Which Tokens to Merge? Diffusion Dynamics for Efficient Image Generation

SeungJu Cha **Ye-Chan Kim** **HyunGee Kim** **Sungho Koh** **Dong-Jin Kim**†

Hanyang University, South Korea

{sju9020, dpcksdl78, khjiiii2002, ksh000906, djdkim}@hanyang.ac.kr

† Corresponding author

Abstract

Token merging has emerged as an effective training-free strategy for accelerating diffusion models by reducing redundant token computation while minimizing generation degradation. A key challenge is determining which tokens can be safely merged and which should be preserved for subsequent computation. Existing method addresses this by preserving prompt-aligned foreground regions, thereby protecting semantically relevant structures. However, this text-relevance criterion overlooks background objects, local details that are not explicitly mentioned in the prompt, leading to degraded structural coherence. In this work, we propose **D**ynamics-**a**ware **T**oken **Me**rging **(DaTMe)**, a training-free framework that grounds token merging in the underlying diffusion dynamics. Our key insight is that token importance should reflect whether a token remains unresolved during denoising and whether it carries structural information. To this end, DaTMe constructs a dynamics-aware importance map from the Tweedie denoised estimate. The map combines a temporal signal, measuring inter-step changes in the predicted clean image, with a structural signal, capturing local high-frequency content such as boundaries and fine details. Using this map, DaTMe adaptively assigns tokens into merging roles under a given compression ratio: important tokens are preserved as anchors, while tokens that are both temporally stabilized and structurally homogeneous are selected as safe merge candidates. Experiments on PixArt-α and FLUX demonstrate that DaTMe maintains visual fidelity across the entire image, highlighting diffusion dynamics as an effective token-level merging decision in efficient generation. Our code is available at `https://github.com/SeungJuCha/DaTMe`.

1 Introduction

Diffusion models [1, 2] have become a dominant paradigm for high-quality image [3–5] and video synthesis [6–8], owing to their strong fidelity and scalability. However, their iterative denoising process incurs substantial computational cost, as a large number of tokens are repeatedly processed at every inference step [9]. To address this, token merging [10] has emerged as an effective training-free acceleration strategy that groups similar tokens, merges them into representative tokens, and uses the merged representations in subsequent layers [9–13]. For example, early token merging methods such as ToMeSD [10] first construct local merge candidate groups in a content-agnostic manner, and then form merge pairs simply based on feature similarity within each group. Consequently, content-relevant tokens can be selected as merge sources, causing local details to be absorbed and leading to visual distortion [9]. This raises a central challenge: *which tokens should be absorbed through merging*, and *which must be retained* as anchors for subsequent computation.

Prior work [9] leverages classifier-free guidance (CFG) [14] as a token importance signal, identifying tokens strongly steered by the text condition through the discrepancy between conditional and

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

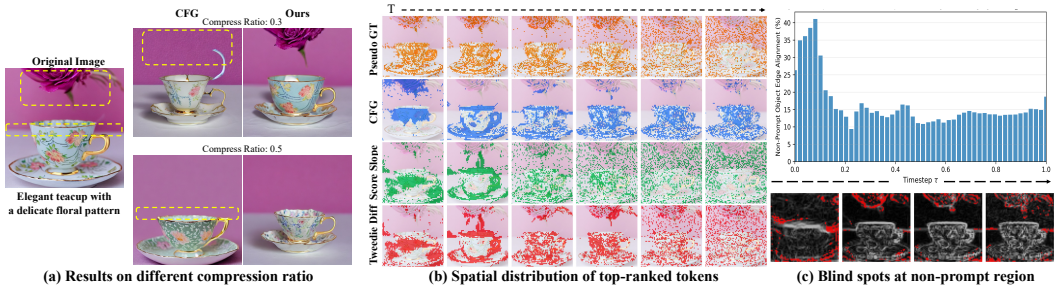

Figure 1: **(a) Comparison under different compression ratios.** CFG-based merging [9] often overlooks background details and structures, whereas our method better preserves background integrity. **(b) Spatial distribution of top-ranked tokens over denoising timesteps.** CFG-based importance remains concentrated on the foreground object with limited temporal variation. In contrast, our score-based signals (Score Slope and Tweedie Diff) progressively shift from object regions to fine-grained spatial details, better matching the pseudo-GT pattern. Here, pseudo-GT is defined as the difference between the final denoised latent and the Tweedie estimate at each timestep. **(c) Non-prompt region blind spot.** Low-ranked tokens under the CFG-based signal consistently overlap with structurally salient regions in text-irrelevant areas (red).

unconditional noise predictions as anchors to be retained. However, as shown in Fig. 1(a), CFG-based merging tends to underprotect non-prompt regions of the image, losing fine details and structural coherence. We attribute this to CFG's prompt-centric bias: by focusing exclusively on text-conditioned influence, CFG overlooks two key aspects of how diffusion models generate images.

First, CFG fails to capture **time-varying token importance**. Diffusion generation follows a coarse-to-fine progression [15, 16], where tokens undergoing large inter-step changes represent regions in which the model is actively injecting details [17]. In this context, prematurely merging such still-evolving tokens directly attenuates their ongoing refinement by averaging their features with neighboring tokens. Crucially, the spatial distribution of these actively evolving tokens is time-varying: while it remains concentrated on foreground objects in early steps, it progressively expands toward broader spatial regions, including non-prompt areas, at later timesteps (Fig. 1(b)). High-ranked tokens in CFG importance signal, however, remain anchored to foreground objects throughout the entire timestep, confining its search for important tokens only to prompt-aligned regions. As a result, CFG fails to identify the very tokens that should be protected: actively evolving tokens outside the foreground are prematurely merged into nearby or similar tokens, suppressing their distinct updates.

Second, CFG is insensitive to **prompt-independent structural importance**. To produce natural and coherent images, diffusion models actively generate scene elements that are not directly mentioned in the prompt such as background objects and supporting structures [1, 2]. Although these prompt-independent structures contribute little to text guidance, they carry essential structural information such as boundaries between homogeneous regions, edges of background objects, and fine textures that anchor visual coherence. However, because CFG only measures text-conditioned discrepancy, structural boundaries in such non-prompt regions tend to receive low scores across all timesteps (Fig. 1(c)), making them vulnerable to merging into neighboring homogeneous areas. Taken together, these observations indicate that token importance should not be defined solely by text relevance, but by diffusion dynamics: *how unresolved a token remains* and *how much structural information it carries* at the current denoising step [18].

To faithfully capture diffusion dynamics, we define two complementary signals that quantify the spatiotemporal dynamics of each token during denoising. For temporal dynamics, we start from the score function [19], which quantifies how each token is updated at the current denoising step. In particular, we project this score into clean-image space via the Tweedie denoised estimate [20, 21], and use its inter-step difference as the *temporal signal*, assigning high values to tokens that are unresolved. Unlike CFG, this score-derived signal reflects the actual update applied to every token, regardless of prompt relevance, providing an image-wide view of generative activity (Fig. 1(b)). For the *structural signal*, we measure the residual between the Tweedie estimate and its coarse-scale smoothed version, which acts as a high-pass filter highlighting structural components throughout the image. This captures prompt-independent structural boundaries, including weakly conditioned

structures that CFG often underestimates. By jointly modeling temporal evolution and structural saliency, we obtain a dynamics-aware token importance signal that guides merging decisions.

Building on this dynamics-aware importance, we propose **D**ynamics-**a**ware **T**oken **Me**rging **(DaTMe)**, a training-free token merging framework for diffusion transformers [2, 5, 22]. At each denoising step, DaTMe assigns token roles at the intersection of the temporal and structural signals. High-scoring tokens are assigned as constraint tokens and excluded from merging as dynamically or structurally important regions, while low-scoring tokens are assigned as source tokens from relatively stable and homogeneous regions that can be safely merged. To meet the target compression budget, DaTMe supplements the source set with grid-based spatial sampling, preventing local collapse. The resulting source tokens are then merged only within nearby regions via locality-constrained matching, built upon [9]. We primarily evaluate DaTMe on PixArt-α and FLUX, where it outperforms prior methods while maintaining visual fidelity across the entire image. Our contributions are as follows:

- We introduce a dynamics-aware importance map derived from the Tweedie estimate, jointly capturing the temporal evolution and structural saliency of each token.
- We propose **DaTMe**, a training-free token merging framework that leverages the proposed temporal and spatial importance signals to identify safely mergeable tokens while preserving structurally critical regions.
- We validate DaTMe on DiT (PixArt-α) and MM-DiT (FLUX), showing strong preservation of image-wide details and state-of-the-art performance across most metrics.

2 Related Works

Token Reduction in Large Models. Token reduction is a widely used strategy for improving Transformer efficiency, mainly through token pruning or token merging [10, 23–26]. Representative methods such as ToMe [27] use training-free bipartite matching, while later approaches refine token selection using attention or frequency-domain cues [25, 26]. However, most of these methods target single-forward-pass architectures with relatively static token importance. Diffusion models, by contrast, update token representations over many steps, motivating diffusion-specific approaches such as ToMeSD [10].

Token Merging in Diffusion Models. A wide range of methods have been proposed to accelerate diffusion models, including improved samplers [28–32] and weight quantization [33, 34]. However, such approaches often require additional training or optimization, introducing non-trivial overhead. Training-free methods instead reduce the number of tokens processed during inference through pruning [18, 35], caching [36–38], or merging [9–13]. Among them, ToMeSD [10] extends token merging in ViTs [27] to diffusion models by grouping tokens based on feature similarity. Several subsequent methods [12, 13] leverage cross-attention maps to identify important tokens from semantically or structurally salient regions, aiming not merely to reduce computation but also to preserve informative representations during token reduction [39]. More recently, [9] uses classifier-free guidance (CFG) as a saliency signal without additional attention computation, but because CFG primarily reflects alignment with the conditioning signal, it fails to capture the spatially heterogeneous and temporally evolving dynamics of the full image. In contrast, our method leverages the Tweedie denoised estimate to construct a dynamics-aware importance map that captures both temporal evolution and spatial structure across the image.

3 Method

Our goal is to design a token merging strategy for efficient image generation by accounting for diffusion dynamics. To this end, we first define a **dynamics-aware importance signal** that characterizes token importance from two complementary perspectives: **temporal dynamics** and **structural saliency**. Based on this importance map, we assign tokens into functional roles for merging, distinguishing safe merge sources and structurally important constraint tokens. We then perform locality-constrained token merging, together with grid-based source completion to satisfy the target compression ratio. In the following, we first motivate our Tweedie-based spatiotemporal importance signals (Sec. 3.2), and then describe how they are used in our proposed DaTMe framework (Sec. 3.3).

