

# ToDo: Token Downsampling for Efficient Generation of High-Resolution Images

Ethan Smith, Nayan Saxena, Aninda Saha

Leonardo AI Research Lab, North Sydney, NSW, Australia

{ethan, nayan.saxena, aninda}@leonardo.ai

## Abstract

Attention has been a crucial component in the success of image diffusion models, however, their quadratic computational complexity limits the sizes of images we can process within reasonable time and memory constraints. This paper investigates the importance of dense attention in generative image models, which often contain redundant features, making them suitable for sparser attention mechanisms. We propose a novel training-free method ToDo that relies on token downsampling of key and value tokens to accelerate Stable Diffusion inference by up to 2x for common sizes and up to 4.5x or more for high resolutions like  $2048 \times 2048$ . We demonstrate that our approach outperforms previous methods in balancing efficient throughput and fidelity.

## 1 Introduction

Transformers, and their key component, attention, have been integral to the success and improvements in generative models in recent years [Vaswani *et al.*, 2023]. Their global receptive fields, ability to compute dynamically based on input context, and large capacities have made them useful building blocks across many tasks [Khan *et al.*, 2022]. The main drawback of Transformer architectures is their quadratic scaling of computational complexity with sequence length, affecting both time and memory requirements. When looking to generate a Stable Diffusion image at  $2048 \times 2048$  resolution, the attention maps of the largest U-Net blocks incur a memory cost of approximately 69 GB in half-precision, calculated as  $(1 \text{ batch} \times 8 \text{ heads} \times (256^2 \text{ tokens})^2 \times 2 \text{ bytes})$ . This exceeds the capabilities of most consumer GPUs [Zhuang *et al.*, 2023]. Specialized kernels, such as those used in Flash Attention, have greatly improved speed and reduced memory costs [Dao *et al.*, 2022], however, challenges due to the unfavorable quadratic scaling with sequence length are persistent.

In the quest for computational efficiency, the concept of sparse attention has gained traction. Methods like Token Merging (ToMe) [Bolya *et al.*, 2023] and its application in latent image diffusion models [Bolya and Hoffman, 2023] have reduced the computation time required by condensing tokens with high similarity, thereby retaining the essence of

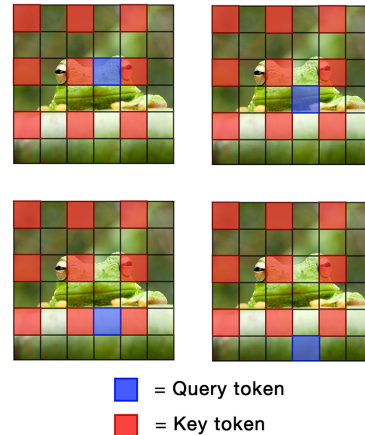


Figure 1: A visualization of our method. From a given latent or image, we subsample positions on the grid in a strided fashion for the keys and values used in attention maintaining the full set of query tokens. Link to demo video is [here](#).

the information with fewer tokens. Similarly, approaches like Neighborhood Attention [Hassani *et al.*, 2023] and Focal Transformers [Yang *et al.*, 2021] have introduced mechanisms where query tokens attend only to a select neighborhood, balancing the trade-off between receptive field and computational load. These strategies aim to efficiently approximate the attention mechanism’s output. While performant, these methods typically require training-time modifications to be successful, incurring significant logistical overheads to leverage their optimizations.

Complementing the sparse attention frameworks, attention approximation methods offer an alternative avenue by exploiting mathematical properties to simplify the attention operation. Techniques ranging from replacing the softmax with more computationally friendly nonlinearities [Chen *et al.*, 2020], to fully linearizing attention [Katharopoulos *et al.*, 2020], and leveraging the kernel trick for dimensionality reduction [Choromanski *et al.*, 2022], have been explored to approximate attention efficiently but are also generally required to be trained into the model.

