

What Kind of Visual Tokens Do We Need? Training-Free Visual Token Pruning for Multi-Modal Large Language Models from the Perspective of Graph

Yutao Jiang^{*1}, Qiong Wu^{*12}, Wenhao Lin¹², Wei Yu¹², Yiyi Zhou^{†12}

¹Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University, 361005, P.R. China.

²Institute of Artificial Intelligence, Xiamen University, 361005, P.R. China.
{yutaojiang, qiong, wenhaolin, weiyu}@stu.xmu.edu.cn, zhoyiyi@xmu.edu.cn

Abstract

Recent *Multimodal Large Language Models* (MLLMs) often use a large number of visual tokens to compensate their visual shortcoming, leading to excessive computation and obvious visual redundancy. In this paper, we investigate what kind of visual tokens are needed for MLLMs, and reveal that both foreground and background tokens are critical for MLLMs given the varying difficulties of examples. Based on this observation, we propose a *graph*-based method towards training-free visual token pruning, termed *G-Prune*. In particular, *G-Prune* regards visual tokens as nodes, and construct their connections based on their semantic similarities. Afterwards, the information flow is propagated via weighted links, and the most important tokens after iterations are kept for MLLMs, which can be front or background. To validate *G-Prune*, we apply it to a recent MLLM called LLaVA-NeXT, and conduct extensive experiments on a set of benchmarks. The experiment results show that *G-Prune* can greatly reduce computation overhead while retaining high performance on both coarse- and fine-grained tasks. For instance, *G-Prune* can reduce 63.57% FLOPs of LLaVA-NeXT on VQA2.0 and TextVQA with only 0.95% and 2.34% accuracy drops, respectively.

Code — <https://github.com/jytmelon/G-Prune>

Introduction

Recently, extending *large language models* (LLMs) to more modalities has become a research hotspot (Achiam et al. 2023; Bai et al. 2023a; Touvron et al. 2023b; Chen et al. 2024b; Guo et al. 2024), *i.e.*, multimodal LLM (MLLMs). For vision-language (VL) tasks, the most common paradigm is to directly project the extracted visual features into the semantic space of LLMs (Li et al. 2023a; Team et al. 2023; Achiam et al. 2023; Liu et al. 2024b,a; Luo et al. 2024) as the input tokens. Despite effective, these MLLMs (Team et al. 2023; Achiam et al. 2023; Liu et al. 2024b,a) often suffer from more severe visual hallucinations, and perform worse on granular VL tasks like TextVQA (Singh et al. 2019).

A straightforward solution to remedy the visual shortcoming of MLLMs is to increase image resolution, *i.e.*,

^{*}These authors contributed equally.

[†]Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

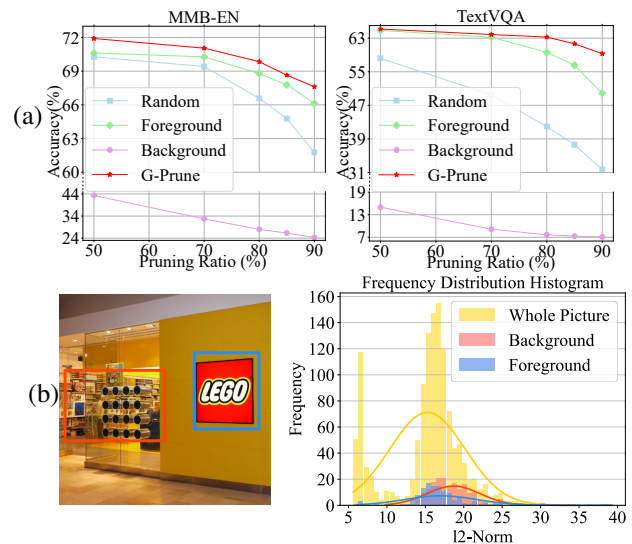


Figure 1: (a) Random pruning, Foreground-preserving pruning, Background-preserving pruning and *G-Prune* performance curves in MMBench and TextVQA, (b) The frequency distribution of l_2 -Norm for the entire image, as well as its specific background and foreground areas.

using more visual tokens, thereby making vision matters in multimodal reasoning. For instance, LLaVA-NeXT (Liu et al. 2024a) partitions images into multiple regions and directly concatenates the visual tokens mapped from each region. InternVL (Chen et al. 2024b) presets a variety of high-resolution image inputs with different aspect ratios. Although simple, this solution does improve the visual reasoning of MLLMs to a large extent, achieving new SOTA performances on a set of benchmarks. However, a counterpart is the prohibitively expensive computation. For instance, LLaVA-NeXT uses 2880 visual tokens, and require about 18.52 TFLOPs computational for each inference.

In this case, a question rises, *do we really need so many visual tokens for MLLMs?* The large image tokens can provide more detailed visual semantics for multi-modal reasoning. However, these visual tokens are often processed by the well pre-trained image encoders, *e.g.*, ViT (Dosovitskiy et al. 2020), and have a large receptive field, *i.e.*, represent-

ing much more information than its actual area. In this case, heavy redundancy does exist in these image tokens (Pan et al. 2021). Thus, randomly dropping a certain number of visual tokens has almost no impact on the performance on MM-bench (Goyal et al. 2017), a widely used MLLM benchmark for common scenarios, also shown in Fig.1-a. However, on the granular VL task, *e.g.* TextVQA (Singh et al. 2019), randomly removing a larger number of visual tokens will lead to a drastic decline in performance, indicating the great loss of key information, as shown in Fig.1-a. Conclusively, the large number of visual tokens are apparently redundant, but what visual tokens to preserve remains an challenge given the varying difficulties of examples.