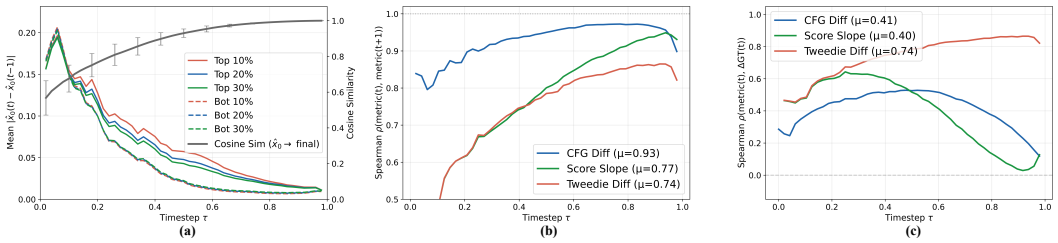

Figure 2: **(a) Cosine similarity and temporal trajectories of top- and bottom-ranked tokens.** The Tweedie estimate progressively approaches the final latent over denoising. Meanwhile, tokens with high pseudo-GT magnitude $|x_0 - \hat{x}_0(t)|$ exhibit consistently larger changes than low-ranked tokens. This indicates that the pseudo-GT effectively separates active tokens from converged ones. **(b) Step-wise autocorrelation of token-importance signals.** Computed between consecutive timesteps, lower values indicate more dynamic importance. The proposed Tweedie-based signal exhibits lower autocorrelation than CFG and score slope. **(c) Correlation with the pseudo-GT change magnitude.** The proposed Tweedie difference remains better aligned with actual token changes across timesteps than CFG or score slope. τ denotes the normalized timestep, scaled to the range $[0, 1]$.

3.1 Preliminaries

Score Function. In diffusion models [1–3], the forward process gradually corrupts a clean image x_0 by adding Gaussian noise over $t \in [0, T]$ timesteps:

$$x_t = \alpha_t x_0 + \sigma_t \epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \tag{1}$$

where α_t and σ_t denote the signal and noise scales at timestep t, respectively, satisfying $\alpha_t^2 + \sigma_t^2 = 1$. The reverse process aims to recover x_0 from x_T by iteratively denoising, guided by the score function $\nabla_{x_t} \log p_t(x_t)$, which represents the gradient of the log-density of the noisy distribution $p_t(x_t)$ with respect to x_t [19]. Intuitively, the score function defines a vector field that points toward regions of higher data likelihood, guiding the noisy sample toward the clean data manifold throughout the denoising process. In practice, the score function is estimated via a neural network ϵ_θ trained to predict the noise ϵ added at each timestep. The score is then obtained from the predicted noise as:

$$s(t) = \nabla_{x_t} \log p_t(x_t) \approx -\frac{1}{\sigma_t} \epsilon_\theta(x_t, t, c), \tag{2}$$

where c denotes the conditioning signal (e.g., text prompt).

Tweedie's Estimation. Tweedie's formula [21] provides a closed-form expression for the posterior mean of the clean sample x_0 given the noisy observation x_t:

$$\hat{x}_0(t) = \mathbb{E}[x_0 \mid x_t] = \frac{x_t + \sigma_t^2 \nabla_{x_t} \log p_t(x_t)}{\alpha_t}. \tag{3}$$

Substituting the learned score estimator from Eq. 2, this becomes:

$$\hat{x}_0(t) \approx \frac{x_t - \sigma_t \, \epsilon_\theta(x_t, t, c)}{\alpha_t}, \tag{4}$$

which gives a denoised estimate of x_0 at timestep t, termed the *Tweedie denoised estimate*. Unlike the noisy sample x_t, which is directly modulated by the noise level σ_t, $\hat{x}_0(t)$ provides an estimate in the predicted clean image space [40], offering a noise-normalized view of the current generation state.

3.2 Spatiotemporal Tweedie Importance Signal

Theoretical Comparison of Temporal Signals. To characterize token-level evolution during denoising, we compare two temporal signals derived from the score function [19]: the raw score slope, $|s(t) - s(t-1)|$, and the Tweedie difference, $|\hat{x}_0(t) - \hat{x}_0(t-1)|$. As derived in Appendix A, the score slope can be decomposed into two components:

$$s(t) - s(t-1) = \underbrace{\frac{\epsilon_\theta(t-1) - \epsilon_\theta(t)}{\sigma_{t-1}}}_{\text{difference term}} + \underbrace{\epsilon_\theta(t) \left(\frac{1}{\sigma_{t-1}} - \frac{1}{\sigma_t} \right)}_{\text{scale term}}. \tag{5}$$

The first term reflects changes in the model prediction across timesteps and is therefore associated with actual token-level evolution. In contrast, the second term arises purely from the predefined global noise schedule σ_t rather than changes in the model prediction itself. As denoising approaches the late steps ($\sigma_t \to 0$), the inverse difference $\left|\frac{1}{\sigma_{t-1}} - \frac{1}{\sigma_t}\right|$ grows rapidly. This suggests that the scale term can increasingly dominate the score slope in later stages, making the resulting signal less reliable. To obtain a cleaner temporal measure, we instead project the score into predicted clean-image space via the Tweedie estimate $\hat{x}_0(t)$ [41] and consider its inter-step residual, $\hat{x}_0(t) - \hat{x}_0(t-1)$. By algebraic decomposition (Appendix A), this residual reduces to:

$$\hat{x}_0(t) - \hat{x}_0(t-1) = \frac{\sigma_{t-1}}{\alpha_{t-1}}\big(\epsilon_\theta(t-1) - \epsilon_\theta(t)\big), \tag{6}$$

which removes the scale term that appears in the raw score slope and retains only the difference term update. Furthermore, the update is scaled by $\frac{\sigma_{t-1}}{\alpha_{t-1}}$, and this factor decreases as denoising concludes, naturally attenuating the signal in late steps. These properties make the Tweedie difference a more stable temporal signal for measuring token evolution.

Empirical Validation of Temporal Signals. We validate the theoretical analysis above through two empirical tests: temporal autocorrelation and correlation with a pseudo-GT change measure. We first examine temporal autocorrelation to assess whether each signal remains responsive to token-wise refinement over denoising steps (Fig. 2(b)). The CFG-based importance [9] exhibits consistently high autocorrelation across all timesteps, confirming its static nature and its bias toward fixed, prompt-aligned foreground regions, as previously observed in Fig. 1(b). In contrast, both score-based signals, namely the raw score slope and the Tweedie difference, initially show noticeably lower autocorrelation, indicating greater sensitivity to image-wide refinement than CFG. However, temporal autocorrelation alone does not naturally imply that a signal reflects token-level convergence.

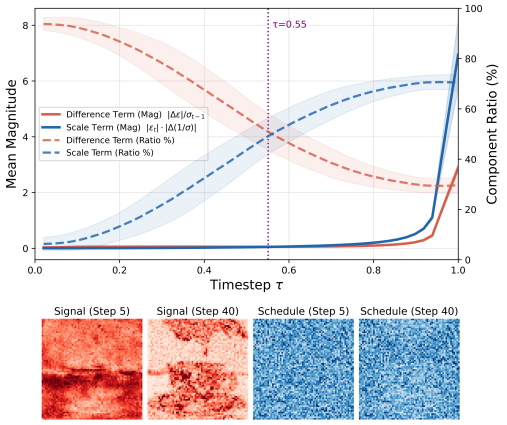

Figure 3: **Difference- and Scale- components of the score slope.** (Top) Mean magnitude and contribution ratio of the two terms: the scale term becomes dominant in late steps. (Bottom) Spatial distributions: the difference term captures structured token-level changes, whereas the scale term remains spatially unstructured.

To evaluate whether a signal faithfully reflects token-level convergence, we measure its correlation with a reference signal constructed from the per-timestep Tweedie estimate $\hat{x}_0(t)$, which we treat as a proxy for the clean image following [41]. We then define the deviation between the final clean latent x_0 and $\hat{x}_0(t)$, denoted as $|x_0 - \hat{x}_0(t)|$, as a pseudo-GT measure of how much each token is still expected to change. As shown in Fig. 2(a), the cosine similarity between $\hat{x}_0(t)$ and the final denoised latent steadily approaches 1.0 over time, indicating that $\hat{x}_0(t)$ progressively approximates the final generation state. Moreover, when tokens are partitioned by their pseudo-GT magnitude, high-ranked tokens consistently exhibit larger changes across denoising steps than low-ranked ones, suggesting that this pseudo-GT provides a meaningful per-token reference for remaining change.

Using this pseudo-GT, a clear difference between the two score-based signals emerges (Fig. 2(c)). While both signals show high correlation with the pseudo-GT in the early stages, the correlation of the raw score slope drops sharply around normalized timestep 0.4–0.6. Consistent with our theoretical decomposition, this degradation closely matches the range where the scale term empirically overtakes the difference term (Fig. 3). At the same time, the autocorrelation of the score slope rises (Fig. 2(b)), indicating increasingly static behavior and reduced sensitivity to token-wise refinement. In contrast, by removing scale term interference, the Tweedie difference maintains both a higher correlation with the pseudo-GT and a consistently lower temporal autocorrelation throughout denoising. Taken together, these results suggest that the Tweedie difference provides a more reliable temporal signal for tracking token evolution than either the raw score slope or CFG.

Tweedie Importance Map. Building on this analysis, we define the temporal importance map as the inter-step residual in predicted clean-image space:

$$\mathcal{M}_{temp}^{(t)} = \left|\hat{x}_0(t) - \hat{x}_0(t-1)\right|, \tag{7}$$

which quantifies how much the predicted clean structure is still evolving at each spatial location. This temporal map captures when each token is actively changing during denoising.

However, as previously observed in Fig. 1(c), CFG also fails to protect structurally salient tokens in non-prompt regions—those carrying essential boundary and texture information. Under CFG-based importance, such tokens consistently receive low scores throughout denoising, leading them to be mistakenly treated as safe merge candidates and causing object boundaries to be diluted (Fig. 8).

To capture this structural property, we introduce a spatial axis derived from the same Tweedie estimate. Specifically, we use the spatial axis to distinguish structurally salient but temporally stable tokens from genuinely mergeable homogeneous regions. To this end, we define the spatial axis as the residual between $\hat{x}_0(t)$ and its coarse-scale approximation obtained via interpolation-based transformation [42]:

$$\mathcal{M}_{spat}^{(t)} = \left|\hat{x}_0(t) - \mathrm{UpSample}_{\times s}\Big(\mathrm{DownSample}_{\times s}\big(\hat{x}_0(t)\big)\Big)\right|, \tag{8}$$

where s denotes the spatial scale used to construct the coarse approximation. This spatial map assigns high responses to contours, boundaries, and fine textures, while giving low responses to locally homogeneous regions. By combining the temporal and spatial axes, we obtain an importance signal that reflects both how tokens evolve over time and how much structural information they carry at each denoising step.

3.3 Dynamics-aware Token Merging (DaTMe)

At each timestep t, given a target compression ratio r, we denote the full token set as $\mathcal{S}_{tot}$ with $N = |\mathcal{S}_{tot}|$, and aim to merge $n_{tgt} = \lfloor rN \rfloor$ tokens. While feature-similarity-based merging strategies work well in U-Net-based architectures [1, 3], directly applying them to Diffusion Transformers (DiTs) [2, 4] often causes severe structural distortion [9, 12]. To prevent such distortion, we propose Dynamics-aware Token Merging (DaTMe), which performs three steps: (1) assigning token roles based on the proposed spatiotemporal importance maps, (2) completing the source budget in a spatially distributed manner, and (3) performing locality-constrained spatial matching. A detailed step-by-step algorithm of DaTMe is provided in the Appendix B.

