

Investigating Length Bias and Robustness in LVLMs for Multiple-Choice Question Answering

Md. Atabuzzaman Hani Alomari Chris Thomas

Department of Computer Science

Virginia Tech

{atabuzzaman, hani, christhomas}@vt.edu

Abstract

Large Vision-Language Models (LVLMs) demonstrate strong performance on multiple-choice question answering (MCQA), yet their reliability under biased conditions remains underexplored. We investigate length bias, the systematic preference for longer answer options, through controlled benchmarks spanning fine-grained visual domains. LVLMs consistently favor longer options even when visually inconsistent, with accuracy dropping 2–38% when correct answers are shorter and increasing 1–70% when longer. This bias persists across model scales. Through layer-wise mechanistic analysis, we reveal that length bias emerges from late-layer suppression of visual evidence in favor of language-side length priors. Blurry-image experiments show that bias intensifies when visual evidence weakens, while abstention tests reveal that LVLMs struggle to refuse incorrect predictions. We propose and evaluate multiple inference-time mitigation strategies; while these provide partial improvements, none achieves deployable bias mitigation without relying on oracle information unavailable at inference time. These results underscore that inference-time interventions are insufficient, pointing to the need for architectural solutions and bias-aware training.¹

1 Introduction

Large Vision-Language Models (LVLMs) (Liu et al., 2024, 2023; Li et al., 2023; Dai et al., 2023; Deitke et al., 2025; Bai et al., 2025b; Chen et al., 2025b; OpenAI, 2024; Chen et al., 2025a) have achieved impressive results across multimodal tasks including visual question answering (VQA), captioning, and fine-grained visual understanding (Kim and Ji, 2024; Atabuzzaman et al., 2025b). However, recent evidence suggests these models exhibit systematic biases that compromise their reliability in structured decision-making tasks. In

¹Code and Dataset: https://github.com/Atabuzzaman/Length_bias_robustness




Figure 1: **Overview of our framework** for investigating length bias in LVLMs. We create length-controlled MCQA variants to systematically expose length bias, conduct layer-wise analysis revealing late-layer suppression of visual evidence, perform diagnostic experiments using abstention and blurry images to localize bias sources, and evaluate mitigation strategies like ELED, caption-guided MCQA and Yes/No QA.

large language models (LLMs), biases in multiple-choice settings include preferences for specific answer positions, verbose options, and particular formatting patterns (Pezeshkpour and Hruschka, 2024; Zheng et al., 2023; Zhao et al., 2025). While early work suggests LVLMs may inherit or amplify such tendencies (Adila et al., 2024; Zong et al., 2024; Atabuzzaman et al., 2025a), the extent and mechanisms of these biases remain underexplored.

MCQA provides an ideal testbed for investigating these biases. By constraining outputs to a fixed set of options, MCQA enables controlled analysis of whether models base predictions on genuine visual-semantic understanding or rely on superficial heuristics. Recent work has identified arbitrary token and position biases in LVLM responses (Adila et al., 2024; Atabuzzaman et al., 2025a), suggesting reliance on textual patterns rather than grounded reasoning. However, a comprehensive understanding of systematic preferences in LVLMs, particularly length bias and its interaction with visual reasoning, remains limited.

In this work, we conduct a systematic investigation of **length bias in LVLMs**: the tendency to

prefer longer answer options regardless of visual evidence or correctness. Through carefully controlled benchmarks across five fine-grained visual domains, we demonstrate that LVLMs exhibit severe length bias, with accuracy dropping 2–38% when correct answers are shorter and increasing 1–70% when longer. We further conduct layer-wise probing analysis, revealing that while early layers encode sufficient visual evidence for correct answers, late decoder layers systematically suppress this evidence when longer options are present. This demonstrates that bias operates through late-layer override of visual evidence by language-side priors rather than initial visual encoding failures. This bias persists across model scales from 3B to 32B parameters and affects even state-of-the-art proprietary models, indicating deeply entrenched priors that extend beyond simple capacity limitations.

Beyond identifying length bias, we probe LVLM robustness under challenging conditions. To investigate whether bias originates primarily from visual processing or language-side reasoning, we conduct blurry-image experiments and find that length bias *intensifies* under degraded visual input. This reveals that language-side heuristics become the dominant fallback when visual evidence is weak or ambiguous. Abstention tests using "None of the Above" options reveal poor abstention capability, with performance collapsing to 3–33.5% when models should refuse to answer.

We propose and evaluate multiple mitigation strategies inspired by human reasoning, including (1) Early-Layer Ensemble Decoding (ELED), which combines predictions from early and late layers through adaptive decoding, (2) caption-guided answering with explicit visual grounding, (3) independent Yes/No evaluation of each option, and (4) selection by rejection, where models identify and discard incorrect options. While these approaches provide partial improvements, none achieves deployable bias mitigation in practice, revealing that inference-time interventions are insufficient as a practical solution. This suggests that architectural-level changes and bias-aware training are needed to fully address this bias.

We make the following contributions:

- We introduce controlled MCQA benchmarks across five fine-grained visual domains that systematically expose length bias in LVLMs.
- We localize length bias to late decoder layers, demonstrating systematic override of visual evidence by language-side priors.

- We conduct diagnostic analyses showing that bias intensifies under degraded visual input and LVLMs exhibit poor abstention capability.
- We propose and evaluate multiple mitigation strategies (ELED, caption-guided MCQA, Yes/No Question-Answering, selection-by-rejection), revealing that inference-time methods are insufficient in practice.

2 Related Work

Recent research has identified systematic biases in LLMs that affect their reliability in MCQA tasks, showing that models consistently favor certain tokens, positions, or options presented early in the sequence (Zheng et al., 2023; Pezeshkpour and Hruschka, 2024; Wang et al., 2025). Robinson and Wingate (2023) showed that multiple-choice question (MCQ) presentation format influences model behavior, with performance varying significantly between prompts with explicit option identifiers versus cloze-style formats. Research suggests models struggle to effectively map option identifiers to their corresponding content (Xue et al., 2024) or exhibit positional sensitivity combined with model uncertainty (Pezeshkpour and Hruschka, 2024), with bias manifesting particularly strongly in knowledge-intensive scenarios (Li et al., 2024). Mitigation strategies include neuron pruning (Yang et al., 2024), elimination of biased feed-forward network (FFN) components and attention heads (Zhou et al., 2024), probability calibration (Wei et al., 2024), majority voting (Guda et al., 2025), and causal debiasing (Yang et al., 2025). Ma and Du (2023) proposed a process-of-elimination approach that discards low-probability incorrect options before masking them to derive final predictions. Similarly, Fu et al. (2025) introduced an exclusion-of-thought method combined with chain-of-thought prompting to reduce cognitive load through systematic option elimination in MCQA tasks. Zhao et al. (2025) showed that LLMs are biased toward verbose options and that fine-tuning fails to mitigate this length bias in MCQA.

Recent work has explored bias mitigation in LVLMs through various approaches, including data augmentation (Gokhale et al., 2020), model editing (Cheng et al., 2023; Wang et al., 2024), and output post-processing (Wang et al., 2021; Zhang et al., 2025). Approaches to address language priors include output probability calibration for classification tasks (Zhang et al., 2025) and causal graph-

based interventions on both visual and textual inputs for VQA (Chen et al., 2024). Researchers have explored activation-level interventions to reduce bias (Adila et al., 2024) and developed logit-based correction methods to address selection bias across varying task difficulty (Atabuzzaman et al., 2025a). However, no prior work has systematically examined length bias in LVLM-based MCQA through controlled experiments that isolate option length while maintaining visual accuracy. We address this gap through layer-wise mechanistic analysis, diagnostic experiments with degraded visual inputs and abstention tests, and evaluation of multiple mitigation strategies.

3 Benchmark Construction

Base Dataset Collection. We build our length-controlled MCQA benchmark using hard-difficulty questions from the fine-grained visual MCQA benchmarks of Atabuzzaman et al. (2025a). This dataset covers five fine-grained visual domains: CUB-200-2011 (200 bird species) (Wah et al., 2011), Stanford Dogs (120 dog breeds) (Khosla et al., 2011), FGVC Aircraft (102 aircraft variants) (Maji et al., 2013), Stanford Cars (196 car models) (Krause et al., 2013), and Food-101 (101 food categories) (Bossard et al., 2014). Each base MCQ includes a question, four answer options representing different fine-grained classes within the same domain (e.g., different bird species or car models), and a single correct answer corresponding to the class depicted in the image.

Length-Controlled Variant Generation. To systematically investigate LVLMs’ length bias, we use GPT-4o (OpenAI, 2024) to generate four length-controlled variants for each base MCQ while preserving semantic content and class structure. The base datasets already include two versions of each MCQ: one **with class names** explicitly mentioned (e.g., "This image shows a Yellow-headed Blackbird which has...") and one **without class names** that provides only visual descriptions (e.g., "This image shows a bird which has a black body, bright yellow head..."). We maintain both versions across all variants to analyze how bias changes with explicit class information. The with class name version examines whether domain knowledge can mitigate length bias, while the without class name version forces reliance on visual grounding, potentially increasing susceptibility to bias. We use carefully designed prompts to ensure

consistency and validity across generated samples. We create the following four variants:

Equal Long: All four options have approximately the same length and contain detailed descriptions of the image class. This serves as a balanced baseline without length bias.

Equal Short: All options are short with approximately the same length (roughly half of Equal Long) and contain 2-3 visually identifiable attributes. This tests model performance with minimal descriptive content.

Shorter Correct: The correct answer is short (matching the Equal Short length) and includes 2-3 visually identifiable attributes that distinguish it from other options, while all three incorrect options are long (matching the Equal Long length). This configuration tests whether models exhibit bias against selecting shorter options.

Longer Correct: The correct answer is long (matching the Equal Long length) and includes multiple visually identifiable attributes that distinguish it from other options, while all three incorrect options are short (matching the Equal Short length). This configuration tests whether models exhibit bias toward selecting longer options.