Previous research (Rao et al. 2021; Liang et al. 2022; Xu et al. 2022; Yang et al. 2022) often focuses on the foreground information of the given images, *i.e.*, the main objects, but we note that both the foreground and background tokens are important for MLLMs. To explain, for a conventional image recognition task, the foreground object mainly responds to the label of this image, and only a small number of foreground tokens can be capable of recognition. In contrast, for MLLM tasks, some details are often questioned. Meanwhile, we computed the l_2 -Norm frequency distribution histograms for both the foreground and background, and found that their distributions have significant overlap, as shown in Fig.1-b. Nevertheless, how to recognize the important tokens for different images is still challenging in practice.

In this paper, we propose a graph-based method towards the training-free visual token pruning of MLLMs, termed G-Prune. In particular, we consider the visual tokens as graph nodes, and construct their connections according to their feature distances. Afterwards, information propagation is executed among nodes with an iterative algorithm to update the importance scores. Lastly, the most important tokens can be selected for MLLMs, which could be either foreground or background ones. In this way, we can select the most representative visual tokens for MLLMs, thereby greatly reducing the sequence length and computation complexity.

To validate G-Prune, we apply it to LLaVA-NeXT(Liu et al. 2024a), and conduct extensive experiments on eight competitive MLLM benchmarks, including VQA2.0 (Goyal et al. 2017), GQA (Hudson and Manning 2019), ChartQA (Masry et al. 2022), TextVQA (Singh et al. 2019), DocVQA (Mathew et al. 2020), POPE (Li et al. 2023b), MME (Fu et al. 2023) and MMB (Liu et al. 2023). The experimental results show that G-Prune significantly reduces computational costs while retaining high performance across various bench-marks. For instance, G-Prune can reduce the computation of LLaVA-NeXT on MME by 63.50% without performance drops. Moreover, even on granular tasks like TextVQA, G-Prune can also achieve a higher accuracy than the compared pruning methods while dropping a larger number of tokens, *e.g.*, 59.31 v.s. 39.02 (ToMe) on TextVQA while dropping 90% tokens.

Conclusively, the contribution of this paper is three-fold:

- We investigate the importance of different types of visual tokens for MLLMs, and show that both foreground and background ones are critical.

- We propose a graph-based method towards the training-free visual token pruning of MLLMs, which can mine the significant tokens via iterative information propagation.
- G-Prune can reduce the visual tokens of MLLMs to a large extent while retaining high performance on different VL benchmarks, even on the text-oriented benchmarks, such as TextVQA and DocVQA.

Related Work

Multimodal Large Language Models. Driven by the remarkable success of *large language models* (LLMs)(Touvron et al. 2023a,b; Meta 2024; Team et al. 2023; Reid et al. 2024; Achiam et al. 2023; Brown et al. 2020) in managing text-only tasks with demonstrated exceptional capabilities, the field of *multimodal large language models* (MLLMs)(Team et al. 2023; Achiam et al. 2023; Liu et al. 2024b,a; Luo et al. 2024; Zhang et al. 2024) is also attracting increasing scholarly interest. The core principle underlying the extension of LLMs to MLLMs entails transforming visual information into a sequence of tokens. These tokens are then amalgamated with textual tokens to create a unified input for processing by LLMs. Some approaches employ a series of learnable tokens to dynamically aggregate information from image sequences, which are then amalgamated with text tokens for processing within LLMs. For instance, BLIP-2 (Li et al. 2023a) introduces the QFormer module to dynamically schedule information aggregation from the outputs of the visual backbone. Similarly, Qwen-VL (Bai et al. 2023b) employs cross-attention mechanism to optimize the processing of information from ViT outputs. Although these methods have demonstrated success and incur only limited additional computational overhead, the inadvertent loss of visual information during the process limits the upper limit of such methods (Yao et al. 2024). To achieve this, other methods map visual information into the feature space of LLMs using a simple Multi-Layer Perceptron (MLP) layer. For example, MINI-GPT4(Zhu et al. 2023) applies a projection layer to map visual features into the semantic space of the LLM. Similarly, LLaVA (Liu et al. 2024b) follows this paradigm but further enhances it with a meticulously devised training strategy. For higher resolution, LLaVA-NeXT (Liu et al. 2024a) partitions images into multiple regions, and ViT is used to obtain visual features from upsampled images. InternVL (Chen et al. 2024b) presets a variety of high-resolution image inputs with different aspect ratios, and ViT is used to extract features from images of a specific size. These methods effectively preserve visual information by fully retaining the visual tokens. While a large amount of redundant information is fed into LLMs. To this end, it is necessary to filter out redundant tokens efficiently and effectively for the applications of MLLMs.

Token Pruning. Token pruning is a widely applied method for Transformer-based networks (Vaswani 2017; Devlin et al. 2018; Dosovitskiy et al. 2020), which reduces the less important tokens to speed up inference. In the aspect of MLLM, existing methods mainly use the calculation process of MLLM to locate redundant visual content. For instance, FastV (Chen et al. 2024a) uses the first two layers as the key to multi-

