

D²Pruner: Debaised Importance and Structural Diversity for MLLM Token Pruning

Evelyn Zhang^{1,2*}, Fufu Yu¹, Aoqi Wu³, Zichen Wen², Ke Yan¹, Shouhong Ding¹, Biqing Qi⁴, Linfeng Zhang^{2†}

¹Tencent YouTu Lab

²School of Artificial Intelligence, Shanghai Jiao Tong University

³School of Computer Science and Technology, Tongji University

⁴Shanghai AI Laboratory
zhanglinfeng@sjtu.edu.cn

Abstract

Processing long visual token sequences poses a significant computational burden on Multimodal Large Language Models (MLLMs). While token pruning offers a path to acceleration, we find that current methods, while adequate for general understanding, catastrophically fail on fine-grained localization tasks. We attribute this failure to the inherent flaws of the two prevailing strategies: importance-based methods suffer from a strong positional bias, an inherent model artifact that distracts from semantic content, while diversity-based methods exhibit structural blindness, disregarding the user’s prompt and spatial redundancy. To address this, we introduce D²Pruner, a framework that rectifies these issues by uniquely combining debaised importance with a structural pruning mechanism. Our method first secures a core set of the most critical tokens as pivots based on a debaised attention score. It then performs a Maximal Independent Set (MIS) selection on the remaining tokens, which are modeled on a hybrid graph where edges signify spatial proximity and semantic similarity. This process iteratively preserves the most important and available token while removing its neighbors, ensuring that the supplementary tokens are chosen to maximize importance and diversity. Extensive experiments demonstrate that D²Pruner achieves exceptional efficiency and fidelity.

Code —

<https://github.com/EvelynZhang-epiclab/D2Pruner>

Introduction

The computational demands of Multimodal Large Language Models (MLLMs) severely constrain their visual understanding capabilities. The quadratic complexity of self-attention and massive KV cache memory impose a critical bottleneck, particularly with high-resolution inputs, hindering efficient inference and making acceleration imperative.

To mitigate this, token pruning has emerged as a promising strategy for reducing the input sequence length. While

*This work was done during Evelyn Zhang’s internship at Tencent YouTu Lab.

†corresponding author

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

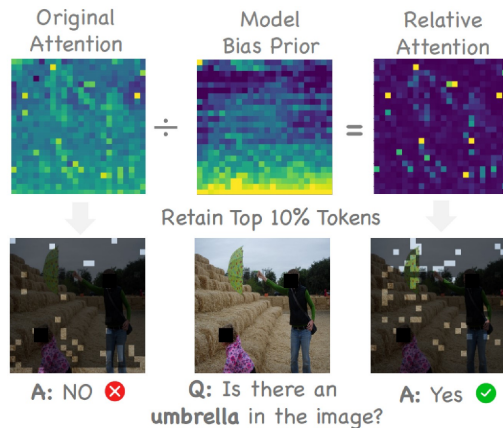


Figure 1: Illustration of our debiasing solution.

existing methods show moderate success on general understanding and reasoning tasks, they catastrophically fail on fine-grained tasks that demand precise spatial awareness and localization. Our investigation pinpoints two critical, previously overlooked failure modes in these approaches:

Position Bias in Importance-based Methods. These methods (Chen et al. 2024b) estimate the importance of visual tokens by leveraging the attention scores assigned to them by the last token. However, such scores are inherently vulnerable to position bias, which can distort the true importance distribution. Prior studies (Wen et al. 2025a; Zhang et al. 2024a) have reported that LLaVA consistently exhibits a content-agnostic emphasis on the bottom region of the image. Building on this observation, we conduct a broader analysis across several representative MLLMs and reveal that position bias is not limited to LLaVA but manifests in diverse forms. Specifically, we average the attention maps of over 1,000 randomly sampled images from the COCO dataset (Lin et al. 2014) to uncover these patterns. As highlighted in Fig.1, model bias causes the naive attention-based method (e.g., FastV) to fail by pruning salient objects like the umbrella. Our relative attention, which normalizes by this model bias prior, rectifies the pruning decision.

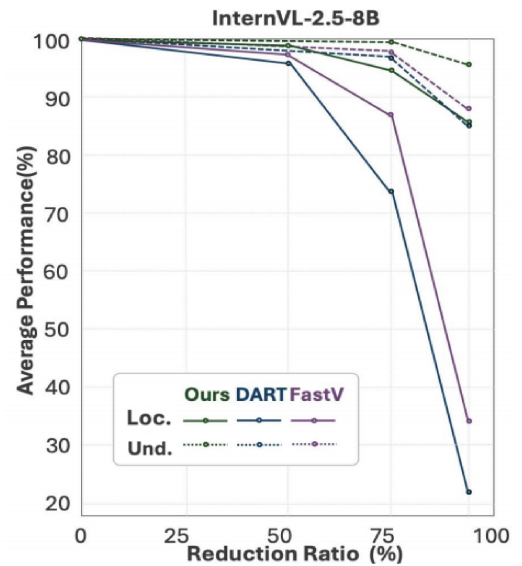
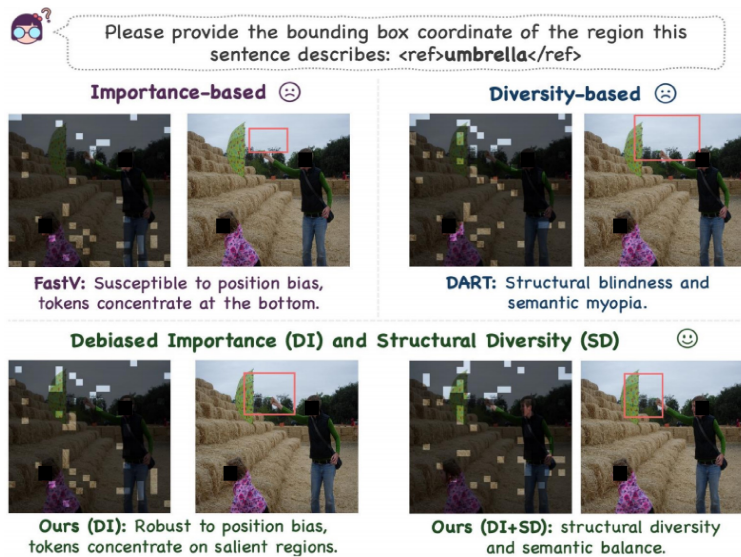


Figure 2: **Qualitative and Quantitative Analysis of D²Pruner against Prior Methods.** *Left:* Visual comparison on localization tasks. Importance-based method, FastV, is misled by positional bias, concentrating tokens on the image bottom, while Diversity-based method, DART, suffers from structural blindness and semantic myopia. In contrast, our Debiased Importance (DI) component successfully overcomes positional bias. Our full model, D²Pruner (DI+SD), jointly addresses both semantic and spatial redundancy, resulting in a more balanced and informative token selection that enhances localization performance. *Right:* Quantitative evaluation of average performance on general understanding and fine-grained localization tasks.