Building upon these works and aiming to address the pre-training requirement, we propose a novel post-hoc method for

accelerating inference, which we refer to as Token Downsampling (ToDo). Our approach, ToDo, is inspired by the observation that adjacent pixels in images tend to exhibit similar values to their neighbors. Hence, we employ a downsampling technique to reduce tokens, akin to grid-based subsampling in image processing. Compared to prior method ToMe [Bolya and Hoffman, 2023], our method not only simplifies the merging process but also significantly reduces computational overhead, as it eliminates the need for exhaustive similarity calculations. In summary, our main contributions are:

- A training-free method that can accelerate inference for Stable Diffusion up to 4.5x faster, beating previous methods in balancing throughput and fidelity.
- An in-depth analysis of attention features within the U-Net, and hypotheses on why attention can be approximated sparsely without substantially hurting fidelity.

## 2 Methods

### 2.1 Background

**Diffusion Models for Image Generation** The diffusion model [Song and Ermon, 2019] employs a U-Net architecture [Ronneberger *et al.*, 2015] with transformer-based blocks that utilize self-attention layers [Rombach *et al.*, 2021]. This setup flattens spatial dimensions into a series of tokens, which are then fed through multiple transformer blocks to predict the denoised image.

**Original Token Merging Scheme** In the original ToMe [Bolya *et al.*, 2023] framework, tokens are categorized into source (src) and destination (dst) sets. The merging process involves identifying the  $r$  most similar tokens from the src set and merging them into the dst set, effectively reducing the total token count by  $r$ . This merging is defined as  $x_{\text{merged}} = \frac{1}{r} \sum_{i=1}^r x_i$  where  $x_i$  represents individual tokens to be merged.

Overall, the original ToMe method is predicated on the reduction of computational load through merging of similar tokens prior to being input to attention layers. This process involves the computation of a similarity matrix, where tokens exhibiting the highest similarity are merged. Subsequently, the unmerging process aims to redistribute the merged token information back to the original token locations. This approach, however, introduces two critical bottlenecks:

- **Computational Complexity:** The similarity matrix calculation,  $\mathcal{O}(n^2)$  complexity, is costly in itself, especially when required at every step of the process.
- **Quality Degradation:** The merge-unmerge cycle inherent to ToMe can lead to significant loss of image detail, particularly at higher merging ratios.

### 2.2 Training Free Enhancements

Our proposed token downsampling (ToDo) methodology extends the original ToMe approach, addressing its computational bottlenecks and quality degradation issues when applied to Stable Diffusion models. We introduce two principal modifications with ToDo: an optimized token merging method based on spatial contiguity and a refined attention mechanism that mitigates the need for unmerging.

**Optimized Merging Through Spatial Contiguity** We introduce a novel token merging strategy that leverages the inherent spatial contiguity of image tokens. Recognizing that tokens in close spatial proximity exhibit higher similarity, thus providing a basis for merging without the extensive computation of pairwise similarities. Therefore, we employ a downsampling function  $D(\cdot)$  using the Nearest-Neighbor algorithm [Bankman, 2008]. We note this approach is akin to strided convolutions, as shown in Figure 1. Formally, let  $T = \{t_1, t_2 \dots t_n\}$  denote the original set of image tokens arranged in a two-dimensional grid reflecting their spatial relationships. The proposed downsampling operation,  $D$  is applied to  $T$  to yield a reduced set of merged tokens  $T'$ , as such:

$$T' = D(T) = \{D(t_1), D(t_2) \dots D(t_m)\}, \text{ where } m < n$$

This enhancement mitigates the computational overhead associated with the pairwise similarity computation inherent in ToMe. By leveraging the assumption that spatially adjacent tokens are more likely to be similar, we bypass the need for  $\mathcal{O}(n^2)$  similarity calculations, instead employing a more computationally efficient  $\mathcal{O}(n)$  downsampling operation.

**Enhanced Attention Mechanism with Downsampling** To mitigate the information loss inherent to the unmerging process in conventional token merging approaches, we introduce a refinement to the attention mechanism within the transformer architecture [Vaswani *et al.*, 2023]. This modification entails the application of the downsampling operation  $D(\cdot)$  to the keys,  $K$ , and values  $V$  of the attention mechanism while preserving the original queries  $Q$ . The modified attention function can be mathematically articulated as follows, with  $d_k$  denoting the dimensionality of the keys, ensuring proper scaling within the softmax operation.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{Q \cdot D(K)^T}{\sqrt{d_k}}\right) \cdot D(V)$$

This refinement ensures that the integrity of the queries is preserved, thereby maintaining the fidelity of the attention process while reducing the dimensionality of the matrices involved in the attention computation.

