inception Score↑
5	45.07	181.69
10	45.02	182.50
15	44.96	181.63

(b) Ablations on sample budget K .

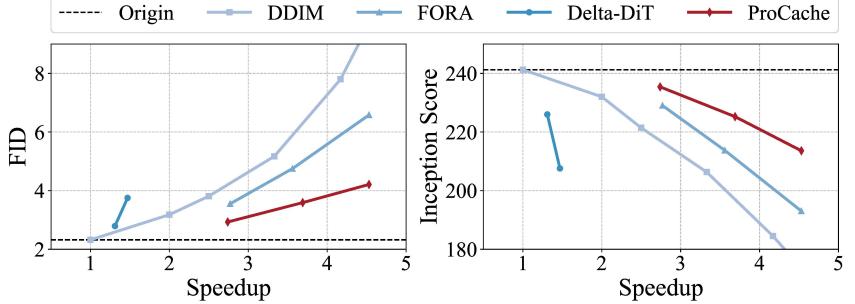
r	FLOPs (T)↓	FID↓	Inception Score↑
50%	9.138	45.32	184.18
65%	8.911	45.33	184.09
75%	8.594	45.31	183.61
90%	8.344	45.38	182.46

(c) Ablations on proportion r of compute blocks.

Table 4: Ablation studies of our ProCache on ImageNet with DiT-XL/2.



(a) Performance stability across 5 run trials.



(b) Comparison of speedup ratio versus quality metrics (FID and Inception Score).

Figure 7: Robustness and efficiency–quality trade-off analysis of DiT-XL/2 on ImageNet-1k (a) and ImageNet-50k (b).

4.3 Qualitative Analysis

Text-to-image generation. We evaluated ProCache in high-resolution text-to-image generation tasks. As illustrated in Figure 5 for FLUX.1-dev, the visual comparison demonstrates ProCache’s ability to maintain perceptual quality while achieving substantial acceleration.

Class-conditional image generation. In Figure 6, we present image generation samples from ProCache. Compared with the original image, ProCache maintains perceptual quality indistinguishable from the original model at $2.90\times$ acceleration, with imperceptible losses in fine detail. We put more qualitative results in the Appendix.

4.4 Further Experiments

Effect of Different Components. In Table 4a, the results demonstrate the effectiveness of ProCache’s two core components. First, the searched caching pattern improves both FID and sFID without any computational overhead, maintaining identical speed to the default baseline while achieving better generation quality. Second, the selective computation mechanism further enhances performance with minimal additional cost, the slight reduction in speed ($2.90\times$ vs. $2.93\times$), yet yields measurable improvements.

Ablations on Sampling Budget. In Table 4b, we present ablations on sampling budget K . As K increases from 5 to 15, FID scores exhibit an improvement, indicating that larger sampling budget enables discovery of higher-quality activation schedules. Notably, even with $K = 5$, the lowest sampling budget and the least performant setting, our method still surpasses existing approaches. Given its superior computational efficiency, we adopt $K = 5$ as the default budget.

Effect of computing #blocks. To study this, we evaluate

performance at different depths (50–90% of total blocks) on DiT-XL/2. In Table 4c, 75% of blocks achieves the lowest FID and a competitive Inception Score. Going deeper (*e.g.*, 50%) brings no FID improvement and lowers Inception Score, indicating diminishing returns beyond a threshold. To maximize speedup, we use 25% in main experiments—the minimal depth that preserves quality.

Robustness of Caching Pattern Search. Our constraint-aware cache pattern search in Sec. 3.2 involves randomness. To verify its robustness in practical use, we conducted five run trials on the ImageNet-1K. In Figure 7a, the variations in both FID and Inception Score across runs are minimal. This indicates that our algorithm always finds promising cache pattern, maintaining reliable generative performance regardless of the random sampling variation.

High-acceleration Ratios Performance. To comprehensively evaluate the trade-off between speedup and model accuracy, we conducted experiments on ImageNet-50k using DDIM with varying steps, Δ -DiT, and FORA. In Figure 7b, our ProCache outperforms prior caching methods, which suffer from feature error accumulation. Our ProCache reduces quality degradation by 56.2% and achieves robust performance in acceleration regimes exceeding $4.53\times$.