Structural Blindness in Diversity-based Methods. Previous methods (Wen et al. 2025b; Alvar et al. 2025) prune tokens by assessing diversity purely in the feature space, typically via cosine similarity. This strategy operates on a flawed premise: that visual information can be processed as a one-dimensional sequence, after flattening the original 2D grid. This action strips away all topological information, forcing the model to treat tokens as an unstructured set and rendering it blind to their crucial spatial relationships. This motivates our approach to adopt a representation that explicitly preserves the 2D structure of images by modeling tokens and their semantic-spatial relationships as a graph.

To address these challenges, we introduce **D²Pruner**, a pruning framework built on the principles of **Debiased Importance (DI)** and **Structural Diversity (SD)**. It first secures a set of the most critical tokens as pivots based on the debiased attention score. It then performs a greedy Maximal Independent Set (MIS) selection on the remaining tokens, which are modeled on a hybrid graph where edges signify both spatial proximity and semantic similarity. This process iteratively preserves the most important available token while pruning its neighbors, ensuring the supplementary tokens are chosen to maximize importance and structural diversity.

Leveraging its dual-stage design, D²Pruner not only excels on general understanding tasks but demonstrates unparalleled strength in fine-grained localization, as illustrated in Fig.2. On challenging localization benchmarks (Kazemzadeh et al. 2014; Mao et al. 2016), our method maintains 85.7% performance on InternVL-2.5-8B (Chen et al. 2024c) even at a 90% pruning rate—surpassing prior diversity-based and importance-

based methods by a massive 63.5 and 51.2 points, respectively. This is complemented by strong general performance, retaining 95.5% of LLaVA-Next-7B’s capabilities. Critically, these results are achieved with remarkable efficiency: D²Pruner delivers a 5× prefill acceleration and a 7.6× reduction in KV cache size, while the pruning process itself adds a negligible overhead of just 0.8% to the forward pass. In summary, our contributions are as follows.

- We systematically analyze position bias in MLLMs, characterize its patterns, and propose an effective mitigation method for token importance estimation.
- We introduce D²Pruner, a unified framework that models spatial-semantic token relationships via hybrid graphs to enable informed pruning decisions.
- Extensive validation shows D²Pruner achieves superior performance with significant efficiency gains, particularly excelling on fine-grained localization tasks.

Related Work

Multimodal Large Language Models

The dominant architecture for Multimodal Large Language Models (MLLMs), pioneered by LLaVA (Liu et al. 2023), connects a pre-trained vision encoder to a Large Language Model (LLM) via a lightweight projection module. This paradigm has spurred rapid advancements, with successors like LLaVA-NeXT (Liu et al. 2024b), Qwen-VL (Bai et al. 2023), and InternVL (Chen et al. 2024d) pushing the state-of-the-art. These models achieve superior performance by leveraging more powerful foundation models, scaling up

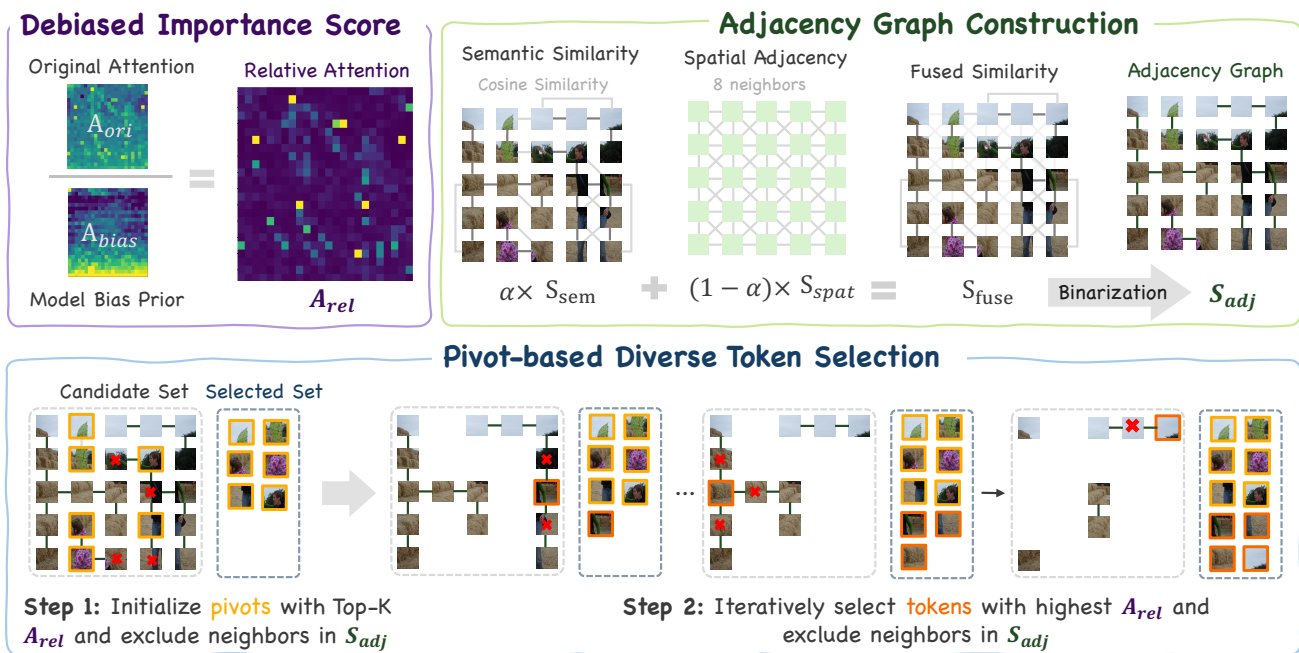


Figure 3: **Overall pipeline for our proposed D²Pruner.** (a) *Debiased Importance Score*: The original attention map A_{ori} is adjusted by subtracting the model bias prior A_{bias} , resulting in the relative attention map A_{rel} . (b) *Adjacency Graph Construction*: The semantic similarity and spatial adjacency are fused to create the adjacency graph S_{adj} . (c) *Pivot-based Diverse Token Selection*: In Step 1, pivots are initialized using the top-K values from A_{rel} , excluding neighbors as defined by S_{adj} ; In Step 2, tokens are iteratively selected based on the highest A_{rel} values, with neighbors excluded according to the adjacency graph.

training data, and enhancing their capability to process high-resolution images. However, pursuing higher performance via high-resolution inputs exacts a heavy computational toll. Visual tokens increase quadratically with resolution, creating a severe bottleneck in the LLM’s self-attention layers.