## 3 Experiments

**Experimental Setup** For our empirical evaluation, we employ the finetuned DreamshaperV7 model [Luo *et al.*, 2023], noted for its superior handling of larger image dimensions which are central to this study. All experiments are conducted on a single A6000 GPU, utilizing float16 precision and flash attention [Dao *et al.*, 2022] for inference as this has become the norm for many users. We use the DDIM sampler [Song *et al.*, 2020] with 50 diffusion steps and a guidance scale of 7.5 [Team, 2024]. Each experiment involves averaging 10 generations comparing ToDo against ToMe with baseline referring to standard generations without token merging. The resolutions benchmarked include:  $1024 \times 1024$ ,  $1536 \times 1536$  and  $2048 \times 2048$  across two token merging ratios, 0.75 and 0.89 which denotes the proportion of tokens removed. This is equivalent to 2x and 3x downsample respectively. For the comparison in Figure 2 we also use a merge ratio of 0.9375 for the  $2048 \times 2048$  images, equivalent to a 4x downsample.

**Image Quality and Throughput** To assess the fidelity and detail preservation of generated images, we utilized Mean Squared Error (MSE) to quantify each method’s deviation from the baseline, and High Pass Filter (HPF) a standard for evaluating image sharpness and texture preservation [Gonzalez, 2009]. Our analysis, substantiated by Figure 2 and Table 1, demonstrates that our method not only closely mirrors the baseline in terms of MSE but also maintains comparable HPF values, underscoring its efficiency in retaining image features while ensuring higher throughput, as depicted in Figure 3.



Figure 2: Qualitative comparison of attention methods with: 25% of tokens at  $1024 \times 1024$ , 11% at  $1536 \times 1536$ , and 6% at  $2048 \times 2048$ , maintaining a consistent token count of 4096 post-merging.

Method	Merge Ratio	MSE	HPF
Baseline	-	-	4.846
ToMe	0.75	$2.686 \times 10^{-2}$	4.022
	0.89	$2.671 \times 10^{-2}$	4.003
ToDo (ours)	0.75	$6.247 \times 10^{-3}$	<b>4.887</b>
	0.89	$9.207 \times 10^{-3}$	<b>4.733</b>

Table 1: Metrics from various attention methods, averaged over 10 generations of different prompts at  $1536 \times 1536$  resolution. MSE denotes the mean squared error relative to the baseline, while HPF represents the mean absolute magnitude post-high pass filtering.

**Latent Feature Redundancy** We investigated latent feature redundancy in the Stable Diffusion U-Net, assessing similarity among adjacent latent features. By extracting latent representations at various stages and noise levels, we constructed cosine similarity matrices, focusing on the proportion of tokens with top-3 similarities within a  $3 \times 3$  area, and the highest, mean, and lowest similarities within  $3 \times 3$  and  $5 \times 5$  areas. We observed high similarity among neighboring tokens within the hidden features and notable trends as seen in Figure 4. Similarity trends varied across different depths without a distinct pattern, possibly due to the increas-

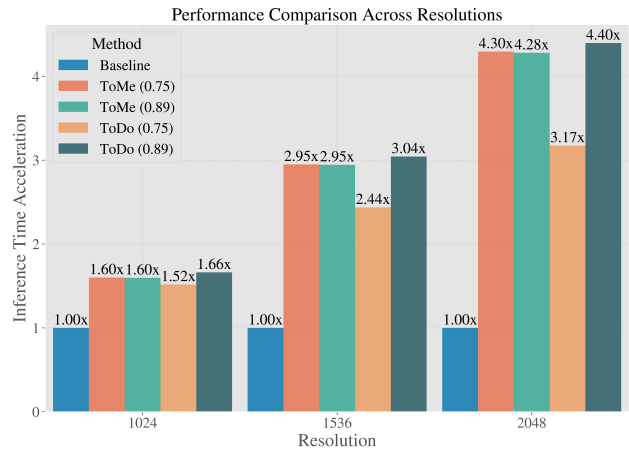


Figure 3: Inference time acceleration across resolutions using attention methods at various merge ratios. Bars represent the relative performance increase against the baseline.

ing spatial compression and consequent reduction in information redundancy with values diminishing as the denoising progresses, likely because diffusion models initially generate broad details and later refine them.

