

Chain-of-Thought Compression Should Not Be Blind: V-Skip for Efficient Multimodal Reasoning via Dual-Path Anchoring

Dongxu Zhang^{1,2,*} Yiding Sun^{1,2,*} Cheng Tan^{4,*} Wenbiao Yan⁵
Ning Yang^{3,†} Jihua Zhu^{1,2,†} Haijun Zhang⁶

¹School of Software Engineering, Xi'an Jiaotong University

²State Key Laboratory of Integrated Services Networks, Xidian University

³Institute of Automation, Chinese Academy of Sciences

⁴Shanghai Artificial Intelligence Laboratory

⁵Harbin Institute of Technology, Shenzhen

⁶University of Science and Technology Beijing

Abstract

While Chain-of-Thought (CoT) reasoning significantly enhances the performance of Multimodal Large Language Models (MLLMs), its autoregressive nature incurs prohibitive latency constraints. Current efforts to mitigate this via token compression often fail by blindly applying text-centric metrics to multimodal contexts. We identify a critical failure mode termed Visual Amnesia, where linguistically redundant tokens are erroneously pruned, leading to hallucinations. To address this, we introduce V-Skip that reformulates token pruning as a Visual-Anchored Information Bottleneck (VA-IB) optimization problem. V-Skip employs a dual-path gating mechanism that weighs token importance through both linguistic surprisal and cross-modal attention flow, effectively rescuing visually salient anchors. Extensive experiments on Qwen2-VL and Llama-3.2 families demonstrate that V-Skip achieves a $2.9\times$ speedup with negligible accuracy loss. Specifically, it preserves fine-grained visual details, outperforming other baselines over 30% on the DocVQA.

1 Introduction

The integration of vision and language has catalyzed a paradigm shift in Artificial Intelligence, culminating in the emergence of Multimodal Large Language Models (MLLMs) (Wang et al., 2025). By extending the Chain-of-Thought (CoT) reasoning capability from text to vision, models such as LLaVA (Li et al., 2024) and Qwen-VL (Wang et al., 2024b) have demonstrated remarkable proficiency in decomposing complex visual queries into step-by-step rationales. This reasoning process significantly enhances performance and interpretability.

However, because of the autoregressive nature of Transformer architectures, generating extensive rea-

soning chains introduces severe latency and computational overhead (Shazeer, 2019; Dao et al., 2022). The linear growth of the Key-Value (KV) cache restricts batch sizes and limits throughput, rendering long chain visual reasoning impractical (Kwon et al., 2023). Empirically, these generated sequences often contain linguistic fillers (Jiang et al., 2023) that support syntactic fluency yet offer minimal information gain for the final prediction.

To mitigate this inefficiency, token compression techniques such as TokenSkip (Xia et al., 2025) or LLMLingua-2 (Pan et al., 2024) have been proposed. While effective in text contexts, we argue that directly applying these methods to multimodal tasks precipitates a critical failure mode we term Visual Amnesia. This phenomenon severs the connection to the input image and induces object hallucinations (Li et al., 2023). Text-centric pruning relies on linguistic surprisal. If a token is statistically predictable from the preceding text, it is deemed redundant. However, in multimodal reasoning, token importance is bimodal. For instance, a token like red may be linguistically predictable given apple, yet it serves as an essential visual anchor. As illustrated in Figure 1, standard algorithms rely solely on linguistic priors and discard such anchors without considering visual context, severing the connection to the input image and inducing object hallucinations (Rohrbach et al., 2018).

In this paper, we introduce V-Skip, a novel method designed to compress multimodal CoT sequences while strictly preserving visual grounding. We reformulate the compression task as a Visual-Anchored Information Bottleneck (VA-IB) optimization problem. Unlike existing methods that rely solely on language probability (Zhang et al., 2025), V-Skip employs a dual-path scoring mechanism that interrogates both linguistic surprisal and cross-modal attention flow. This allows the model

*Equal contribution.

†Corresponding authors.

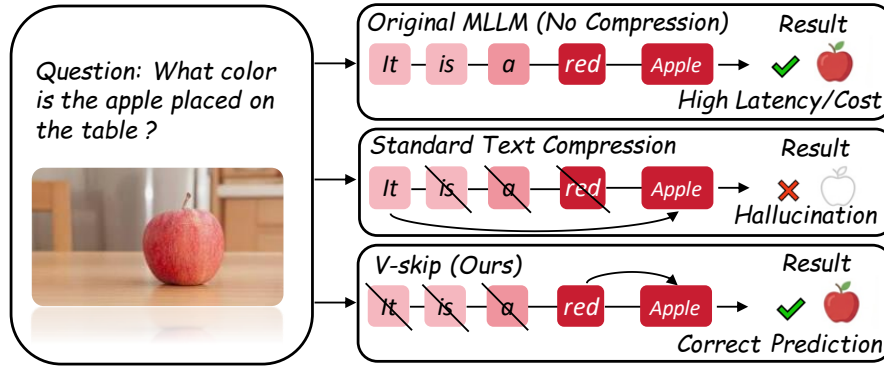


Figure 1: Comparison of token compression paradigms. **(Top)** The original MLLMs retains all tokens, ensuring accuracy but incurring high computational costs. **(Middle)** Standard text-centric compression aggressively prunes the adjective red due to its high linguistic probability given apple, detaching the reasoning from the visual ground truth and leading to hallucination. **(Bottom)** V-Skip utilizes a dual-path scoring mechanism to identify and rescue visually salient tokens while compressing redundant linguistic fillers.

to identify and rescue tokens that are linguistically redundant but visually salient. Furthermore, to avoid the latency overhead of computing attention maps, we treat compression not as an online search but as an offline policy learning problem. We distill the V-Skip pruning policy into the base model via Low-Rank Adaptation (LoRA) (Hu et al., 2022), resulting in an efficient reasoner that generates concise, visually grounded rationales without requiring explicit scoring or online filtering during inference.

Our contributions are summarized as follows:

- We identify Visual Amnesia as a fundamental failure mode in multimodal CoT compression, where text-centric pruning inadvertently discards tokens essential for visual grounding.
- The proposed VA-IB framework reformulates multimodal compression as an information theoretic objective to balance linguistic efficiency with cross-modal grounding.
- We introduce V-Skip, a novel framework utilizing a dual-path scoring mechanism to selectively preserve visual anchors, distilled into a lightweight adapter for efficient inference.
- Extensive experiments on Qwen2-VL demonstrate that V-Skip achieves a $2.9\times$ speedup with minimal accuracy trade-offs. It outperforms other baselines over 30% on DocVQA.

2 Related Work

2.1 Multimodal Large Language Models

Recent advancements in MLLMs have bridged the gap between visual perception and linguis-

tic intelligence (Lin et al., 2025, 2026; Zhang et al., 2026b). Foundational models such as Llama (Dubey et al., 2024), LLaVA (Li et al., 2024), and Qwen-VL (Wang et al., 2024b) project visual features into the embedding space of LLMs, enabling them to process interleaved image-text inputs.

On one hand, text-centric methods like Tree-of-Thoughts (Yao et al., 2023) and Graph-of-Thoughts (Besta et al., 2024) explore non-linear reasoning topologies. Parallel to these explorations, multimodal strategies focus heavily on cross-modal alignment and rationale quality. For instance, PICa (Yang et al., 2022) converts visual information into captions to prompt GPT-3 for reasoning. Multimodal-CoT (Zhang et al., 2023) adopts a two-stage method that fine tunes models using human annotated rationales. Other works like Prompt-CoT (Yao et al., 2024) and T-SciQ (Wang et al., 2024a) leverage prompt engineering and teacher signals to guide the generation of reasoning paths.