Visual Token Compression

Visual token pruning, aiming to ease computational burdens, primarily employs either importance-based or diversity-based methods to reduce token count.

(1) Importance-Based Pruning. Early and prominent methods operate on the principle of importance-based pruning, where tokens with the highest scores are retained. Some methods (Chen et al. 2024b; Zhang et al. 2024b; Xing et al. 2025) leverage the attention scores assigned to visual tokens by the last token within the language model to evaluate the importance of each visual token. However, such methods suffer from positional bias, often favoring tokens at the bottom of image regardless of content (Wen et al. 2025a; Zhang et al. 2024a). Other methods (Shang et al. 2024; Yang et al. 2025; Zhang et al. 2024a) rely on the attention scores assigned to visual tokens by class token in the vision encoder, suffering from instruction-agnostic token selection.

(2) Diversity-Based Pruning. Recent studies have moved beyond importance-based token pruning, highlighting the benefits of reducing redundancy through diversity-based strategies. Approaches such as DivPrune (Alvar et al. 2025) formulate token selection as a diversity maximization prob-

lem, ensuring that the retained tokens are highly representative of the original set. Similarly, DART (Wen et al. 2025b) demonstrates that importance scores may not be reliable indicators for pruning and instead proposes a duplication-aware method, which selects tokens with minimal overlap to a set of pivots. But, these methods only consider semantic redundancy and often exhibit *structural blindness* by ignoring spatial relationships among tokens. They also fail to incorporate user instructions, limiting their adaptability.

Methodology

In this section, we present **D²Pruner**, a training-free framework designed to accelerate the inference of MLLMs. D²Prune is built on two key principles: debiased importance and structural diversity, as illustrated in Fig. 3.

Preliminary and Problem Formulation

An MLLM typically comprises a vision encoder \mathcal{E}_v and a Large Language Model (LLM) \mathcal{L} . Given an image I , \mathcal{E}_v produces a sequence of N visual tokens, $V = \{v_1, \dots, v_N\}$, where each $v_i \in \mathbb{R}^d$. These are prepended to M text tokens T to form the input sequence $X = [V; T]$ for the LLM. The self-attention mechanism’s quadratic complexity, $O((N + M)^2)$, makes the large number of visual tokens a primary computational bottleneck.

Token pruning aims to mitigate this cost by selecting an informative subset of visual tokens. We perform this pruning dynamically within the LLM’s architecture. Let $H^{(k)}$

be the sequence of hidden states at the output of the k -th transformer layer. We apply our pruner at a pre-determined depth, after layer $K - 1$. The pruner operates on the visual component of the hidden states, $H_V^{(K-1)}$, to produce a condensed set $H'_V{}^{(K-1)}$ of size $n = r \cdot N$ for a given pruning ratio r . This pruned set is then concatenated with the untouched text hidden states, $H_T^{(K-1)}$, to form the input $H^{(K-1)} = [H'_V{}^{(K-1)}; H_T^{(K-1)}]$ for all subsequent layers $k \geq K$. We aim to select a subset of hidden states $H_V^{(K-1)} = \{h_1, \dots, h_N\}$ that is maximally informative and representative. We define our objective as finding the selection mask \mathbf{s}^* that maximizes a composite score function \mathcal{F} :

$$\mathbf{s}^* = \arg \max_{\mathbf{s}} \mathcal{F}(H_V^{(K-1)}, \mathbf{s}) \quad \text{subject to} \quad \sum_{i=1}^N s_i = n \quad (1)$$

The score function \mathcal{F} is designed to explicitly reward both individual token importance and overall subset diversity:

$$\mathcal{F}(H_V^{(K-1)}, \mathbf{s}) = \underbrace{\sum_{i=1}^N s_i \cdot \mathcal{I}(h_i)}_{\text{Token Importance}} + \lambda \cdot \underbrace{\text{Div}(\{h_i | s_i = 1\})}_{\text{Subset Diversity}} \quad (2)$$

where $\mathcal{I}(h_i)$ is the debiased Importance score of token h_i , quantifying its intrinsic value. $\text{Div}(\cdot)$ is the structural diversity function, measuring the structural comprehensiveness and low redundancy of the selected subset. λ is a hyperparameter balancing the trade-off between importance and diversity. The following sections will detail our methods for computing $\mathcal{I}(h_i)$ and for designing an efficient selection algorithm that jointly optimizes this objective.

Debiased Importance Score

The first term in our objective function (Eq. 2), Total Importance, relies on computing a robust and reliable importance score $\mathcal{I}(h_i)$ for each visual token’s hidden state h_i . A naive reliance on raw attention scores is prone to architectural artifacts, such as positional biases, where tokens at certain positions (e.g., the end of a sequence) receive artificially high attention regardless of their semantic content. We leverage a novel *relative attention score*, which is specifically designed to disentangle true semantic salience from systemic positional bias. We define this score as follows:

First, for a given image-text input, we extract the raw attention weights directed from the final text token to each of the N visual tokens at layer $K - 1$. Let this vector of content-specific attention scores be $\mathcal{A}_{ori} \in \mathbb{R}^N$, where $\mathcal{A}_{ori}(i)$ reflects the model’s focus on visual token h_i for the current task. Second, to isolate the inherent positional bias, we conduct a one-time calibration to compute a positional bias prior, denoted as $\mathcal{A}_{bias} \in \mathbb{R}^N$. This prior is obtained by feeding the model 1,000 randomly selected images from the COCO dataset, paired with a generic, non-specific prompt (e.g., "Please describe the provided image."). The final attention map is then averaged over all 1,000 images. The resulting attention scores, $\mathcal{A}_{bias}(i)$, capture the model’s default attention distribution across the visual token positions when no meaningful visual information is present.