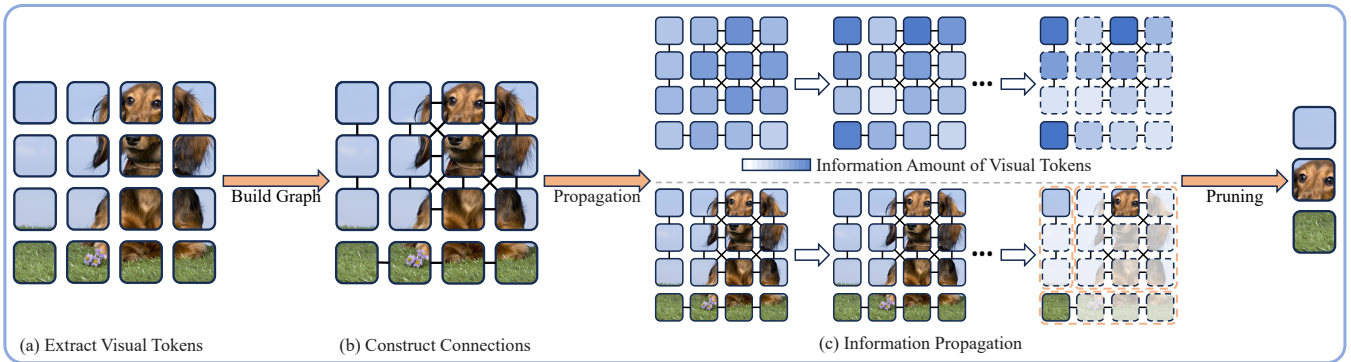


Figure 2: The overview of G-Prune. G-Prune aims to find out the important visual tokens for MLLMs, thereby reducing the computation complexity. In practice, G-Prune regards all visual tokens as graph nodes, and build their connections based on their semantic similarities. Afterwards, an iterative algorithm is performed to propagate information among nodes and upgrade the importance scores of visual tokens. After iterations, we can select the top- k tokens for MLLMs, which could be both foreground and background ones.

modal information exchange and directly removes most of the visual tokens with lower attention weights after the second layer. Similarly, VTW (Lin et al. 2024) directly remove all visual tokens after a certain layer. Some research efforts have also focused on pruning both textual and visual tokens simultaneously. For instance, PuMer (Cao, Paranjape, and Hajishirzi 2023) reduces both textual and visual tokens in vision-language models by applying text-informed image pruning and modality-aware token merging. Moreover, MADTP (Cao et al. 2024) leverages the Multi-modality Alignment Guidance (MAG) module to align image and text features, and the Dynamic Token Pruning (DTP) module to adaptively prune tokens in both modalities at each layer. The computational process involving MLLM will inevitably generate a lot of computational overhead. To this end, directly adopting the pruning method for ViT may be an effective approach. With additional module, DynamicViT (Rao et al. 2021) reduces the number of tokens according to the probabilities generated by a small MLP layers. Besides, some training-free methods, e.g., EViT (Liang et al. 2022) and Evo-ViT (Xu et al. 2022) calculate the similarity between the CLS token and other tokens, and remove or merge tokens with low similarity. BAT (Long et al. 2023) combines token importance with matching and clustering techniques to effectively preserve both the most discriminative and diverse tokens. These methods can effectively help us reduce computational overhead in ViT. However, the results of highlighting the foreground while ignoring the background make them possibly not directly applicable to MLLM. More deeply bound to the transformer computation process, Zero-TPrune (Wang, Dedhia, and Jha 2024) initializes the importance of each token and continuously updates the importance of each token based on the attention matrix during the calculation process, and only the tokens with higher importance are retained. Despite effectiveness, some existing methods take foreground as the main consideration ignores the requirement of the background in MLLMs. Others need to be integrated into the calculation process of MLLM, of which trials are expensive for MLLMs. In this paper, we focus on a graph-based

training-free method to filter out the redundant tokens while maintaining all objects by determining the most representative token of each object.

Method

In this paper, we propose a graph-based method towards the training-free visual token pruning of MLLMs, termed *G-Prune*. The principle of G-Prune is to select the most representative tokens from the perspective of graph, which could be either foreground or background ones.

In practice, G-Prune considers the visual tokens as graph nodes and builds their connections based on the feature similarity. Afterwards, an iterative algorithm is executed to propagate information among nodes, based on which the most important ones emerge.

Concretely, given a set of visual tokens $\mathbf{X} \in \mathbb{R}^{N \times d}$, where N and d represent the number and dimension of visual tokens, we first construct the graph according to node-wise feature distances:

$$\mathbf{A}_{ij} = \begin{cases} \cos(\mathbf{X}_i, \mathbf{X}_j), & \cos(\mathbf{X}_i, \mathbf{X}_j) \geq s, \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where $\cos(\cdot, \cdot)$ represents the *cosine* similarity between two tokens. s is a threshold for the connection. Then we apply the softmax function to normalize each row of \mathbf{A} . In this way, the tokens with similar semantics are connected, e.g., the background ones or the ones of the same object.

Afterwards, we iterate information propagation to find out the most representative token for each connected component. We first initialize the amount of information of visual tokens according to their l_2 -Norm:

$$\mathbf{S}_i^{(0)} = \sqrt{\sum_{j=1}^d \mathbf{X}_{ij}^2}, \quad (2)$$

where \mathbf{X}_{ij} denotes the j -th component of the i -th visual token. And $\mathbf{S} \in \mathbb{R}^{1 \times N}$ denotes the initial score for each token. Then, the score of each token at the t -th step is obtained by