Despite these improvements in accuracy, the resulting extensive reasoning chains often introduce significant computational overhead (Zhang et al., 2023). In contrast to these generation centric methods, our work focuses on the efficient utilization of reasoning chains, pruning redundant computations while preserving the accurate prediction.

2.2 Efficient LLM Reasoning & Compression

Efforts to mitigate the computational demands of LLMs have focused on system level optimizations. Techniques such as quantization (Dettmers et al., 2022), structured compression (Zhu et al., 2025),

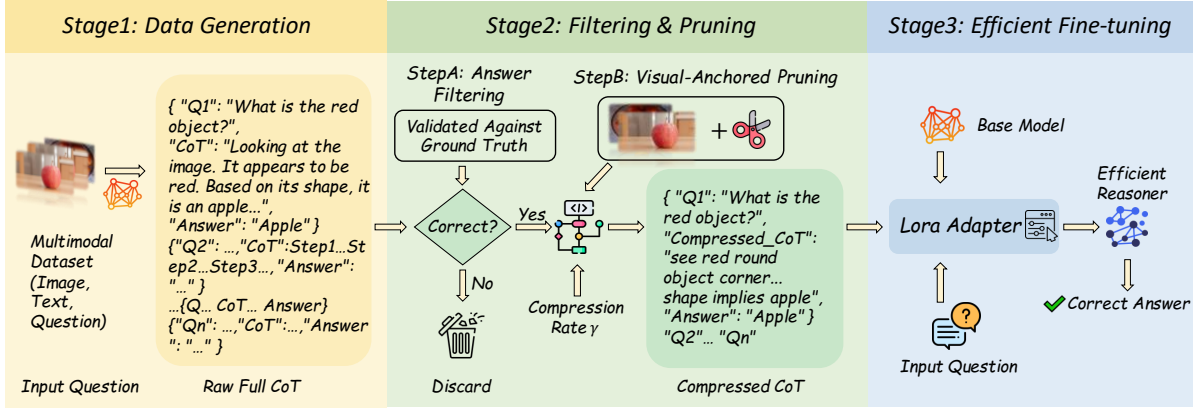


Figure 2: The V-Skip Training Pipeline. Our method automates the construction of efficient multimodal reasoners through three distinct stages. (1) Data Generation, a VLM generates comprehensive CoT rationales for input image-question pairs. (2) Filtering & Pruning, Step A filters out rationales that lead to incorrect answers compared to the ground truth. Step B applies our Visual-Anchored Pruning mechanism (detailed in Figure 3), which selectively removes redundant tokens based on a target compression rate γ , producing concise Compressed_CoT sequences. (3) Efficient Fine-tuning, the resulting high-quality compressed dataset is used to supervise a Base Model via LoRA, yielding an Efficient Reasoner capable of generating succinct and visually grounded reasoning paths directly.

and KV cache optimization strategies like PagedAttention (Kwon et al., 2023; Zhang et al., 2026a) effectively reduce memory footprints. However, these methods do not address the fundamental bottleneck of sequence length in autoregressive generation.

To directly tackle the computational complexity, token pruning and compression have emerged as effective strategies (Chuang et al., 2024; Yang et al., 2025). Methods like TokenSkip (Xia et al., 2025) accelerate inference by structurally bypassing layers for redundant tokens. Methods like AS-CoT (Zhang et al., 2025) and LLMLingua-2 (Pan et al., 2024) leverage perplexity and entropy distributions to identify and prune tokens with low linguistic information gain, reducing the generation of unnecessary intermediate steps.

Despite their efficacy in NLP, compression metrics based on linguistic probability often discard critical visual anchors. V-Skip distinguishes itself as the first framework to explicitly incorporate Visual Anchoring into the pruning metric.

3 Method

We introduce V-Skip, a method designed to compress multimodal CoT sequences. Departing from standard text-centric pruning, we theoretically reformulate the compression task as a VA-IB optimization. This formulation guides a dual-path selection mechanism that distinguishes between linguistic redundancy and visual necessity. The

overall pipeline is illustrated in Figure 2.

3.1 Preliminaries and Problem Formulation

We formalize the reasoning process of a MLLM as a probabilistic generation task. Let $\mathcal{I} = (\mathcal{V}, \mathcal{Q})$ denote the multimodal input context, comprising the visual input \mathcal{V} and textual query \mathcal{Q} . The model generates a reasoning chain $\mathcal{C} = \{c_1, \dots, c_T\}$, where c_t represents the t -th reasoning token in a sequence of total length T . The joint distribution is factorized autoregressively as:

$$P_\theta(\mathcal{A}, \mathcal{C} | \mathcal{I}) = P_\theta(\mathcal{A} | \mathcal{C}, \mathcal{I}) \prod_{t=1}^T P_\theta(c_t | \mathcal{C}_{<t}, \mathcal{I}), \quad (1)$$

where $P_\theta(\cdot)$ denotes the conditional probability distribution parameterized by the MLLM with weights θ . The objective of token compression is to identify a sparse subsequence $\hat{\mathcal{C}} \subset \mathcal{C}$ with length $|\hat{\mathcal{C}}| \ll |\mathcal{C}|$ that maximizes the predictive likelihood of the answer \mathcal{A} while minimizing inference latency.

Standard text pruning methods typically operate by minimizing the negative log-likelihood of the sequence based on linguistic priors. However, directly applying this unimodal metric to MLLMs precipitates a critical failure mode we term Visual Amnesia. In a multimodal context, visually grounded tokens often exhibit high linguistic perplexity because they cannot be inferred from the preceding text alone. As a result, probability-based pruners misidentify these high-entropy visual anchors as redundancy. This results in compressed

chains that are fluent but factually ungrounded.

3.2 Visual-Anchored Information Bottleneck

To mitigate Visual Amnesia, we reformulate the compression task through the lens of the Information Bottleneck (IB) principle (Tishby et al., 2000; Alemi et al., 2016). Building upon this principle, we posit that a valid compressed rationale $\hat{\mathcal{C}}$ must satisfy two distinct information-theoretic criteria. The first is sufficiency, which ensures the retention of adequate semantic content to predict the answer \mathcal{A} . The second is grounding, which necessitates maintaining high mutual dependence with the visual input \mathcal{V} to mitigate hallucination risks.

To operationalize these criteria, we propose the VA-IB objective. Unlike standard text compression which minimizes length based solely on linguistic priors, VA-IB explicitly imposes a visual constraint. We formulate the optimization as maximizing the joint utility subject to a length budget:

$$\begin{aligned} \max_{\hat{\mathcal{C}}} & \left[\underbrace{I(\hat{\mathcal{C}}; \mathcal{A})}_{\text{Sufficiency}} + \lambda \underbrace{I(\hat{\mathcal{C}}; \mathcal{V} | \mathcal{Q})}_{\text{Anchoring}} \right] \\ \text{s.t.} & \quad |\hat{\mathcal{C}}| \approx \gamma |\mathcal{C}|, \end{aligned} \quad (2)$$

where λ controls the strength of visual grounding and γ is the target compression ratio. The core lies in the anchoring term $I(\hat{\mathcal{C}}; \mathcal{V} | \mathcal{Q})$. To elucidate its physical meaning, we decompose this conditional mutual information using Shannon entropy:

$$I(\hat{\mathcal{C}}; \mathcal{V} | \mathcal{Q}) = H(\hat{\mathcal{C}} | \mathcal{Q}) - H(\hat{\mathcal{C}} | \mathcal{V}, \mathcal{Q}), \quad (3)$$

where $H(\hat{\mathcal{C}} | \mathcal{Q})$ represents Textual Uncertainty and $H(\hat{\mathcal{C}} | \mathcal{V}, \mathcal{Q})$ measures Multimodal Uncertainty. Maximizing Eq. (3) effectively selects tokens that are highly unpredictable from text alone but become deterministic given the image.