5 Conclusion

In this work, we addressed the challenge of high computational costs associated with DiTs in generative modeling by proposed ProCache, a training-free dynamic feature caching framework. ProCache achieves $2.90\times$ speedup with minimal impact on generation quality, outperforming existing methods. Our results highlight its potential to enable efficient real-time deployment of DiTs.

Acknowledgments

This work was partially supported by Postdoctoral Fellowship Program of CPSF (GZC20251043) and the Young Scholar Project of Pazhou Lab (No.PZL2021KF0021).

References

Bolya, D.; and Hoffman, J. 2023. Token Merging for Fast Stable Diffusion. arXiv:2303.17604.

Brooks, T.; Peebles, B.; Holmes, C.; DePue, W.; Guo, Y.; Jing, L.; Schnurr, D.; Taylor, J.; Luhman, T.; Luhman, E.; Ng, C.; Wang, R.; and Ramesh, A. 2024. Video generation models as world simulators.

Chen, J.; Yu, J.; Ge, C.; Yao, L.; Xie, E.; Wu, Y.; Wang, Z.; Kwok, J.; Luo, P.; Lu, H.; and Li, Z. 2023. PixArt- α : Fast Training of Diffusion Transformer for Photorealistic Text-to-Image Synthesis. arXiv:2310.00426.

Chen, P.; Shen, M.; Ye, P.; Cao, J.; Tu, C.; Bouganis, C.-S.; Zhao, Y.; and Chen, T. 2024. Δ -DiT: A Training-Free Acceleration Method Tailored for Diffusion Transformers. arXiv:2406.01125.

Dao, T.; Fu, D. Y.; Ermon, S.; Rudra, A.; and Ré, C. 2022. FlashAttention: Fast and memory-efficient exact attention with IO-awareness. *arXiv preprint arXiv:2205.14135*.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.

Fang, G.; Li, K.; Ma, X.; and Wang, X. 2024. TinyFusion: Diffusion Transformers Learned Shallow. arXiv:2412.01199.

Fang, G.; Ma, X.; and Wang, X. 2023. Structural Pruning for Diffusion Models. arXiv:2305.10924.

Hessel, J.; Holtzman, A.; Forbes, M.; Bras, R. L.; and Choi, Y. 2022. CLIPScore: A Reference-free Evaluation Metric for Image Captioning. arXiv:2104.08718.

Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2018. GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. arXiv:1706.08500.

Labs, B. F. 2024. FLUX. <https://github.com/black-forest-labs/flux>.

Li, S.; Hu, T.; van de Weijer, J.; Shahbaz Khan, F.; Liu, T.; Li, L.; Yang, S.; Wang, Y.; Cheng, M.-M.; et al. 2024. Faster diffusion: Rethinking the role of the encoder for diffusion model inference. *Advances in Neural Information Processing Systems*, 37: 85203–85240.

Li, Y.; Wang, H.; Jin, Q.; Hu, J.; Chemerys, P.; Fu, Y.; Wang, Y.; Tulyakov, S.; and Ren, J. 2023. SnapFusion: Text-to-Image Diffusion Model on Mobile Devices within Two Seconds. arXiv:2306.00980.

Lin, T.-Y.; Maire, M.; Belongie, S.; Bourdev, L.; Girshick, R.; Hays, J.; Perona, P.; Ramanan, D.; Zitnick, C. L.; and Dollár, P. 2015. Microsoft COCO: Common Objects in Context. arXiv:1405.0312.

Lipman, Y.; Chen, R. T. Q.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. arXiv:2210.02747.

Liu, F.; Zhang, S.; Wang, X.; Wei, Y.; Qiu, H.; Zhao, Y.; Zhang, Y.; Ye, Q.; and Wan, F. 2025. Timestep Embedding Tells: It’s Time to Cache for Video Diffusion Model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 7353–7363.

Liu, J.; Geddes, J.; Guo, Z.; Jiang, H.; and Nandwana, M. K. 2024. SmoothCache: A Universal Inference Acceleration Technique for Diffusion Transformers. arXiv:2411.10510.