$$\mathbf{S}^{(t)} = \mathbf{S}^{(0)} (\mathbf{A}')^t. \quad (3)$$

Model	Method	Pruning Ratio	General VQA		MLLM benchmarks			Text-oriented VQA			Throughput (samples/sec)
			GQA	VQA2.0	MME	POPE	MMB-EN	ChartQA	DocVQA	TextVQA	
LLaVA-NeXT-8B	Baseline	0%	65.38	82.70	1587.72	87.84	72.08	69.28	78.22	65.41	2.49
	Random	50%	64.95	81.61	1605.43	86.47	70.27	56.00	61.54	58.21	4.03(1.62×)
		70%	64.22	80.17	1576.33	84.98	69.42	44.48	48.42	49.48	5.23(2.10×)
		90%	60.55	74.23	1475.07	79.63	61.77	26.72	27.59	31.73	6.92(2.78×)
	ToMe	50%	65.07	81.82	1566.60	87.56	70.88	68.20	73.53	59.07	1.00(0.40×)
		70%	64.07	80.56	1564.36	87.33	68.21	62.88	65.28	52.19	1.11(0.45×)
		90%	59.72	76.31	1453.13	84.29	61.77	41.04	41.04	38.36	1.21(0.48×)
	FastV	50%	65.11	82.51	1604.14	87.51	71.82	67.60	73.92	65.15	3.97(1.59×)
		70%	64.34	81.83	1607.83	87.08	68.05	62.80	66.67	63.08	4.82(1.94×)
		90%	60.20	77.21	1488.16	83.01	68.30	39.56	42.57	53.53	5.90(2.37×)
	G-Prune(Ours)	50%	65.34	82.54	1623.81	87.76	71.91	67.72	75.62	65.17	4.03(1.62×)
		70%	64.37	81.91	1609.87	87.69	70.19	65.04	70.72	63.87	5.21(2.09×)
		90%	61.40	77.51	1456.14	84.49	67.27	51.72	48.94	59.31	6.91(2.77×)

Table 1: Comparison with SOTA methods under different FLOPs ratios for three strategies of benchmarks. The best results for each FLOPs ratios are marked in **bold**.




Figure 3: Comparison between our G-Prune and other compression methods for the LLaVA-NeXT model tested on the TextVQA, DocVQA, POPE and GQA benchmarks.

Each object has a token that best represents it, and we gather scores for all tokens belonging to that object. However, since objects contain different numbers of tokens, their representative tokens can vary significantly. This variation makes it challenging to directly identify the most representative token. To mitigate the influence of connected components size, we apply the degree of each token to normalize the scores. The degree $\mathbf{D} \in \mathbb{R}^N$ of tokens can be calculated by

$$\mathbf{D}_i = \sum_{j=1}^N \mathbf{1}(A_{ij} > 0), \quad (4)$$

where $\mathbf{1}(\cdot)$ indicates an indicator function, which returns 1 when the condition is met. Subsequently, the normalized score of each token can be defined as

$$\mathbf{S}'_i^{(t)} = \mathbf{S}_i^{(t)} / \mathbf{D}_i. \quad (5)$$

Here, $\mathbf{S}'_i^{(t)}$ is the normalized representativeness of each token in its object. Thanks to the normalization operation, we can use a unified standard to directly obtain the most representative token of each object. Finally, the index of the retained

tokens can be selected by

$$I_k = \arg \max_{I \subset \{1, 2, \dots, n\} | |I|=k} \sum_{i \in I} \mathbf{S}'_i^{(t)}. \quad (6)$$

We find the top k tokens by token score $\mathbf{S}'_i^{(t)}$. In this way, we can retain the objects from both the foreground and the background, while only the most representative token is retained to represent each object. The detailed algorithm of G-Prune is described in Alg. 1.

Experiments

Datasets and Metrics

We first evaluate G-Prune for two common general benchmarks, *i.e.*, VQA2.0 (Goyal et al. 2017), GQA (Hudson and Manning 2019). Furthermore, for fine-grained tasks, we evaluate OncePrue for text-oriented VQA benchmarks, including TextVQA (Singh et al. 2019), DocVQA (Mathew et al. 2020) and ChartQA (Masry et al. 2022). We also evaluate G-Prune on three emerging multimodal benchmarks for MLLMs, including POPE (Li et al. 2023b), MME (Fu et al. 2023) and MMB (Liu et al. 2023). These benchmarks tend

Algorithm 1: G-Prune: Token Reduction via Graph-based Information Propagation

Input: Visual tokens $\mathbf{X} \in \mathbb{R}^{N \times d}$, Similarity threshold s , Number of tokens to retain k , Number of iterations t

Output: Indices of selected representative tokens I_k

```

1: Construct  $\mathbf{A} \in \mathbb{R}^{N \times N}$  according to Eq. 1:
2: for  $i, j = 1$  to  $N$  do
3:    $\mathbf{A}_{ij} \leftarrow \max(0, \cos(\mathbf{X}_i, \mathbf{X}_j) - s)$ 
4: end for
5: Iterate information propagation:
6: for  $t = 1$  to  $t$  do
7:    $\mathbf{S}^{(t)} \leftarrow \mathbf{S}^{(t-1)} \cdot (\mathbf{A}^t)$ 
8: end for
9: Calculate degree of each token  $\mathbf{D} \in \mathbb{R}^N$ :
10: for  $i = 1$  to  $N$  do
11:    $\mathbf{D}_i \leftarrow \sum_{j=1}^N \mathbf{1}(\mathbf{A}_{ij} > 0)$ 
12: end for
13: Normalize token scores  $\mathbf{S}'^{(t)} \in \mathbb{R}^N$ :
14: for  $i = 1$  to  $N$  do
15:    $\mathbf{S}'_i \leftarrow \mathbf{S}_i^{(t)} / \mathbf{D}_i$ 
16: end for
17: Select top- $k$  representative tokens:
18:  $I_k \leftarrow \arg \max_{I \subset \{1, 2, \dots, N\} | |I|=k} \sum_{i \in I} \mathbf{S}'_i^{(t)}$ 
19: return  $I_k$ 

```

l_2 -Norm	Graph	GQA	POPE	TextVQA	Avg.	Avg. TFLOPs
\times	\times	64.22	84.98	49.48	66.23	6.76
\checkmark	\times	64.30	87.48	53.63	68.47	6.76
\times	\checkmark	64.23	87.46	64.01	71.90	6.76
\checkmark	\checkmark	64.37	87.69	64.05	72.04	6.76
Baseline		65.38	87.84	65.41	72.88	18.52

Table 2: Ablation study of each component in G-Prune.

to be more challenging than traditional vision-language evaluations, aiming to assess different aspects of MLLMs, like detailed reasoning and OCR.