3.3 Visual Anchoring Score via Attention

Optimizing the VA-IB objective requires evaluating the mutual information $I(c_t; \mathcal{V} | \mathcal{Q})$ for each token. Since direct computation over high-dimensional latent spaces is intractable, we derive an efficient proxy by leveraging the model’s internal attention mechanism, detailed in Figure 3.

In MLLMs trained with contrastive objectives, cross-modal attention weights serve as a variational lower bound for the pointwise mutual information between textual and visual representations. Specifically, the attention weight from a text token c_t to

a visual patch v is proportional to their density ratio (Oord et al., 2018). This theoretical alignment implies that the cumulative attention mass directed toward the image manifests the degree of cross-modal dependence. We interpret the total attention flow as an empirical measure of a token’s grounding necessity. Based on this insight, we instantiate a dual-path scoring mechanism that evaluates linguistic redundancy and visual necessity.

Textual Path S_{text} To evaluate the linguistic redundancy of a generated token c_t , we quantify its information content using self-information. Formally, we define the textual score as the negative log-likelihood of the token conditioned on the multimodal context:

$$S_{\text{text}}(c_t) = -\log P_{\theta}(c_t | c_{<t}, \mathcal{Q}, \mathcal{V}), \quad (4)$$

where $P_{\theta}(\cdot)$ is consistent with the formulation in Eq. (1). Intuitively, tokens with high predictability (e.g., functional words like *is* or *the*) yield high probabilities and consequently low S_{text} scores. These tokens provide minimal information gain and are considered candidates for pruning, provided they do not serve as visual anchors.

Visual Path S_{vis} Simultaneously, we measure the visual grounding by analyzing the attention flow from the generated text back to the visual context. Let $A^{(\ell, h)} \in \mathbb{R}^{T \times T}$ denote the attention weight matrix for layer ℓ and head h . Specifically, $A_{t,k}^{(\ell, h)}$ represents the attention weight from the query at position t to the key at position k . We define the total *visual attention mass* $\Omega_t^{(\ell, h)}$ as the sum of attention weights directed towards the set of visual patch indices \mathcal{I}_{img} :

$$\Omega_t^{(\ell, h)} = \sum_{k \in \mathcal{I}_{\text{img}}} A_{t,k}^{(\ell, h)}. \quad (5)$$

A high value of Ω_t indicates that the generation at step t is heavily dependent on visual features, marking the corresponding token as a potential visual anchor. However, recent mechanistic interpretability studies suggest that visual-text semantic alignment is not uniform. It typically peaks in the middle layers and is handled by sparse (Ethayarajh, 2019), specialized induction heads rather than being distributed evenly (Olsson et al., 2022). Simple averaging would dilute this signal with noise from syntactic heads. Therefore, we define the final Visual Anchoring Score (VAS) for the t -th token by aggregating $\Omega_t^{(\ell, h)}$ over a salient subset of layers

$\mathcal{L}_{\text{focus}}$ and employing a max-pooling strategy over heads \mathcal{H} to isolate the most confident grounding signal. This is formalized as:

$$S_{\text{vis}}(t) = \frac{1}{|\mathcal{L}_{\text{focus}}|} \sum_{\ell \in \mathcal{L}_{\text{focus}}} \max_{h \in \mathcal{H}} \left(\Omega_t^{(\ell, h)} \right). \quad (6)$$

By defining $S_{\text{vis}}(t)$ in this manner, we effectively capture the Multimodal Uncertainty reduction described in Eq. (3). This ensures that tokens acting as visual anchors are identified and preserved regardless of their linguistic probability.

3.4 The V-Skip Gating Mechanism

We integrate the established dual utility metrics into the V-Skip Gate to construct a binary retention mask $M = \{m_1, \dots, m_T\}$, where $m_t \in \{0, 1\}$, identifying the optimal subsequence $\hat{\mathcal{C}}$ that satisfies the VA-IB objective defined in Eq.(2). Unlike conventional pruning methods that rely on a single confidence threshold, V-Skip employs a Union-of-Saliency strategy.

This design is motivated by the complementary nature of our dual-path scores. S_{text} identifies tokens essential for linguistic coherence, while S_{vis} identifies tokens essential for visual grounding. Consequently, a token should be retained if it is salient in either modality. Formally, given a target compression rate γ , we dynamically calibrate two thresholds τ_{text} and τ_{vis} to ensure that the cardinality of their boolean union precisely satisfies the target retention budget $(1 - \gamma)$. The retention mask m_t is determined by the disjunction of two boolean conditions, formulated as:

$$m_t = \mathbb{I}(S_{\text{text}} \geq \tau_{\text{text}}) \vee \mathbb{I}(S_{\text{vis}} \geq \tau_{\text{vis}}), \quad (7)$$

where $\mathbb{I}(\cdot)$ is the indicator function. The final compressed chain is obtained by filtering the sequence with this mask. By allowing visual saliency to override linguistic redundancy, this mechanism explicitly prevents the pruning of tokens that are linguistically predictable (low S_{text}) but visually essential (high S_{vis}). This ensures that critical visual descriptors, which are often treated as noise by unimodal compressors, are preserved, thereby maintaining the factual integrity of the reasoning chain.

3.5 Efficient Inference via V-Skip Distillation

To eliminate the runtime overhead of computing dual-path scores, we distill the V-Skip policy into the model parameters via efficient instruction tuning. We first construct a distilled dataset $\mathcal{D}_{\text{distill}}$

offline by applying the retention mask (Eq. (7)) to raw reasoning chains. Subsequently, we fine-tune the model using LoRA (Hu et al., 2022) under standard autoregressive supervision. This optimization allows the model to internalize the pruning logic, enabling it to generate concise and visually anchored rationales without explicit attention analysis during inference.

4 Experiments

4.1 Experimental Setup

Models and Architectures To evaluate V-Skip across varying scales and architectures, we employ the Qwen2-VL-Instruct series (encompassing the 2B, 7B, and 72B variants) (Wang et al., 2024b) alongside Llama-3.2-11B-Vision-Instruct (Dubey et al., 2024). We implement V-Skip by applying LoRA to the attention layers while freezing the pre-trained visual and language backbones. All experiments are conducted on a server equipped with 8 NVIDIA RTX 3090 GPUs.

Datasets and Metrics To evaluate the efficacy of V-Skip, we employ two distinct benchmarks. We utilize the MMMU dataset to assess complex multidisciplinary reasoning (Yue et al., 2024). Furthermore, we leverage the DocVQA benchmark to strictly evaluate fine-grained OCR and spatial grounding abilities (Mathew et al., 2021). We report Accuracy, End-to-End Latency, and the Actual Compression Ratio (ActRatio). Detailed experimental configurations and dataset splits are provided in the Appendix B.

Baselines To strictly evaluate the contribution of our visual-anchored mechanism, we compare V-Skip against both text-centric and multimodal compression strategies. For text-centric methods, we employ ASCoT (Zhang et al., 2025) and LLMLingua-2 (Pan et al., 2024), which rely solely on linguistic probability distributions to identify redundancy. For multimodal comparison, we implement a Visual-Attention Pruning as a lower bound for performance stability (Wang et al., 2024b).

4.2 Main Results

The quantitative comparison on Qwen2-VL-7B detailed in Table 1. Text-centric baselines suffer significant performance degradation on the fine-grained DocVQA task. Notably, LLMLingua-2 drops over 53% at $\gamma = 0.5$. In contrast, V-Skip effectively preserves visual anchors, retaining 83.7% ANLS and outperforming LLMLingua-2 by 45.2%.