The final relative attention score, \mathcal{A}_{rel} , is then calculated by normalizing the content-specific attention by this positional bias prior:

$$\mathcal{A}_{rel} = \frac{\mathcal{A}_{ori}}{\mathcal{A}_{bias} + \epsilon} \quad (3)$$

where ϵ is a small constant (e.g., 10^{-7}) to ensure numerical stability. Our method also supports dynamic input resolutions. Specifically, when the number of visual tokens varies (e.g., due to different image resolutions or patch sizes), we resize the positional bias prior \mathcal{A}_{bias} to match the size of \mathcal{A}_{ori} using interpolation. This ensures \mathcal{A}_{rel} remains consistent and applicable across varying input resolutions.

Adjacency Graph Construction

Let the set of visual token features be denoted by $H_V = \{h_1, h_2, \dots, h_N\}$, where $h_i \in \mathbb{R}^D$ is the D -dimensional feature vector for the i -th token. These N tokens correspond to a spatial grid of patches of size $h \times w$, such that $N = h \times w$. Our goal is to compute an adjacency matrix $\mathcal{S} \in \{0, 1\}^{N \times N}$ that defines the connectivity of the graph.

To capture the content-based relationships between tokens, we compute the pairwise cosine similarity. The semantic similarity $S_{sem}(i, j)$ between two token features h_i and h_j is calculated as:

$$S_{sem}(i, j) = \frac{h_i \cdot h_j}{\|h_i\| \|h_j\|} \quad (4)$$

This computation is performed for all pairs of tokens, yielding a dense semantic similarity matrix $\mathbf{S}_{sem} \in \mathbb{R}^{N \times N}$. Then, we apply min-max normalization to the semantic similarity matrix for normalization:

$$\hat{\mathbf{S}}_{sem} = \frac{\mathbf{S}_{sem} - \min(\mathbf{S}_{sem})}{\max(\mathbf{S}_{sem}) - \min(\mathbf{S}_{sem})} \quad (5)$$

To re-introduce the spatial structure, we define a spatial proximity matrix \mathbf{S}_{spat} based on adjacency in the original image grid. An edge is considered to exist between two tokens if their corresponding patches are immediate neighbors. We employ an 8-connectivity rule. Formally, the spatial proximity matrix $\mathbf{S}_{spat} \in \{0, 1\}^{N \times N}$ is defined as:

$$S_{spat}(i, j) = \begin{cases} 1 & \text{if } j \text{ is in the 8-neighborhood of } i \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

This results in a sparse, binary matrix that exclusively encodes the local spatial topology.

Subsequently, the two normalized matrices are fused via a weighted sum to produce a final similarity matrix \mathbf{S}_{fused} :

$$\mathbf{S}_{fused} = \alpha \cdot \hat{\mathbf{S}}_{sem} + (1 - \alpha) \cdot \mathbf{S}_{spat} \quad (7)$$

where $\alpha \in [0, 1]$ is a hyperparameter that balances the contribution of semantic content versus spatial structure.

Finally, the fused similarity matrix \mathbf{S}_{fused} is converted into a binary adjacency matrix \mathcal{S} by applying a similarity threshold, θ_{sim} . An edge is created between two nodes if their fused similarity exceeds this threshold. The elements of the final adjacency matrix \mathcal{S} are defined as:

$$S(i, j) = \begin{cases} 1 & \text{if } S_{fused}(i, j) > \theta_{sim} \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