To ensure high-quality variants that accurately test for length bias without introducing confounding factors, we manually inspect a sample of generated variants and enforce strict prompting constraints for semantic preservation, visual consistency, and structural uniformity (detailed in Appendix A.1). Additionally, for the Dogs dataset, we include samples from easy and hard difficulty categories to provide broader coverage of task complexity. **Our final datasets consist of 5,976 MCQs** across five domains with four length-controlled variants and two class-name settings per question.

4 Experimental Setup

We evaluate length bias, robustness, and mitigation strategies using both open-source and proprietary LVLMs. Open-source models include Qwen3-VL-30B-A3B-Instruct (Bai et al., 2025a), Qwen2.5-VL-3B-Instruct, 7B-Instruct, and 32B-Instruct (referred to as Qwen3-VL-30B, Qwen2.5-VL-3B, Qwen2.5-VL-7B, and Qwen2.5-VL-32B) (Bai et al., 2025b), InternVL_2.5-8B (Chen et al., 2025b), LLaVA-v1.5-7B, and LLaVA-v1.5-13B (Liu et al., 2023). Proprietary models include OpenAI’s GPT-4o, o4-mini (with reasoning) (OpenAI, 2024), and Claude Sonnet-




Figure 2: **Accuracy (%) of LVLMs across five datasets under different option-length conditions**, with and without class names. The vertical line at "Equal Long" indicates baseline performance for each LVLm-dataset pair.

4.5 (used as C-Sonnet-4.5) (with reasoning) (Anthropic, 2025). For each MCQ, we use 5 images per class for open-source models and 3 images per class for proprietary models due to API cost constraints. For example, in the CUB dataset (200 classes) with class names, we evaluate 1,000 images (200×5) per variant for open-source models and 600 images (200×3) for proprietary models. We apply this protocol across all experiments.

Dataset	Model	Equal Long	Equal Short	Shorter Correct	Longer Correct
CUB (W/o)	GPT-4o	71.17	62.83	45.17	92.83
	o4-mini	71.67	66.00	44.33	91.00
	C-Sonnet-4.5	65.33	54.50	54.50	80.00
Dogs (W/o)	GPT-4o	86.11	75.56	55.83	97.78
	o4-mini	86.11	73.61	46.94	96.67
	C-Sonnet-4.5	68.06	69.17	57.50	86.94
CUB (W.)	GPT-4o	89.50	86.33	80.33	95.83
	o4-mini	81.67	81.33	73.33	89.83
	C-Sonnet-4.5	72.17	72.00	72.67	85.17
Dogs (W.)	GPT-4o	85.28	85.28	78.06	94.72
	o4-mini	80.56	80.00	70.28	90.56
	C-Sonnet-4.5	68.06	70.83	65.83	85.83

Table 1: **Accuracy (%) of advanced proprietary LVLms** on CUB and Dogs datasets with (W.) and without (W/o) class names in the answer options.

5 Length Bias in LVLms

To investigate the presence of systematic response preferences in LVLms, we analyze how answer selection varies with the relative length of candidate options in MCQA settings. As shown in Figure 2 (detailed results in Table 8, Appendix A.8), all evaluated LVLms exhibit measurable length bias, though its magnitude varies by model scale, dataset, and presence of class names within the options.

LVLms Overwhelmingly Prefer Longer Options, Even When Incorrect. Across all LVLms and datasets, models overwhelmingly prefer longer options, even when incorrect. This is most appar-

ent in the *Shorter Correct* condition, where accuracy drops sharply. For example, LLaVA-v1.5-13B drops by 24.10 percentage points (pp) on CUB (without class names) and 10.83% on Dogs (with class names), while Qwen2.5-VL-7B suffers losses of 38.30% on CUB and 15.10% on Cars. Conversely, in the *Longer Correct* setting, accuracy improves dramatically, often by 40–70%, indicating reliance on textual length rather than visual evidence.

Class Names Reduce but Do Not Eliminate Length Bias. Including class names in candidate options reduces length bias for all LVLms. On CUB, the performance drop in "Shorter Correct" reduces from 38.30% to 25.40% for Qwen2.5-VL-7B and from 24.10% to 9.20% for LLaVA-v1.5-13B when class names are included. However, length bias persists across all datasets even with class names, indicating that when visual distinctions are subtle, models rely on surface-level textual patterns rather than visual evidence.

Scaling Models Improves Accuracy but Does Not Reduce Length Bias. Comparing model families across parameter scales reveals that increasing model size consistently improves accuracy, particularly in Equal Long settings. On Cars (without class names), Qwen2.5-VL improves from 64.80% (3B) to 73.16% (32B) under "Equal Long" options. However, even the strongest models experience substantial drops when the correct option is shorter, and both LLaVA variants show 50–70% performance boosts in "Longer Correct" conditions. This indicates scaling primarily enhances discriminative ability under balanced conditions but does not address reliance on option length as a heuristic.

Proprietary and Advanced Models Demonstrate Reduced but Persistent Length Bias. Tables 1 and 8 show that advanced models (GPT-4o, o4-

mini, Claude Sonnet-4.5, Qwen3-VL-30B) substantially outperform smaller open-source models while exhibiting reduced length bias. With class names, GPT-4o achieves 89.50% on CUB "Equal Long" compared to 76.10% for Qwen2.5-VL-7B and maintains higher "Shorter Correct" accuracy (80.33% vs 50.70%). The gap between "Shorter Correct" and "Longer Correct" narrows significantly: GPT-4o shows 15.50 points on CUB and Claude Sonnet-4.5 shows 12.50 points, compared to 43.20 points for Qwen2.5-VL-7B. However, even these advanced models exhibit persistent length preference, with "Longer Correct" consistently outperforming "Shorter Correct," indicating that scaling and advanced training reduce but do not eliminate length bias.

Length Bias Reflects a Learned Statistical Prior.

The persistence of length bias across models, datasets, and scaling indicates it reflects a learned statistical prior rather than a decoding artifact. Option length functions as an implicit confidence cue that often outweighs visual evidence when distinctions are subtle, revealing that models rely on surface-level heuristics under uncertainty.




Figure 3: **Layer-wise preference for long distractors in the Shorter Correct setting using QWEN2.5-VL-7B-INSTRUCT.** Mean preference $\Delta\ell = P(\text{long incorrect}) - P(\text{short correct})$ across 200 CUB examples via logit-lens decoding of hidden states (95% CI shaded). Early layers show near-zero systematic preference for long options; mid layers exhibit a positive shift, indicating emergence of length bias in internal representations prior to decoding. Dashed line: *stable inversion layer* ($L = 14$), the first layer at which the mean $\Delta\ell$ exceeds zero for three consecutive layers. Dotted line: *median commitment layer* ($L = 24$), beyond which the predicted answer remains unchanged.

6 Where Length Bias Emerges in LVLMM Computation

To understand *where* length bias arises inside LVLMMs, we perform layer-wise analysis using

logit-lens decoding: at each layer, we apply the language model head (`lm_head`) to the hidden state, restricted to the option-identifier tokens, to obtain a probability distribution over answer options in QWEN2.5-VL-7B. At each layer, we compute the model’s probability assigned to a fixed long distractor (the longest of the three incorrect options) and to the short correct answer, tracking their difference $\Delta\ell = P(\text{long incorrect}) - P(\text{short correct})$.

As shown in Figure 3, early layers exhibit near-zero or slightly negative values of $\Delta\ell$, indicating that initial representations do not favor longer options. Beginning around the stable inversion layer ($L = 14$), the mean $\Delta\ell$ becomes consistently positive, revealing a *semantic inversion* where long distractors overtake the short correct option within the model’s internal computation on average. This transition occurs well before final decoding, demonstrating that length bias is introduced at the representation level rather than at final generation.

A representative individual example illustrates this progression at the single-example level: early layers favor the short correct option, but at layer 7, predictions first switch to a long distractor. By layer 11, the long distractor dominates consistently, and after layer 25, the model reaches a *commitment point* where the predicted answer no longer changes. The final layer assigns substantially higher probability to the long distractor ($P = 0.65$) than to the short correct option ($P = 0.31$), demonstrating clear length-biased failure despite early evidence favoring the correct answer. Notably, $\Delta\ell$ decreases after the commitment layer, reflecting a shift to probability sharpening rather than option re-evaluation. Section 7 analyzes this late-layer conflict resolution process in detail, demonstrating systematic override of visual evidence.

7 Vision–Language Conflict Resolution in Late Layers

To understand the mechanism underlying late-layer override of visual evidence, we analyze how competing signals evolve across all decoder layers. We use prediction correctness as an operational proxy for alignment with visual evidence. For 200 examples from the Shorter Correct split, we extract probabilities at each layer by applying the language model head to hidden states, tracking the short correct answer versus the average probability of each long distractor (mean of three long options). Both signals begin near the unbiased baseline of 0.25,




Figure 4: **Vision-language conflict resolution in late layers.** Layer-wise probabilities for the short correct answer (green) versus average probability of each long distractor (red), both starting at unbiased baseline 0.25 (dotted line). **(a)** In length-biased errors ($N=144$, 72%), late layers (20–28, shaded) amplify long distractors to 0.31 while suppressing the correct answer to 0.10—each long option becomes $3.1\times$ more likely than correct. **(b)** When predictions align with ground truth (visual wins) ($N=56$, 28%), the correct answer amplifies to 0.73 while each long distractor suppresses to 0.06 ($8.0\times$ preference ratio). Shaded regions: 95% confidence intervals. Despite stronger amplification capacity for correct answers (+0.48 vs +0.06), the model utilizes this in only 28% of cases.

enabling direct comparison of their evolution.