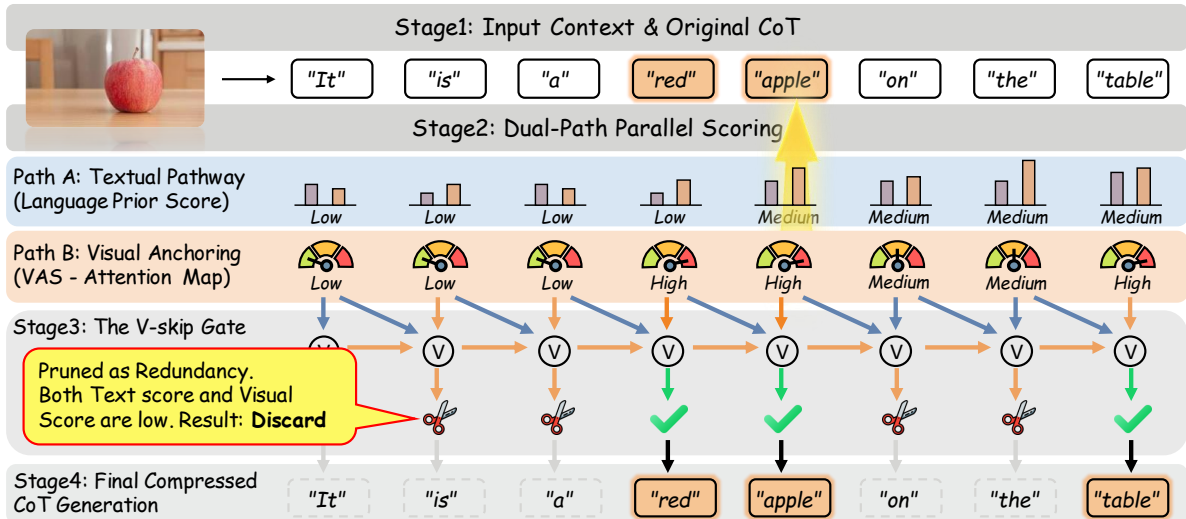


Figure 3: The Visual-Anchored Token Pruning Mechanism. The pruning process evaluates each generated token through a dual-pathway scoring framework. The Textual Pathway quantifies the linguistic necessity of a token by identifying redundant functional fillers, such as “it” or “is” based on their language prior scores. The Visual Pathway computes the VAS by aggregating cross-modal attention weights, which quantifies the grounding strength of a token relative to specific image regions. Finally, the V-Skip Gate applies a union-of-saliency operation to determine token retention. A token is pruned only if it is diagnosed as both linguistically redundant and visually irrelevant.

Similarly, on MMMU, V-Skip limits accuracy loss to 5.9%, whereas baselines degrade by over 20%. In terms of efficiency, V-Skip achieves a $1.8\times$ speedup on DocVQA (2.71s). Crucially, it processes faster than LLMLingua-2 (2.93s) despite generating longer sequences, confirming the runtime advantage of our zero-overhead LoRA distillation over calculation-heavy online filtering.

To verify the scalability of our method, we extend our evaluation to the Llama-3.2-11B-Vision-Instruct model, as illustrated in Figure 4. The V-Skip method maintains a flat performance curve across varying compression ratios for both MMMU and DocVQA tasks. Specifically, even under aggressive pruning ($\gamma = 0.5$), Llama-3.2 with V-Skip retains over 91% of its original performance on both benchmarks. This demonstrates that the principle of Visual Anchoring is model-agnostic and essential for robust multimodal compression.

Robustness Increases with Model Size We further investigate the scalability of V-Skip across the Qwen2-VL family (2B, 7B, and Large) on the complex MMMU benchmark. As detailed in Table 2, we observe a distinct Positive Scaling Law regarding pruning robustness. The lightweight 2B model exhibits high sensitivity to token removal, suffering a significant performance degradation of 8.5%. We attribute this to the fragility of reasoning chains in parameter constrained models, where




Figure 4: Performance robustness of V-Skip on Llama-3.2-11B-Vision-Instruct across varying compression ratios. The dual-axis plot displays the impact of the target compression ratio γ (x-axis) on model performance. The left axis (blue circle) reports the Accuracy (%) for the MMMU, while the right axis (green square) indicates the ANLS Score for the DocVQA.

information is densely packed and lacks redundancy. Removing intermediate steps in such concise chains often leads to logical fractures. Conversely, the 72B model demonstrates superior resilience. Despite removing 50% of the tokens, the accuracy drop is limited to just 3.2%, effectively halving the loss observed in the 7B baseline. This result suggests that as model capacity increases, the generated CoT becomes inherently more verbose and redundant.

Table 1: Main results of Qwen2-VL-7B-Instruct on MMMU and DocVQA datasets. We report Accuracy (%) for MMMU and ANLS (%) for DocVQA. To evaluate efficiency, we detail the average number of generated reasoning tokens (**Tokens**), end-to-end inference latency (**Latency**), and the actual compression ratio (**ActRatio**). The symbol γ denotes the target compression ratio set for the pruning algorithms. \downarrow indicates the performance drop compared to the original full model.

| Methods | Ratio γ | MMMU | | | | DocVQA | | | |
|-----------------|----------------|--------------------------|--------|-------------|----------|--------------------------|--------|-------------|----------|
| | | Acc | Tokens | Latency (s) | ActRatio | ANLS | Tokens | Latency (s) | ActRatio |
| Original (Full) | - | 54.1 | 245.0 | 6.42 | - | 91.6 | 189.0 | 4.87 | - |
| Truncation | 0.9 | 50.8(3.3 \downarrow) | 220.5 | 5.84 | 0.90 | 84.2(7.4 \downarrow) | 170.1 | 4.52 | 0.90 |
| | 0.7 | 44.5(9.6 \downarrow) | 171.5 | 4.51 | 0.70 | 71.5(20.1 \downarrow) | 132.3 | 3.51 | 0.70 |
| | 0.5 | 38.5(15.6 \downarrow) | 122.5 | 3.23 | 0.50 | 62.5(29.1 \downarrow) | 94.5 | 2.57 | 0.50 |
| LLMLingua-2 | 0.9 | 49.5(4.6 \downarrow) | 223.2 | 6.03 | 0.91 | 78.4(13.2 \downarrow) | 173.8 | 4.74 | 0.92 |
| | 0.7 | 40.2(13.9 \downarrow) | 166.6 | 4.97 | 0.68 | 55.6(36.0 \downarrow) | 130.4 | 3.81 | 0.69 |
| | 0.5 | 32.4(21.7 \downarrow) | 115.1 | 3.73 | 0.47 | 38.5(53.1 \downarrow) | 88.8 | 2.93 | 0.47 |
| ASCoT | 0.9 | 50.1(4.0 \downarrow) | 218.0 | 6.12 | 0.89 | 79.8(11.8 \downarrow) | 168.2 | 4.78 | 0.89 |
| | 0.7 | 41.8(12.3 \downarrow) | 176.4 | 5.20 | 0.72 | 58.2(33.4 \downarrow) | 136.1 | 4.03 | 0.72 |
| | 0.5 | 33.1(21.0 \downarrow) | 124.9 | 3.91 | 0.51 | 40.2(51.4 \downarrow) | 90.7 | 3.10 | 0.48 |
| V-Skip (Ours) | 1.0 | 54.1(0.0 \downarrow) | 245.0 | 6.54 | 1.00 | 91.6(0.0 \downarrow) | 189.0 | 5.09 | 1.00 |
| | 0.9 | 53.6(0.5 \downarrow) | 227.8 | 5.89 | 0.93 | 90.8(0.8 \downarrow) | 172.1 | 4.61 | 0.91 |
| | 0.8 | 52.8(1.3 \downarrow) | 193.6 | 5.24 | 0.79 | 89.5(2.1 \downarrow) | 156.9 | 4.08 | 0.83 |
| | 0.7 | 51.5(2.6 \downarrow) | 173.4 | 4.65 | 0.71 | 87.9(3.7 \downarrow) | 132.4 | 3.68 | 0.70 |
| | 0.6 | 50.1(4.0 \downarrow) | 151.9 | 4.05 | 0.62 | 85.8(5.8 \downarrow) | 115.3 | 3.19 | 0.61 |
| | 0.5 | 48.2(5.9 \downarrow) | 120.1 | 3.49 | 0.49 | 83.7(7.9 \downarrow) | 98.4 | 2.71 | 0.52 |