Implementation Details

We introduce G-Prune, a training-free, plug-and-play method that seamlessly integrates into existing MLLMs. We implement G-Prune on widely-used open-source MLLM, LLaVA-NeXT-8B (Liu et al. 2024a) following its default settings. We use the evaluation toolkit LMMs-Eval (Li et al. 2024) to evaluate the performance of these MLLMs across various datasets. For specific implementations, we configure 5 iterations with a similarity threshold of 0.5.

Quantitative Experiments

Comparison with state-of-the-art methods. In Tab. 1, we compare the effect of G-Prune with existing pruning methods on LLaVA-NeXT, including ToMe (Bolya et al. 2022), FastV (Chen et al. 2024a), and the random pruning. From the Tab. 1, we can first observe that, overall, our method achieves significant advantages in the above three types of data sets.

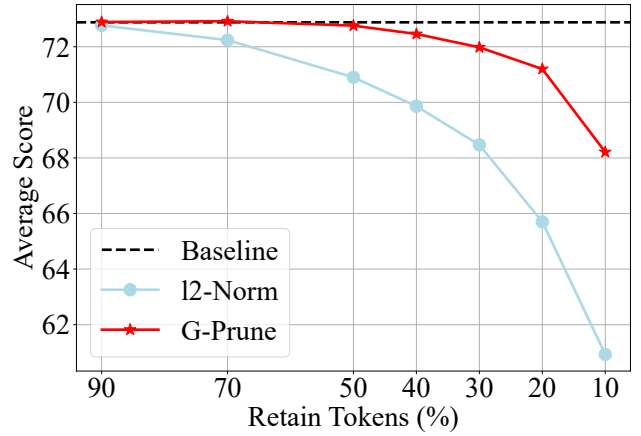


Figure 4: Comparison between our G-Prune method and l_2 -Norm based pruning method. The result is based on the average performance across GQA, POPE and TextVQA.

Iteration t	Similarity s	GQA	POPE	TextVQA	Avg.
5	0.3	64.29	87.48	63.61	71.79
5	0.5	64.37	87.69	63.87	71.98
5	0.7	64.04	87.78	64.14	71.99
5	0.9	63.87	87.22	54.20	68.43
1	0.5	64.26	87.73	63.84	71.94
5	0.5	64.37	87.69	63.87	71.98
10	0.5	64.37	87.60	63.90	71.96
50	0.5	64.05	87.34	63.85	71.75

Table 3: Ablation study on the impact of hyper-parameters: Iteration num t and Similarity threshold s .

On the general VQA tasks, the proposed G-Prune method has significant advantages at all pruning ratios, *e.g.*, G-Prune improves performance by 0.88%-1.58% for VQA2.0 compared to ToMe. For the MLLM benchmarks, the proposed G-Prune improves performance on MME by 2.27% while using only 50% of the visual tokens. On the POPE dataset, our method consistently outperforms the baseline across all pruning ratios. As for text-oriented VQA tasks, the proposed G-Prune can better maintain the performance compared to other methods. When the pruning ratio is increased to 90%, we can observe that the performance of previous methods, *e.g.*, ToMe, has a significant decline, especially on text-oriented VQA benchmarks, *i.e.*, the performance drops by 40.35% on TextVQA and 26.02% on ChartVQA. In terms of ToMe, its complex indexing leads to large latency in MLLMs. In contrast, the proposed G-Prune can better maintain the performance, *e.g.*, G-Prune improves the performance by 51.99% for TextVQA benchmark compared to ToMe when pruning 90% tokens. These results well validate the effectiveness of our G-Prune method in pruning visual tokens in the MLLM.

For a more fine-grained comparison, we plot the computational cost-performance curve in Fig. 3. From the curve, we can see that in the initial stage of reducing computa-

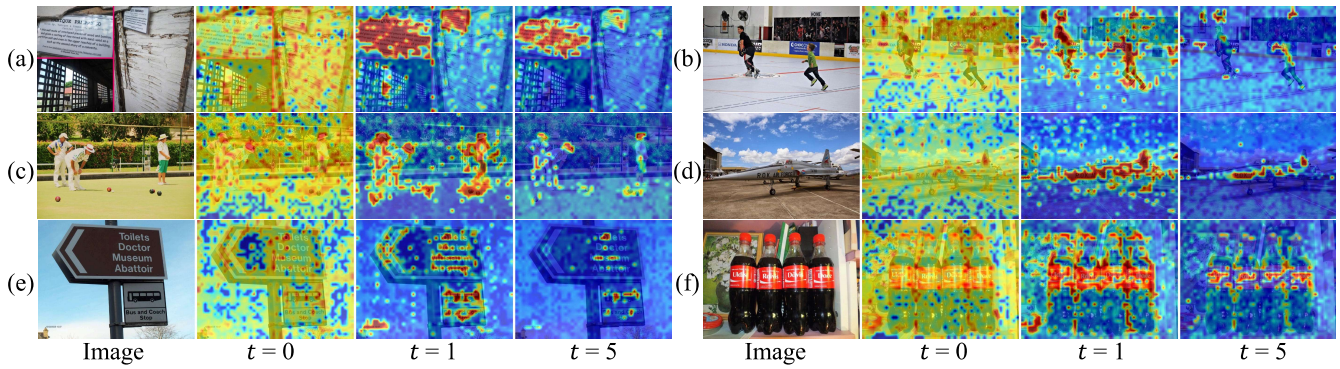


Figure 5: The visualization of information propagation in G-Prune across different iteration num t . The heatmaps demonstrates the score of each visual token on LLaVA-NeXT for TextVQA.