Method	GQA	MMB	MMB-CN	MME	POPE	SQA	VQA ^{V2}	VQA ^{Text}	VizWiz	Avg.
LLaVA-1.5-7B	<i>Upper Bound, 576 Tokens (100%)</i>									
Vanilla	61.9	64.7	58.1	1862	85.9	69.5	78.5	58.2	50.0	100%
LLaVA-1.5-7B	<i>Retain 192 Tokens (↓ 66.7%)</i>									
ToMe (ICLR23)	54.3	60.5	-	1563	72.4	65.2	68.0	52.1	-	88.5%
FastV (ECCV24)	52.7	61.2	57.0	1612	64.8	67.3	67.1	52.5	50.8	91.2%
HiRED (AAAI25)	58.7	62.8	54.7	1737	82.8	68.4	74.9	47.4	50.1	91.5%
FitPrune (AAAI25)	60.4	63.3	56.4	1831	83.4	67.8	-	57.4	50.9	98.2%
LLaVA-PruMerge (ICCV25)	54.3	59.6	52.9	1632	71.3	67.9	70.6	54.3	50.1	90.8%
SparseVLM (ICML25)	57.6	62.5	53.7	1721	83.6	69.1	75.6	56.1	50.5	96.3%
PDrop (CVPR25)	57.1	63.2	56.8	1766	82.3	68.8	75.1	56.1	51.1	96.7%
VisionZip (CVPR25)	59.3	63.0	-	1783	85.3	68.9	77.4	57.3	-	97.8%
DART (Arxiv25)	60.0	63.6	57.0	1856	82.8	69.8	76.7	57.4	51.2	98.8%
D ² Pruner (Ours)	60.8	64.3	57.1	1879	85.6	70.0	78.0	58.2	51.4	99.9%
LLaVA-1.5-7B	<i>Retain 128 Tokens (↓ 77.8%)</i>									
ToMe (ICLR23)	52.4	53.3	-	1343	62.8	59.6	63.0	49.1	-	80.4%
FastV (ECCV24)	49.6	56.1	56.4	1490	59.6	60.2	61.8	50.6	51.3	86.4%
HiRED (AAAI25)	57.2	61.5	53.6	1710	79.8	68.1	73.4	46.1	51.3	90.2%
FitPrune (AAAI25)	58.5	62.7	56.2	1776	77.9	68.0	-	55.7	51.7	96.4%
LLaVA-PruMerge (ICCV25)	53.3	58.1	51.7	1554	67.2	67.1	68.8	54.3	50.3	88.8%
SparseVLM (ICML25)	56.0	60.0	51.1	1696	80.5	67.1	73.8	54.9	51.4	93.8%
PDrop (CVPR25)	56.0	61.1	56.6	1644	82.3	68.3	72.9	55.1	51.0	95.1%
VisionZip (CVPR25)	57.6	62.0	-	1763	83.2	68.9	75.6	56.8	-	96.2%
DART (Arxiv25)	58.7	63.2	57.5	1840	80.1	69.1	75.9	56.4	51.7	98.0%
D ² Pruner (Ours)	60.3	63.7	56.5	1850	85.1	69.3	77.4	58.0	51.8	99.2%
LLaVA-1.5-7B	<i>Retain 64 Tokens (↓ 88.9%)</i>									
ToMe (ICLR23)	48.6	43.7	-	1138	52.5	50.0	57.1	45.3	-	70.1%
FastV (ECCV24)	46.1	48.0	52.7	1256	48.0	51.1	55.0	47.8	50.8	77.3%
HiRED (AAAI25)	54.6	60.2	51.4	1599	73.6	68.2	69.7	44.2	50.2	87.0%
FitPrune (AAAI25)	52.3	58.5	49.7	1556	60.9	68.0	-	51.2	51.1	87.8%
LLaVA-PruMerge (ICCV25)	51.9	55.3	49.1	1549	65.3	68.1	67.4	54.0	50.1	87.4%
SparseVLM (ICML25)	52.7	56.2	46.1	1505	75.1	62.2	68.2	51.8	50.1	84.6%
PDrop (CVPR25)	41.9	33.3	50.5	1092	55.9	68.6	69.2	45.9	50.7	78.1%
VisionZip (CVPR25)	55.1	60.1	-	1690	77.0	69.0	72.4	55.5	-	92.7%
DART (Arxiv25)	55.9	60.6	53.2	1765	73.9	69.8	72.4	54.4	51.6	93.7%
D ² Pruner (Ours)	57.9	61.9	55.6	1823	82.4	70.0	74.6	56.1	52.2	97.3%
LLaVA-Next-7B	<i>Upper Bound, 2880 Tokens (100%)</i>									
Vanilla	64.2	67.4	60.6	1851	86.5	70.1	81.8	64.9	57.6	100%
LLaVA-Next-7B	<i>Retain 320 Tokens (↓ 88.9%)</i>									
FastV (ECCV24)	55.9	61.6	51.9	1661	71.7	62.8	71.9	55.7	53.1	86.4%
HiRED (AAAI25)	59.3	64.2	55.9	1690	83.3	66.7	75.7	58.8	54.2	91.8%
LLaVA-PruMerge (ICCV25)	53.6	61.3	55.3	1534	60.8	66.4	69.7	50.6	54.0	79.9%
SparseVLM (ICML25)	56.1	60.6	54.5	1533	82.4	66.1	71.5	58.4	52.0	85.9%
PDrop (CVPR25)	56.4	63.4	56.2	1663	77.6	67.5	73.5	54.4	54.1	86.8%
FasterVLM (Arxiv24)	56.9	61.6	53.5	1701	83.6	66.5	74.0	56.5	52.6	89.8%
GlobalCom ² (Arxiv25)	57.1	61.8	53.4	1698	83.8	67.4	76.7	57.2	54.6	90.3%
DART (Arxiv25)	61.7	64.2	58.2	1710	84.1	68.4	75.7	58.7	56.1	93.9%
D ² Pruner (Ours)	62.3	64.9	57.2	1717	84.6	69.0	78.9	58.9	54.8	95.5%

Table 1: Performance comparisons on LLaVA-1.5-7B and LLaVA-Next-7B across several understanding benchmarks.

The resulting matrix \mathcal{S} represents the adjacency structure of the graph \mathcal{G} , where an edge signifies that two tokens are both semantically similar and spatially proximal.

Pivot-based Diverse Token Selection

To construct a token subset of size n that is both representative and non-redundant, we introduce a novel selection algorithm named Pivot-based Diverse Token Selection. Its core mechanism is inspired by greedy approaches to the Maximal Independent Set (MIS) problem: it iteratively selects a high-value token and disqualifies its topological neighbors. This method leverages pre-computed importance scores $\mathcal{I}(h_i)$ to gauge the value of each token. Instead of merely selecting the top-scoring tokens, which could be highly redundant, our algorithm actively enforces diversity through a two-stage topological exclusion mechanism. First, based on a pre-

defined pivot ratio r_{pivot} , it selects $n_p = \lfloor n \cdot r_{pivot} \rfloor$ high-importance "pivots" with the highest scores, followed by an iterative expansion to choose the remaining $n - n_p$ tokens while excluding topological neighbors to ensure diversity.

Experiments

Experimental Settings

Models and Baselines. We apply D²Prune to four popular MLLMs with different architectures to evaluate its general effectiveness. Specifically, we follow previous work in this field to compare performance on LLaVA-1.5-7B (Liu et al. 2024a) and LLaVA-NeXT-7B (Liu et al. 2024b). Furthermore, we also present experimental results on the recent Qwen-2.5-VL-7B (Bai et al. 2025) and InternVL-2.5-8B (Chen et al. 2024c). We compare the performance of our approach with multiple token reduction methods:

Method	Qwen2.5-VL-7B							InternVL-2.5-8B								
	AI2D	ChartQA	DocVQA	MME	MMStar	POPE	TextVQA	Avg.	MME	POPE	Nocaps	OKVQA	VizWiz	Filckr30k	VQAv2	Avg.
	Upper Bound, All Tokens (100%)							Upper Bound, All Tokens (100%)								
Vanilla	74.1	77.4	91.0	2240	58.4	87.7	76.23	100%	2349	90.49	1.16	64.56	63.51	95.99	81.70	100%
	Retain 25% Tokens in Average (↓ 75%)							Retain 25% Tokens in Average (↓ 75%)								
FastV	67.49	63.68	71.26	2073	49.26	81.54	74.65	88.49%	2251	89.76	1.14	63.89	62.91	91.54	79.46	97.57%
DART	62.37	49.44	52.99	2086	50.58	82.53	64.74	80.72%	2252	89.70	1.11	63.83	62.66	89.33	79.47	96.94%
DivPrune	69.95	58.84	73.88	2097	52.56	83.98	70.43	89.05%	2189	89.91	1.11	63.46	62.29	89.29	79.23	96.37%
D ² Pruner	71.83	68.40	81.14	2190	54.50	85.22	74.64	94.38%	2316	90.18	1.16	64.23	63.56	94.50	80.13	99.19%
	Retain 10% Tokens in Average (↓ 90%)							Retain 10% Tokens in Average (↓ 90%)								
FastV	57.71	38.64	36.85	1692	37.4	65.52	66.40	67.10%	1999	85.05	1.02	59.99	59.65	75.76	69.65	88.31%
DART	58.55	33.04	30.1	1853	42.08	71.29	48.31	64.89%	2057	81.53	0.90	59.94	60.33	66.64	69.91	85.45%
DivPrune	62.73	39.28	49.34	1909	45.10	77.97	58.79	74.01%	2081	88.29	1.00	61.63	60.56	78.50	74.96	90.96%
D ² Pruner	64.73	51.96	57.54	2004	48.33	78.34	70.54	81.69%	2196	88.23	1.17	62.48	63.61	90.05	76.92	96.67%