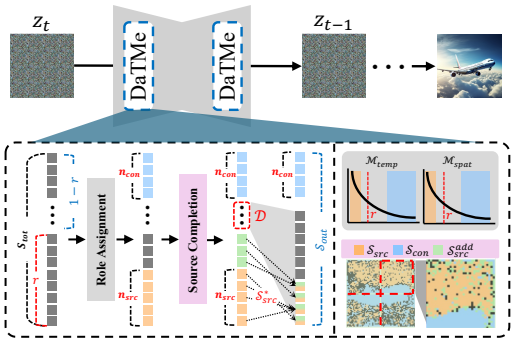

Figure 4: **Proposed DaTMe pipeline.** Using the spatiotemporal axes, DaTMe assigns token roles and performs merging within diffusion layers.

Dynamics-Aware Role Assignment. Rather than imposing a hand-crafted cutoff for what should be considered important, DaTMe adaptively assigns token roles at each denoising step based on their temporal and spatial responses under the given compression ratio r. To prevent quality degradation caused by merging temporally evolving or structurally salient tokens, we first define a conservative source set $\mathcal{S}_{src} \subset \mathcal{S}_{tot}$ as the intersection of the lowest-r fractions of tokens ranked by $\mathcal{M}_{temp}$ and $\mathcal{M}_{spat}$. Conversely, we define a constraint set, $\mathcal{S}_{con} \subset \mathcal{S}_{tot}$, as the intersection of the top-$(1-r)$ fractions of tokens on both axes, preserving tokens that are both temporally unstable and structurally salient. By preserving $\mathcal{S}_{con}$ throughout merging, we use them as structural anchors that keep critical boundaries from dissolving into surrounding regions.

Spatially Distributed Source Completion. Because $\mathcal{S}_{src}$ is defined by the strict intersection of two conditions, its size $n_{src} = |\mathcal{S}_{src}|$ is typically smaller than the target budget n_{tgt} (i.e., $n_{src} < n_{tgt}$). We therefore select an additional set $\mathcal{S}_{src}^{add}$ from the remaining unassigned tokens to satisfy $|\mathcal{S}_{src} \cup \mathcal{S}_{src}^{add}| = n_{tgt}$. To avoid concentrating these additional source tokens in a specific region—which would lead to local information loss after merging—we adopt a grid-based spatial sampling strategy. Specifically, we reshape the token sequence into a spatial feature map, partition



Figure 5: **Qualitative comparison under compression ratio 0.5 on PixArt-α.** DaTMe better preserves structural details and background content.



Figure 6: **Qualitative comparison under compression ratio 0.5 on FLUX.** DaTMe better preserves structural details with lower quality degradation.

it into P non-overlapping grid cells, and uniformly sample $(n_{tgt} - n_{src})/P$ tokens from each cell. The final source set is then given by $\mathcal{S}^*_{src} = \mathcal{S}_{src} \cup \mathcal{S}^{add}_{src}$.

Locality-Constrained Routing. The remaining tokens, those not assigned to $\mathcal{S}^*_{src}$ or $\mathcal{S}_{con}$, form the destination pool $\mathcal{D} \subset \mathcal{S}_{tot}$, into which $\mathcal{S}^*_{src}$ is merged via bipartite soft matching. Because DiTs process tokens as a sequence without inherent spatial inductive bias [43], relying purely on feature similarity can collapse spatially distant tokens and distort local spatial structure. We therefore use spatial adjacency as the matching criterion [9], pairing each source token with its spatially nearest destination and obtaining $\mathcal{S}_{merged} = \mathrm{Merge}(\mathcal{D}, \mathcal{S}^*_{src})$. The final output token set is $\mathcal{S}_{out} = \mathcal{S}_{merged} \cup \mathcal{S}_{con}$, with $|\mathcal{S}_{out}| = N - n_{tgt}$. By excluding $\mathcal{S}_{con}$ from both pools and restricting merging to local neighborhoods, DaTMe compensates for the weak spatial inductive bias of DiTs and preserves visual fidelity throughout merging.

4 Experiments

Implementation Details and Evaluation Metrics. We primarily conduct experiments on PixArt-α [5] using 10k image-prompt pairs sampled from the COCO 2014 validation set [44]. We set the interpolation scale to $s = 4$, and for grid-based source completion, we use $P = 4$ grid cells. We evaluate three compression ratios: 0.3, 0.5, and 0.7. Following [9], we apply DaTMe to the middle-range self-attention layers of the DiT [2] backbone. To further evaluate scalability on a larger DiT model, we apply DaTMe to FLUX [22]. We use FLUX.1-dev and additionally apply DaTMe from the middle block to the final block within the joint-attention blocks using the same hyperparameters. More details are in the supplementary materials. For evaluation, we follow prior work [9, 10, 12, 13] and report CLIP-T and FID. In addition, we report HPSv2 [45], which complements CLIP-T by measuring preference-oriented prompt faithfulness and perceptual quality. We also report SSIM, PSNR, and LPIPS to measure image-level fidelity and detail preservation relative to the original output. We further report the Aesthetic score using Aesthetic Predictor V2.5, a SigLIP-based reference-free image aesthetic predictor. All experiments are conducted on a single RTX PRO 6000 Blackwell.

Table 1: **Text alignment and efficiency scores on PixArt-α.** DaTMe achieves the best overall quality-efficiency trade-off, preserving strong text alignment with minimal degradation under token compression.

	CLIP-T↑	HPSv2↑	FID↓	Latency(s)↓	TFLOPs↓
PixArt-α (Base) [5]	31.2485	29.2780	42.1001	1.9417	183.45
Compression ratio 0.3					
ToMeSD [10]	30.5735	24.1699	61.5229	1.6917	172.514
CFG [9]	<u>30.8801</u>	<u>27.7375</u>	<u>44.6070</u>	1.7458	172.514
DaTMe (Ours)	**31.1375**	**28.5591**	**42.4474**	1.8030	175.631
Compression ratio 0.5					
ToMeSD	29.5178	21.3334	94.8901	1.5918	165.201
CFG	<u>30.4645</u>	<u>25.8345</u>	<u>52.2912</u>	1.6445	165.201
DaTMe (Ours)	**31.0142**	**27.5174**	**43.4155**	1.6884	170.410
Compression ratio 0.7					
ToMeSD	28.2805	18.9797	123.0139	1.5356	157.897
CFG	<u>29.4543</u>	<u>22.6156</u>	<u>70.7062</u>	1.5761	157.897
DaTMe (Ours)	**30.6259**	**25.2393**	**47.0519**	1.6426	165.193

Table 2: **Image fidelity score on PixArt-α.** DaTMe achieves stronger structural preservation and image-level fidelity under token compression.

	SSIM↑	PSNR↑	LPIPS↓	Aes.↑
Compression ratio 0.3				
ToMeSD [10]	<u>0.7306</u>	<u>18.9748</u>	0.2433	5.9906
CFG [9]	0.7137	17.8157	<u>0.2369</u>	<u>6.1521</u>
DaTMe (Ours)	**0.8112**	**20.5136**	**0.1501**	**6.2014**
Compression ratio 0.5				
ToMeSD	<u>0.6737</u>	<u>17.7161</u>	<u>0.3151</u>	5.7739
CFG	0.6331	16.2768	0.3254	<u>6.1135</u>
DaTMe (Ours)	**0.7300**	**18.4277**	**0.2258**	**6.1695**
Compression ratio 0.7				
ToMeSD	<u>0.6301</u>	<u>16.8537</u>	<u>0.3682</u>	5.5809
CFG	0.5695	15.2003	0.4129	<u>5.9559</u>
DaTMe (Ours)	**0.6581**	**16.9048**	**0.3078**	**6.0544**

Table 3: **Text alignment and efficiency score on FLUX.** DaTMe achieves stronger text alignment with comparable efficiency under token compression. † denotes our reproduced results.

	CLIP-T↑	HPSv2↑	FID↓	Latency(s)↓	TFLOPs↓
FLUX.1-dev (Base) [22]	31.2274	29.6927	30.6141	14.601	2098
Compression ratio 0.3					
ToMeSD† [10]	<u>31.3254</u>	<u>28.9339</u>	**30.4576**	13.866	2046
ToMA [12]	31.1787	28.5153	<u>30.6520</u>	13.790	1971
DaTMe (Ours)	**31.3524**	**29.5042**	30.9258	13.772	2033
Compression ratio 0.5					
ToMeSD†	31.1852	<u>28.3595</u>	**30.8491**	13.278	2000
ToMA	<u>31.2036</u>	27.7566	<u>30.9883</u>	13.288	1895
DaTMe (Ours)	**31.2621**	**28.9266**	31.3423	13.271	1990
Compression ratio 0.7					
ToMeSD†	31.2045	<u>27.6455</u>	**31.0033**	12.728	1954
ToMA	<u>31.2191</u>	26.4796	35.0541	12.775	1793
DaTMe (Ours)	**31.2415**	**27.9093**	<u>31.1585</u>	12.737	1946

Table 4: **Image fidelity score on FLUX.** DaTMe achieves stronger structural preservation and image fidelity.

	SSIM↑	PSNR↑	LPIPS↓	Aes.↑
Compression ratio 0.3				
ToMeSD† [10]	0.6402	17.5070	0.2616	5.7626
ToMA [12]	<u>0.7415</u>	<u>20.6450</u>	<u>0.1844</u>	<u>6.0210</u>
DaTMe (Ours)	**0.7737**	**21.6352**	**0.1694**	**6.0318**
Compression ratio 0.5				
ToMeSD†	0.6166	16.9059	0.2932	5.6665
ToMA	<u>0.7009</u>	<u>20.0375</u>	<u>0.2305</u>	<u>5.9030</u>
DaTMe (Ours)	**0.7087**	**20.2021**	**0.2176**	**6.0191**
Compression ratio 0.7				
ToMeSD†	0.6032	16.6127	0.3181	5.5739
ToMA	<u>0.6280</u>	<u>17.6972</u>	**0.2835**	<u>5.6971</u>
DaTMe (Ours)	**0.6378**	**17.7655**	<u>0.2877</u>	**5.9147**

4.1 Main Experiments

Qualitative Results. In Fig. 5 and Fig. 6, we compare qualitative results at a compression ratio of 0.5. On PixArt-α, directly merging tokens based on feature similarity (ToMeSD) severely disrupts spatial structure, often leading to obvious structural distortion. Although the CFG-based method preserves foreground structures that are explicitly described in the prompt, it shows noticeably weaker preservation in background and other weakly conditioned regions, as highlighted by the red boxes. This qualitatively supports our analysis that CFG-based importance is strongly biased toward prompt-conditioned content and is therefore less effective at preserving the full image context.