Liu, W.; and Zhang, S. Q. 2024. HQ-DiT: Efficient Diffusion Transformer with FP4 Hybrid Quantization. arXiv:2405.19751.

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:2206.00927.

Lu, C.; Zhou, Y.; Bao, F.; Chen, J.; Li, C.; and Zhu, J. 2023. DPM-Solver++: Fast Solver for Guided Sampling of Diffusion Probabilistic Models. arXiv:2211.01095.

Ma, X.; Fang, G.; Mi, M. B.; and Wang, X. 2024a. Learning-to-Cache: Accelerating Diffusion Transformer via Layer Caching. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

Ma, X.; Fang, G.; and Wang, X. 2023. DeepCache: Accelerating Diffusion Models for Free. arXiv:2312.00858.

Ma, X.; Wang, Y.; Jia, G.; Chen, X.; Liu, Z.; Li, Y.-F.; Chen, C.; and Qiao, Y. 2024b. Latte: Latent Diffusion Transformer for Video Generation. arXiv:2401.03048.

Norooshi, M.; Hadji, I.; Martinez, B.; Bulat, A.; and Tzimiropoulos, G. 2024. You Only Need One Step: Fast Super-Resolution with Stable Diffusion via Scale Distillation. arXiv:2401.17258.

Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4195–4205.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. arXiv:2112.10752.

Salimans, T.; and Ho, J. 2022. Progressive Distillation for Fast Sampling of Diffusion Models. arXiv:2202.00512.

Selvaraju, P.; Ding, T.; Chen, T.; Zharkov, I.; and Liang, L. 2024. FORA: Fast-Forward Caching in Diffusion Transformer Acceleration. arXiv:2407.01425.

Shang, Y.; Yuan, Z.; Xie, B.; Wu, B.; and Yan, Y. 2023. Post-training Quantization on Diffusion Models. arXiv:2211.15736.

So, J.; Lee, J.; Ahn, D.; Kim, H.; and Park, E. 2023. Temporal Dynamic Quantization for Diffusion Models. arXiv:2306.02316.

Song, J.; Meng, C.; and Ermon, S. 2022. Denoising Diffusion Implicit Models. arXiv:2010.02502.

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: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. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 15903–15935.

Yu, J.; Xu, Y.; Koh, J. Y.; Luong, T.; Baid, G.; Wang, Z.; Vasudevan, V.; Ku, A.; Yang, Y.; Ayan, B. K.; Hutchinson, B.; Han, W.; Parekh, Z.; Li, X.; Zhang, H.; Baldridge, J.; and Wu, Y. 2022. Scaling Autoregressive Models for Content-Rich Text-to-Image Generation. arXiv:2206.10789.

Yue, Z.; Liao, K.; and Loy, C. C. 2025. Arbitrary-steps image super-resolution via diffusion inversion. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 23153–23163.

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, W.; Wang, H.; Zhou, J.; and Lu, J. 2024. DC-Solver: Improving Predictor-Corrector Diffusion Sampler via Dynamic Compensation. arXiv:2409.03755.

Zhao, X.; Jin, X.; Wang, K.; and You, Y. 2025. Real-Time Video Generation with Pyramid Attention Broadcast. In *The Thirteenth International Conference on Learning Representations*.

Zheng, Z.; Peng, X.; Yang, T.; Shen, C.; Li, S.; Liu, H.; Zhou, Y.; Li, T.; and You, Y. 2024. Open-sora: Democratizing efficient video production for all. *arXiv preprint arXiv:2412.20404*.

Zou, C.; Liu, X.; Liu, T.; Huang, S.; and Zhang, L. 2025. Accelerating Diffusion Transformers with Token-wise Feature Caching. In *The Thirteenth International Conference on Learning Representations*.

Appendix of “ProCache: Constraint-Aware Feature Caching with Selective Computation for Diffusion Transformer Acceleration”

In the supplementary, we provide more experimental results and more details on the ProCache method. We organize our supplementary as follows.

- In Section A, we provide more implementation details for evaluation and model configuration.
- In Section B, we provide more details about the proposed ProCache.
- In Section C, we provide additional experimental results.
- In Section D, we show more visualization results for qualitative analysis.
- In Section E, we provide the final pattern searched by ProCache for different models.