Figure 4 reveals distinct conflict resolution patterns based on prediction outcomes. In length-biased errors (72%, $N=144$), late layers (20–28) amplify each long distractor to 0.31 while suppressing the correct answer to 0.10—each long option becomes $3.1\times$ more likely than correct. When predictions align with ground truth (28%, $N=56$), the pattern reverses dramatically: the correct answer amplifies to 0.73 while each long distractor suppresses to 0.06, yielding an $8.0\times$ preference ratio. Critically, the correct answer demonstrates stronger amplification capacity (+0.48 above baseline) than length bias (+0.06), yet this capacity is utilized in only 28% of cases. This demonstrates that late layers perform winner-take-all conflict resolution but systematically fail to privilege signals aligned with visual evidence, revealing a failure mode where the same architectural components inconsistently resolve vision-language conflicts.

8 Vision vs. Language Contributions to Length Bias

To causally test whether the conflict resolution failure in Section 7 stems from insufficient visual information versus language-side heuristics, we degrade input images using Gaussian blur ($\sigma = 12$), reducing fine-grained details while preserving coarse structure. We test whether degraded visual input (1) intensifies length bias in standard MCQA, and (2) triggers appropriate abstention when a "None of the Above" option is available.

Blurry Images Degrade Performance with

Dataset	Setting	Model	Equal Long	Equal Short	Shorter Correct	Longer Correct
CUB (W/o)	w/o NoA	Qwen2.5-VL-7B	35.70	34.80	11.70	83.70
		InternVL_2.5-8B	32.10	31.90	21.90	74.30
	NoA w/C	Qwen2.5-VL-7B	25.20	18.60	11.60	66.40
		InternVL_2.5-8B	2.30	4.40	1.80	10.80
Dogs (W/o)	w/o NoA	Qwen2.5-VL-7B	34.83	36.83	14.50	82.67
		InternVL_2.5-8B	32.50	40.00	26.17	76.67
	NoA w/C	Qwen2.5-VL-7B	24.00	18.50	11.33	62.67
		InternVL_2.5-8B	2.83	5.50	2.00	11.83
CUB (W.)	w/o NoA	Qwen2.5-VL-7B	39.30	41.50	16.00	85.70
		InternVL_2.5-8B	30.90	34.60	18.20	80.10
	NoA w/C	Qwen2.5-VL-7B	20.50	17.00	10.00	60.80
		InternVL_2.5-8B	0.70	2.70	0.90	10.20
Dogs (W.)	w/o NoA	Qwen2.5-VL-7B	35.83	36.83	16.50	80.50
		InternVL_2.5-8B	34.33	39.67	18.00	76.67
	NoA w/C	Qwen2.5-VL-7B	22.00	15.50	8.83	57.17
		InternVL_2.5-8B	3.00	4.67	1.67	12.67

Table 2: **Accuracy (%) of LVLMs on blurry image evaluation** across CUB and Dogs datasets with (W.) and without (W/o) class names in options. "w/o NoA" denotes standard MCQ format without "None of the Above" options. "NoA w/C" includes "E. None of the above" while the correct answer remains present among A–D, testing over-abstention. For each setting, **bold** marks the **lowest** accuracy between the two models.

Mixed Abstention Behavior. Table 2 shows that blurred images substantially reduce accuracy and intensify length bias. For Qwen2.5-VL-7B on CUB without class names, performance drops from 64.90% to 35.70% in "Equal Long," and the gap between "Shorter Correct" (11.70%) and "Longer Correct" (83.70%) widens to 72 points versus 66.80 under clear images. Models show divergent abstention behavior: InternVL_2.5-8B appropriately selects NoA at high rates (85–97%), while Qwen2.5-VL-7B severely under-abstains (5–10%), continuing to rely on length heuristics despite degraded visual evidence.

Visual Degradation Accelerates Length Bias

Emergence. Layer-wise analysis under visual degradation (Appendix A.2) reveals that the median override layer (Figure 8) shifts from $L = 19$ (clear images) to $L = 11$ (blurred images), demonstrating that stronger visual evidence delays—but does not prevent—bias emergence. Attention analysis (Appendix A.3) further shows that long incorrect options query visual features more than short correct answers, a difference partly attributable to their greater token length. Despite this increased visual attention, models still select options with visually inconsistent attributes, revealing systematic override of visual evidence. Blur accelerates this override by weakening visual evidence that contradicts incorrect attributes.

Length Bias Is Language-Driven with Architecture-Dependent Robustness. The intensification of length bias under degraded visual input reveals that language-side heuristics dominate when visual evidence weakens. While length bias is fundamentally language-driven, abstention capability varies by architecture (InternVL: 85–97% NoA; Qwen: 5–10%), with implications for safe deployment where input quality cannot be guaranteed.

9 Mitigation Methods

Having localized length bias to late decoder layers, we evaluate whether it can be mitigated at inference time, considering the following approaches.

9.1 Output-Level Calibration Baselines

We adapt two widely used calibration techniques from the LLM MCQA literature as baselines, and are the first to systematically evaluate them for length bias in LVLM-based MCQA. Both operate on the model’s output distribution over options, without accessing intermediate representations.

Pointwise Mutual Information (PMI)-based scoring. Following Holtzman et al. (2021), we reweight each option by its unconditional likelihood to counter surface-form competition. For an option o_k , we score

$$\text{PMI}(o_k) = \log P(o_k | I, Q) - \log P(o_k), \quad (1)$$

where $P(o_k | I, Q)$ is the conditional probability given image I and question Q , and $P(o_k)$ is the option’s prior probability estimated without visual grounding. This penalizes options that are intrinsically likely regardless of the image, partially offsetting the advantage of longer, more elaborate options.

Dataset	Method	Shorter Correct	Longer Correct	Avg. Gain
CUB (W/o)	Baseline	26.6	93.4	–
	PMI	46.5	74.5	0.5
	Len. Norm ($\alpha=1.0$)	41.9	85.8	3.85
	Len. Norm ($\alpha=0.1$)	30.5	92.2	1.35
Dogs (W/o)	Baseline	33.2	90.8	–
	PMI	48.5	73.5	-1.0
	Len. Norm ($\alpha=1.0$)	48.5	82.3	3.4
	Len. Norm ($\alpha=0.1$)	36.3	90.0	1.15
CUB (W.)	Baseline	50.7	93.9	–
	PMI	69.1	75.7	0.1
	Len. Norm ($\alpha=1.0$)	64.7	89.2	4.65
	Len. Norm ($\alpha=0.1$)	53.5	93.3	1.1
Dogs (W.)	Baseline	58.8	93.3	–
	PMI	69.7	78.5	-1.95
	Len. Norm ($\alpha=1.0$)	67.0	86.7	0.8
	Len. Norm ($\alpha=0.1$)	58.7	92.3	-0.55

Table 3: Comparison of output-level calibration baselines on Qwen2.5-VL-7B across CUB and Dogs, with (W.) and without (W/o) class names. **Avg. Gain** is the mean change over the Baseline across the Shorter Correct and Longer Correct splits. Both PMI and length normalization improve Shorter Correct at the cost of degrading Longer Correct, yielding small or negative average gains.

Length normalization. We adjust conditional scores for token length by normalizing the log-probability of each option by its length raised to a power α :

$$\text{score}(o_k) = \frac{\log P(o_k | I, Q)}{|o_k|^\alpha}, \quad (2)$$

where $|o_k|$ is the token length of option o_k and α controls the strength of the correction. We report two settings, $\alpha = 1.0$ (full normalization) and $\alpha = 0.1$ (mild normalization), to characterize the method’s sensitivity to this hyperparameter.

9.2 Early-Layer Ensemble Decoding (ELED)

Our layer-wise analysis reveals that length bias emerges progressively through the network rather than being encoded initially. For Qwen2.5-VL-7B, the mean $\Delta\ell$ trajectory reaches its stable inversion layer at $L = 14$, with a median commitment layer of $L = 24$ (Section 6). For InternVL_2.5-8B, bias remains near-zero through layer 16 before explosive growth from layers 17–28. This temporal separation suggests a natural mitigation: decode from earlier layers before bias dominates.

We first test whether decoding from pre-inversion layers can counteract length bias using a single fixed rule applied uniformly, with no information about which option length is correct. This rule averages logits over early layers (5–12 for Qwen2.5-VL-7B) and defers to the early-layer pre-

Algorithm 1 Split-Aware ELED

Require: Hidden states $\{h_\ell\}_{\ell=0}^{L_{\max}}$, option texts $\{o_1, \dots, o_K\}$ with corresponding option-identifier token ids $\{t_1, \dots, t_K\}$, language model head $\text{lm_head}(\cdot)$, early layers $\mathcal{L}_{\text{early}}$, threshold $\tau = 1.2$, oracle split label $s \in \{\text{shorter correct, longer correct, equal}\}$

Ensure: Prediction $\hat{y}$

```

1:  $z_{\text{early}} \leftarrow \text{lm\_head}\left(\frac{1}{|\mathcal{L}_{\text{early}}|} \sum_{\ell \in \mathcal{L}_{\text{early}}} h_\ell\right)$ 
2:  $z_{\text{final}} \leftarrow \text{lm\_head}(h_{L_{\max}})$ 
3:  $p_{\text{early}} \leftarrow \text{softmax}([z_{\text{early}}]_{t_1}, \dots, [z_{\text{early}}]_{t_K})$ 
4:  $p_{\text{final}} \leftarrow \text{softmax}([z_{\text{final}}]_{t_1}, \dots, [z_{\text{final}}]_{t_K})$ 
5:  $\hat{y}_{\text{early}} \leftarrow \arg \max p_{\text{early}}, \hat{y}_{\text{final}} \leftarrow \arg \max p_{\text{final}}$ 
6: if  $\hat{y}_{\text{early}} = \hat{y}_{\text{final}}$  then
7:   return  $\hat{y}_{\text{final}}$   $\triangleright$  High confidence, no intervention
8: end if
9:  $\ell_{\text{early}} \leftarrow \text{length}(o_{\hat{y}_{\text{early}}}), \ell_{\text{final}} \leftarrow \text{length}(o_{\hat{y}_{\text{final}}})$ 
10: if  $s = \text{shorter correct}$  and  $\ell_{\text{final}} > \tau \cdot \ell_{\text{early}}$  then
11:   return  $\hat{y}_{\text{early}}$   $\triangleright$  Length bias detected; correct toward short
12: else if  $s = \text{longer correct}$  and  $\ell_{\text{early}} > \tau \cdot \ell_{\text{final}}$  then
13:   return  $\hat{y}_{\text{early}}$   $\triangleright$  Anti-length bias detected; correct toward long
14: else
15:   return  $\hat{y}_{\text{final}}$   $\triangleright$  Default to mature representations
16: end if
```

diction whenever it disagrees with the final layer in a length-correlated way. Applied uniformly, this rule improves accuracy on Shorter Correct from 26.6% to 45.5% but substantially degrades Longer Correct from 93.4% to 24.7%: the same rule that overrides spuriously long distractors also overrides genuinely long correct answers, since early layers cannot distinguish the two without additional information. This mirrors the one-directional trade-off shown by the output-level calibration baselines (Table 3), motivating a split-aware gating mechanism that applies the correction selectively.