On FLUX, existing methods generally preserve the coarse prompt semantics, but still introduce noticeable deviations from the original generation. ToMeSD tends to alter the original composition, object geometry, or viewpoint, as seen in the relative arrangement of the cats, the fire hydrant structure, and the airplane viewpoint. ToMA shows similar structural drift while additionally introducing noisy visual degradation that weakens overall image quality. In contrast, DaTMe more faithfully preserves the original spatial configuration and object structure while maintaining the intended prompt semantics. These results demonstrate that DaTMe reduces both structural deviation and visual degradation across different architectures. Additional results under various compression ratios are in Appendix C.

Quantitative Results. In Tab. 1 and Tab. 2, we compare DaTMe with prior methods on PixArt-α [5] under various compression ratios. Although DaTMe is slightly slower than prior methods in terms of latency, the prior methods suffer from substantial quality degradation as the compression ratio

Table 5: **Ablation on temporal and spatial axes.** We compare using only the temporal axis, only the spatial axis, both axes, and both axes with score slope instead of Tweedie difference.

Setting	CLIP-T↑	SSIM↑	PSNR↑	FID↓
CFG [9]	30.8801	0.7137	17.8157	44.6070
Only $\mathcal{M}_{\text{temp}}$	31.0355	0.8109	20.5021	<u>42.5269</u>
Only $\mathcal{M}_{\text{spat}}$	<u>31.1318</u>	<u>0.8110</u>	<u>20.5029</u>	42.5736
Ours (score slope)	31.1166	0.8028	20.3944	42.9554
Ours	**31.1375**	**0.8112**	**20.5136**	**42.4474**

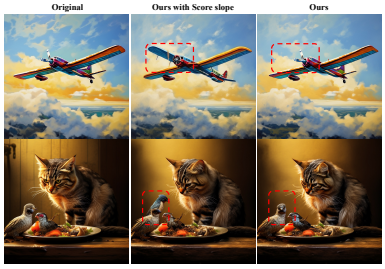

Figure 7: **Qualitative comparison difference between score slope and Tweedie difference.** Using the score slope introduces noticeable artifacts.

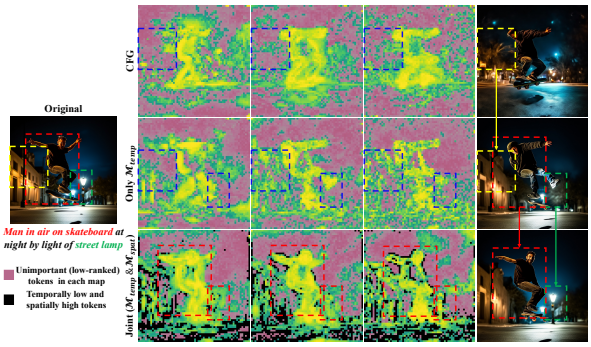

Figure 8: **Effect of different importance signals.** Jointly using temporal and spatial signals, the proposed importance map preserves both structurally important object regions and unconditioned background content.

Table 6: **Ablation on constraint-set definition $\mathcal{S}_{\text{con}}$.** We compare different choices for selecting constraint tokens.

Setting	CLIP-T↑	SSIM↑	PSNR↑	LPIPS↓	FID↓
None	**31.1419**	0.8037	20.3736	0.1601	43.7692
$\mathcal{M}_{\text{temp}}\uparrow \cap \mathcal{M}_{\text{spat}}\uparrow$	<u>31.1375</u>	**0.8112**	**20.5136**	**0.1501**	**42.4474**
$\mathcal{M}_{\text{temp}}\downarrow \cap \mathcal{M}_{\text{spat}}\uparrow$	31.1372	<u>0.8076</u>	<u>20.5027</u>	<u>0.1569</u>	<u>43.6605</u>
$\mathcal{M}_{\text{temp}}\uparrow \cap \mathcal{M}_{\text{spat}}\downarrow$	31.1250	0.8005	20.3459	0.1641	43.8062

increases, particularly in FID. In contrast, DaTMe remains much more robust under compression, consistently preserving image quality and achieving the most favorable overall quality-efficiency trade-off. Notably, DaTMe also attains substantially stronger structural preservation, which indicates that the proposed spatiotemporal importance signal effectively prevents structurally important tokens from being merged prematurely (Tab. 2).

In the FLUX results, CFG-based token selection is not applicable because FLUX does not employ CFG, whereas DaTMe remains directly applicable despite the architectural difference. Across all compression ratios, DaTMe consistently achieves the strongest CLIP-T and HPSv2 scores while maintaining comparable inference efficiency (Tab. 3). Although ToMeSD preserves reasonable text–image alignment, its substantially lower SSIM and PSNR and higher LPIPS indicate larger structural deviations from the original generation. ToMA better preserves the original structure than ToMeSD, but still exhibits noticeable visual degradation, particularly in fine details and overall image quality. In contrast, DaTMe achieves the strongest overall image fidelity, consistently improving SSIM, PSNR, and aesthetic quality while maintaining strong prompt alignment (Tab. 4). These quantitative trends are consistent with Fig. 6, where DaTMe more faithfully preserves the original image structure.

4.2 Additional Analysis

We conduct various ablation studies on PixArt-α. Additional experiments are in Appendix C.

Importance of Temporal and Spatial Signals. In Tab. 5, we ablate the two components of our importance map by using only the temporal axis, only the spatial axis, or both. Using both axes yields the best overall performance, confirming that temporal evolution and structural saliency provide complementary cues for token selection, whereas either axis alone is insufficient to fully characterize mergeability. We further replace our temporal signal with the raw score slope to evaluate its effectiveness in the actual merging process. As shown in Tab. 5, this substitution consistently degrades all evaluation metrics, and Fig. 7 shows that it also introduces noticeable visual artifacts. This is consistent with our earlier analysis that the score slope becomes increasingly contaminated by schedule-driven effects in later denoising steps, making it less reliable for token merging.

Table 7: **Effectiveness of the proposed importance signal on PixArt-α at $r = 0.3$.** "Merging" indicates the merging strategy adopted from either [9] or ours, while "Signal" indicates whether token importance is determined by CFG [9] or our proposed importance map.

Merging	Signal	CLIP-T↑	HPSv2↑	SSIM↑	PSNR↑	LPIPS↓
CFG [9]	CFG [9]	30.8801	27.7375	0.7137	17.8157	0.2369
CFG [9]	Ours	<u>31.0545</u>	<u>27.9366</u>	<u>0.7644</u>	<u>20.0917</u>	<u>0.1915</u>
Ours	Ours	**31.1375**	**28.5591**	**0.8112**	**20.5136**	**0.1501**

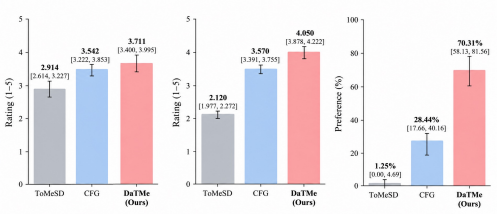

Figure 9: **Human preference study on PixArt-α at $r = 0.3$.** DaTMe receives the highest ratings for text alignment and visual quality, as well as the highest overall preference.

Fig. 8 further illustrates how different importance signals affect region preservation. The CFG-based signal [9] is strongly biased toward prompt-aligned foreground regions and therefore neglects unmentioned background content. Using only the temporal signal alleviates this issue by preserving evolving background regions (yellow box), but it still misses structurally important boundaries (black regions), leading to structural distortions (red box). By jointly modeling temporal and spatial signals, our importance map preserves both evolving background content and critical object boundaries, resulting in more faithful image structure and higher overall visual quality.

Constraint Token Selection. In Tab. 6, we ablate different definitions of the constraint set. The best results are obtained when the constraint tokens are selected from regions with both high temporal and high spatial responses. In particular, including spatially high-response tokens improves structural preservation, indicating that the spatial axis captures fine object details and structurally salient regions. By contrast, using spatially low-response tokens degrades performance because such tokens are biased toward homogeneous regions and therefore fail to protect meaningful object structure.

Effectiveness of the Importance Signal. To isolate the contribution of our importance signal, we keep the entire merging pipeline of [9] unchanged and replace only its CFG-based importance with ours. Since [9] requires a single token ranking, we combine our temporal and spatial rankings using rank-sum. As shown in Tab. 7, replacing only the CFG-based signal with ours improves CLIP-T, HPSv2, and all fidelity metrics, demonstrating that the proposed dynamics-aware importance more effectively identifies tokens to preserve or merge. Using the full DaTMe merging strategy further improves all metrics, indicating that the role-assignment and spatial merging design provides additional complementary benefits.

Preference Study. We evaluate perceptual quality at a compression ratio of $r = 0.3$ using human judgments. For 20 randomly sampled PixArt-α cases, 32 evaluators compare outputs from ToMeSD, CFG [9], and DaTMe without access to the uncompressed reference. Evaluators independently rate text alignment and visual quality, and select their overall preferred result. As shown in Fig. 9, DaTMe achieves the highest ratings for both text alignment and visual quality, while also being the most preferred method overall.

5 Conclusion

We presented **DaTMe**, a diffusion-dynamics-aware token merging framework for efficient image generation. By leveraging the Tweedie denoised estimate, DaTMe defines token importance along two complementary axes—temporal evolution and structural saliency—that reflect the image generation process. This allows DaTMe to identify safely mergeable tokens while preserving structurally critical regions, achieving strong image fidelity and favorable quality-efficiency trade-offs across DiT-based diffusion models. Our work shows that grounding token merging in the underlying generation dynamics is a key step toward faithful and efficient diffusion inference.

Acknowledgments and Disclosure of Funding

This work was partly supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No. RS-2025-

25422680, Metacognitive AGI Framework) and the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korean government (MSIT) (No.RS-2025-02215122, Development and Demonstration of Lightweight AI Model for Smart Homes).

References

[1] Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695 (2022).

[2] Peebles, W. & Xie, S. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205 (2023).

[3] Podell, D. *et al.* Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952* (2023).

[4] Esser, P. *et al.* Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning* (2024).

[5] Chen, J. *et al.* Pixart-α: Fast training of diffusion transformer for photorealistic text-to-image synthesis (2023). URL `https://arxiv.org/abs/2310.00426`. `2310.00426`.

[6] Li, J. *et al.* T2v-turbo-v2: Enhancing video generation model post-training through data, reward, and conditional guidance design. *arXiv preprint arXiv:2410.05677* (2024).

[7] Kong, W. *et al.* Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603* (2024).

[8] Polyak, A. *et al.* Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720* (2024).

[9] Wu, H., Xu, J., Le, H. & Samaras, D. Importance-based token merging for efficient image and video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 4983–4995 (2025).

[10] Bolya, D. & Hoffman, J. Token merging for fast stable diffusion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4599–4603 (2023).

[11] Zhang, E., Tang, J., Ning, X. & Zhang, L. Training-free and hardware-friendly acceleration for diffusion models via similarity-based token pruning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, 9878–9886 (2025).

[12] Lu, W., Zheng, S., Xia, Y. & Wang, S. Toma: Token merge with attention for diffusion models. *arXiv preprint arXiv:2509.10918* (2025).

[13] Fang, H. *et al.* Attend to not attended: Structure-then-detail token merging for post-training dit acceleration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 18083–18092 (2025).

[14] Ho, J. & Salimans, T. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022).