A More Implementation Details

A.1 More Details of Evaluation

For DiT-XL/2, we performed class-conditional generation on the full 1,000-class ImageNet dataset (Deng et al. 2009), synthesizing 50,000 images at 256×256 resolution. Model optimization was guided by FID-1k (Heusel et al. 2018) scores, with final performance evaluated using FID-50k. Secondary evaluation metrics included sFID, Inception Score, and Precision & Recall metrics. The PixArt- α text-to-image experiments employed 30,000 randomly selected captions from the COCO-2017 (Lin et al. 2015) validation set, generating corresponding images at 256×256 resolution. We conducted parameter optimization using FID-3k and performed final quality assessment with FID-30k. Image-text alignment was quantified using CLIP Score (Hessel et al. 2022) between generated images and input prompts. For the FLUX family, we generated high-resolution 1024×1024 images on the PartiPrompts dataset (Yu et al. 2022) containing 1,632 text prompts. We use Image Reward (Xu et al. 2023), a metric that better suited to measuring human preference, to evaluate the generation quality. All experiments were implemented in Python 3.9 with PyTorch 2.4.0. Computational resources included two NVIDIA RTX 3090 (24GB) GPUs for DiT-XL/2 and PixArt- α experiments, and a single NVIDIA A800 (80GB) GPU for FLUX models evaluations.

A.2 More Details of Hyper-parameters in Different Diffusion Models

The ProCache framework employs six key hyper-parameters that govern its dynamic caching and selective computation strategies: K denotes the sampling budget for candidate pattern generation, B represents the activation budget constraining total computational steps, p specifies the token selection ratio for importance-based token computation, v^{\min} and v^{\max} define the minimum and maximum reuse intervals for feature caching, and r indicates the proportion of transformer blocks requiring computation updates. These hyper-parameters collectively balance acceleration efficiency with generation quality by adapting to the non-uniform temporal dynamics of diffusion transformers, with their optimal configurations varying across different model architectures as detailed in Table A.

Model	K	B	p	v^{\min}	v^{\max}	r
DiT-XL/2	5	17	7%	2	5	75%
PixArt- α	5	7	30%	2	3	50%
FLUX.1-dev	5	17	30%	2	3	50%
FLUX.1-schnell	-	-	30%	-	-	50%

Table A: Hyperparameter settings on different models. In FLUX.1-schnell, we abandoned the pattern search strategy.

B More Details of our ProCache

B.1 Constraint-Aware Caching Pattern Search

The constraint-aware caching pattern search framework lies at the core of ProCache, enabling the discovery of optimal non-uniform caching schedules that align with the temporal dynamics of diffusion transformers. As formally presented in Algorithm A, our framework employs a constrained random sampling strategy to efficiently explore the vast space of possible caching patterns while respecting three critical constraints identified in Section 3.2 of the main paper.

The algorithm begins by initializing an empty candidate set \mathcal{C} and iteratively generates random binary sequences $\mathbf{s} \in \{0, 1\}^T$ where each element s_t indicates whether computation (1) or caching (0) occurs at timestep t . For each candidate sequence,

Algorithm A: Constraint-Aware Caching Pattern Search

T Total denoising steps
 K Sampling quota
Require: B Maximum number of activation steps
 v^{\min} Minimum reuse interval
 v^{\max} Maximum reuse interval
 M Maximum number of attempts

Ensure: $\mathcal{S} \subseteq \{0, 1\}^T$: Set of K feasible cache-activation scheduling strategies

Initialize an empty set $\mathcal{C} \leftarrow \emptyset$
 Set attempt counter $m \leftarrow 0$
while $|\mathcal{C}| < K$ and $m < M$ **do**
 Generate a random binary sequence $\mathbf{s} \in \{0, 1\}^T$ via uniform stochastic sampling
 $m \leftarrow m + 1$
if $\sum_{t=1}^T s_t > B$ **then**
 Discard \mathbf{s} due to budget violation
 Continue to next iteration
end if
 Extract activation indices $\mathcal{A} = \{t \in [1, T] \mid s_t = 1\}$, sorted in increasing order
for $i = 1$ to $|\mathcal{A}| - 1$ **do**
 Compute inter-activation interval $v_i = \mathcal{A}_{i+1} - \mathcal{A}_i - 1$
if $v_i < v^{\min}$ or $v_i > v^{\max}$ **then**
 Discard \mathbf{s} due to interval constraint violation
 Break and continue to next iteration
end if
end for
 Add \mathbf{s} to candidate set: $\mathcal{C} \leftarrow \mathcal{C} \cup \{\mathbf{s}\}$
end while
return $\mathcal{S} \leftarrow$ top- K distinct strategies from \mathcal{C}