Method. Split-aware ELED averages logits across early layers and applies gating conditioned on an oracle split label s indicating whether the correct answer is expected to be shorter or longer than the distractors (Algorithm 1). For Qwen2.5-VL-7B, we use layers 5–12 (before the stable inversion layer at $L = 14$); for InternVL_2.5-8B, we use layers 5–16 (before explosive onset at $L = 17$). Conditioning on s allows the gate to apply the early-layer correction only in the direction indicated by the split, avoiding the one-directional trade-off observed above. Since s is not available at inference time in deployment, we report these results as an oracle upper bound characterizing the ceiling achievable if the gating direction were known.

Results. Table 4 presents split-aware ELED results on CUB and Stanford Dogs with and without class names. On *Shorter Correct*, split-aware

Dataset	Model	Shorter Correct	Longer Correct	Δ
CUB (W/o)	Qwen2.5-VL-7B	48.40 (+21.80)	95.00 (+1.60)	46.60 (-20.20)
	InternVL_2.5-8B	52.10 (+14.30)	90.40 (+2.10)	38.30 (-12.20)
Dogs (W/o)	Qwen2.5-VL-7B	46.49 (+13.32)	91.47 (+0.64)	44.98 (-12.68)
	InternVL_2.5-8B	61.17 (+23.00)	81.67 (+2.00)	20.50 (-21.00)
CUB (W.)	Qwen2.5-VL-7B	61.50 (+10.80)	94.70 (+0.80)	33.20 (-10.00)
	InternVL_2.5-8B	52.20 (+13.00)	90.40 (+2.10)	38.20 (-10.90)
Dogs (W.)	Qwen2.5-VL-7B	64.05 (+5.22)	94.31 (+0.98)	30.26 (-4.24)
	InternVL_2.5-8B	57.83 (+14.16)	81.83 (+2.50)	24.00 (-11.66)

Table 4: **Split-aware ELED performance** for length bias mitigation, using oracle split information s (Algorithm 1; see Section 9.2 for details). Values in parentheses show changes over baseline.

ELED achieves substantial improvements: +21.8% (Qwen/CUB W/o), +14.3% (InternVL/CUB W/o), +13.3% (Qwen/Dogs W/o), and +23.0% (InternVL/Dogs W/o). Critically, *Longer Correct* performance also improves (+0.6–2.5%): conditioning the gate on the oracle split label s allows early layers to correct in the appropriate direction for each split, rather than applying a single rule uniformly (cf. the motivation analysis above). The overall bias gap ($\Delta = \text{Longer} - \text{Shorter}$) reduces substantially: from 66.8% to 46.6% on Qwen/CUB (30.2% relative reduction) and from 50.5% to 38.3% on InternVL/CUB (24.2% relative reduction). For Equal-length conditions (Equal Long, Equal Short), the gating rarely triggers since all options have similar lengths, maintaining baseline performance. As Table 4 reflects an oracle upper bound, these results characterize the best achievable outcome given accurate gating direction, rather than a deployable inference-time method.

Validation of mechanistic insights. Split-aware ELED’s effectiveness validates our layer-wise analysis findings. First, substantial improvements on *Shorter Correct* confirm that early layers (5–12 for Qwen, 5–16 for InternVL) encode less biased representations, directly validating our logit-lens measurements showing near-zero Δ_ℓ in these layers. Second, improvements emerging specifically when gating detects length-correlated disagreements demonstrate that the observed bias corresponds to functionally relevant decision changes, not measurement artifacts. Third, consistent gains across architectures indicate this is a general property of current LVLM training, not model-specific.

Comparison with other mitigation strategies. Beyond ELED, we evaluate three additional strategies inspired by human reasoning: caption-guided MCQA (Appendix A.5), Yes/No question answering (Appendix A.6), and MCQA with selection by rejection (Appendix A.7).




Figure 5: Accuracy (%) of LVLMs on "None of the Above" (NoA) evaluation settings across CUB datasets, reported both with and without class names. "w/o NoA" denotes the original MCQ format without a NoA option. "NoA w/C" appends a fifth option, "E. None of the above", while keeping the correct answer present among A–D (over-abstention test). "NoA wo/C" removes the correct answer and uses "D. None of the above." as the only valid choice (forced-abstention test).

10 Evaluating Abstention: When LVLMs Should Refuse to Answer

To probe LVLm robustness beyond length bias, we introduce a "None of the Above" (NoA) evaluation setting that tests models' ability to abstain when no valid option is present. This is essential for safe decision-making in high-stakes applications where overconfident guessing can be harmful. We design two complementary NoA configurations to distinguish between false abstention (unnecessarily choosing NoA) and failure to abstain (guessing when abstention is appropriate):

Appended-NoA (Over-Abstention Test): We keep the original four options (A–D) unchanged and add a fifth option, "E. None of the above." The correct answer remains present among A–D. This evaluates whether models unnecessarily abstain when a valid answer is available, revealing tendencies toward excessive caution or uncertainty.

Removed-Correct (Forced-Abstention Test): We remove the correct answer from the candidate set and shuffle the remaining incorrect options into positions A–C, while "D. None of the above" becomes the only valid response. This assesses whether models can correctly abstain when appropriate rather than guessing incorrectly.

Results. We evaluate LVLm abstention capability under three settings. Figure 5 presents representative results; comprehensive results are provided in Appendix A.9 (Figure 11 and Table 9).

Models Exhibit Over-Abstention When NoA Is Added to Valid Options. When we append a NoA option while the correct answer remains avail-

able (NoA w/C), models show measurable over-abstention. For example, InternVL_2.5-8B drops from 61.4% to 55.7% on CUB, indicating inappropriate selection of the abstention option.

Models Struggle to Abstain When No Correct Answer Exists. When the correct answer is removed (NoA wo/C), performance collapses. Without class names, Qwen2.5-VL-7B falls to 3.2% on CUB and 3.83% on Dogs, indicating a near-total failure to select the abstention option even when it is the only valid answer, while InternVL_2.5-8B achieves 17.7% and 33.5%, respectively. With class names, performance improves to 20.0% and 11.33% but remains far below acceptable levels, indicating models default to guessing even when semantic grounding cues are available.

Models Handle Abstention Reliably Only in Trivial Cases. In Dogs (Easy), where distractors come from unrelated domains, both models achieve near-perfect abstention (100%). This demonstrates that abstention failures are tied to fine-grained ambiguity rather than fundamental incapacity, revealing a critical robustness gap for safe deployment in challenging scenarios.

11 Conclusion

We present a systematic analysis of length bias and robustness in LVLMs for MCQA across five fine-grained visual domains. LVLMs consistently prefer longer answer options, with accuracy dropping substantially when correct answers are shorter and increasing dramatically when longer. This bias persists across model scales and state-of-the-art proprietary models. Through layer-wise mechanistic analysis, we reveal that length bias emerges from late-layer suppression of visual evidence in favor of language-side priors, where models systematically fail to privilege visual signals despite stronger amplification capacity for correct answers. Blurry-image experiments show that degraded visual input accelerates this override, confirming that bias stems from language-side heuristics dominating when visual evidence weakens. Abstention tests reveal poor capability, with performance collapsing when models should refuse to answer. We propose and evaluate multiple inference-time mitigation strategies; while these provide partial improvements, none achieves deployable mitigation without relying on oracle information. This underscores the need for architectural solutions and bias-aware training.

Limitations

Our work focuses specifically on length bias and robustness in MCQA settings, not addressing other biases such as positional bias, modality imbalance, or cultural bias in LVLMs. Our benchmark is constructed from fine-grained visual classification MCQA across five domains; while this setting enables controlled isolation of length effects, direct empirical validation on broader VQA task families (e.g., open-ended VQA, compositional reasoning, or non-classification visual question answering) remains important future work, and the extent to which our findings generalize to these settings is an open question. The mitigation approaches we explore are inference-time strategies that do not modify model parameters. More effective solutions may require training-time methods such as bias-aware fine-tuning, including parameter-efficient approaches like Low-Rank Adaptation (LoRA), adapters, and prefix tuning, or architectural modifications that explicitly address length-based heuristics. Finally, our study examines a broad set of LVLM architectures spanning open-source models from 3B to 32B parameters, a mixture-of-experts model (Qwen3-VL-30B-A3B), and proprietary frontier models (GPT-4o, o4-mini, Claude Sonnet-4.5); continuous evaluation remains important as newer architectures continue to emerge.

Acknowledgments

We acknowledge Advanced Research Computing (ARC) at Virginia Tech for providing the computational resources and technical support that contributed to the results reported in this paper. We thank the reviewers for their constructive feedback, which helped improve this paper.