[15] Hertz, A. *et al.* Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626* (2022).

[16] Hwang, J., Park, Y.-H. & Jo, J. Resolution chromatography of diffusion models. *arXiv preprint arXiv:2401.10247* (2023).

[17] Cao, Y., Zhao, Z., Patras, I. & Gong, S. Temporal score analysis for understanding and correcting diffusion artifacts. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 7707–7716 (2025).

[18] Zhao, W. *et al.* Dynamic diffusion transformer. *arXiv preprint arXiv:2410.03456* (2024).

[19] Song, Y. *et al.* Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456* (2020).

[20] Robbins, H. E. An empirical bayes approach to statistics. In *Breakthroughs in Statistics: Foundations and basic theory*, 388–394 (Springer, 1992).

[21] Efron, B. Tweedie's formula and selection bias. *Journal of the American Statistical Association* **106**, 1602–1614 (2011).

[22] Labs, B. F. Flux. `https://github.com/black-forest-labs/flux` (2024).

[23] Rao, Y. *et al.* Dynamicvit: Efficient vision transformers with dynamic token sparsification. *Advances in neural information processing systems* **34**, 13937–13949 (2021).

[24] Chen, M. *et al.* Diffrate: Differentiable compression rate for efficient vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 17164–17174 (2023).

[25] Yang, L. *et al.* Libra-merging: Importance-redundancy and pruning-merging trade-off for acceleration plug-in in large vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9402–9412 (2025).

[26] Zhang, X. *et al.* Tf-atm: Training-free adaptive token merging. In *Proceedings of the 33rd ACM International Conference on Multimedia*, 7172–7180 (2025).

[27] Bolya, D. *et al.* Token merging: Your vit but faster. *arXiv preprint arXiv:2210.09461* (2022).

[28] Luo, S., Tan, Y., Huang, L., Li, J. & Zhao, H. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378* (2023).

[29] Salimans, T. & Ho, J. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512* (2022).

[30] Meng, C. *et al.* On distillation of guided diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14297–14306 (2023).

[31] Song, J., Meng, C. & Ermon, S. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020).

[32] Song, Y., Dhariwal, P., Chen, M. & Sutskever, I. Consistency models (2023).

[33] Li, X. *et al.* Q-diffusion: Quantizing diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17535–17545 (2023).

[34] Shang, Y., Yuan, Z., Xie, B., Wu, B. & Yan, Y. Post-training quantization on diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1972–1981 (2023).

[35] Kim, M., Gao, S., Hsu, Y.-C., Shen, Y. & Jin, H. Token fusion: Bridging the gap between token pruning and token merging. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1383–1392 (2024).

[36] Ma, X., Fang, G. & Wang, X. Deepcache: Accelerating diffusion models for free. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 15762–15772 (2024).

[37] Wimbauer, F. *et al.* Cache me if you can: Accelerating diffusion models through block caching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6211–6220 (2024).

[38] Liu, H. *et al.* Faster diffusion via temporal attention decomposition. *arXiv preprint arXiv:2404.02747* (2024).

[39] Kong, Z. *et al.* Token reduction should go beyond efficiency in generative models–from vision, language to multimodality. *arXiv preprint arXiv:2505.18227* (2025).

[40] Rout, L. *et al.* Beyond first-order tweedie: Solving inverse problems using latent diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9472–9481 (2024).

[41] Rout, L. *et al.* Solving linear inverse problems provably via posterior sampling with latent diffusion models. *Advances in Neural Information Processing Systems* **36**, 49960–49990 (2023).

[42] Koh, S., Cha, S., Oh, H., Lee, K. & Kim, D.-J. Scalediff: Higher-resolution image synthesis via efficient and model-agnostic diffusion. *arXiv preprint arXiv:2510.25818* (2025).

[43] d'Ascoli, S. *et al.* Convit: Improving vision transformers with soft convolutional inductive biases. In *International conference on machine learning*, 2286–2296 (PMLR, 2021).

[44] Lin, T.-Y. *et al.* Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755 (Springer, 2014).

[45] Wu, X. *et al.* Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341* (2023).

[46] Ho, J., Jain, A. & Abbeel, P. Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020).

[47] Karras, T., Aittala, M., Aila, T. & Laine, S. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems* **35**, 26565–26577 (2022).

[48] Luo, C. Understanding diffusion models: A unified perspective (2022). 2208.11970.

A Formal Decomposition of Score Slope Instability

We provide a formal analysis of why the score slope fails to reliably capture per-token convergence in late denoising steps.

Decomposition of score slope. Recall that the score function is estimated as:

$$s_t \approx -\frac{1}{\sigma_t}\epsilon_\theta(x_t, t, c). \tag{9}$$

To understand the temporal variation of the score, we analyze the score slope between consecutive timesteps. For brevity, let $\epsilon_\theta(t)$ denote $\epsilon_\theta(x_t, t, c)$. The score slope is:

$$s_t - s_{t-1} = \frac{\epsilon_\theta(t-1)}{\sigma_{t-1}} - \frac{\epsilon_\theta(t)}{\sigma_t}. \tag{10}$$

By adding and subtracting the term $\frac{\epsilon_\theta(t)}{\sigma_{t-1}}$,

$$s_t - s_{t-1} = \frac{\epsilon_\theta(t-1)}{\sigma_{t-1}} - \frac{\epsilon_\theta(t)}{\sigma_{t-1}} + \frac{\epsilon_\theta(t)}{\sigma_{t-1}} - \frac{\epsilon_\theta(t)}{\sigma_t}, \tag{11}$$

we can decouple this slope into two distinct components:

$$s_t - s_{t-1} = \underbrace{\frac{\epsilon_\theta(t-1) - \epsilon_\theta(t)}{\sigma_{t-1}}}_{\text{difference term}} + \underbrace{\epsilon_\theta(t)\left(\frac{1}{\sigma_{t-1}} - \frac{1}{\sigma_t}\right)}_{\text{scale term}}. \tag{12}$$

This decomposition reveals a limitation. The first term captures the actual change in the model's noise prediction, reflecting semantic token-level evolution. However, the second term is entirely dictated by the predefined noise schedule σ_t, scaled by the current noise prediction $\epsilon_\theta(t)$.

Behavior in late denoising steps. As generation progresses toward the final steps, the noise scale σ_t becomes extremely small. Because σ_t decays rapidly, the inverse difference $\left|\frac{1}{\sigma_{t-1}} - \frac{1}{\sigma_t}\right|$ grows disproportionately large. Meanwhile, since ϵ_θ is trained to match samples from $\mathcal{N}(0, I)$ [46], its norm is concentrated near $\sqrt{D}$ in expectation, where D is the dimensionality [47]. Consequently, the scale term in Eq. (12) significantly amplifies the raw noise prediction $\epsilon_\theta(t)$, eventually overshadowing the difference term. Therefore, in later denoising stages, the score slope $s_t - s_{t-1}$ is heavily dominated by the noise schedule scaling rather than actual changes in each token, which is consistent with the high autocorrelation and loss of discriminative power observed in Fig. 2.

Decomposition of Tweedie difference. To understand why the Tweedie difference avoids this scale-induced amplification, we first define the Tweedie denoised estimate. Under variance-preserving diffusion (VP diffusion) [19, 46], the noisy state is $x_t = \alpha_t x_0 + \sigma_t \epsilon$, where the signal and noise schedules satisfy $\alpha_t^2 + \sigma_t^2 = 1$. The posterior mean estimate of the clean data is [21]:

$$\hat{x}_0(t) = \mathbb{E}[x_0 \mid x_t] \approx \frac{x_t - \sigma_t \epsilon_\theta(t)}{\alpha_t}, \tag{13}$$

where the approximation arises from using the learned $\epsilon_\theta(t)$ in place of the true noise. This clean-space projection is embedded within the deterministic update rule (e.g., DDIM [31]). The latent transition from timestep t to $t-1$ is:

$$x_{t-1} = \alpha_{t-1} \underbrace{\left(\frac{x_t - \sigma_t \epsilon_\theta(t)}{\alpha_t}\right)}_{\text{Tweedie estimate } \hat{x}_0(t)} + \sigma_{t-1}\epsilon_\theta(t). \tag{14}$$

Since the parenthesized term is exactly $\hat{x}_0(t)$, we can write:

$$x_{t-1} = \alpha_{t-1}\hat{x}_0(t) + \sigma_{t-1}\epsilon_\theta(t). \tag{15}$$

We decompose the inter-step difference by substituting this x_{t-1} into the Tweedie estimate at the next timestep, $\hat{x}_0(t-1) = \frac{x_{t-1} - \sigma_{t-1}\epsilon_\theta(t-1)}{\alpha_{t-1}}$:

$$\hat{x}_0(t-1) = \frac{\big(\alpha_{t-1}\hat{x}_0(t) + \sigma_{t-1}\epsilon_\theta(t)\big) - \sigma_{t-1}\epsilon_\theta(t-1)}{\alpha_{t-1}}. \tag{16}$$

Simplifying yields:

$$\hat{x}_0(t-1) = \hat{x}_0(t) + \frac{\sigma_{t-1}}{\alpha_{t-1}}\big(\epsilon_\theta(t) - \epsilon_\theta(t-1)\big). \tag{17}$$

Rearranging, the Tweedie inter-step difference reduces to:

$$\hat{x}_0(t) - \hat{x}_0(t-1) = \frac{\sigma_{t-1}}{\alpha_{t-1}}\big(\epsilon_\theta(t-1) - \epsilon_\theta(t)\big). \tag{18}$$

Comparing Eq. (18) with Eq. (12) reveals a difference in how the noise schedule enters each formulation. First, the problematic scale term—the $1/\sigma_t$ scaling that overshadowed the genuine signal—no longer dominates the slope. The Tweedie difference cleanly isolates the difference term change, $\epsilon_\theta(t-1) - \epsilon_\theta(t)$. Second, instead of amplifying this change with a rapidly growing $1/\sigma_{t-1}$ multiplier, the Tweedie difference scales it by $\frac{\sigma_{t-1}}{\alpha_{t-1}}$ (closely related to the inverse signal-to-noise ratio [48]). As $t \to 0$, we have $\sigma_{t-1} \to 0$ and $\alpha_{t-1} \to 1$, so $\frac{\sigma_{t-1}}{\alpha_{t-1}} \to 0$. This decays the variation as the network's prediction stabilizes, ensuring that the Tweedie difference provides a stable metric for convergence without scale-induced amplification.

Extension to flow matching. The Tweedie formulation extends naturally to flow matching architectures such as FLUX [22]. While VP diffusion satisfies $\alpha_t^2 + \sigma_t^2 = 1$, FLUX adopts the *rectified flow* schedule $\alpha_t = 1 - t,\ \sigma_t = t$, which does not satisfy the VP constraint. Both nonetheless share the general form $x_t = \alpha_t x_0 + \sigma_t \epsilon$, allowing the Tweedie derivation to proceed with FLUX-specific schedule values. Since rectified flow uses continuous time $t \in [0, 1]$, we denote the step size between consecutive sampling times by Δt.