Table 2: Performance comparisons on Qwen2.5-VL-7B and InternVL-2.5-8B across several understanding benchmarks

Method	LLaVA-1.5-7B							InternVL-2.5-8B										
	RefCOCO		RefCOCO+		RefCOCOg	Average		RefCOCO		RefCOCO+		RefCOCOg	Average					
	val	testA	testB	val	testA	testB	val	test	val	testA	testB	val	test	val	test			
	Upper Bound, All Tokens (100%)							Upper Bound, All Tokens (100%)										
Upper Bound	69.29	77.07	59.88	57.37	69.56	45.53	65.36	65.43	100.0%	90.16	94.54	85.98	85.05	91.56	78.77	87.03	87.72	100.0%
	Retain 75% Tokens in Average (↓ 25%)							Retain 50% Tokens in Average (↓ 50%)										
FastV	65.94	73.47	57.61	54.93	65.86	43.1	61.83	61.84	95.11%	88.07	92.15	83.77	82.68	88.91	75.95	85.01	85.63	97.33%
DivPrune	43.78	48.56	38.98	35.23	41.83	27.45	39.89	39.02	61.74%	87.06	91.32	81.77	81.44	88.47	74.37	83.95	84.42	95.97%
DART	65.54	73.01	57.74	54.17	64.81	44.00	62.38	61.84	95.11%	86.15	91.32	81.73	80.91	87.72	74.13	83.37	84.1	95.49%
D ² Pruner	68.37	76.15	59.33	56.91	68.06	44.22	64.07	63.03	98.14%	88.97	93.44	84.73	83.42	90.57	77.71	85.85	86.03	98.43%
	Retain 50% Tokens in Average (↓ 50%)							Retain 25% Tokens in Average (↓ 75%)										
FastV	56.36	63.27	48.32	46.13	55.34	35.96	52.96	51.6	80.39%	78.17	83.38	72.91	71.68	77.98	65.04	74.96	76.37	85.64%
DivPrune	20.11	24.02	18.41	16.1	20.47	12.78	18.22	18.02	28.99%	73.28	78.56	68.44	67.75	75.22	61.44	70.73	71.5	80.96%
DART	53.54	60.62	47.22	44.79	53.89	36.69	51.41	51.17	78.49%	66.83	70.94	62.65	61.54	67.69	56.99	66.83	67.56	74.33%
D ² Pruner	62.7	70.64	54.29	51.53	62.57	40.7	58.44	57.81	90.00%	86.09	89.78	82.08	79.8	85.85	74.56	82.23	83.27	94.71%
	Retain 40% Tokens in Average (↓ 60%)							Retain 10% Tokens in Average (↓ 90%)										
FastV	49.10	56.19	42.73	40.47	48.13	31.72	46.26	44.99	70.76%	31.98	36.63	28.13	27.1	31.52	23.6	32.27	31.26	34.50%
DivPrune	12.80	14.77	12.68	10.19	12.00	9.16	12.32	11.60	18.82%	36.99	39.12	35.72	32.87	35.61	30.84	35.19	36.94	40.65%
DART	45.69	52.87	39.02	37.67	45.95	30.31	44.36	43.80	66.87%	19.51	20.35	20.22	16.78	17.99	17.37	20.94	21.93	22.15%
D ² Pruner	59.59	68.07	52.22	48.56	59.87	38.23	57.03	54.63	85.89%	78.21	82.69	73.99	72.05	78.03	65.56	75.12	75.02	85.68%

Table 3: Performance comparisons on LLaVA-1.5-7B and InternVL2.5-8B across localization benchmarks.

ToMe (Bolya et al. 2022), FastV (Chen et al. 2024b), HiRED (Arif et al. 2025), FitPrune (Ye et al. 2025), LLaVA-PruMerge (Shang et al. 2024), SparseVLM (Zhang et al. 2024b), HiRED (Arif et al. 2025), PyramidDrop (Xing et al. 2025), VisionZip (Yang et al. 2025), FasterVLM (Zhang et al. 2024a), GlobalCom³ (Liu et al. 2025a) DivPrune (Alvar et al. 2025) and DART (Wen et al. 2025b).

Benchmarks. We perform comprehensive evaluations across general understanding benchmarks, which include visual understanding and reasoning datasets such as GQA (Hudson and Manning 2019), ScienceQA (Lu et al. 2022), VQAv2 (Goyal et al. 2017), TextVQA (Singh et al. 2019), and VizWiz (Bigham et al. 2010), as well as multi-modal reasoning benchmarks like MMBench (Liu et al. 2025b), MMBench-CN (Liu et al. 2025b), MME (Fu et al. 2023), POPE (Li et al. 2023), AI2D (Kembhavi et al. 2016),

ChartQA (Masry et al. 2022), DocVQA (Mathew, Karatzas, and Jawahar 2021), OKVQA (Marino et al. 2019), MMStar (Chen et al. 2024a), Filckr30k (Young et al. 2014) and Nocaps (Agrawal et al. 2019). Our experiments also extend to more challenging referring grounding tasks, utilizing RefCOCO (Kazemzadeh et al. 2014), RefCOCO+ (Kazemzadeh et al. 2014), and RefCOCOg (Mao et al. 2016).

Implementation Details. We follow the default inference settings provided in the official codebase for each MLLM. For our method’s hyperparameters, pruning is initiated at layer $K = 2$ for all scenarios except for the LLaVA-based localization task, where K is set to 5. The weighting parameter α is assigned a value of 1 for understanding tasks and 0.5 for localization tasks. By default, r_{pivot} is set to 0.7 and θ_{sim} is set to 0.8.