References

- Dyah Adila, Shuai Zhang, Boran Han, and Bernie Wang. 2024. Discovering bias in latent space: An unsupervised debiasing approach. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235, pages 246–261. PMLR.
- Anthropic. 2025. [Claude sonnet 4.5 system card](#). Technical report, Anthropic.
- Md Atabuzzaman, Ali Asgarov, and Chris Thomas. 2025a. Benchmarking and mitigating MCQA selection bias of large vision-language models. *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 33548–33562.
- Md Atabuzzaman, Andrew Zhang, and Chris Thomas. 2025b. Zero-shot fine-grained image classification using large vision-language models. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 23569–23582.
- Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhi-fang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, and 45 others. 2025a. [Qwen3-vl technical report](#). Preprint, arXiv:2511.21631.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025b. [Qwen2.5-vl technical report](#). Preprint, arXiv:2502.13923.
- Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. 2014. Food-101 – mining discriminative components with random forests. In *Computer Vision – ECCV 2014*, pages 446–461. Springer International Publishing.
- Meiqi Chen, Yixin Cao, Yan Zhang, and Chaochao Lu. 2024. Quantifying and mitigating unimodal biases in multimodal large language models: A causal perspective. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16449–16469.
- Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. 2025a. [Janus-pro: Unified multimodal understanding and generation with data and model scaling](#). Preprint, arXiv:2501.17811.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, Lixin Gu, Xuehui Wang, Qingyun Li, Yiming Ren, Zixuan Chen, Jiapeng Luo, Jiahao Wang, Tan Jiang, Bo Wang, and 23 others. 2025b. [Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling](#). Preprint, arXiv:2412.05271.
- Siyuan Cheng, Bozhong Tian, Qingbin Liu, Xi Chen, Yongheng Wang, Huajun Chen, and Ningyu Zhang. 2023. Can we edit multimodal large language models? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13877–13888.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in Neural Information Processing Systems*, 37.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza

- Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, and 1 others. 2025. Molmo and Pixmo: Open weights and open data for state-of-the-art vision-language models. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 91–104. IEEE.
- Qihang Fu, Yongbin Qin, Ruizhang Huang, Yanping Chen, Yulin Zhou, and Lintao Long. 2025. Exclusion of thought: Mitigating cognitive load in large language models for enhanced reasoning in multiple-choice tasks. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21673–21686.
- Tejas Gokhale, Pratyay Banerjee, Chitta Baral, and Yezhou Yang. 2020. Mutant: A training paradigm for out-of-distribution generalization in visual question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 878–892.
- Blessed Guda, Lawrence Francis, Gabriel Zencha Ashungafac, Carlee Joe-Wong, and Moise Busogi. 2025. Tiny: Rethinking selection bias in llms: Quantification and mitigation using efficient majority voting. In *ICLR Workshop: Quantify Uncertainty and Hallucination in Foundation Models: The Next Frontier in Reliable AI*.
- Ari Holtzman, Peter West, Vered Shwartz, Yejin Choi, and Luke Zettlemoyer. 2021. Surface form competition: Why the highest probability answer isn’t always right. In *Proceedings of the 2021 conference on empirical methods in natural language processing*, pages 7038–7051.
- Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. 2011. Novel dataset for fine-grained image categorization: Stanford dogs. In *Proc. CVPR workshop on fine-grained visual categorization (FGVC)*, volume 2.
- Jeonghwan Kim and Heng Ji. 2024. Finer: Investigating and enhancing fine-grained visual concept recognition in large vision language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6187–6207.
- Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Wangyue Li, Liangzhi Li, Tong Xiang, Xiao Liu, Wei Deng, and Noa Garcia. 2024. Can multiple-choice questions really be useful in detecting the abilities of llms? In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2819–2834.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36.
- Chenkai Ma and Xinya Du. 2023. Poe: Process of elimination for multiple choice reasoning. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4487–4496.
- Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. 2013. Fine-grained visual classification of aircraft. *arXiv preprint arXiv:1306.5151*.
- OpenAI. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Pouya Pezeshkpour and Estevam Hruschka. 2024. Large language models sensitivity to the order of options in multiple-choice questions. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2006–2017.
- Joshua Robinson and David Wingate. 2023. Leveraging large language models for multiple choice question answering. In *The Eleventh International Conference on Learning Representations (ICLR)*.
- Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. 2011. The caltech-ucsd birds-200-2011 dataset. *California Institute of Technology*.
- Haochun Wang, Sendong Zhao, Zewen Qiang, Nuwa Xi, Bing Qin, and Ting Liu. 2025. LLMs may perform mcqa by selecting the least incorrect option. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 5852–5862.
- Jialu Wang, Yang Liu, and Xin Wang. 2021. Are gender-neutral queries really gender-neutral? mitigating gender bias in image search. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1995–2008.
- Zecheng Wang, Xinye Li, Zhanyue Qin, Chunshan Li, Zhiying Tu, Dianhui Chu, and Dianbo Sui. 2024. Can we debias multimodal large language models via model editing? In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 3219–3228.
- Sheng-Lun Wei, Cheng-Kuang Wu, Hen-Hsen Huang, and Hsin-Hsi Chen. 2024. Unveiling selection biases: Exploring order and token sensitivity in large language models. In *Findings of the Association for*

Computational Linguistics ACL 2024, pages 5598–5621.

Mengge Xue, Zhenyu Hu, Liqun Liu, Kuo Liao, Honglin Han, Meng Zhao, Chengguo Yin, and 1 others. 2024. Strengthened symbol binding makes large language models reliable multiple-choice selectors. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4331–4344.

Nakyeong Yang, Taegwan Kang, Stanley Jungkyu Choi, Honglak Lee, and Kyomin Jung. 2024. Mitigating biases for instruction-following language models via bias neurons elimination. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9061–9073.

Zhen Yang, Ping Jian, and Chengzhi Li. 2025. Option symbol matters: Investigating and mitigating multiple-choice option symbol bias of large language models. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1902–1917.

YiFan Zhang, Yang Shi, Weichen Yu, Qingsong Wen, Xue Wang, Wenjing Yang, Zhang Zhang, Liang Wang, and Rong Jin. 2025. Debiasing multimodal large language models via penalization of language priors. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 4232–4241.

Guangxiang Zhao, Saier Hu, Xiaoqi Jian, Wu Jinzhu, Yuhao Wu, Lin Sun, and Xiangzheng Zhang. 2025. Large language models badly generalize across option length, problem types, and irrelevant noun replacements. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 26825–26834.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. *Advances in Neural Information Processing Systems*, 36:46595–46623.

Hanzhang Zhou, Zijian Feng, Zixiao Zhu, Junlang Qian, and Kezhi Mao. 2024. Unibias: Unveiling and mitigating llm bias through internal attention and fnn manipulation. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

Yongshuo Zong, Tingyang Yu, Ruchika Chavhan, Bingchen Zhao, and Timothy Hospedales. 2024. Fool your (vision and) language model with embarrassingly simple permutations. In *International Conference on Machine Learning*, pages 62892–62913. PMLR.

A Appendix

This section contains the details of the following topics.

- Dataset and Dataset Generation Constraints (Appendix A.1)
- Effect of Visual Degradation on Conflict Resolution (Appendix A.2)
- Architectural Localization via Attention Analysis (Appendix A.3)
- Caption-Guided MCQA Prompt (Appendix A.4)
- Caption-Guided MCQA (Appendix A.5)
- Yes/No Question Answering (Appendix A.6)
- MCQA with Selection by Rejection (Appendix A.7)
- LVLMS Performance Under Different Option-Length Conditions (Appendix A.8)
- Results for the Abstention (NoA) Settings (Appendix A.9)
- Position Bias and Answer Distribution Analysis (Appendix A.10)

A.1 Dataset and Dataset Generation Constraints




Figure 6: Example of normal and blurred image of German Shepherd.

We enforced the following generation constraints to ensure semantic and visual fidelity of the length-controlled variants.

Semantic preservation: We rephrase options

only to adjust length; the factual meaning and fine-grained class identity remain unchanged.

Visual consistency: We ground all modifications in visual attributes present in the image. We introduce no information absent from the base question or image.

Structural consistency: We ensure all options within a variant share the same fixed prefix (e.g., "This image shows..." or "This is..."), with length variations applied only to the descriptive content following the prefix.

Natural phrasing: We use natural, grammatically correct language to avoid awkward constructions that might signal correctness independent of content.

Class structure preservation: Each option continues to represent a distinct fine-grained class, maintaining the original task difficulty.

Verification of Dataset Fidelity. To verify that GPT-4o rewrites preserved semantic correctness and visual fidelity, we conducted structured manual inspection of stratified samples from every dataset. For each rewritten option, we verified that: (1) all class-defining visual attributes were preserved, (2) no contradictory or hallucinated attributes were introduced, and (3) the class identity remained unambiguous relative to distractors. All examined samples satisfied these criteria.

We further validated fidelity through two indirect measures. First, the **Equal Long** and **Equal Short** variants yield broadly comparable accuracies for most models and domains (see Table 8), with a mean absolute difference of approximately 5 points; notable exceptions include LLaVA-v1.5 variants on Food, where gaps reach up to 28 points, likely reflecting model-specific sensitivity to phrasing length rather than dataset construction artifacts. Critically, even these largest baseline gaps remain substantially smaller than the length bias effect itself: for example, GPT-4o’s Equal Long–Equal Short gap on CUB is 8.34 points, compared to a 47.66-point gap between Shorter Correct and Longer Correct (Table 1). This indicates that residual variation from the rewriting process is minor relative to the systematic bias our benchmark is designed to measure. Second, if GPT-4o had introduced systematic bias during rewriting, GPT-4o’s evaluation performance would differ noticeably from models uninvolved in dataset creation. Instead, GPT-4o exhibits length bias patterns closely matching o4-mini and Claude Sonnet-4.5 (Table 1),

confirming that the rewritten dataset is semantically faithful and unbiased.

Figure 7 illustrates an example from our length-controlled MCQA benchmark. Each variant manipulates the relative length of the correct answer and distractors while preserving semantic content and visual accuracy. In **Equal Long** and **Equal Short**, all options have approximately the same length, serving as balanced baselines. In **Shorter Correct**, the correct answer (D) is concise while distractors (A, B, C) are verbose, testing whether models exhibit bias against selecting shorter options. Conversely, in **Longer Correct**, the correct answer is detailed while distractors are brief, testing preference for longer options. These controlled manipulations enable direct measurement of how answer length influences LVLM selection independent of visual or semantic factors.