4.3 Analysis

We conduct an in-depth analysis to interpret the underlying mechanisms of V-Skip. We first quantify the preservation of fine-grained descriptors through a visual attribute retention analysis. Complementing this quantitative evidence, we examine the information entropy mismatch through a qualitative case study. Ultimately, we link these granular behaviors to broader model safety, evaluating how preserving visual grounding acts as a stabilizer to mitigate hallucination bias.

4.3.1 Visual Attribute Retention Analysis

To investigate the underlying cause of the performance degradation observed in text-only baselines, we hypothesize that text-centric pruning inadvertently discards visual descriptors (*e.g.*, colors, shapes) due to their low linguistic perplexity. To quantify this, we measure the Visual Attribute Retention Rate (VARR) on MMMU using the GQA vocabulary. As shown in Table 3, probability-based baselines suffer significant information loss. Notably, LLMLingua-2 retains only 42.5% of color tokens, treating the majority as redundancy. In contrast, V-Skip effectively rescues these anchors via visual saliency. At the same compression ratio ($\gamma = 0.5$), V-Skip achieves 89.4% retention for colors (+46.9% over LLMLingua-2) and 91.2% for objects. This empirical evidence confirms that

Table 2: Scaling analysis on the MMMU. We evaluate the impact of V-Skip across the Qwen2-VL series models (2B, 7B, and 72B) with a fixed target compression ratio of $\gamma = 0.5$. We reports the accuracy of the uncompressed baseline (Full CoT), the compressed model (V-Skip), and the absolute performance decrease (Δ).

| Model Size | Full CoT | V-Skip | Drop (Δ) |
|-------------|----------|--------|-------------------|
| 2B (Small) | 41.2 | 32.7 | -8.5 |
| 7B (Base) | 54.1 | 48.2 | -5.9 |
| 72B (Large) | 64.5 | 61.3 | -3.2 |

V-Skip successfully anchors the compressed CoT to the visual modality, preventing the reasoning collapse observed in baselines.

4.3.2 Case Study

To empirically validate the mechanism behind the performance divergence observed in Table 1, we present a qualitative case study in Figure 5. The task requires the model to extract a specific invoice total \$45.20. This case exposes a fundamental Information Entropy Mismatch between modalities. From a unimodal linguistic perspective, specific numerical values often appear as high entropy outliers, they lack semantic redundancy with the preceding context. Consequently, probability based pruners like LLMLingua-2 misinterpret these critical tokens as noise, leading to the erroneous pruning of the answer itself.

Table 3: Visual Attribute Retention Rate (VARR) analysis on the MMMU dataset. We categorize tokens within the CoT into specific visual attribute classes (Color, Shape, Object) using the GQA vocabulary. The table reports the percentage of retained tokens for each category at a fixed compression ratio of $\gamma = 0.5$.

| Method | Color | Shape | Object |
|---------------|--------------|--------------|--------------|
| LLMLingua-2 | 42.5% | 55.1% | 64.8% |
| ASCoT | 48.2% | 58.7% | 69.5% |
| V-Skip (Ours) | 89.4% | 86.3% | 91.2% |

V-Skip resolves this conflict by integrating the Visual Score S_{vis} . Although the token “\$45.20” exhibits low linguistic coherence, it triggers a massive activation in the cross-attention layers, indicating a strong dependency on image pixels. By detecting this signal, V-Skip overrides the low textual probability and identifies the token as an indispensable bridge to the visual context. This enables V-Skip to function as a grounding-aware filter, systematically preserving tokens characterized by high visual saliency despite their high linguistic perplexity, which are essential for precise visual reasoning.

4.3.3 Mitigating Hallucination Bias

A critical concern in MLLMs is the tendency to hallucinate non-existent objects. To determine whether token pruning exacerbates this issue, we investigate whether pruning exacerbates object hallucination using the balanced POPE benchmark. As detailed in Table 4, text-centric baselines exhibit a severe Yes-Bias, with positive response rates surging to 66.8% (ASCoT) and 64.5% (LLMLingua-2). While this inflates Recall, it severely degrades Precision to roughly 63.5%, indicating a reliance on parametric priors when visual anchors are stripped. In contrast, V-Skip effectively preserves visual grounding via S_{vis} , maintaining a neutral Yes-Ratio of 51.2% that closely mirrors the uncompressed baseline (50.4%). Consequently, V-Skip achieves a balanced F1 score of 88.9, demonstrating its capability to compress reasoning chains without compromising factual integrity.

4.4 Ablation Experiments

To validate our dual-path design, we decouple the gating components on the MMMU dataset. As shown in Table 5, results reveal a distinct trade-off between linguistic coherence and visual grounding. Relying solely on textual probability (Row B) minimizes syntax errors (1.1%) but sacrifices

Table 4: Hallucination evaluation on the POPE benchmark. We report standard classification metrics (Accuracy, Precision, Recall, F1) and the Yes-Ratio (the proportion of positive responses generated by the model).

| Method | Yes-Ratio | Acc | Prec | Recall | F1 |
|---------------|--------------|-------------|-------------|-------------|-------------|
| Full CoT | 50.4% | 89.2 | 89.8 | 88.5 | 89.1 |
| ASCoT | 66.8% | 73.1 | 63.5 | 98.2 | 77.1 |
| LLMLingua-2 | 64.5% | 75.4 | 66.8 | 96.5 | 78.9 |
| V-Skip (Ours) | 51.2% | 88.9 | 89.1 | 88.7 | 88.9 |

Table 5: Component analysis of the V-Skip gating mechanism on the MMMU dataset. We evaluate the impact of individual scoring components (S_{text} , S_{vis}) and their combinations on model performance. The table reports classification Accuracy, VARR, and Syntax Error rate (Err) at a fixed target compression ratio of $\gamma = 0.5$.

| Pruning Strategy | Acc (%) | VARR (%) \uparrow | Err (%) \downarrow |
|------------------|-------------|---------------------|----------------------|
| (A) Random | 27.5 | 49.8 | 14.5 |
| (B) Text-Only | 32.4 | 58.8 | 1.1 |
| (C) Vision-Only | 40.5 | 92.4 | 11.2 |
| (D) Intersection | 30.1 | 52.4 | 3.5 |
| (E) V-Skip Union | 48.2 | 89.0 | 1.3 |

visual recall (low VARR). Conversely, the Vision-Only method (Row C) maximizes VARR (92.4%) but disrupts syntactic fluency by pruning essential connectives. V-Skip’s Union strategy (Row E) effectively resolves this conflict, achieving the highest accuracy (48.2%) by synergizing linguistic structure with visual anchors. Notably, the Intersection strategy (Row D) underperforms (30.1%) due to overly restrictive filtering, which significantly reduces information recall. Additional sensitivity analyses are provided in Appendix C.