tional overhead, each method has no significant impact on the performance of the model on all benchmarks. As the more FLOPs are reduced, the performance of the model begins to decline rapidly. Specifically, on POPE benchmark, when reducing FLOPs from 1.83×10^{13} to 6.69×10^{12} , there is only 0.09% performance decreased. While reducing FLOPs from 1.83×10^{13} to 3.37×10^{12} , the performance significantly drops by 4.11%. This phenomenon shows that there is obvious redundancy in the visual tokens of MLLM. Therefore, in the initial stage, all methods can achieve token pruning without performance loss. However, as the number of pruning increases, tokens must be selected more carefully to achieve competitive performance. Meanwhile, we can observe that the proposed G-Prune method maintains the great advantages over other methods under different FLOPs reductions. For instance, when reducing 1.53×10^{13} FLOPs for TextVQA benchmark, the proposed G-Prune improves performance by 52.05% compared to ToMe. As for reducing 1.41×10^{13} FLOPs, the proposed G-Prune keep 46.67% advantage than FastV for DocVQA benchmark. These results better validating the advantages of G-Prune for token pruning task in MLLM.

Comparison with Alternatives. In Fig. 4, we compare our G-Prune with the alternative using l_2 -Norm as the metric for LLaVA-NeXT on GQA, POPE and TextVQA benchmarks. The curve reflects the average performance of GQA, POPE and TextVQA under different pruning ratios. Specifically, *baseline* denotes the default MLLM without token pruning. “ l_2 -Norm” indicates retaining the tokens with the largest l_2 -Norm value according to the percentage.

As shown in Fig. 4, it is apparent that when the retention tokens exceed 50%, G-Prune consistently maintains stable performance compared to the baseline method, *i.e.*, G-Prune still retains 99.84% of the original performance. At the same time, “ l_2 -Norm” method drops performance by 2.56%. Based on this, we can find that in the complex scenarios that MLLM needs to deal with, it is impossible to achieve token pruning by relying solely on the amount of information. Moreover, across all evaluated pruning ratios, G-Prune performs better than the method according to l_2 -Norm only. This advantage grows more pronounced as the pruning ratio increases.

For instance, when the pruning ratio is increased to 90%, our G-Prune can improve the performance by 7.28% compared to “ l_2 -Norm” method. It fully shows that it is effective to measure the value of tokens by constructing a similarity graph between tokens in the visual token pruning task. These experiment results fully demonstrate that the proposed graph-based information propagation method in our G-Prune can effectively help us better perform visual token pruning in MLLMs.

Ablation Study. Further, we conduct ablation experiments for each component in Tab. 2 to explore the contributions and impacts of different components of our method. From Tab. 2, we can first observe that when randomly pruning 30% tokens, *i.e.*, the first line in Tab. 2, there is a significant performance drop, *e.g.*, the performance drops by 22.75% for TextVQA dataset. Then, we apply l_2 -Norm to measure the value of each token, and directly prune the tokens with lower l_2 -norm, this approach primarily preserves elements with higher information content (Lu and Zhang 2022), resulting in an average performance of 68.47. Conversely, we employ information propagation by constructing graph, and initial the score of each token to 1. With the better consideration to objects from both foreground and background, this approach improves the performance by 8.56% compared to l_2 -Norm only method. Finally, when applying both two methods, *i.e.*, our G-Prune method, the average performance is improved by 8.77% compared to the random pruning method. These experiment results not only validate the effectiveness of each component in our G-Prune, but also show that every object from the foreground and background should be fully considered during the token pruning of MLLM.

In Tab. 3 we conduct experiment based on the different hyper-parameter settings. We first fix the number of iteration turn t in Eq.3. When the similarity threshold s in Eq.1 gradually increased from 0.3 to 0.9, we can observe that the average performance gradually increase to 2.78% when $s = 0.7$. With the further increment of the similarity threshold s , the performance is dropped by 4.95%. The reason for this phenomenon is that a too low similarity threshold s will cause all regions to be regarded as the same object, while a too high s will cause an object to be split into multiple

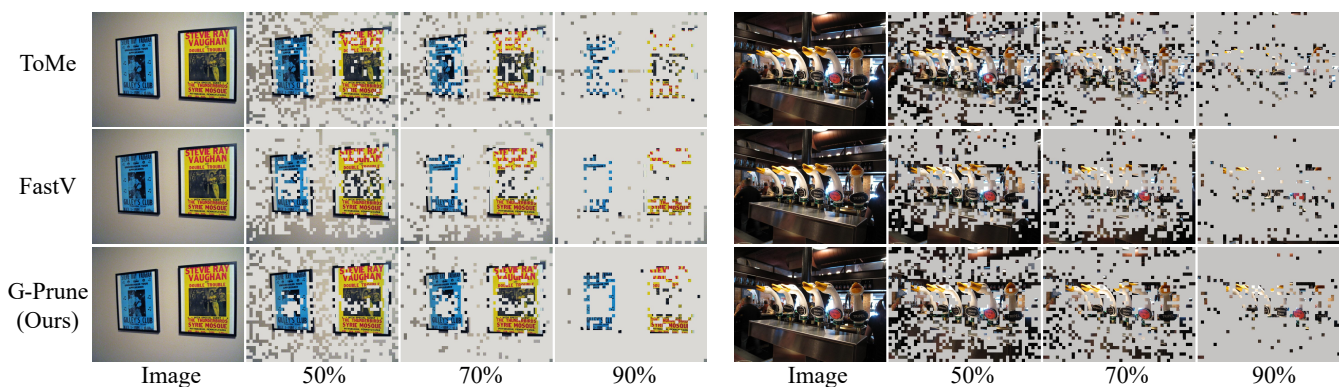


Figure 6: The comparative visualization of ToMe, FastV, and G-Prune on LLaVA-NeXT. G-Prune effectively retains tokens representative of regions with high information content, giving it a significant advantage in fine-grained tasks.