Figure 6 illustrates two German Shepherd examples from the Dogs dataset: one original and one blurred version used for robustness evaluation.

A.2 Effect of Visual Degradation on Conflict Resolution

To causally probe the role of visual evidence in length-biased decision making, we repeat our layer-wise analysis under controlled visual degradation. Specifically, we apply Gaussian blur ($\sigma = 12$) to the input images while keeping the question, answer options, and decoding procedure unchanged. This manipulation weakens fine-grained visual cues without altering linguistic content, allowing us to isolate the effect of visual signal strength on internal conflict resolution dynamics. We define the *override layer* as the first layer where the average long distractor probability exceeds the short correct answer probability for three consecutive layers, indicating stable dominance of length bias over visual evidence. This statistic is distinct from the *stable inversion layer* ($L = 14$) reported in Section 6: the override layer is computed only on length-biased error cases using the averaged probability across all three long distractors, whereas the stable inversion layer is computed across the full Shorter Correct split using a single fixed (longest) distractor.

Figure 8 compares layer-wise probability trajectories for clear versus blurred images on the same set of length-biased error cases. Under clear images, visual evidence delays the dominance of long distractors until late layers (median override layer $L = 19$). When visual inputs are blurred, this




Figure 7: Example multiple-choice questions from the Dogs dataset illustrating four length-controlled variants: **Equal Long**, **Equal Short**, **Shorter Correct**, and **Longer Correct**. The correct answer (D: German Shepherd) remains constant while option lengths vary systematically to isolate the effect of textual length on model predictions.




Figure 8: **Effect of visual degradation on length-bias conflict resolution.** Layer-wise probabilities for the short correct answer (green) and the average long distractor (red) under clear images (left) and Gaussian blur ($\sigma = 12$, right), evaluated on the same set of length-biased error cases. Dashed lines mark the median override layer under each condition. Clear images exhibit greater dispersion in probability trajectories due to heterogeneity in visual evidence strength across examples, while blurring homogenizes visual inputs and produces earlier, more uniform dominance of length bias. Shaded regions denote 95% confidence intervals.

dominance emerges substantially earlier (median override layer $L = 11$), indicating that weakened visual evidence accelerates the resolution of conflict in favor of linguistic priors.

Notably, probability trajectories under clear images exhibit greater dispersion than those under blur. This reflects heterogeneity in visual evidence strength across examples: clear images preserve variation in pose, viewpoint, and fine-grained discriminability, allowing some instances to sustain

visual alignment longer than others. Blurring homogenizes visual inputs across examples, reducing this variability and resulting in more uniform, earlier dominance of length bias. The systematic shift in override timing and reduced dispersion under blur provide causal evidence that late-layer conflict resolution in LVLMs depends on the strength of visual evidence, rather than arising solely from inherited language-model biases.

Together, these results demonstrate that visual

information can delay—but does not reliably prevent—the amplification of length bias. When visual evidence is degraded, linguistic priors dominate earlier in the computation, confirming that length bias in LVLMs emerges from dynamic competition between modalities rather than from static decoder preferences alone.

A.3 Architectural Localization via Attention Analysis




Figure 9: Text-to-vision attention correlation across transformer layers (N=200 examples). We measure the correlation between option length and visual attention at each layer. Late layers (20–28, shaded region) show statistically significant positive correlation (mean $r = 0.063$, $p < 0.001$), indicating that longer options query vision more than shorter options. Despite this increased visual grounding, long incorrect options are preferentially selected, demonstrating that length bias operates through override of visual evidence in the language decoder rather than insufficient visual information retrieval.

To identify whether length bias originates in vision-language fusion or language processing, we analyze text-to-vision attention patterns across all layers (N=200 examples). We extract attention weights from text option tokens to vision tokens, measuring whether longer options differentially query visual information during decision-making.

Late-layer analysis (layers 20–28, Figure 9) reveals a statistically significant positive correlation between option length and visual attention ($r = 0.063$, $p < 0.001$). Critically, long incorrect options attend 9% more to visual features than short correct answers (0.000342 vs 0.000313, $p < 0.001$), yet are still preferentially selected. This demonstrates that length bias does not arise from inadequate visual grounding but rather from the language decoder systematically overriding visual evidence when processing longer text se-

Dataset	Captioner	Model	Equal Long	Equal Short	Shorter Correct	Longer Correct
CUB	Qwen2.5-VL-7B	Qwen2.5-VL-7B	59.30	52.10	41.30	70.50
	Qwen3-VL-30B	Qwen2.5-VL-7B	66.25	55.50	44.25	70.00
	GPT-4o	Qwen2.5-VL-7B	67.00	58.00	48.50	70.75
	InternVL-2.5-8B	InternVL-2.5-8B	64.50	55.10	41.80	74.40
	Qwen3-VL-30B	InternVL-2.5-8B	66.75	57.50	46.00	78.25
	GPT-4o	InternVL-2.5-8B	65.25	57.25	43.50	76.75
Dogs	Qwen2.5-VL-7B	Qwen2.5-VL-7B	57.33	51.50	53.00	61.17
	Qwen3-VL-30B	Qwen2.5-VL-7B	65.00	56.67	58.33	60.83
	GPT-4o	Qwen2.5-VL-7B	67.08	53.33	56.67	59.58
	InternVL-2.5-8B	InternVL-2.5-8B	56.33	52.33	42.33	68.83
	Qwen3-VL-30B	InternVL-2.5-8B	61.67	59.58	50.42	71.67
	GPT-4o	InternVL-2.5-8B	60.83	57.50	47.50	68.75

Table 5: Accuracy (%) of LVLMs on **Caption-Guided MCQA** across CUB and Dogs without class names in the answer options.

quences.

This localization explains why caption-guided MCQA (Section A.5) provides only partial mitigation: explicit visual descriptions cannot overcome bias that operates independently of visual evidence strength. The finding indicates that effective solutions must target the language decoder through architectural modifications or length-normalized scoring mechanisms.

A.4 Caption-Guided MCQA Prompt

For caption-guided MCQA experiments, we use the following prompt template (10) to guide the answering model to ground its reasoning in both the image and the generated caption:

Caption-Guided MCQA Prompt

You are given a multiple-choice question about an image.
Image Description: {image_description}
Use BOTH the image and the description above to answer the question. Return ONLY a single option ID from ('A', 'B', 'C', 'D') of the correct answer. Do not include explanations or extra text. {formatted_prompt}
Answer: (choose one from A, B, C, D)

Figure 10: Prompt template used for caption-guided MCQA.

where {image_description} is replaced with the generated caption from the captioning model, and {formatted_prompt} contains the question and answer options.

A.5 Caption-Guided MCQA

To investigate whether bias arises from visual perception or decision-making, we introduce **Caption-Conditioned MCQA (CC-MCQA)** inspired by

human problem-solving: first forming an internal description of the scene, then comparing it against candidate answers. This framework decouples perception from answer selection through two sequential stages, serving as both a mitigation strategy and diagnostic tool.

Caption Generation (Perception Stage). Given an image I , a captioning model M_{cap} generates a fine-grained description:

$$C = M_{\text{cap}}(I, p_{\text{cap}}),$$

where p_{cap} encourages grounding of discriminative attributes (e.g., color, pattern, texture) while discouraging speculative language.

Caption-Guided Answer Selection (Reasoning Stage). The caption C is provided to an answering model M_{ans} with the question Q and candidate options $\{O_1, \dots, O_n\}$:

$$\hat{y} = M_{\text{ans}}(I, Q, \{O_1, \dots, O_n\}, C, p_{\text{ans}}),$$

where p_{ans} guides the model to ground reasoning in the caption rather than option length. This disentangles perception errors (from M_{cap}) from decision biases (from M_{ans}).

We evaluate two variants: **Self-Conditioned** ($M_{\text{cap}} = M_{\text{ans}}$) and **Decoupled** ($M_{\text{cap}} \neq M_{\text{ans}}$). If performance improves with accurate captions, biases originate from perception errors; if performance degrades despite accurate captions, biases stem from decision-making heuristics.

Results. Table 5 presents Caption-Guided MCQA performance without class names. We compare against baseline (Table 8) to assess whether explicit visual grounding mitigates length bias.

Caption Guidance Substantially Reduces Length Bias. Caption-guided MCQA with GPT-4o substantially narrows the gap between "Shorter Correct" and "Longer Correct." For Qwen2.5-VL-7B on CUB, the baseline gap of 66.80 points (26.60% vs 93.40%) reduces to 22.25 points (48.50% vs 70.75%), a 66.7% bias reduction. The decrease in "Longer Correct" (93.40% to 70.75%) represents correction of artificially inflated accuracy, while "Shorter Correct" improves substantially (26.60% to 48.50%), demonstrating that visual grounding overcomes bias against concise options.

External Captions Outperform Self-Conditioned Approaches. Using stronger

Dataset	Model	Equal Long	Equal Short	Shorter Correct	Longer Correct
CUB (W/o)	Qwen2.5-VL-7B	58.40	51.30	62.80	45.70
	InternVL_2.5-8B	59.00	53.50	67.10	44.00
Dogs (W/o)	Qwen2.5-VL-7B	57.17	53.17	53.00	61.33
	InternVL_2.5-8B	52.00	52.00	73.33	32.50
CUB (W.)	Qwen2.5-VL-7B	71.20	68.60	69.10	70.40
	InternVL_2.5-8B	63.60	62.00	53.20	70.10
Dogs (W.)	Qwen2.5-VL-7B	71.83	71.17	77.33	61.83
	InternVL_2.5-8B	55.17	57.33	67.17	44.00

Table 6: Accuracy (%) of LVLMs on "Yes/No" evaluation across CUB and Dogs, reported with (W.) and without (W/o) class names in the answer options.

captioners (GPT-4o, Qwen3-VL-30B) consistently outperforms both self-conditioned and baseline approaches. On CUB "Equal Long," GPT-4o achieves 67.00% versus 59.30% (self-conditioned) and 64.90% (baseline). Self-conditioned captioning often degrades performance, revealing that perception errors compound decision-making biases when the same model handles both stages.