Clean-image estimate under velocity parameterization. Unlike VP models that predict noise ϵ_θ, rectified flow models predict the velocity field $v_\theta(t) \approx \epsilon - x_0$ [22]. From the forward process $x_t = (1-t)x_0 + t\,\epsilon$ and the velocity definition $\epsilon = v_\theta(t) + x_0$, we have:

$$x_t = (1-t)x_0 + t\big(v_\theta(t) + x_0\big) = x_0 + t \cdot v_\theta(t). \tag{19}$$

Solving for x_0 yields the FLUX clean-image estimate:

$$\hat{x}_0(t) = x_t - t \cdot v_\theta(t), \tag{20}$$

which is the velocity-parameterization analog of Eq. (4).

Tweedie inter-step difference under Euler update. For rectified flow, the deterministic Euler step from time t to $t - \Delta t$ is:

$$x_{t-\Delta t} = x_t - \Delta t \cdot v_\theta(t). \tag{21}$$

Applying Eq. (20) at $t - \Delta t$ and substituting Eq. (21):

$$\begin{aligned}\hat{x}_0(t-\Delta t) &= x_{t-\Delta t} - (t-\Delta t)\cdot v_\theta(t-\Delta t) \\ &= \big[x_t - \Delta t \cdot v_\theta(t)\big] - (t-\Delta t)\cdot v_\theta(t-\Delta t).\end{aligned} \tag{22}$$

Subtracting from $\hat{x}_0(t) = x_t - t \cdot v_\theta(t)$:

$$\begin{aligned}\hat{x}_0(t) - \hat{x}_0(t-\Delta t) &= -t\cdot v_\theta(t) + \Delta t \cdot v_\theta(t) + (t-\Delta t)\cdot v_\theta(t-\Delta t) \\ &= -(t-\Delta t)\cdot v_\theta(t) + (t-\Delta t)\cdot v_\theta(t-\Delta t) \\ &= (t-\Delta t)\big(v_\theta(t-\Delta t) - v_\theta(t)\big).\end{aligned} \tag{23}$$

Eq. (23) mirrors Eq. (18): the Tweedie inter-step difference isolates the change in prediction—velocity rather than noise—scaled by the schedule factor $(t - \Delta t)$, which decays to zero as $t \to 0$. The Tweedie difference therefore remains a stable temporal signal under flow matching as well.

B DaTMe Details

DaTMe Process. In Alg. 1, we show the detailed step-by-step process of our proposed DaTMe framework.

Additional Implementation Details. For PixArt-α, we use 20 inference steps with a guidance scale of 4.5. Token merging is applied to transformer layers 9–19 starting from the second denoising step,

Algorithm 1 Dynamics-Aware Token Merging (DaTMe) at timestep t

Require: token features $\mathbf{F}^{(t)}$, temporal map $\mathcal{M}^{(t)}_{\text{temp}}$, spatial map $\mathcal{M}^{(t)}_{\text{spat}}$, compression ratio r, number of grids P

Ensure: merged token features $\mathbf{F}'^{(t)}$

$N \leftarrow H \times W, \quad \mathcal{S}_{\text{tot}} \leftarrow \{1, \dots, N\}, \quad n_{\text{tgt}} \leftarrow \lfloor rN \rfloor$

// 1. Dynamics-Aware Role Assignment

$\mathcal{S}_{\text{con}} \leftarrow \text{TopFrac}_{1-r}\big(\mathcal{M}^{(t)}_{\text{temp}}\big) \cap \text{TopFrac}_{1-r}\big(\mathcal{M}^{(t)}_{\text{spat}}\big)$ ▷ protected tokens

$\mathcal{S}_{\text{src}} \leftarrow \text{BotFrac}_{r}\big(\mathcal{M}^{(t)}_{\text{temp}}\big) \cap \text{BotFrac}_{r}\big(\mathcal{M}^{(t)}_{\text{spat}}\big)$ ▷ safe merge sources

// 2. Spatially Distributed Source Completion

$\mathcal{U} \leftarrow \mathcal{S}_{\text{tot}} \setminus \big(\mathcal{S}_{\text{con}} \cup \mathcal{S}_{\text{src}}\big)$

$\mathcal{S}^{\text{add}}_{\text{src}} \leftarrow \text{GridSample}\Big(\mathcal{U}, n_{\text{tgt}} - |\mathcal{S}_{\text{src}}|, P\Big)$

$\mathcal{S}^{*}_{\text{src}} \leftarrow \mathcal{S}_{\text{src}} \cup \mathcal{S}^{\text{add}}_{\text{src}}$

// 3. Locality-Constrained Spatial Matching

$\mathcal{D} \leftarrow \mathcal{S}_{\text{tot}} \setminus \big(\mathcal{S}_{\text{con}} \cup \mathcal{S}^{*}_{\text{src}}\big)$ ▷ destination tokens

for each $j \in \mathcal{S}^{*}_{\text{src}}$ **do**

  $d(j) \leftarrow \arg\min_{k \in \mathcal{D}} \|\mathbf{p}_j - \mathbf{p}_k\|_2$

end for

$\mathbf{F}'^{(t)} \leftarrow \text{Merge}\Big(\mathbf{F}^{(t)}, \{j \rightarrow d(j)\}_{j \in \mathcal{S}^{*}_{\text{src}}}\Big)$ ▷ $\mathcal{S}_{\text{con}}$ remains unchanged

return $\mathbf{F}'^{(t)}$

since the temporal signal requires the previous Tweedie estimate, which is unavailable at the first step. For FLUX, images are generated using 35 inference steps with a guidance scale of 3.5. Within the FLUX architecture, we apply token merging to blocks 9 through 19 out of the 19 dual-stream blocks, and to blocks 7 through 22 out of the 38 single-stream blocks.

We further evaluate DaTMe on U-Net-based architectures, specifically Stable Diffusion 1.5 and SDXL. For these models, we set the interpolation scale to $s = 2$ and the number of grid cells to $P = 4$. Images are generated using 50 inference steps, with a guidance scale of 7.5 for Stable Diffusion 1.5 and 5.0 for SDXL. Following [9, 10], we apply token merging at the 64×64 latent resolution in both architectures. Unlike in DiT-based models, we do not explicitly construct the constraint set $\mathcal{S}_{\text{con}}$ for U-Net backbones. Because U-Net architectures already possess a strong spatial inductive bias through convolutions, tokens with similar features are naturally organized in nearby spatial regions [43]. Consequently, spatial locality is implicitly preserved. We therefore only identify the source set $\mathcal{S}_{\text{src}}$ and sample the remaining merge candidates without enforcing an additional constraint set. Finally, token merging begins at the same denoising step as in the DiT experiments.

C Additional Experiments

Additional Comparison Results. We compare additional results on PixArt-α under different compression ratios in Fig. 10. While the CFG-based method degrades noticeably at higher compression ratios, DaTMe more faithfully preserves subject details and unconditioned content, such as the woman's hat and the floor decoration beneath the cake (red boxes). On FLUX, existing methods increasingly deviate from the original generation as compression increases (Fig. 11). ToMeSD introduces changes in spatial composition and object structure, while ToMA additionally exhibits noisy visual degradation. In contrast, DaTMe more faithfully preserves the original composition and fine details across compression ratios.

Extended Experiments on U-Net Architecture. We further evaluate DaTMe on U-Net-based diffusion models under different compression ratios. For Stable Diffusion 1.5, Tab. 8 and Tab. 9 show that DaTMe better preserves structural fidelity while maintaining strong text alignment across compression ratios. A similar trend is observed on SDXL in Tab. 10 and Tab. 11, demonstrating the effectiveness of DaTMe beyond transformer-based diffusion architectures.

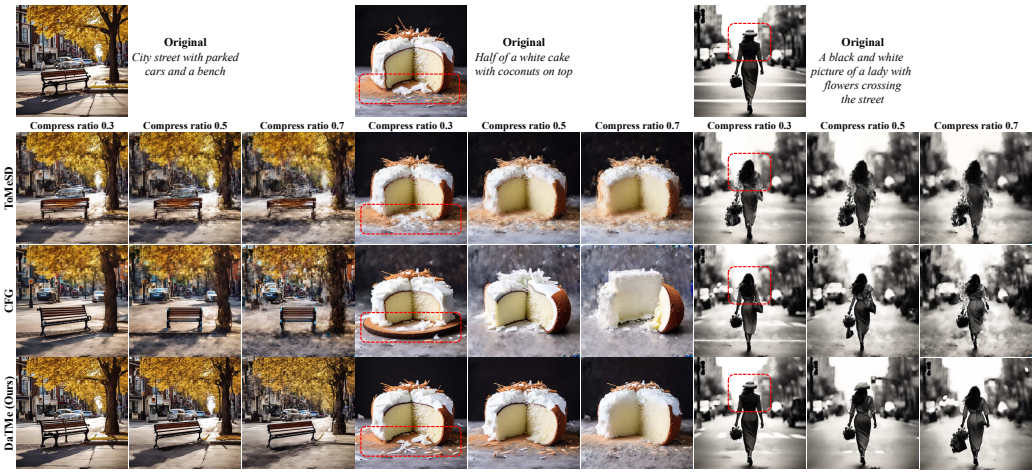

Figure 10: **Additional PixArt-α results under different compression ratios.** Unlike prior methods, DaTMe more faithfully preserves the fine details present in the original image.

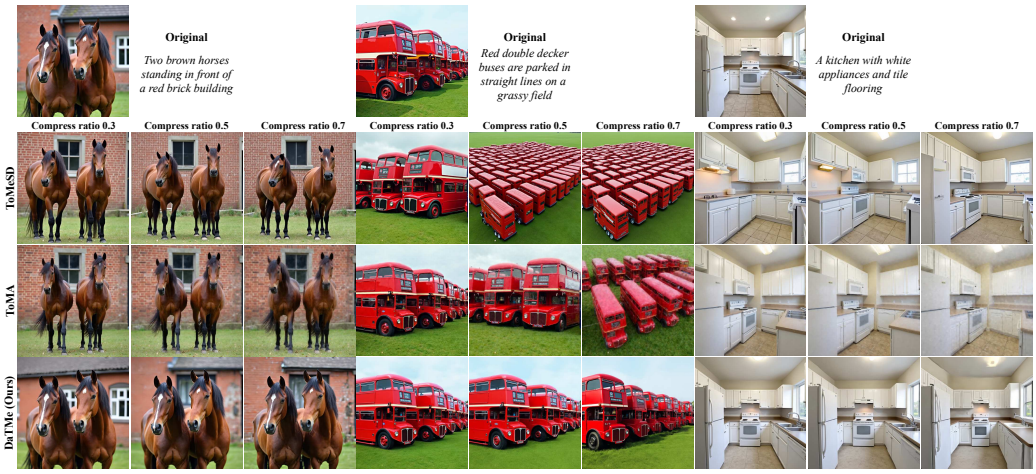

Figure 11: **Additional FLUX results under different compression ratios.** Unlike prior methods, DaTMe more faithfully preserves the fine details present in the original image.

This is further supported by Fig. 12 and Fig. 13: prior methods often fail to preserve the original image structure, particularly in non-prompt regions highlighted by red boxes, whereas DaTMe maintains these regions more faithfully.