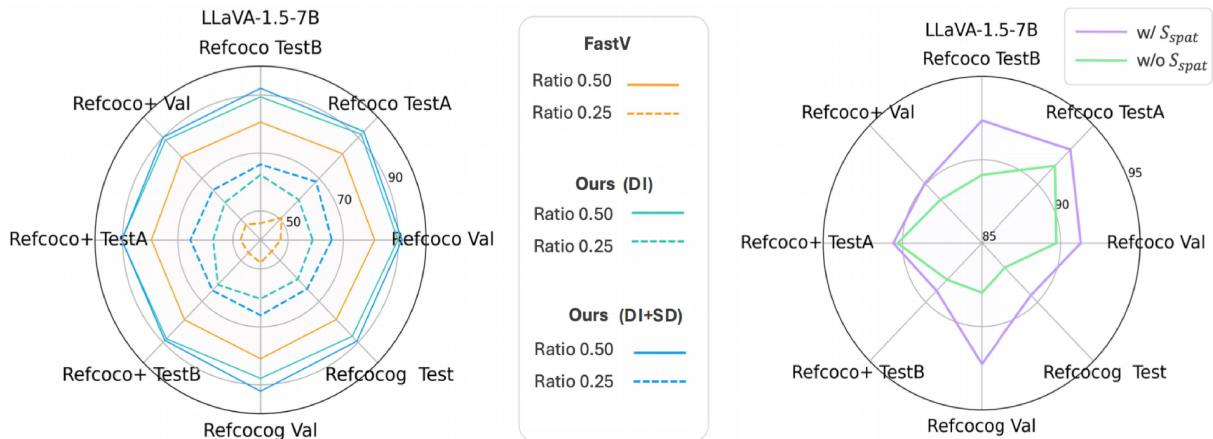


Figure 4: **Ablation study.** *Left:* Effectiveness of Debaised Importance (DI) and Structural Diversity (SD). Our full method (DI+SD) consistently outperforms FastV and the DI-only variant. *Right:* Effectiveness of spatial adjacency. Our method with S_{spat} outperforms the variant without it, demonstrating the benefit of considering spatial distribution for localization tasks.

Main Results

Results on understanding task. We present the performance on understanding tasks in Tab. 1 and Tab. 2. For **LLaVA-1.5-7B**, D²Pruner achieves an impressive average performance of 99.9%, with only 192 tokens retained (\downarrow 66.7%). When further reducing token retention to 64 tokens (\downarrow 77.8%), D²Pruner still maintains a strong performance of 97.3%. For **LLaVA-Next-7B**, the D²Pruner achieves exceptional scalability, achieving an average score of 95.5% with only 11.1% of tokens retained, surpassing all other methods by a notable margin. For **Qwen2.5-VL-7B**, with 90% token reduction, D²Pruner achieves an impressive 81.69%, outperforming DART (64.89%), FastV (67.10%) and DivPrune (74.01%) by notable margins. For **InternVL2.5-8B**, D²Pruner prunes a substantial 75% of tokens while retaining 99.19% of the original performance.

Results on localization task. We further evaluate our method on RefCOCO, RefCOCO+, and RefCOCOg. The results are summarized in Tab. 3. For **LLaVA-1.5-7B**, under a high 60% reduction ratio, our method’s average performance of 85.89% surpasses that of other methods by 15.13-67.07 percentage points. The robustness of our method is particularly evident on **InternVL2.5-8B**. Even when 90% of tokens are pruned, our method maintains a high performance of 85.68%, in stark contrast to competing methods, whose performance drops to a mere 22.1%-40.65%.

Ablation Studies

Effectiveness of DI and SD. We conduct experiments on the localization task with LLaVA-1.5-7B (Liu et al. 2024a), comparing our full method (*Ours (DI+SD)*) against the FastV (Chen et al. 2024b) baseline and a variant using only Debaised Importance (*Ours (DI)*). As shown in Fig. 4 (left), Ours (DI) already surpasses the baseline, demonstrating that correcting for attention bias is effective. The addition of Structural Diversity further improves performance by encouraging diversity among the selected tokens.

Effectiveness of spatial adjacency. Our ablation study underscores the critical role of spatial adjacency. As shown in Fig. 4(right) for LLaVA-1.5-7B, this component consistently enhances localization accuracy by ensuring a uniform spatial distribution, thus preserving vital structural cues post-pruning.

Efficiency Analysis

Our method provides substantial and flexible efficiency gains, as shown in Tab.4 for LLaVA-Next-7B. Retaining just 33.4% of tokens accelerates prefilling by 2.93 \times and cuts FLOPs/KV Cache by over 2.8 \times . Under a more aggressive 11.2% keep ratio of visual tokens, we achieve a 5.09 \times prefill speedup and reduce FLOPs/KV Cache by 6.10 \times /7.63 \times . Notably, the entire pruning process is highly efficient, incurring a negligible overhead of just **0.8%** of a single forward pass.

Methods	Prefilling Time \downarrow (ms/sample)	FLOPs \downarrow (T)	KV Cache \downarrow (MB)	POPE \uparrow (F1-Score)
LLaVA-Next-7B	350 (1.00 \times)	16.9 (1.00 \times)	1512 (1.00 \times)	86.5
+Ours(33.4%)	119 (2.93 \times)	6.0 (2.82 \times)	526 (2.87 \times)	85.4
+Ours(11.2%)	68 (5.09 \times)	2.8 (6.10 \times)	198 (7.63 \times)	84.6

Table 4: Efficiency comparisons on LLaVA-Next-7B.

Conclusion

In this paper, we reveal that importance-based methods suffer from positional bias, while diversity-based methods exhibit structural blindness. To address these issues, we propose D²Pruner, a framework combining Debaised Importance (DI) and Structural Diversity (SD) to retain both critical and diverse tokens. Experimentally, D²Pruner achieves near-original performance (99.2%) on general tasks with a 74.2% FLOPs reduction. Critically, it surpasses prior methods by up to 63.5% on challenging localization tasks, maintaining 85.7% accuracy even when pruning 90% of visual tokens from InternVL-2.5-8B.

Acknowledgments

This research was supported by the Shanghai Science and Technology Program (Grant No. 25ZR1402278).