the algorithm verifies three essential constraints in a specific order of computational efficiency. First, it checks the **budget constraint** by ensuring the total number of activation steps satisfies $\sum_{t=1}^T s_t \leq B$, discarding sequences that violate this fundamental requirement for target acceleration ratio. Second, the algorithm extracts the sorted activation indices $\mathcal{A} = \{t_1, t_2, \dots, t_M\}$ (where $t_1 < t_2 < \dots < t_M$) and verifies the **bounded constraint** by confirming all inter-activation intervals $v_i = t_{i+1} - t_i - 1$ fall within the specified range $[v^{\min}, v^{\max}]$, which prevents both excessive error accumulation from long reuse intervals and unnecessary computation from overly short ones.

Notably, while the **monotonic constraint** ($v_{i+1} \leq v_i$) is described in the main paper as ensuring longer reuse intervals in early stable stages and shorter ones in later dynamic stages, our implementation handles this constraint implicitly through the evaluation phase rather than the sampling phase. This design choice significantly improves sampling efficiency while still capturing the desired temporal characteristics—candidate patterns violating the monotonic principle typically receive lower quality scores during evaluation and are thus naturally filtered out.

The algorithm continues this constrained sampling process until either K valid patterns are collected or the maximum attempt count M is reached, ensuring computational tractability. The resulting candidate set \mathcal{C} is then evaluated on a small representative dataset using quality metrics such as FID, with the top- K performing patterns selected for final deployment. This two-stage approach—constrained sampling followed by lightweight evaluation—strikes an optimal balance between exploration efficiency and solution quality, requiring only approximately one hour on a single GPU to identify the optimal caching pattern for a given diffusion transformer model.

B.2 Selective Computation for Caching Step

Token Selection and Updates. DiT architectures allow processing a variable number of tokens, offering flexibility for efficient inference. Leveraging this property, we first identify a critical subset of tokens—referred to as *base tokens*—and restrict heavy computation to only these tokens, while skipping the rest using cached values from previous timesteps. The cache is then selectively updated across layers.

A key challenge lies in identifying these base tokens efficiently. Prior works (Zhang et al. 2025; Zou et al. 2025) often rely on attention-based token importance scores, e.g., by selecting tokens with high attention weights in self-attention maps. However, such strategies require access to attention scores, which are not available in optimized attention implementations like FlashAttention (Dao et al. 2022) and memory-efficient attention. These implementations trade off intermediate score outputs for speed and memory efficiency, making score-based selection incompatible in practice. To resolve this incompatibility, we adopt

a score-free yet effective proxy for token importance—namely, the ℓ_2 -norm of the value vector in attention layers. Empirically, tokens with higher attention scores often correspond to larger value vector norms. Hence, the norm of the value projection serves as a practical and efficient criterion for importance estimation. Therefore, we define the importance score of token x_i using its corresponding value vector in the attention module: $v_i = \text{ValueProj}(x_i)$.

Computation Injection for Long Cache Sequences. We consider the risk of prolonged cache reuse. Even with deep-layer updates, long contiguous sequences of $s_t = 0$ can lead to irreversible error drift. To prevent this, we inject lightweight partial computations into maximal contiguous zero blocks in s .

Specifically, for every such block $[0, 0, \dots, 0]$, we insert partial updates at regular intervals—specifically, at every second position starting from the second:

$$\begin{aligned} [0] &\rightarrow [0] \\ [0, 0] &\rightarrow [0, \alpha] \\ [0, 0, 0] &\rightarrow [0, \alpha, 0] \\ [0, 0, 0, 0] &\rightarrow [0, \alpha, 0, \alpha] \\ [0, 0, 0, 0, 0] &\rightarrow [0, \alpha, 0, \alpha, 0] \\ [0, 0, 0, 0, 0, 0] &\rightarrow [0, \alpha, 0, \alpha, 0, \alpha] \end{aligned}$$

where α denotes a timestep with partial computation (deep-layer update only). This injection ensures that no cache-only sequence exceeds two consecutive steps without correction, improving stability with minimal overhead.