5 Conclusion

In this work, we identify Visual Amnesia, where text-centric pruning discards tokens essential for visual grounding, as a key efficiency bottleneck in MLLMs. To address this, we propose V-Skip, grounded in the VA-IB principle, which synergizes linguistic surprisal with cross-modal attention flow and distills this policy into a lightweight LoRA adapter. V-Skip achieves a $2.9\times$ speedup with negligible performance loss, outperforming state-of-the-art baselines by over 45% on fine-grained visual tasks while significantly mitigating object hallucination. We hope this work serves as a stepping stone for exploring alignment-based compression in broader modalities, such as dynamic video streams or audio-visual interactions.

Limitations

Despite V-Skip’s robustness, we acknowledge two limitations. First, it depends on the base model’s cross-modal alignment, potentially degrading on architectures with attention. Second, unlike training-free methods, it necessitates offline distillation, incurring additional training overhead.

Acknowledgements

This work was supported in part by the Natural Science Basis Research Plan in Shaanxi Province of China under Grant No. 2025JC-JCQN-091 and the Open Project Program of State Key Laboratory of Integrated Services Networks of Xidian University under Grant. ISN25-25.

References

- Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. 2016. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*.
- Maciej Besta, Nils Blach, Ales Kubicek, Robert Gerstenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and 1 others. 2024. Graph of thoughts: Solving elaborate problems with large language models. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 17682–17690.
- Yu-Neng Chuang, Tianwei Xing, Chia-Yuan Chang, Zirui Liu, Xun Chen, and Xia Hu. 2024. Learning to compress prompt in natural language formats. *arXiv preprint arXiv:2402.18700*.
- Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in neural information processing systems*, 35:16344–16359.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in neural information processing systems*, 35:30318–30332.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv e-prints*, pages arXiv–2407.
- Kawin Ethayarajh. 2019. How contextual are contextualized word representations. *Comparing the geometry of BERT, ELMo, and GPT-2 Embeddings*, 2.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023. Llmlingua: Compressing prompts for accelerated inference of large language models. *arXiv preprint arXiv:2310.05736*.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626.
- Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. 2024. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Liang Lin, Zhihao Xu, Junhao Dong, Jian Zhao, Yuchen Yuan, Guibin Zhang, Miao Yu, Yiming Zhang, Zhengtao Yao, Huahui Yi, and 1 others. 2025. Orthalign: Orthogonal subspace decomposition for non-interfering multi-objective alignment. *arXiv preprint arXiv:2509.24610*.
- Liang Lin, Miao Yu, Kaiwen Luo, Yibo Zhang, Lilan Peng, Dexian Wang, Xuehai Tang, Yuanhe Zhang, Xikang Yang, Zhenhong Zhou, and 1 others. 2026. Hidden in the noise: Unveiling backdoors in audio llms alignment through latent acoustic pattern triggers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 32015–32023.
- Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. 2021. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209.
- Catherine Olsson, Nelson Elhage, Neel Nanda, Nicholas Joseph, Nova DasSarma, Tom Henighan, Ben Mann, Amanda Askell, Yuntao Bai, Anna Chen, and 1 others. 2022. In-context learning and induction heads. *arXiv preprint arXiv:2209.11895*.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Menglin Xia, Xufang Luo, Jue Zhang, Qingwei Lin, Victor Rühle, Yuqing Yang, Chin-Yew Lin, and 1 others. 2024. Llmlingua-2: Data distillation for efficient and faithful task-agnostic prompt compression. *arXiv preprint arXiv:2403.12968*.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. Object hallucination in image captioning. *arXiv preprint arXiv:1809.02156*.

- Noam Shazeer. 2019. Fast transformer decoding: One write-head is all you need. *arXiv preprint arXiv:1911.02150*.
- Naftali Tishby, Fernando C Pereira, and William Bialek. 2000. The information bottleneck method. *arXiv preprint physics/0004057*.
- Lei Wang, Yi Hu, Jiabang He, Xing Xu, Ning Liu, Hui Liu, and Heng Tao Shen. 2024a. T-sciq: Teaching multimodal chain-of-thought reasoning via large language model signals for science question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19162–19170.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, and 1 others. 2024b. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Yaoting Wang, Shengqiong Wu, Yuecheng Zhang, Shuicheng Yan, Ziwei Liu, Jiebo Luo, and Hao Fei. 2025. Multimodal chain-of-thought reasoning: A comprehensive survey. *arXiv preprint arXiv:2503.12605*.
- Heming Xia, Chak Tou Leong, Wenjie Wang, Yongqi Li, and Wenjie Li. 2025. Tokenskip: Controllable chain-of-thought compression in llms. *arXiv preprint arXiv:2502.12067*.
- Ning Yang, Hai Lin, Yibo Liu, Baoliang Tian, Guoqing Liu, and Haijun Zhang. 2025. Token-importance guided direct preference optimization. *arXiv preprint arXiv:2505.19653*.
- Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. 2022. An empirical study of gpt-3 for few-shot knowledge-based vqa. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 3081–3089.
- Junyi Yao, Yijiang Liu, Zhen Dong, Mingfei Guo, Helan Hu, Kurt Keutzer, Li Du, Daquan Zhou, and Shanghang Zhang. 2024. Promptcot: Align prompt distribution via adapted chain-of-thought. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7027–7037.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models, 2023. URL <https://arxiv.org/abs/2305.10601>, 3:1.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567.
- Dongxu Zhang, Hongqiang Lin, Yiding Sun, Pengyu Wang, Qirui Wang, Ning Yang, and Jihua Zhu. 2026a. Not all queries need deep thought: Cofcot for adaptive coarse-to-fine stateful refinement. *arXiv preprint arXiv:2603.08251*.
- Dongxu Zhang, Yiding Sun, Pengcheng Li, Yumou Liu, Hongqiang Lin, Haoran Xu, Xiaoxuan Mu, Liang Lin, Wenbiao Yan, Ning Yang, and 1 others. 2026b. Pointcot: A multi-modal benchmark for explicit 3d geometric reasoning. *arXiv preprint arXiv:2602.23945*.
- Dongxu Zhang, Ning Yang, Jihua Zhu, Jinnan Yang, Miao Xin, and Baoliang Tian. 2025. Ascot: An adaptive self-correction chain-of-thought method for late-stage fragility in llms. *arXiv preprint arXiv:2508.05282*.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2023. Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.
- Peijun Zhu, Ning Yang, Jiayu Wei, Jinghang Wu, and Haijun Zhang. 2025. Breaking the moe llm trilemma: Dynamic expert clustering with structured compression. *arXiv preprint arXiv:2510.02345*.

A More details of our Algorithms

A.1 Algorithms

We summarize the complete V-Skip pipeline in Algorithm 1. The process consists of two sequential phases. For phase 1, we first utilize a frozen Teacher MLLM to generate full CoT rationales. For every token in the sequence, we compute two utility metrics, the textual information score S_{text} derived from linguistic probability, and the visual anchoring score S_{vis} derived from cross-modal attention maps (as defined in Eq. (6)). An Union-of-Saliency gate is then applied. A token is retained if it exceeds the dynamic threshold in either the textual or visual modality. This mechanism explicitly prevents Visual Amnesia by rescuing high entropy visual tokens that would be pruned by text only metrics. For phase 2, to avoid the computational overhead of calculating attention scores at runtime, we distill the pruning policy into a student model. Using the compressed rationales \hat{C} generated in Phase 1 as targets, we fine-tune the Base Student Model via LoRA. This allows the student to internalize the V-Skip logic, enabling it to directly generate concise, visually grounded rationales during inference with zero additional latency.

A.2 Qualitative Analysis: What is Skipped?