References

- Agrawal, H.; Desai, K.; Wang, Y.; Chen, X.; Jain, R.; Johnson, M.; Batra, D.; Parikh, D.; Lee, S.; and Anderson, P. 2019. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, 8948–8957.
- Alvar, S. R.; Singh, G.; Akbari, M.; and Zhang, Y. 2025. Divprune: Diversity-based visual token pruning for large multimodal models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 9392–9401.
- Arif, K. H. I.; Yoon, J.; Nikolopoulos, D. S.; Vandierendonck, H.; John, D.; and Ji, B. 2025. HiRED: Attention-Guided Token Dropping for Efficient Inference of High-Resolution Vision-Language Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 1773–1781.
- Bai, J.; Bai, S.; Yang, S.; Wang, S.; Tan, S.; Wang, P.; Lin, J.; Zhou, C.; and Zhou, J. 2023. Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond. *arXiv preprint arXiv:2308.12966*.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; Zhong, H.; Zhu, Y.; Yang, M.; Li, Z.; Wan, J.; Wang, P.; Ding, W.; Fu, Z.; Xu, Y.; Ye, J.; Zhang, X.; Xie, T.; Cheng, Z.; Zhang, H.; Yang, Z.; Xu, H.; and Lin, J. 2025. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923*.
- Bigham, J. P.; Jayant, C.; Ji, H.; Little, G.; Miller, A.; Miller, R. C.; Miller, R.; Tatarowicz, A.; White, B.; White, S.; et al. 2010. Vizwiz: nearly real-time answers to visual questions. In *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, 333–342.
- 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*.
- Chen, L.; Li, J.; Dong, X.; Zhang, P.; Zang, Y.; Chen, Z.; Duan, H.; Wang, J.; Qiao, Y.; Lin, D.; et al. 2024a. Are we on the right way for evaluating large vision-language models? *Advances in Neural Information Processing Systems*, 37: 27056–27087.
- Chen, L.; Zhao, H.; Liu, T.; Bai, S.; Lin, J.; Zhou, C.; and Chang, B. 2024b. An image is worth 1/2 tokens after layer 2: Plug-and-play inference acceleration for large vision-language models. In *European Conference on Computer Vision*, 19–35. Springer.
- Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Cui, E.; Zhu, J.; Ye, S.; Tian, H.; Liu, Z.; et al. 2024c. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; et al. 2024d. 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.
- Fu, C.; Chen, P.; Shen, Y.; Qin, Y.; Zhang, M.; Lin, X.; Yang, J.; Zheng, X.; Li, K.; Sun, X.; et al. 2023. MME: A Comprehensive Evaluation Benchmark for Multimodal Large Language Models. *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.
- Hudson, D. A.; and Manning, C. D. 2019. GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering. *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kazemzadeh, S.; Ordonez, V.; Matten, M.; and Berg, T. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 787–798.
- Kembhavi, A.; Salvato, M.; Kolve, E.; Seo, M.; Hajishirzi, H.; and Farhadi, A. 2016. A diagram is worth a dozen images. In *European conference on computer vision*, 235–251. Springer.
- Li, Y.; Du, Y.; Zhou, K.; Wang, J.; Zhao, W. X.; and Wen, J.-R. 2023. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2024a. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26296–26306.
- Liu, H.; Li, C.; Li, Y.; Li, B.; Zhang, Y.; Shen, S.; and Lee, Y. J. 2024b. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36: 34892–34916.
- Liu, X.; Wang, Z.; Han, Y.; Wang, Y.; Yuan, J.; Song, J.; Zheng, B.; Zhang, L.; Huang, S.; and Chen, H. 2025a. Global Compression Commander: Plug-and-Play Inference Acceleration for High-Resolution Large Vision-Language Models. *arXiv preprint arXiv:2501.05179*.
- Liu, Y.; Duan, H.; Zhang, Y.; Li, B.; Zhang, S.; Zhao, W.; Yuan, Y.; Wang, J.; He, C.; Liu, Z.; et al. 2025b. Mmbench: Is your multi-modal model an all-around player? In *European Conference on Computer Vision*, 216–233. Springer.
- Lu, P.; Mishra, S.; Xia, T.; Qiu, L.; Chang, K.-W.; Zhu, S.-C.; Tafjord, O.; Clark, P.; and Kalyan, A. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35: 2507–2521.

- Mao, J.; Huang, J.; Toshev, A.; Camburu, O.; Yuille, A. L.; and Murphy, K. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 11–20.
- Marino, K.; Rastegari, M.; Farhadi, A.; and Mottaghi, R. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, 3195–3204.
- 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.; and Jawahar, C. 2021. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 2200–2209.
- Shang, Y.; Cai, M.; Xu, B.; Lee, Y. J.; and Yan, Y. 2024. Llava-prumerge: Adaptive token reduction for efficient large multimodal models. *arXiv preprint arXiv:2403.15388*.
- 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.
- Wen, Z.; Gao, Y.; Li, W.; He, C.; and Zhang, L. 2025a. Token Pruning in Multimodal Large Language Models: Are We Solving the Right Problem? *arXiv preprint arXiv:2502.11501*.
- Wen, Z.; Gao, Y.; Wang, S.; Zhang, J.; Zhang, Q.; Li, W.; He, C.; and Zhang, L. 2025b. Stop looking for important tokens in multimodal language models: Duplication matters more. *arXiv preprint arXiv:2502.11494*.
- Xing, L.; Huang, Q.; Dong, X.; Lu, J.; Zhang, P.; Zang, Y.; Cao, Y.; He, C.; Wang, J.; Wu, F.; et al. 2025. Conical Visual Concentration for Efficient Large Vision-Language Models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 14593–14603.
- Yang, S.; Chen, Y.; Tian, Z.; Wang, C.; Li, J.; Yu, B.; and Jia, J. 2025. Visionzip: Longer is better but not necessary in vision language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 19792–19802.
- Ye, W.; Wu, Q.; Lin, W.; and Zhou, Y. 2025. Fit and prune: Fast and training-free visual token pruning for multi-modal large language models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 22128–22136.
- Young, P.; Lai, A.; Hodosh, M.; and Hockenmaier, J. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the association for computational linguistics*, 2: 67–78.
- Zhang, Q.; Cheng, A.; Lu, M.; Zhuo, Z.; Wang, M.; Cao, J.; Guo, S.; She, Q.; and Zhang, S. 2024a. [CLS] Attention is All You Need for Training-Free Visual Token Pruning: Make VLM Inference Faster. *arXiv e-prints*, arXiv–2412.
- Zhang, Y.; Fan, C.-K.; Ma, J.; Zheng, W.; Huang, T.; Cheng, K.; Gudovskiy, D.; Okuno, T.; Nakata, Y.; Keutzer, K.; et al. 2024b. Sparsevlm: Visual token sparsification for efficient vision-language model inference. *arXiv preprint arXiv:2410.04417*.