Additional Ablation Studies on PixArt-α. Tab. 12 studies the number of spatial grid cells used for grid-based source completion. Increasing the number of grid cells generally improves structure preservation by distributing additional source tokens more evenly across the image, but overly fine grids introduce extra latency. We therefore choose $P = 4$, which achieves the best overall quality while remaining faster than the unmerged PixArt-α baseline. Fig. 15 further visualizes this trend: denser grids better preserve structural details by preventing source tokens from concentrating in a limited region. Tab. 13 ablates the down/up-sampling scale used to construct the spatial axis. We vary the scale from 2 to 16 and set $s = 4$, which achieves the best overall trade-off across quality metrics.

Fig. 14 presents the merging maps over denoising timesteps. Our method better preserves structural details, such as the bowl shape and leaves, thereby yielding higher-fidelity images.

Ablation Studies on SD. We select the hyperparameters for U-Net based model based on Tab. 14, Tab. 15 and Tab. 16.

Table 8: **Text alignment and efficiency score on SD 1.5.**

	CLIP-T↑	HPSv2↑	FID↓	Latency (s)↓	TFLOPs↓
SD 1.5 (Base)	31.2669	25.5643	33.2610	1.0186	70.396
Compression ratio 0.3					
ToMeSD	31.3262	25.1050	32.7718	1.0283	70.744
CFG	31.3453	25.1062	**32.6889**	1.0669	70.744
DaTMe (Ours)	**31.3456**	**25.1185**	33.3034	0.9155	70.849
Compression ratio 0.5					
ToMeSD	31.3212	24.6347	**33.4531**	1.0141	70.408
CFG	**31.3631**	24.6633	33.4947	1.0704	70.408
DaTMe (Ours)	31.3306	**24.6682**	34.3004	0.9138	70.727
Compression ratio 0.7					
ToMeSD	31.2225	23.6237	36.3672	1.0028	70.073
CFG	31.2227	**23.7960**	**34.9584**	1.0507	70.073
DaTMe (Ours)	**31.2550**	23.7728	35.7699	0.9025	70.192

Table 9: **Image fidelity score on SD 1.5.**

	SSIM↑	PSNR↑	LPIPS↓
Compression ratio 0.3			
ToMeSD	0.8245	20.0186	0.1435
CFG	0.8265	20.1034	0.1448
DaTMe (Ours)	**0.8579**	**21.2101**	**0.1171**
Compression ratio 0.5			
ToMeSD	0.7575	18.4332	0.2038
CFG	0.7596	18.4743	0.2036
DaTMe (Ours)	**0.7820**	**19.1236**	**0.1842**
Compression ratio 0.7			
ToMeSD	0.7001	17.2531	0.2648
CFG	0.6804	17.1733	0.2730
DaTMe (Ours)	**0.7035**	**17.5194**	**0.2611**

Table 10: **Text alignment and efficiency score on SDXL.**

	CLIP-T↑	HPSv2↑	FID↓	Latency (s)↓	TFLOPs↓
SDXL (Base)	31.9501	27.0768	32.5454	5.0761	607.849
Compression ratio 0.3					
ToMeSD	31.9603	26.7230	32.4035	5.0839	607.851
CFG	**31.9795**	26.7159	**32.3833**	5.2559	607.851
DaTMe (Ours)	31.9622	**26.9697**	32.3880	4.2480	608.035
Compression ratio 0.5					
ToMeSD	31.9411	26.1445	33.3079	5.0216	605.164
CFG	**32.0154**	26.1695	32.8330	5.2054	605.164
DaTMe (Ours)	31.9604	**26.8040**	**32.4385**	4.2622	607.312
Compression ratio 0.7					
ToMeSD	31.6383	24.4532	39.8369	4.9837	602.481
CFG	**31.9262**	24.9657	36.0965	5.1569	602.481
DaTMe (Ours)	31.8824	**25.1762**	**35.2577**	4.2672	605.930

Table 11: **Image fidelity score on SDXL.**

	SSIM↑	PSNR↑	LPIPS↓
Compression ratio 0.3			
ToMeSD	0.8076	21.1286	0.1715
CFG	0.8078	21.1278	0.1715
DaTMe (Ours)	**0.8413**	**22.2688**	**0.1412**
Compression ratio 0.5			
ToMeSD	0.7458	19.5331	0.2367
CFG	0.7408	19.3888	0.2433
DaTMe (Ours)	**0.7801**	**20.3747**	**0.2007**
Compression ratio 0.7			
ToMeSD	0.6942	18.5884	0.3097
CFG	0.6827	18.1728	0.3194
DaTMe (Ours)	**0.7114**	**18.8536**	**0.2860**

DaTMe Additional Results. In Fig. 16 and Fig. 17, we present additional samples under different compression ratios. DaTMe consistently preserves the structural details of the original generation while maintaining strong text alignment.

D Limitations

Similar to prior token merging methods, DaTMe faces limitations under aggressive token compression. In highly complex scenes with many objects or intricate backgrounds, excessive merging can inevitably remove necessary tokens, leading to degraded generation quality. In addition, the layer-wise roles of diffusion models remain underexplored, especially in DiT-based architectures. Understanding which layers contribute most to structural layout, fine details, or background refinement could further improve token merging by enabling layer-aware compression strategies.



Figure 18: **Limitations.** In highly complex images, aggressive compression adversely affects image quality, leading to loss of structural details and visual coherence.

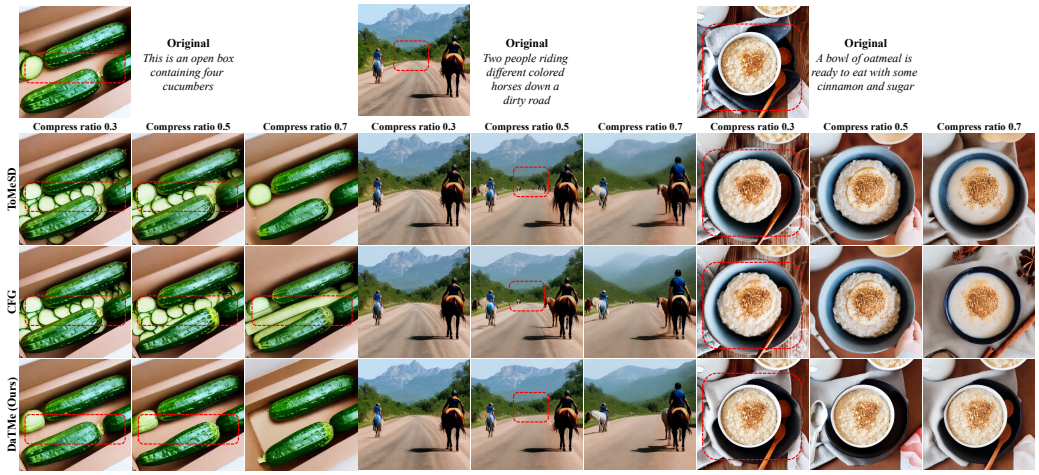

Figure 12: **SD 1.5 results under different compression ratios.** Unlike prior methods, DaTMe more faithfully preserves the contents of the original image.

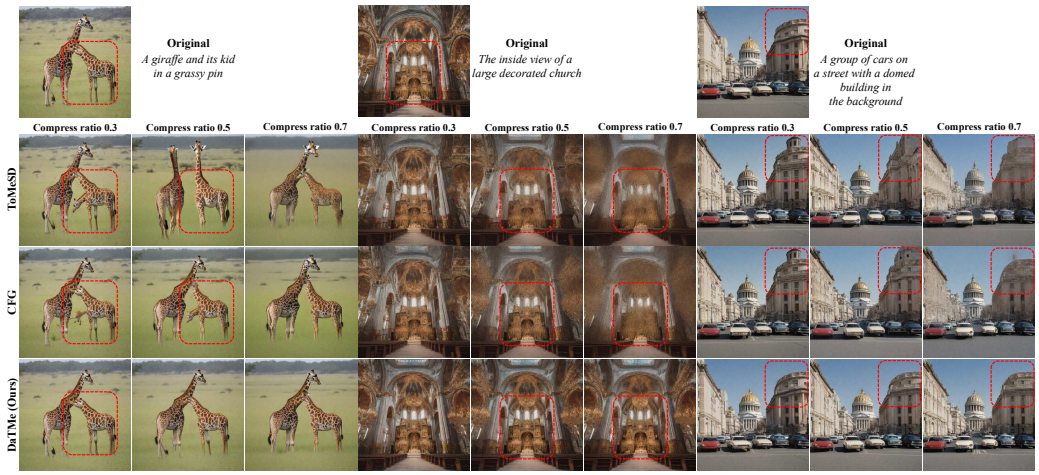

Figure 13: **SDXL results under different compression ratios.** Unlike prior methods, DaTMe more faithfully preserves the contents of the original image especially background details.

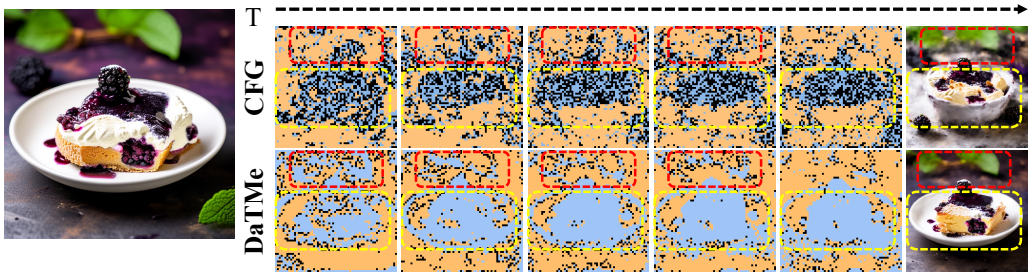

Figure 14: **Difference in merging maps.** Black spots indicate the destinations of merged source tokens. Our method better preserves structurally important details across the entire image.

Table 12: **Ablation on the number of spatial grids in PixArt-α.** We vary the number of grid cells used for grid-based source completion and select $P = 4$ as the best trade-off between latency and performance.

Grid	CLIP-T↑	SSIM↑	PSNR↑	LPIPS↓	FID↓	Latency (s)↓
PixArt-α [5]	31.2485	–	–	–	42.1001	1.9417
CFG [9]	30.8801	0.7137	17.8157	0.2369	44.6070	1.7458
1	**31.1426**	0.8099	20.4707	0.1513	42.5585	**1.7343**$_{-0.2074}$
4	31.1375	0.8112	20.5136	0.1501	42.4474	1.8030$_{-0.1387}$
16	31.1374	0.8128	20.5705	0.1485	**42.3615**	1.9765$_{+0.0348}$
64	31.1307	**0.8140**	**20.6223**	**0.1477**	42.3792	2.4443$_{+0.5026}$

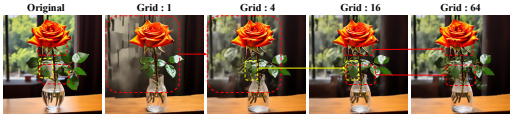

Figure 15: **Visualization of different grid sizes on PixArt-α.** Denser grids better preserve structural details by preventing source tokens from concentrating in a limited region.

