

The Art of Saying “Maybe”: A Conformal Lens for Uncertainty Benchmarking in VLMs

Asif Azad^{1*†} Mohammad Sadat Hossain^{1*} MD Sadik Hossain Shanto^{1*}
M Saifur Rahman¹ Md Rizwan Parvez²

¹Bangladesh University of Engineering and Technology

²Qatar Computing Research Institute

asifazad0178@gmail.com, sadat@cse.buet.ac.bd

shantosadikrglhs@gmail.com, mrahman@cse.buet.ac.bd

mparvez@hbku.edu.qa

Abstract

Vision-Language Models (VLMs) have achieved remarkable progress in complex visual understanding across scientific and reasoning tasks. While performance benchmarking has advanced our understanding of these capabilities, the critical dimension of uncertainty quantification has received insufficient attention. Therefore, unlike prior conformal prediction studies that focused on limited settings, we conduct a comprehensive uncertainty benchmarking study, evaluating 18 state-of-the-art VLMs (open and closed-source) across 6 multimodal datasets with 3 distinct scoring functions. For closed-source models lacking token-level logprob access, we develop and validate instruction-guided likelihood proxies. Our findings demonstrate that larger models consistently exhibit better uncertainty quantification; models that know more also know better what they don’t know. More certain models achieve higher accuracy, while mathematical and reasoning tasks elicit poorer uncertainty performance across all models compared to other domains. This work establishes a foundation for reliable uncertainty evaluation in multimodal systems.

1 Introduction

Recent advances in large vision-language models (VLMs) have led to remarkable progress in complex visual understanding and reasoning across diverse domains such as mathematics (Wang et al., 2024), science (Lu et al., 2022), and medicine (Matos et al., 2024). These models now achieve impressive results on challenging multimodal benchmarks, demonstrating their potential for real-world impact.

Yet, despite these capabilities, significant challenges remain. As VLMs are increasingly deployed in high-stakes domains like medical diagnostics (Li

et al., 2025), educational assessments, and scientific reasoning, the consequences of model failure become critical. While accuracy metrics highlight overall performance, they do not reveal when a model is uncertain or likely to err. In practical applications, especially in sensitive fields like healthcare, an overconfident but incorrect prediction can have severe repercussions. Thus, quantifying and understanding model uncertainty in a computationally efficient way is essential for building reliable and trustworthy VLM systems.

Quantifying uncertainty in VLMs is therefore crucial for building reliable and trustworthy systems, especially in high-stakes domains. While classical approaches such as Bayesian neural networks (Blundell et al., 2015), deep ensembles (Lakshminarayanan et al., 2017), and calibration-based methods (Guo et al., 2017) have been explored for uncertainty estimation mostly in traditional machine learning models, their application to foundation models (e.g., LLMs, VLMs, and multimodal architectures), where parameters often scale to billions or trillions, is limited by computational cost and scalability issues. Conformal prediction, in contrast, offers a computationally feasible, model-agnostic framework with formal statistical guarantees, making it particularly attractive for uncertainty quantification in complex multimodal settings. Prior work has applied conformal prediction to LLMs for benchmarking predictive confidence (Ye et al., 2024), but its utility for VLMs, where uncertainty arises from both visual and textual modalities, remains largely unexplored. This motivates our study, which systematically investigates conformal prediction as a principled approach for uncertainty benchmarking in VLMs across diverse set of tasks.

This study is guided by several core research questions:

1. Do different conformal scoring functions

*Equal contribution.

†Corresponding author.

yield similar efficiency in terms of prediction set size, or do their behaviors diverge across tasks and models?

2. Is there a correlation between model accuracy and the size of conformal prediction sets, indicating calibration quality?
3. How do uncertainty metrics (set size) vary with model scale and architecture?
4. Can this uncertainty quantification approach be applied to black-box proprietary models, provided they expose token-level probabilities?
5. How does uncertainty calibration vary across domains and do certain model families show systematic advantages?