C More Experimental Results

C.1 Demonstration of ProCache’s Flexibility

The flexibility of ProCache is demonstrated in Table B, which presents comprehensive performance across varying computational budgets. Our framework exhibits remarkable adaptability by allowing precise control over the acceleration-quality trade-off through the activation budget parameter B . When configured with $B = 21$, ProCache achieves a $2.74 \times$ speedup while delivering state-of-the-art FID, significantly outperforming existing methods at comparable acceleration ratios. With a more aggressive budget of $B = 17$, the framework maintains exceptional generation quality while achieving a $2.90 \times$ speedup, surpassing all baseline methods in both quantitative metrics and visual fidelity. Most notably, when pushed to an extreme budget of $B = 13$, ProCache still delivers a substantial $3.69 \times$ acceleration with acceptable quality degradation, demonstrating robust performance even in high-acceleration regimes. This tunable characteristic makes ProCache particularly valuable for diverse deployment scenarios—from resource-constrained edge devices requiring maximum acceleration to quality-sensitive applications where minimal degradation is paramount—without requiring any model retraining or architectural modifications.

Method	Latency(s) \downarrow	FLOPs(T) \downarrow	Speed \uparrow	FID \downarrow	sFID \downarrow	Precision \uparrow	Recall \uparrow	Inception Score \uparrow
DDPM-250 steps	49.564	118.68	-	2.31	4.98	0.82	0.58	243.42
DDIM-50 steps	4.549	23.74	1.00 \times	2.43	4.40	0.80	0.59	241.25
DDIM-25 steps	2.263	11.87	2.00 \times	3.18	4.74	0.79	0.58	232.01
DDIM-20 steps	1.862	9.49	2.50 \times	3.81	5.15	0.78	0.58	221.43
DDIM-17 steps	1.563	8.07	2.94 \times	4.58	5.76	0.77	0.56	208.72
Δ-DiT($\mathcal{N} = 3$)	2.572	16.46	1.47 \times	3.75	5.70	0.77	0.54	207.57
FORA ($\mathcal{N} = 3$)	2.191	8.59	2.76 \times	3.88	6.43	0.79	0.54	229.02
ToCa ($\mathcal{N} = 3$)	2.087	10.23	2.32 \times	3.04	5.14	0.79	0.56	230.70
ToCa ($\mathcal{N} = 4$)	2.063	8.73	2.72 \times	3.64	5.55	0.78	0.56	223.25
ProCache ($B = 21$)	2.049	8.66	2.74 \times	2.93	4.94	0.81	0.57	235.40
ProCache ($B = 17$)	1.725	8.18	2.90 \times	2.96	4.93	0.80	0.57	232.85
ProCache ($B = 13$)	1.654	6.42	3.69\times	3.59	6.05	0.80	0.55	225.23

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

C.2 Feasibility of Offline Evaluation

For the ProCache framework, the optimal caching pattern is selected through offline evaluation on a small representative dataset to minimize computational overhead. To validate the feasibility of this approach, we conduct ablation studies on ImageNet where we fix a computational budget and generate 5 random patterns. As shown in Table C, the pattern achieving the best

performance on the small dataset (FID-1k) consistently demonstrates superior performance on the full dataset (FID-50k) as well, confirming that offline evaluation on a small subset reliably identifies patterns that generalize to the complete dataset. This critical finding validates our practical design choice of using lightweight offline evaluation, which significantly reduces the search time from hours to minutes while maintaining selection accuracy and eliminating the need for resource-intensive full-dataset evaluation during the pattern search phase.

Run Trials	FID-1k ↓	FID-50k ↓
1	45.44	3.41
2	45.48	3.42
3	45.64	3.61
4	45.45	3.46
5	45.36	3.39

Table C: Consistency between small-scale (FID-1k) and full-scale (FID-50k) evaluation of caching patterns on ImageNet. The pattern achieving optimal performance on the small dataset consistently generalizes to the complete dataset.