Table 13: **Ablation on the interpolation scale.** We vary the down/up-sampling scale used to construct the spatial axis in PixArt-α.

Scale	CLIP-T↑	SSIM↑	PSNR↑	LPIPS↓	FID↓
2	31.1368	0.8102	20.4854	0.1507	42.5226
4	31.1375	**0.8112**	**20.5136**	**0.1501**	**42.4474**
8	**31.1445**	0.8091	20.4785	0.1522	42.6547
16	31.1259	0.8016	20.2553	0.1596	42.7489

Table 14: **Ablation on constraint token selection in SD 1.5.**

	CLIP-T↑	SSIM↑	PSNR↑	LPIPS↓	FID↓
$\mathcal{M}_{\text{temp}}\uparrow \cap \mathcal{M}_{\text{spat}}\downarrow$	31.3346	0.8402	20.6266	0.1316	**33.1811**
$\mathcal{M}_{\text{temp}}\downarrow \cap \mathcal{M}_{\text{spat}}\uparrow$	**31.3592**	0.8529	21.0306	0.1214	33.4103
$\mathcal{M}_{\text{temp}}\uparrow \cap \mathcal{M}_{\text{spat}}\uparrow$	31.3527	0.8472	20.8985	0.1273	33.3237
None	31.3456	**0.8579**	**21.2101**	**0.1171**	33.3034

Table 15: **Ablation on the interpolation scale in SD 1.5.**

Scale	CLIP-T↑	SSIM↑	PSNR↑	LPIPS↓	FID↓
2	31.3456	**0.8579**	**21.2101**	**0.1171**	**33.3034**
4	**31.3543**	0.8569	21.1812	0.1182	33.4557
8	31.3468	0.8541	21.1044	0.1209	33.4293
16	31.3422	0.8514	21.0544	0.1237	33.4424

Table 16: **Ablation on the number of spatial grids in SD 1.5.**

Grid	CLIP-T↑	SSIM↑	PSNR↑	LPIPS↓	FID↓
1	**31.3529**	0.8571	21.1678	0.1178	33.3703
4	31.3456	**0.8579**	**21.2101**	**0.1171**	**33.3034**
16	31.3443	0.8576	21.1988	0.1176	33.3631
64	31.3365	0.8570	21.1814	0.1180	33.3841

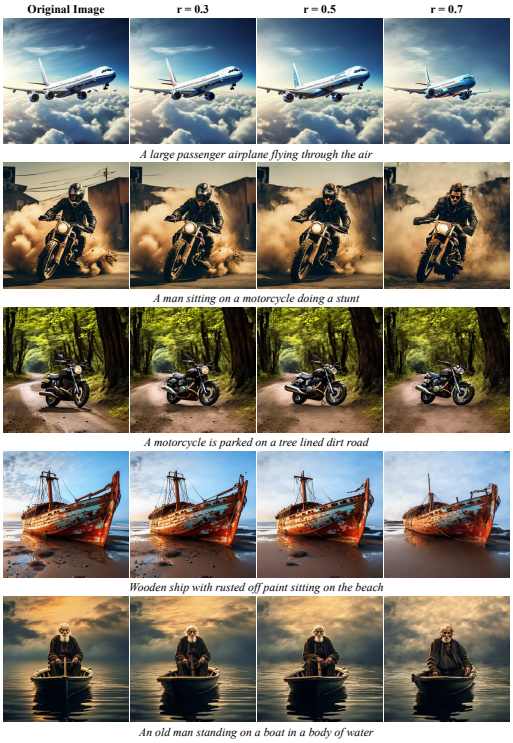

Figure 16: **Additional PixArt-α results.**

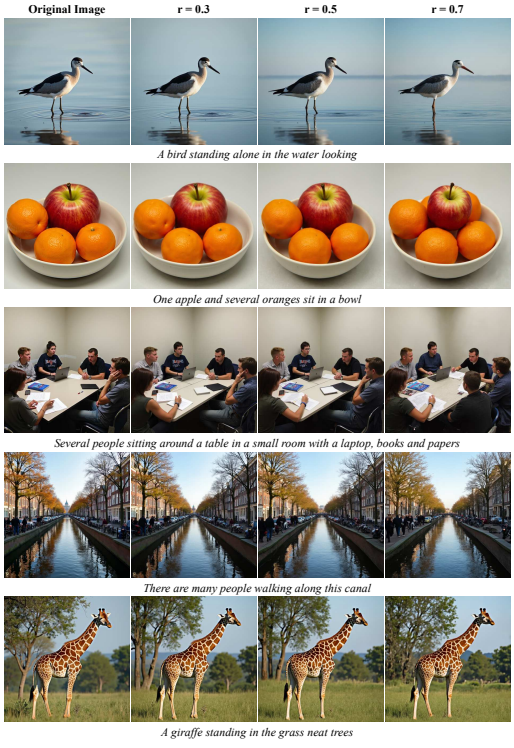

Figure 17: **Additional FLUX results.**

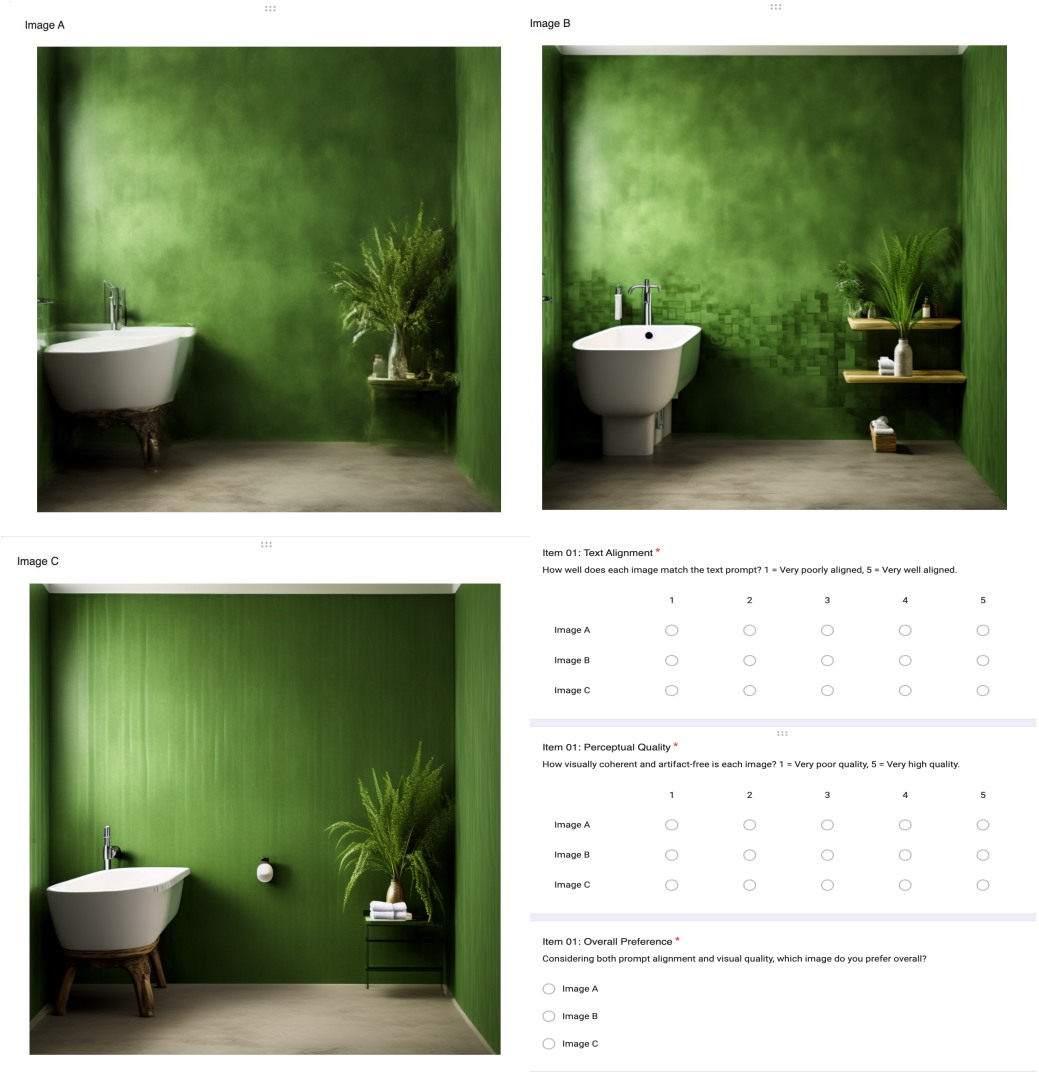

Figure 19: **Example questionnaire for the human preference study.** Each evaluation sample includes three questions.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The main claims regarding the effectiveness of the proposed spatiotemporal importance map are clearly stated in Section 1 and are fully supported by the quantitative and qualitative results presented in Section 4. We do not overstate our claims beyond the evaluated architectures.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We discuss the limitations of our method in the supplementary material.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper does not introduce formal mathematical theorems or lemmas that require proofs. The equations provided in Section 3 and Appendix A are algebraic derivations and decompositions.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: All implementation details necessary to reproduce our experiments, including hyperparameters and model configurations, are explicitly provided in Section 4 and Appendix B.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The source code and implementation details are publicly available at `https://github.com/SeungJuCha/DaTMe`.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: The experimental settings necessary to understand and reproduce the results are presented in the core paper (Section 4). The full comprehensive details, including exact hyperparameters and model configurations, are provided in Appendix B.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: Following standard practices, we report exact metrics (FID, CLIP-T) computed over a large, fixed set of samples. Because the inference process under these conditions is deterministic and the sample size is large enough to ensure stability, variance across multiple runs is negligible and error bars are not typically reported.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: The exact hardware specifications and computational metrics are explicitly detailed in Section 4 and the supplementary materials.

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Our research follows the NeurIPS Code of Ethics. The human preference study involved 32 evaluators who assessed generated images for text alignment, visual quality, and overall preference.

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [No]

    Justification: Our work proposes a generic inference acceleration method via token merging rather than a specific end-user application. Thus, it does not introduce new negative societal impacts beyond those already inherent to the baseline diffusion models.

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: Our paper proposes a training-free token merging framework for existing, publicly available diffusion model.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All existing assets used in our experiments, including the COCO 2014 dataset and pre-trained models (PixArt-α, FLUX, Stable Diffusion 1.5, SDXL), are properly cited in the main text.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [N/A]

    Justification: The paper does not introduce or release any new assets such as datasets or pre-trained foundation models that require structured documentation.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [Yes]

    Justification: We conducted a human preference study with 32 participants evaluating 20 randomly sampled cases. The supplementary material provides the complete evaluation instructions, three evaluation questions, and an example screenshot of the questionnaire. No compensation was provided to participants.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A]

    Justification: We conducted a minimal-risk human preference study in which 32 evaluators rated generated images.

16. **Declaration of LLM usage**

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

    Answer: [N/A]

    Justification: Our proposed methodology does not utilize LLMs. LLMs were not used as a component of the core research.