objects unnecessarily. When fix the similarity threshold s to 0.5, we can observe that the best performance occurred with 5 iterations, *i.e.*, the average performance is achieved by 71.98. It is worth noting that the performance difference of our G-Prune is very limited under different iteration numbers. Specifically, the difference between the highest and the lowest is only 0.32%. It is because that the propagation of information will reach equilibrium after a certain number of rounds, *i.e.*, the amount of information flowing in and out of each token will have no change. The above experiment not only validate the effectiveness of our G-Prune method, but also shows that our G-Prune method has good robustness.

Qualitative Experiments

To further delve into the intrinsic mechanisms of our method, we visualize the information propagation in Fig. 5, the heatmap represents the score of tokens. The tokens in the red areas have higher scores. In the initial stage, multiple tokens have high scores for the same object. This also shows that there is a lot of redundant information. As the information continues to propagate, we can find that the information gradually gathers in several tokens from both foreground and background. This shows that our G-Prune can effectively retain representative tokens for objects. These experiment results clearly prove the effectiveness of our G-Prune in visual token pruning. Then we visualize the results on the LLaVA-NeXT using the TextVQA dataset, pruning 50%, 70%, and 90% of the tokens to explore the pruning mechanisms of our method. We also compare our G-Prune method with ToMe and FastV. As shown in Fig. 6, we first focus on the performance of our method under different pruning ratios. When keep 50% retain tokens, the foreground occupies a large proportion, while the background area is preserved uniformly. This phenomenon shows that our G-Prune can effectively preserve information for each area. With the decrease of retain tokens, we can observe that more tokens will be retained in areas with complex textures, whether from the background or the foreground. At the same time, in each area with high similarity, a certain number of tokens will be retained. Subsequently, we conducted a visual comparison

of our method with others. For instance, in the left image, even with a pruning ratio of 50%, ToMe and FastV begin to lose fine-grained details, such as text on a wall image. This situation is even more serious at a 90% pruning ratio. Our method not only preserves more detailed information but also significantly reduces redundancy by targeting areas with higher similarity. On the other hand, from the right image, we can find out that the visual tokens from the background are completely pruned. This may cause MLLM to be unable to solve some tasks involving background information. Above all, with the decrease of retain tokens, we can observe that more tokens will be retained in areas with complex textures, whether from the background or the foreground. At the same time, in each area with high similarity, a certain number of tokens will be retained. Under this paradigm, our G-Prune method can better preserve the information in the original image. These visualizations prove that G-Prune not only effectively maintains representative visual tokens, but also well alleviates redundancy.

Conclusion

In this paper, we introduced a novel, training-free visual token pruning method for multimodal large language models (MLLMs) named G-Prune. Our G-Prune method does not need to intervene in the reasoning process of MLLM, saving a lot of computational overhead. Meanwhile, the proposed G-Prune method considers objects in the foreground and background at the same time, fully meeting the complex problems that MLLM needs to deal with. Specifically, our approach constructs a graph based on token similarities, iteratively propagating information to identify and retain the most representative tokens for objects from both foreground and background. Extensive experiments across various benchmarks prove superiority of G-Prune over existing pruning methods, maintaining a great balance between computational efficiency and performance. For instance, G-Prune achieves a 63.57% FLOPs reduction on LLaVA-NeXT for TextVQA while only incurring a 2% decrease in accuracy. These findings suggest that G-Prune significantly advances the efficiency and scalability of MLLMs.

Acknowledgments

This work was supported by the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No. U21B2037, No. U22B2051, No. U23A20383, No. U21A20472, No. 62176222, No. 62176223, No. 62176226, No. 62072386, No. 62072387, No. 62072389, No. 62002305 and No. 62272401), and the Natural Science Foundation of Fujian Province of China (No. 2021J06003, No.2022J06001).