To provide an intuitive understanding of the pruning dynamics, Figure 6 visualizes the token-level compression trajectory for the DocVQA invoice case discussed in Section 4.3.2.

We compare the retention masks generated by LLMingua-2 and V-Skip at a compression ratio of $\gamma = 0.5$. Deleted tokens are marked with strikethroughs.

LLMingua-2 successfully removes linguistic fillers, but disastrously removes the answer \$45.20 due to its high linguistic perplexity. This results in an incoherent statement. V-Skip demonstrates the ability to learn semantic shortcuts. It skips the same linguistic fillers as the baseline but selectively retains the high entropy visual token \$45.20. The resulting compressed CoT (Find total look bottom right... total amount is \$45.20) retains the minimal sufficient statistics required for the final prediction.

This visualization confirms that V-Skip operates not merely as a text compressor, but as a grounding-aware filter that aligns latent reasoning with visual evidence.

Algorithm 1 The V-Skip Framework: From Dual-Path Pruning to Efficient Distillation

Require: Dataset \mathcal{D} , Teacher MLLM θ_{teach} , Base Student θ_{student} , Target Rate γ

Ensure: Efficient Student Model θ_{student}

```
1: // Phase 1: Data Construction via Dual-Path Pruning
2: Initialize training set  $\mathcal{D}_{\text{train}} \leftarrow \emptyset$ 
3: for each sample  $(\mathcal{V}, \mathcal{Q}) \in \mathcal{D}$  do
4:    $\mathcal{C} \leftarrow \text{GenerateCoT}(\theta_{\text{teach}}, \mathcal{V}, \mathcal{Q})$ 
5:   // Calculate Scores per Token  $c_t$ 
6:    $S_{\text{text}} \leftarrow -\log P(c_t | c_{<t})$ 
7:    $S_{\text{vis}} \leftarrow \text{AggAttention}(c_t \rightarrow \mathcal{V})$ 
8:   // Dynamic Thresholding
9:    $\tau_{\text{text}}, \tau_{\text{vis}} \leftarrow \text{CalibrateGate}((1 - \gamma), \{S_{\text{text}}\}, \{S_{\text{vis}}\})$ 
10:  // Apply Union Gate
11:   $\hat{\mathcal{C}} \leftarrow \emptyset$ 
12:  for  $t = 1$  to  $\text{Length}(\mathcal{C})$  do
13:    if  $S_{\text{text}}^{(t)} \geq \tau_{\text{text}} \vee S_{\text{vis}}^{(t)} \geq \tau_{\text{vis}}$  then
14:      Append  $c_t$  to  $\hat{\mathcal{C}}$ 
15:    end if
16:  end for
17:  Add  $(\mathcal{V}, \mathcal{Q}, \hat{\mathcal{C}})$  to  $\mathcal{D}_{\text{train}}$ 
18: end for
19: // Phase 2: LoRA Fine-tuning
20: Inject LoRA adapters into  $\theta_{\text{student}}$ 
21: while not converged do
22:   Sample batch from  $\mathcal{D}_{\text{train}}$ 
23:    $\mathcal{L} \leftarrow -\log P_{\theta_{\text{student}}}(\hat{\mathcal{C}} | \mathcal{V}, \mathcal{Q})$ 
24:   Update parameters via SGD
25: end while
26: return  $\theta_{\text{student}}$ 
```

B Implementation Details

B.1 Datasets

We first utilize the MMMU benchmark, a massive multi-discipline evaluation suite comprising collegiate-level tasks across disciplines such as Art, Science, and Engineering. Unlike simple recognition tasks, MMMU demands deep multimodal reasoning and expert-level knowledge application. We select this dataset to rigorously assess whether the compressed rationales generated by V-Skip maintain the logical coherence and semantic integrity required for complex problem solving. Complementing this, we leverage the DocVQA benchmark, which consists of diverse document images (e.g., invoices, forms, reports) paired with questions requiring precise Optical Character Recogni-

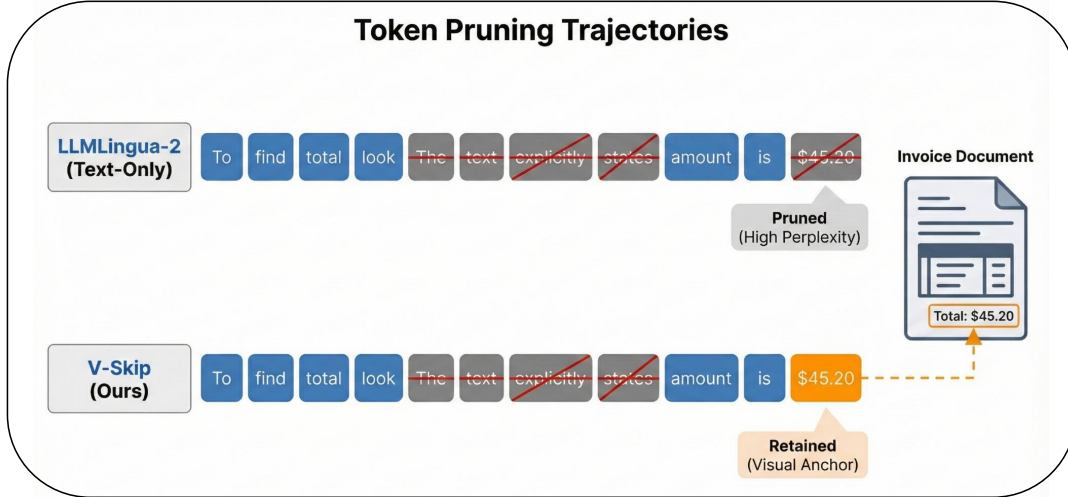


Figure 5: Qualitative comparison on DocVQA illustrating the Information Entropy Mismatch. V-Skip successfully retains the high-surprisal token \$45.20 by leveraging visual attention, whereas text-only pruning fails.

tion (OCR) and spatial layout understanding. In DocVQA, correct answers often hinge on high entropy visual tokens that appear statistically random to language models. To measure performance and controllability, we report Accuracy, End-to-End Latency, and the ActRatio.

B.2 Detailed Evaluation Metrics

To provide a comprehensive assessment of V-Skip’s performance, we report three key efficiency metrics alongside task accuracy. Here, we provide the precise definitions and measurement protocols for each metric reported in Table 1.

Tokens This metric quantifies the verbosity of the reasoning chain. For a given dataset \mathcal{D} , let $\hat{\mathcal{C}}_i$ denote the generated Chain-of-Thought sequence for the i -th sample. The Average Generated Tokens, L_{avg} , is defined as:

$$L_{\text{avg}} = \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} \text{Length}(\hat{\mathcal{C}}_i), \quad (8)$$

where $\text{Length}(\cdot)$ counts the number of tokens decoded by the model, excluding the input prompt. This metric serves as a proxy for the information density of the reasoning process. A significantly lower L_{avg} with maintained accuracy indicates a more concise and efficient reasoning path.

Latency We report the wall-clock inference time to strictly evaluate the real-world efficiency of our method. Unlike theoretical complexity measures (FLOPs), this metric captures the total computational cost, including method specific overheads.

ActRatio While we set a target compression ratio γ (e.g., 0.5 implies retaining 50% of tokens) as a hyperparameter, the realized compression often varies due to the distinct pruning logic of each algorithm. We define the ActRatio as:

$$\text{ActRatio} = \frac{\sum \text{Length}(\hat{\mathcal{C}}_{\text{compressed}})}{\sum \text{Length}(\mathcal{C}_{\text{original}})}. \quad (9)$$

The ActRatio is dynamic. Consequently, if an image requires extensive visual grounding, V-Skip may yield an ActRatio slightly higher than γ to rescue critical visual tokens. Conversely, for simple images, it may compress more aggressively. We report this metric to transparently demonstrate the distinct behavior of our visual-anchored mechanism compared to rigid text-only truncation.