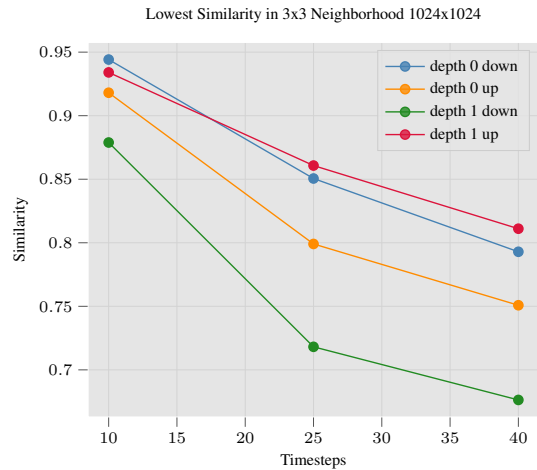


Figure 4: Lowest cosine similarity between tokens in a  $3 \times 3$  area across diffusion timesteps and U-Net locations extracted from 10 generations of different prompts at  $1024 \times 1024$ . Timesteps out of 50 indicate noise reduction; Depth 0 is initial resolution, Depth 1 is after 2x downsampling. Up/down denotes encoder/decoder blocks.

## 4 Conclusion

We demonstrate that our approach ToDo is capable of maintaining the balance between efficient throughput and fidelity, especially in high-frequency components. We also show that nearby features within the U-Net might be redundant and postulate that our method can benefit other attention based generative image models, especially those operating on a large number of tokens. Future work can explore the differentiability of our method, and leverage it to efficiently finetune Stable Diffusion at previously unseen larger image dimensions.

## References

- [Bankman, 2008] Isaac Bankman. *Handbook of medical image processing and analysis*. Elsevier, 2008.
- [Bolya and Hoffman, 2023] Daniel Bolya and Judy Hoffman. Token merging for fast stable diffusion. *CVPR Workshop on Efficient Deep Learning for Computer Vision*, 2023.
- [Bolya *et al.*, 2023] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your ViT but faster. In *International Conference on Learning Representations*, 2023.
- [Chen *et al.*, 2020] Dengsheng Chen, Jun Li, and Kai Xu. Arelu: Attention-based rectified linear unit, 2020.
- [Choromanski *et al.*, 2022] Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, David Belanger, Lucy Colwell, and Adrian Weller. Rethinking attention with performers, 2022.
- [Dao *et al.*, 2022] Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems*, 2022.
- [Gonzalez, 2009] Rafael C Gonzalez. *Digital image processing*. Pearson education india, 2009.
- [Hassani *et al.*, 2023] Ali Hassani, Steven Walton, Jiachen Li, Shen Li, and Humphrey Shi. Neighborhood attention transformer, 2023.
- [Katharopoulos *et al.*, 2020] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention, 2020.
- [Khan *et al.*, 2022] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *ACM Computing Surveys*, 54(10s):1–41, January 2022.
- [Luo *et al.*, 2023] Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023.
- [Rombach *et al.*, 2021] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021.
- [Ronneberger *et al.*, 2015] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [Song and Ermon, 2019] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.
- [Song *et al.*, 2020] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv:2010.02502*, October 2020.
- [Team, 2024] Huggingface Diffusers Team. Speed up inference, 2024.
- [Vaswani *et al.*, 2023] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need, 2023.
- [Yang *et al.*, 2021] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers. *CoRR*, abs/2107.00641, 2021.
- [Zhuang *et al.*, 2023] Bohan Zhuang, Jing Liu, Zizheng Pan, Haoyu He, Yuetian Weng, and Chunhua Shen. A survey on efficient training of transformers. *arXiv preprint arXiv:2302.01107*, 2023.