Residual Bias Reveals Multi-Stage Architecture. Despite substantial improvements, a 22.25 point gap persists even with GPT-4o captions, revealing that length preference operates at multiple processing stages. While explicit visual grounding addresses perceptual biases, decision-making biases in textual weighting remain partially intact.

A.6 Yes/No Question Answering (QA)

We transform the multiple-choice task into binary Yes/No questions. For each option, we ask the model whether it correctly answers the question and select the option with the highest probability assigned to "Yes" as the predicted answer. This independent evaluation removes direct comparisons between options during selection.

Table 6 presents results with and without class names. The Yes/No format substantially improves "Shorter Correct" performance: for Qwen2.5-VL-7B on CUB with class names, accuracy reaches 69.10% compared to 26.60% baseline (Table 8), nearly achieving parity with "Longer Correct" (70.40%). However, performance degrades in balanced settings, with "Equal Short" dropping from 61.00% to 51.30%. While independent option evaluation mitigates length bias by removing direct length comparisons, eliminating comparative context hinders fine-grained discrimination when length is not a confounding factor.

Dataset	Model	Equal Long	Equal Short	Shorter Correct	Longer Correct
CUB (W/o)	Qwen2.5-VL-7B	52.80	52.40	31.40	81.80
	InternVL_2.5-8B	56.80	49.80	33.10	82.90
Dogs (W/o)	Qwen2.5-VL-7B	54.17	53.83	41.50	73.50
	InternVL_2.5-8B	54.00	56.33	32.67	76.67
CUB (W.)	Qwen2.5-VL-7B	66.70	63.40	50.00	84.50
	InternVL_2.5-8B	57.90	56.00	33.40	88.10
Dogs (W.)	Qwen2.5-VL-7B	69.83	70.17	60.17	89.67
	InternVL_2.5-8B	58.67	52.83	38.17	78.00

Table 7: Accuracy (%) of LVLMs on "**selection by rejection**" across CUB and Dogs, with (W.) and without (W/o) class names in the answer options.

A.7 MCQA with Selection by Rejection

In multiple-choice reasoning, humans often approach difficult questions by first eliminating clearly incorrect options before selecting a final answer. Inspired by this process, we design a *selection-by-rejection* strategy to test whether explicit elimination can mitigate the length bias.

Instead of prompting the model to select a single option directly, we instruct it to *reject* three candidates, inferring the final prediction as the surviving option. We hypothesize that explicitly discarding incorrect candidates might reduce reliance on superficial heuristics such as option length and instead focus on genuine grounding in the image.

Selection-by-Rejection Degrades Performance and Reinforces Length Bias. Table 7 shows that this strategy consistently degrades performance across all models. Comparing with Table 8, we observe substantial accuracy drops. On CUB without class names, Qwen2.5-VL-7B drops from 64.90% to 52.80% in "Equal Long" and from 93.40% to 81.80% in "Longer Correct." Although LVLMs produce syntactically valid rejections, the eliminations are not reliably aligned with correctness. Further analysis reveals that generated justifications are typically plausible-sounding but not systematically grounded in visual or textual evidence. Unlike humans, who perform elimination as a structured logical process, LVLMs treat rejection as another generative task influenced by the same shallow heuristics. This suggests that the rejection procedure does not reduce length bias but instead reinforces the preference for visually inconsistent longer options.

A.8 LVLM Performance Under Different Option-Length Conditions

Table 8 presents comprehensive results across all evaluated models and datasets under four length-controlled conditions. The Δ column quantifies length bias magnitude, showing substantial bias

across all models: without class names, Δ ranges from 8.52 (InternVL_2.5-8B on Food) to 93.50 (LLaVA-v1.5-7B on Dogs). Performance consistently drops in "Shorter Correct" conditions (negative values in parentheses) and increases in "Longer Correct" conditions (positive values), with drops of 2–38 points and gains of 1–70 points depending on model and dataset. Class name inclusion mitigates bias, reducing Δ values substantially: for Qwen2.5-VL-7B on CUB, Δ decreases from 66.80 to 43.20. However, even with class names, significant bias persists (e.g., LLaVA-v1.5-13B on CUB: $\Delta = 57.30$), indicating that length preference is deeply entrenched. Notably, the Food dataset shows lower bias across all models, suggesting that domain characteristics influence susceptibility to length-based heuristics.

A.9 Results for the Abstention (NoA) Settings




Figure 11: Accuracy (%) of LVLMs on "None of the Above" (NoA) evaluation settings across Dogs datasets, reported both with and without class names in the answer options. "w/o NoA" denotes the original MCQ format without a NoA option. "NoA w/C" appends a fifth option, "E. None of the above", while keeping the correct answer present among A–D (over-abstention test). "NoA wo/C" removes the correct answer and uses "D. None of the above." as the only valid choice (forced-abstention test).

Figure 11 and Table 9 summarize LVLM performance under different abstention settings. We observe several consistent patterns.

Without Class Names, Abstention Performance Collapses Completely. Without class names, LVLMs fail almost completely under abstention tests. Qwen2.5-VL-7B drops to 3.2% on CUB and 3.83% on Dogs under NoA wo/C, while InternVL_2.5-8B falls to 17.7% and 33.5%, respec-

Class Name	Dataset	Model	Equal Long	Equal Short	Shorter Correct	Longer Correct	Δ
Without	Aircraft	LLaVA-v1.5-7B	32.86	34.00 (+1.14)	18.29 (-14.57)	84.57 (+51.71)	66.28
		LLaVA-v1.5-13B	39.71	35.71 (-4.00)	18.86 (-20.85)	87.43 (+47.72)	68.57
		Qwen2.5-VL-7B	54.86	43.71 (-11.15)	21.43 (-33.43)	84.00 (+29.14)	62.57
		InternVL_2.5-8B	39.71	36.29 (-3.42)	32.29 (-7.42)	57.43 (+17.72)	25.14
	Cars	LLaVA-v1.5-7B	26.73	32.65 (+5.92)	12.35 (-14.38)	94.29 (+67.56)	81.94
		LLaVA-v1.5-13B	40.71	44.08 (+3.37)	16.94 (-23.77)	93.16 (+52.45)	76.22
		Qwen2.5-VL-3B	64.80	55.51 (-9.29)	35.82 (-28.98)	85.61 (+20.81)	49.79
		Qwen2.5-VL-7B	68.88	59.59 (-9.29)	38.57 (-30.31)	89.69 (+20.81)	51.12
		Qwen2.5-VL-32B	73.16	60.71 (-12.45)	42.86 (-30.30)	92.65 (+19.49)	49.79
		InternVL_2.5-8B	56.94	47.45 (-9.49)	37.55 (-19.39)	72.65 (+15.71)	35.10
	CUB	LLaVA-v1.5-7B	26.70	31.20 (+4.50)	8.90 (-17.80)	96.80 (+70.10)	87.90
		LLaVA-v1.5-13B	39.30	42.60 (+3.30)	15.20 (-24.10)	93.30 (+54.00)	78.10
		Qwen2.5-VL-3B	60.60	58.40 (-2.20)	38.10 (-22.50)	81.40 (+20.80)	43.30
		Qwen2.5-VL-7B	64.90	61.00 (-3.90)	26.60 (-38.30)	93.40 (+28.50)	66.80
		Qwen2.5-VL-32B	68.00	61.30 (-6.70)	37.20 (-30.80)	93.30 (+25.30)	56.10
		Qwen3-VL-30B	72.30	63.20 (-9.10)	47.30 (-25.00)	85.00 (+12.70)	37.70
		InternVL_2.5-8B	62.50	57.10 (-5.40)	37.80 (-24.70)	88.30 (+25.80)	50.50
	Dogs	LLaVA-v1.5-7B	30.00	33.67 (+3.67)	5.67 (-24.33)	99.17 (+69.17)	93.50
		LLaVA-v1.5-13B	36.67	42.33 (+5.66)	14.00 (-22.67)	95.67 (+59.00)	81.67
		Qwen2.5-VL-7B	62.83	58.50 (-4.33)	33.17 (-29.66)	90.83 (+28.00)	57.66
		Qwen3-VL-30B	71.17	64.00 (-7.17)	49.50 (-21.67)	89.17 (+18.00)	39.67
		InternVL_2.5-8B	57.50	58.00 (+0.50)	38.17 (-19.33)	79.67 (+22.17)	41.50
	Food	LLaVA-v1.5-7B	33.47	53.07 (+19.60)	27.72 (-5.75)	93.07 (+59.60)	65.35
		LLaVA-v1.5-13B	49.50	65.15 (+15.65)	42.57 (-6.93)	95.84 (+46.34)	53.27
		Qwen2.5-VL-7B	86.73	86.93 (+0.20)	78.61 (-8.12)	92.08 (+5.35)	13.47
		InternVL_2.5-8B	80.79	76.63 (-4.16)	75.84 (-4.95)	84.36 (+3.57)	8.52
With	Aircraft	LLaVA-v1.5-7B	19.43	21.71 (+2.28)	12.86 (-6.57)	54.57 (+35.14)	41.71
		LLaVA-v1.5-13B	21.43	24.00 (+2.57)	10.57 (-10.86)	75.14 (+53.71)	64.57
		Qwen2.5-VL-7B	70.29	75.14 (+4.85)	45.43 (-24.86)	96.00 (+25.71)	50.57
		InternVL_2.5-8B	54.29	56.57 (+2.28)	31.43 (-22.86)	81.14 (+26.85)	49.71
	Cars	LLaVA-v1.5-7B	25.82	32.35 (+6.53)	19.18 (-6.64)	79.08 (+53.26)	59.90
		LLaVA-v1.5-13B	31.84	40.92 (+9.08)	26.84 (-5.00)	80.41 (+48.57)	53.57
		Qwen2.5-VL-3B	76.73	75.41 (-1.32)	62.76 (-13.97)	89.59 (+12.86)	26.83
		Qwen2.5-VL-7B	82.24	84.49 (+2.25)	67.14 (-15.10)	94.90 (+12.66)	27.76
		Qwen2.5-VL-32B	85.10	82.96 (-2.14)	74.39 (-10.71)	91.73 (+6.63)	17.34
		InternVL_2.5-8B	56.12	57.04 (+0.92)	44.80 (-11.32)	75.51 (+19.39)	30.71
	CUB	LLaVA-v1.5-7B	27.50	28.70 (+1.20)	23.40 (-4.10)	63.50 (+36.00)	40.10
		LLaVA-v1.5-13B	28.30	33.70 (+5.40)	19.10 (-9.20)	76.40 (+48.10)	57.30
		Qwen2.5-VL-3B	67.50	70.90 (+3.40)	52.50 (-15.00)	87.30 (+19.80)	34.80
		Qwen2.5-VL-7B	76.10	74.30 (-1.80)	50.70 (-25.40)	93.90 (+17.80)	43.20
		Qwen2.5-VL-32B	76.40	73.60 (-2.80)	61.00 (-15.40)	88.60 (+12.20)	27.60
		Qwen3-VL-30B	77.80	75.20 (-2.60)	61.90 (-15.90)	88.50 (+10.70)	26.60
		InternVL_2.5-8B	61.40	60.00 (-1.40)	39.20 (-22.20)	88.30 (+26.90)	49.10
	Dogs	LLaVA-v1.5-7B	26.67	31.83 (+5.16)	16.00 (-10.67)	76.67 (+50.00)	60.67
		LLaVA-v1.5-13B	30.83	37.67 (+6.84)	20.00 (-10.83)	76.83 (+46.00)	56.83
		Qwen2.5-VL-7B	76.33	78.00 (+1.67)	58.83 (-17.50)	93.33 (+17.00)	34.50
		Qwen3-VL-30B	77.67	77.67 (+0.00)	67.67 (-10.00)	88.50 (+10.83)	20.83
		InternVL_2.5-8B	59.33	57.67 (-1.66)	43.67 (-15.66)	79.33 (+20.00)	35.66
	Food	LLaVA-v1.5-7B	30.10	58.22 (+28.12)	32.87 (+2.77)	89.90 (+59.80)	57.03
		LLaVA-v1.5-13B	58.42	71.88 (+13.46)	56.63 (-1.79)	96.44 (+38.02)	39.81
		Qwen2.5-VL-7B	88.12	89.31 (+1.19)	84.36 (-3.76)	93.66 (+5.54)	9.30
		InternVL_2.5-8B	85.74	83.76 (-1.98)	81.39 (-4.35)	86.34 (+0.60)	4.95