Acc for MMMU The MMMU benchmark comprises multidisciplinary tasks requiring expert level reasoning, formatted primarily as multiple choice questions. We report the standard Top-1 Accuracy:

$$\text{Acc} = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\hat{y}_i = y_i), \quad (10)$$

where N is the total number of evaluation samples, \hat{y}_i is the model’s predicted option, y_i is the ground truth, and $\mathbb{I}(\cdot)$ is the indicator function.

ANLS for DocVQA For the DocVQA dataset, where answers are often extracted text spans susceptible to minor Optical Character Recognition (OCR) variations, exact match accuracy is overly penalizing. Following standard evaluation protocols for document intelligence, we employ the ANLS metric.

ANLS measures the edit distance between the predicted text sequence and the ground truth, normalized by the sequence length. For a prediction P and ground truth G , the score $s(P, G)$ is defined as:

$$\text{Let } M = \max(|P|, |G|).$$

$$s(P, G) = \begin{cases} 1 - \frac{d_L(P, G)}{M}, & \text{if } d_L(P, G) < \tau \cdot M \\ 0, & \text{otherwise} \end{cases} \quad (11)$$

where $d_L(\cdot)$ denotes the Levenshtein distance, $|\cdot|$ denotes the string length, and $\tau = 0.5$ is the threshold penalty. The final ANLS reported in Table 1 is the mean score over all samples, scaled to a percentage (0 – 100).

B.3 Dataset Splits and Evaluation

For all benchmarking experiments reported in the main text, we strictly utilize the official validation splits of both MMMU and DocVQA. This protocol ensures a consistent and fair comparison with baseline methods while preventing any potential test set leakage.

B.4 Training Configurations

We implement the V-Skip distillation process using LoRA to efficiently fine-tune the attention mechanisms of the pre-trained MLLMs. This allows the model to internalize the pruning policy with minimal computational overhead. We set the LoRA rank $r = 16$ and the scaling factor $\alpha = 32$. To fully capture the attention-based pruning logic, we apply the low-rank adapters to all linear projection layers within the self-attention modules (W_q, W_k, W_v, W_o). The LoRA dropout rate is set to 0.05 to prevent overfitting during the distillation phase.

The models are optimized using AdamW with a learning rate of 2×10^{-4} and a cosine learning rate scheduler featuring a warm-up ratio of 0.03. We fine-tune the student models for a total of 3 epochs, which we found empirically sufficient for the model to learn the V-Skip policy without degrading its generation quality. We utilize a global batch size of 16 to ensure stable convergence on our $8 \times \text{RTX } 3090$ GPU infrastructure.

C Additional Ablation Studies

In this section, we provide supplementary experiments to further characterize the robustness and implementation choices of V-Skip.

C.1 Sensitivity to Visual Layer Selection

The calculation of the Visual Anchoring Score S_{vis} relies on aggregating cross-attention weights. However, interpretability studies suggest that visual text alignment is stratified across the Transformer depth. To verify this hypothesis and identify the optimal signal source, we conduct a sensitivity analysis on the MMMU dataset, as summarized in Table 6.

Relying on shallow layers results in the 41.5% performance drop. These layers primarily encode low level visual features and basic syntactic information. Consequently, the attention maps here are too noisy to reflect high level semantic importance, leading to the erroneous pruning of critical reasoning tokens. The final layers also yield suboptimal results (46.8%). While deep layers capture rich semantics, they suffer from abstraction drift as the representation methods the output layer, it becomes highly specialized for the next token prediction objective, diluting the direct correspondence with specific image patches. The intermediate layers serve as the optimal source, achieving the peak accuracy of 48.2%. This confirms that the semantic bridging between visual embeddings and textual concepts is strongest in the middle of the network. By aggregating attention from this range, V-Skip captures the most robust alignment signals for visual grounding.

C.2 Hyperparameter Robustness: LoRA Rank Analysis

V-Skip utilizes LoRA to efficiently fine-tune the gating mechanism. To determine the optimal adapter size, we conduct a sensitivity analysis on the LoRA rank (r) ranging from 4 to 64, evaluating both classification accuracy on the MMMU dataset and computational efficiency (VRAM usage).

Table 7 illustrates the trade-off between model capacity and resource consumption. At $r = 4$, the model exhibits a performance deficit, reaching only 46.5% accuracy. This suggests that the extremely constrained parameter budget (0.04%) creates a representational bottleneck, preventing the adapter from effectively learning the complex cross-modal alignment patterns necessary for precise pruning. Increasing the rank beyond 16 yields diminishing returns. While scaling r to 32 or 64 marginally improves accuracy to 48.3% (+0.1%), it comes at the cost of doubled parameter counts and increased memory overhead. Consequently, we identify $r = 16$ as the Pareto-optimal configuration. It achieves near peak performance (48.2%)

Table 6: Sensitivity to visual layer selection on the MMMU dataset. We evaluate the impact of deriving the Visual Anchoring Score S_{vis} from different ranges of Transformer layers. The table reports the model accuracy and the performance deviation (Δ) relative to the Middle Layers setting, conducted at a fixed compression ratio of $\gamma = 0.5$.

| Layer Selection | Depth Range | Acc (%) | Drop (Δ) |
|-----------------|------------------|-------------|-------------------|
| Early Layers | 0% – 25% | 41.5 | –6.7 |
| Late Layers | 75% – 100% | 46.8 | –1.4 |
| All Layers | 0% – 100% | 47.1 | –1.1 |
| Middle (Ours) | 25% - 75% | 48.2 | – |

Table 7: Impact of LoRA Rank (r) on performance and efficiency. We evaluate the sensitivity of the model to varying adapter sizes on the MMMU dataset. The table reports the proportion of trainable parameters, classification accuracy, and peak training VRAM usage at a fixed compression ratio of $\gamma = 0.5$.

| LoRA Rank (r) | Trainable Params | Acc (%) | Training VRAM |
|-------------------|------------------|-------------|----------------|
| 4 | 0.04% | 46.5 | 18.2 GB |
| 8 | 0.08% | 47.6 | 18.3 GB |
| 16 (Ours) | 0.16% | 48.2 | 18.5 GB |
| 32 | 0.32% | 48.3 | 18.9 GB |
| 64 | 0.64% | 48.3 | 19.6 GB |

Qualitative Comparison on DocVQA

Question: What is the invoice total shown at the bottom?

Image Context: An invoice image containing the text Total: \$45.20 in small print.

Original CoT:

To find the total, look at the bottom right corner. The text explicitly states that the total amount of the invoice is **\$45.20**. Therefore, the answer is \$45.20.

LLMLingua-2 (Text-Only Pruning):

To find total look bottom right corner. ~~The text explicitly states that total amount of the invoice is \$45.20. Therefore answer is \$45.20.~~

→ Prediction: "Total amount is." (Incorrect)

V-Skip (Ours):

Find total look bottom right. ~~The text explicitly states that total amount of the invoice is \$45.20. Therefore answer \$45.20.~~

→ Prediction: "\$45.20" (Correct)

Figure 6: Qualitative Comparison. Black strikethrough text denotes pruned tokens. Note that LLMLingua-2 prunes the crucial value \$45.20 because numbers often have high perplexity in isolation. V-Skip retains it because the Visual Anchoring Score is high.

while maintaining a minimal memory footprint (18.5 GB), demonstrating that V-Skip is highly parameter efficient and does not require extensive fine-tuning parameters to function effectively.