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023a. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023b. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Bolya, D.; Fu, C.-Y.; Dai, X.; Zhang, P.; Feichtenhofer, C.; and Hoffman, J. 2022. Token merging: Your vit but faster. *arXiv preprint arXiv:2210.09461*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Cao, J.; Ye, P.; Li, S.; Yu, C.; Tang, Y.; Lu, J.; and Chen, T. 2024. MADTP: Multimodal Alignment-Guided Dynamic Token Pruning for Accelerating Vision-Language Transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15710–15719.
- Cao, Q.; Paranjape, B.; and Hajishirzi, H. 2023. PuMer: Pruning and merging tokens for efficient vision language models. *arXiv preprint arXiv:2305.17530*.
- Chen, L.; Zhao, H.; Liu, T.; Bai, S.; Lin, J.; Zhou, C.; and Chang, B. 2024a. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. *arXiv preprint arXiv:2403.06764*.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24185–24198.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Fu, C.; Chen, P.; Shen, Y.; Qin, Y.; Zhang, M.; Lin, X.; Yang, J.; Zheng, X.; Li, K.; Sun, X.; Wu, Y.; and Ji, R. 2023. MME: A Comprehensive Evaluation Benchmark for Multimodal Large Language Models. *arXiv preprint arXiv:2306.13394*.
- Goyal, Y.; Khot, T.; Summers-Stay, D.; Batra, D.; and Parikh, D. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6904–6913.
- Guo, D.; Zhu, Q.; Yang, D.; Xie, Z.; Dong, K.; Zhang, W.; Chen, G.; Bi, X.; Wu, Y.; Li, Y.; et al. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming—The Rise of Code Intelligence. *arXiv preprint arXiv:2401.14196*.
- Hudson, D. A.; and Manning, C. D. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6700–6709.
- Li, B.; Zhang, P.; Zhang, K.; Pu, F.; Du, X.; Dong, Y.; Liu, H.; Zhang, Y.; Zhang, G.; Li, C.; et al. 2024. Lmms-eval: Accelerating the development of large multimodal models.
- Li, J.; Li, D.; Savarese, S.; and Hoi, S. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, 19730–19742. PMLR.
- Li, Y.; Du, Y.; Zhou, K.; Wang, J.; Zhao, W. X.; and Wen, J.-R. 2023b. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Liang, Y.; Ge, C.; Tong, Z.; Song, Y.; Wang, J.; and Xie, P. 2022. Not all patches are what you need: Expediting vision transformers via token reorganizations. *arXiv preprint arXiv:2202.07800*.
- Lin, Z.; Lin, M.; Lin, L.; and Ji, R. 2024. Boosting Multimodal Large Language Models with Visual Tokens Withdrawal for Rapid Inference. *arXiv preprint arXiv:2405.05803*.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024a. Llava-next: Improved reasoning, ocr, and world knowledge.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2024b. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Liu, Y.; Duan, H.; Zhang, Y.; Li, B.; Zhang, S.; Zhao, W.; Yuan, Y.; Wang, J.; He, C.; Liu, Z.; et al. 2023. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*.
- Long, S.; Zhao, Z.; Pi, J.; Wang, S.; and Wang, J. 2023. Beyond attentive tokens: Incorporating token importance and diversity for efficient vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10334–10343.
- Lu, Y.; and Zhang, J. 2022. Norm-based noisy corpora filtering and refurbishing in neural machine translation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 5414–5425.

- Luo, G.; Zhou, Y.; Ren, T.; Chen, S.; Sun, X.; and Ji, R. 2024. Cheap and quick: Efficient vision-language instruction tuning for large language models. *Advances in Neural Information Processing Systems*, 36.
- Masry, A.; Long, D. X.; Tan, J. Q.; Joty, S.; and Hoque, E. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*.
- Mathew, M.; Karatzas, D.; Manmatha, R.; and Jawahar, C. V. 2020. DocVQA: A Dataset for VQA on Document Images. *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*.
- Meta, A. 2024. Introducing meta llama 3: The most capable openly available llm to date. *Meta AI*.
- Pan, B.; Panda, R.; Jiang, Y.; Wang, Z.; Feris, R.; and Oliva, A. 2021. IA-RED2: Interpretability-Aware Redundancy Reduction for Vision Transformers. *Advances in Neural Information Processing Systems*, 34: 24898–24911.
- Rao, Y.; Zhao, W.; Liu, B.; Lu, J.; Zhou, J.; and Hsieh, C.-J. 2021. Dynamicvit: Efficient vision transformers with dynamic token sparsification. *Advances in neural information processing systems*, 34: 13937–13949.
- Reid, M.; Savinov, N.; Teplyashin, D.; Lepikhin, D.; Lillcrap, T.; Alayrac, J.-b.; Soricut, R.; Lazaridou, A.; Firat, O.; Schrittwieser, J.; et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Singh, A.; Natarajan, V.; Shah, M.; Jiang, Y.; Chen, X.; Batra, D.; Parikh, D.; and Rohrbach, M. 2019. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8317–8326.
- Team, G.; Anil, R.; Borgeaud, S.; Wu, Y.; Alayrac, J.-B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A. M.; Hauth, A.; et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Vaswani, A. 2017. Attention is All You Need. *arXiv preprint arXiv:1706.03762*.
- Wang, H.; Dedhia, B.; and Jha, N. K. 2024. Zero-TPrune: Zero-shot token pruning through leveraging of the attention graph in pre-trained transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16070–16079.
- Xu, Y.; Zhang, Z.; Zhang, M.; Sheng, K.; Li, K.; Dong, W.; Zhang, L.; Xu, C.; and Sun, X. 2022. Evo-vit: Slow-fast token evolution for dynamic vision transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2964–2972.
- Yang, J.; Li, C.; Dai, X.; and Gao, J. 2022. Focal modulation networks. *Advances in Neural Information Processing Systems*, 35: 4203–4217.
- Yao, L.; Li, L.; Ren, S.; Wang, L.; Liu, Y.; Sun, X.; and Hou, L. 2024. DeCo: Decoupling Token Compression from Semantic Abstraction in Multimodal Large Language Models. *arXiv preprint arXiv:2405.20985*.
- Zhang, P.; Dong, X.; Zang, Y.; Cao, Y.; Qian, R.; Chen, L.; Guo, Q.; Duan, H.; Wang, B.; Ouyang, L.; Zhang, S.; Zhang, W.; Li, Y.; Gao, Y.; Sun, P.; Zhang, X.; Li, W.; Li, J.; Wang, W.; Yan, H.; He, C.; Zhang, X.; Chen, K.; Dai, J.; Qiao, Y.; Lin, D.; and Wang, J. 2024. InternLM-XComposer-2.5: A Versatile Large Vision Language Model Supporting Long-Contextual Input and Output. *arXiv preprint arXiv:2407.03320*.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.