C.3 Computational Efficiency of Pattern Search

The computational efficiency of our constraint-aware pattern search algorithm is critical for practical deployment. We evaluate its performance on a standard CPU environment to demonstrate its lightweight nature. When configured with the default sampling quota $K = 5$, the algorithm completes the search process in approximately 0.001 seconds on an AMD Ryzen 7 5800H processor (3.2GHz). This near-instantaneous execution time confirms that the pattern search introduces negligible overhead to the overall inference pipeline. As expected, the search duration scales linearly with the maximum attempt count M , following the relationship $t \propto M$, while remaining well below one second even for significantly larger M values.

To further investigate the relationship between search effort and solution space exploration, we conduct experiments with varying maximum attempt counts while fixing other parameters ($B = 17$, denoising steps $T = 50$, $v^{\min} = 2$, $v^{\max} = 5$). Table D presents the number of valid patterns discovered across different M values, with the sampling quota set to $K = 1000$ to maximize solution discovery without artificial constraints. Notably, the results reveal that the constrained search space contains only 30 valid patterns for this configuration. The search process rapidly converges, finding 20 patterns with just 10^3 attempts and reaching the complete set of 30 patterns by 10^6 attempts. Crucially, increasing M beyond 10^6 yields no additional valid patterns, as evidenced by identical counts at $M = 10^6$ and $M = 10^7$. This convergence behavior confirms that our default setting of $M = 10^6$ provides complete exploration of the feasible solution space while maintaining computational efficiency, making higher attempt counts unnecessary for practical implementation.

Maximum Attempts (M)	Valid Patterns Found	Search Time (s)
10^3	20	0.008
10^4	25	0.086
10^5	27	0.856
10^6	30	8.572
10^7	30	86.07

Table D: Relationship between maximum attempt count and the number of valid patterns discovered under fixed constraints ($B = 17$, $T = 50$, $v^{\min} = 2$, $v^{\max} = 5$).

D More Visualization Results

To further illustrate the qualitative improvements of our method, we present visualization examples on Pixart- α and DiT (Figure B, Figure C, Figure D). With similar speed-up ratios compared with FORA (Selvaraju et al. 2024) ($1.96 \times$), our method produces more realistic, detailed images and better alignment with original images and text prompts. For example, under the prompt *A young man dressed in ancient Chinese clothing, asian people, has a handsome face and wears a white robe.*”, the image generated by FORA exhibit significant visual degradation, affecting both the main subject and fine details. In contrast, ProCache preserves details effectively, retaining even subtle elements such as the blue light adjacent to the samurai’s body. These results showcase the superior fidelity and consistency of our method in generating high-quality outputs across diverse scenarios.

E Visualization of Searched Caching Pattern

The searched caching patterns for PixArt- α , DiT-XL/2, and FLUX.1-dev are shown in Figure A. Unlike uniform caching strategies that apply fixed intervals throughout the denoising process, ProCache discovers non-uniform patterns that align with

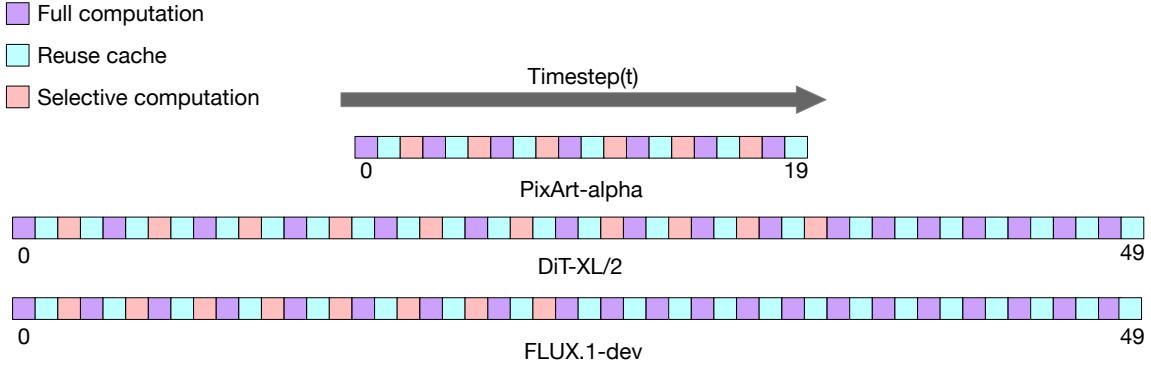


Figure A: Visualization of the final caching patterns searched by ProCache for different models.