Table 8: Accuracy (%) of LVLMS across five datasets under different option-length conditions, with and without class names in answer choices. Results show LVLMS performance on "Equal Long," "Equal Short," "Shorter Correct," and "Longer Correct" variants. Values in parentheses indicate differences from the "Equal Long" baseline for each model-dataset pair. Δ (Longer Correct - Shorter Correct) quantifies length bias magnitude. The best result for each setting and dataset is highlighted in bold.

tively. Such near-random behavior suggests that without semantic anchors, models cannot recognize when no valid answer is available and instead produce arbitrary guesses. This underscores the

importance of class names as grounding cues for abstention, not just classification.

Abstention Reveals Critical Robustness Failures in LVLMS. While LVLMS handle abstention re-

liably in trivial cases, they perform poorly under challenging conditions. They exhibit both false abstention (choosing NoA when the correct answer is present) and failure to abstain (guessing when abstention is correct). This highlights abstention as a critical but underexplored aspect of LVLM robustness with direct implications for safe deployment in real-world applications.

A.10 Position Bias and Answer Distribution Analysis

To ensure that our observed length bias is not confounded by position bias, we provide detailed analysis of answer position distributions and our experimental design controls.

A.10.1 Answer Position Distribution Across Datasets

Our length-controlled benchmarks inherit the answer position distributions from the underlying fine-grained MCQA datasets of [Atabuzzaman et al. \(2025a\)](#). Table 10 presents the distribution of correct answers across positions A, B, C, and D for each dataset. While distributions show some natural variation, no position exhibits systematic overrepresentation that could confound our length manipulation experiments.

A.10.2 Position-Controlled Length Manipulation

Critically, our length-controlled variant design explicitly neutralizes position effects. For every question, the correct answer remains in the *same position* (A, B, C, or D) across all four variants: Equal Long, Equal Short, Shorter Correct, and Longer Correct.

For example, consider a question where the correct answer originally appears at position C:

- **Equal Long:** Position C contains a detailed correct answer (all options long)
- **Equal Short:** Position C contains a brief correct answer (all options short)
- **Shorter Correct:** Position C contains a *short* correct answer (distractors long)
- **Longer Correct:** Position C contains a *long* correct answer (distractors short)

Because position is held constant while answer length is systematically varied, any performance differences across these variants cannot be attributed to position bias. Length is the only manipulated factor. This within-question design en-

sures that our results isolate length effects from positional preferences.

A.10.3 Class-Name-Only Ablation

To further verify that content variation does not confound length effects, we conducted additional ablations where only class names differ between conditions. Table 11 presents results where the baseline (Class-name-only) provides all options as bare class names, the Shorter Correct condition adds detailed descriptions only to distractors, and the Longer Correct condition adds detailed descriptions only to the correct answer.

Results demonstrate consistent accuracy shifts driven solely by length manipulation. For instance, Qwen2.5-VL-7B on CUB drops from 67.0% to 53.3% when the correct answer is short and rises to 93.2% when it is long—a 39.9 percentage point gap attributable entirely to length. These results confirm that position bias does not explain our findings.

Class Name	Dataset	Setting	Model	Equal Long	Equal Short	Shorter Correct	Longer Correct
Without	CUB	w/o NoA	Qwen2.5-VL-7B	64.90	61.00	26.60	93.40
			InternVL_2.5-8B	62.50	57.10	37.80	88.30
		NoA w/C	Qwen2.5-VL-7B	64.10	60.20	28.10	93.30
			InternVL_2.5-8B	59.30	54.40	34.70	87.40
	Dogs	NoA wo/C	Qwen2.5-VL-7B	3.20	5.80	3.20	5.40
			InternVL_2.5-8B	17.70	21.00	18.00	20.00
		w/o NoA	Qwen2.5-VL-7B	62.83	58.50	33.17	90.83
			InternVL_2.5-8B	57.50	58.00	38.17	79.67
	Dogs	NoA w/C	Qwen2.5-VL-7B	63.00	57.67	34.00	91.33
			InternVL_2.5-8B	48.83	53.83	31.33	76.33
		NoA wo/C	Qwen2.5-VL-7B	3.83	9.17	4.50	8.00
			InternVL_2.5-8B	33.50	35.17	33.50	31.17
With	CUB	w/o NoA	Qwen2.5-VL-7B	76.10	74.30	50.70	93.90
			InternVL_2.5-8B	61.40	60.00	39.20	88.30
		NoA w/C	Qwen2.5-VL-7B	75.00	73.20	50.70	93.50
			InternVL_2.5-8B	55.70	54.30	33.90	86.00
	Dogs	NoA wo/C	Qwen2.5-VL-7B	20.00	20.50	12.20	19.70
			InternVL_2.5-8B	23.00	29.70	23.10	30.10
		w/o NoA	Qwen2.5-VL-7B	76.33	78.00	58.83	93.33
			InternVL_2.5-8B	59.33	57.67	43.67	79.33
	Dogs	NoA w/C	Qwen2.5-VL-7B	76.33	77.17	59.00	92.50
			InternVL_2.5-8B	50.67	49.17	36.00	73.83
		NoA wo/C	Qwen2.5-VL-7B	11.33	17.17	11.00	17.00
			InternVL_2.5-8B	32.50	30.33	32.33	30.67
	Dogs (Easy)	w/o NoA	Qwen2.5-VL-7B	100.0	100.0	100.0	100.0
			InternVL_2.5-8B	100.0	100.0	100.0	100.0
		NoA w/C	Qwen2.5-VL-7B	99.17	98.33	97.67	99.50
			InternVL_2.5-8B	94.67	91.83	92.83	95.33
	Dogs (Easy)	NoA wo/C	Qwen2.5-VL-7B	100.0	100.0	100.0	100.0
			InternVL_2.5-8B	100.0	100.0	100.0	100.0

Table 9: Accuracy (%) of LVLMs on "None of the Above" (NoA) evaluation settings across CUB and Dogs datasets, reported both with and without class names in the answer options. "w/o NoA" denotes the original MCQ format without a NoA option. "NoA w/C" appends a fifth option, "E. None of the above", while keeping the correct answer present among A–D (over-abstention test). "NoA wo/C" removes the correct answer and uses "D. None of the above." as the only valid choice (forced-abstention test). In the "NoA wo/C" setting, removing the correct answer from "Shorter Correct" leaves only long distractors, making its option pool structurally similar to "Equal Long"; likewise, removing the correct answer from "Longer Correct" leaves only short distractors, similar to "Equal Short". Accordingly, LVLM performance under "NoA wo/C" for these two conditions closely aligns with, though does not exactly match, "Equal Long" and "Equal Short", respectively.

Dataset	A	B	C	D
Cars	245	245	240	250
CUB	260	215	250	275
Dogs	170	145	160	125
Food	140	125	100	140

Table 10: Distribution of correct answers across positions A, B, C, and D for each dataset. Distributions are inherited from base benchmarks and remain unchanged across all length-controlled variants.

Model	Dataset	Class name only	Shorter Correct	Longer Correct
InternVL-2.5-8B	Dogs	54.17	45.33	79.17
InternVL-2.5-8B	CUB	47.10	48.60	89.60
Qwen2.5-VL-7B	CUB	67.00	53.30	93.20
Qwen2.5-VL-7B	Dogs	75.83	60.17	93.33

Table 11: Accuracy (%) under class-name-only ablations. All conditions use identical answer positions, varying only the length of the correct option versus distractors.