the temporal dynamics of diffusion transformers. Specifically, the patterns exhibit longer reuse intervals in early and middle denoising stages where feature changes remain minimal, and progressively shorter intervals in later stages where features evolve rapidly. This adaptive scheduling effectively balances computational efficiency with generation quality by allocating more computation resources to critical time periods while maximizing reuse during stable phases. ProCache achieves superior quality-speed trade-offs compared to uniform caching strategies.

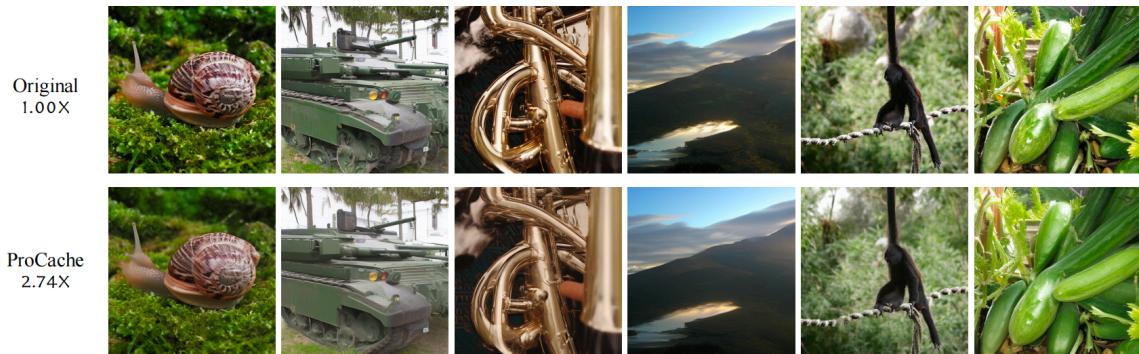


Figure B: Visualization examples of DiT-XL/2.



Figure C: Visualization examples for different acceleration methods on Pixart- α .

Original
1.00X

ProCache
1.96X

FORA
1.98X

"A bucket bag made of blue suede. The bag is decorated with intricate golden paisley patterns."



"A modern architectural building with large glass windows, situated on a cliff."



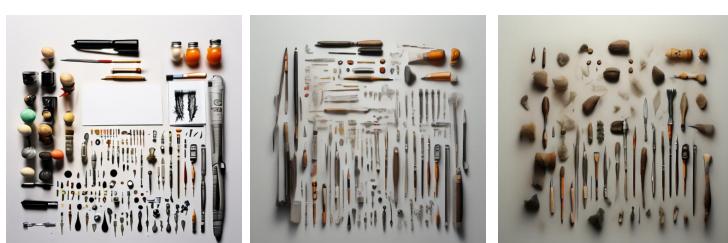
"A small cactus with a happy face in the Sahara desert."



"An astronaut sits in a diner, eating fries, in a cinematic scene."



"knolling of a drawing tools for painter."



"A professional portrait photo of an anthropomorphic cat wearing fancy gentleman hat and jacket."



"The image features a woman wearing a red shirt with an icon. She appears to be posing for the camera."

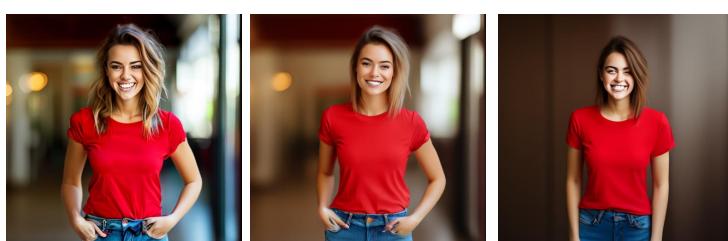


Figure D: Visualization examples for different acceleration methods on Pixart- α .