Our evaluation spans a suite of carefully chosen datasets- MMMU (Yue et al., 2024a), MMMU-Pro (Yue et al., 2024b), AI2D (Kembhavi et al., 2016), MathVision (Wang et al., 2024, 2025), ScienceQA (Lu et al., 2022), and WorldMedQAV (Matos et al., 2024)- each probing distinct aspects of visual and scientific understanding. We systematically compare multiple scoring functions within the conformal framework to provide a comprehensive analysis of uncertainty in VLMs.

Our findings reveal patterns of uncertainty that correlate not only with accuracy but also with task modality and semantic complexity, offering deeper insights into when and why VLMs hesitate.

2 Related Works

Uncertainty quantification has long been central to machine learning, especially for risk-sensitive decisions (Abdar et al., 2021). Classic methods include Bayesian neural networks (Blundell et al., 2015), deep ensembles (Lakshminarayanan et al., 2017), and calibration techniques (Guo et al., 2017). These work well in low dimensions but can be computationally heavy or lack expressiveness for deep multimodal models.

Conformal prediction provides robust, statistical guarantees and applies across domains (Zhou et al., 2025). It’s distribution-free, model-agnostic, and efficient, making it ideal for large-scale models. Recent efforts have used it for LLMs (Angelopoulos and Bates, 2021; Ye et al., 2024) and VLMs (Kostumov et al., 2024), delivering coverage via

prediction sets. Still, past work focused on text-only models or simpler benchmarks with outdated VLMs. In contrast, while Kostumov et al. (2024) mainly examined accuracy–uncertainty alignment and scaling effects, our study investigates a broader set of five research questions covering scoring efficiency, calibration–accuracy correlations, architecture and scale influences, black-box model applicability, and domain-specific trends. We also evaluate a larger and more diverse suite of state-of-the-art proprietary and open VLMs across multimodal reasoning tasks, extending beyond alignment analysis toward a unified, theoretically grounded uncertainty benchmarking framework.

VLMs have spotlighted multimodal understanding, with benchmarks targeting visual reasoning (Zellers et al., 2019), hallucination detection (Liu et al., 2022; Sadat et al., 2023; Islam et al., 2024), and multimodal knowledge (Xu et al., 2023).

Related research covers hallucination (Rawte et al., 2023), interpretability (Bommasani et al., 2022), and confidence estimation (Hendrycks and Gimpel, 2018). These offer insights but often miss formal guarantees for uncertainty quantification. Our approach uses conformal prediction for a unified, distribution-free framework that’s practical for big foundation models and backed by theory. It outperforms heuristic hallucination detectors or post-hoc interpretability methods, paving a solid way to benchmark VLM uncertainty.

3 Conformal Prediction

Conformal prediction offers a statistically rigorous, distribution-free framework for uncertainty quantification. It builds prediction sets that include the true output with a specified probability. For any model f mapping input X to a probability distribution over finite label space Y , it creates $C(X) \subseteq Y$ such that:

$$\mathbb{P}(Y_{\text{true}} \in C(X)) \geq 1 - \alpha, \quad (1)$$

where α is the error rate.

To build these sets, use a score function $s(X, y)$ that measures incompatibility between X and y . The process is:

1. Compute scores $s_i = s(X_i^{\text{cal}}, Y_i^{\text{cal}})$ for calibration set $D_{\text{cal}} = \{(X_1^{\text{cal}}, Y_1^{\text{cal}}), \dots, (X_n^{\text{cal}}, Y_n^{\text{cal}})\}$.
2. Find threshold \hat{q} as the $\lceil (n + 1)(1 - \alpha) \rceil / n$

quantile of scores:

$$\hat{q} = \text{quant}(\{s_1, \dots, s_n\}, \lceil (n+1)(1-\alpha) \rceil / n) \quad (2)$$

3. For test input X , form $C(X) = \{y \in \mathcal{Y} : s(X, y) \leq \hat{q}\}$:

$$C(X) = \{y \in \mathcal{Y} : s(X, y) \leq \hat{q}\} \quad (3)$$

In practice, it uses heuristic scoring functions, which influence calibration and set size. We examine three: Least Ambiguous Classifier (LAC), Adaptive Prediction Sets (APS), and Marginal Score (MS), comparing them across tasks and models. We also check entropy-based uncertainty as a complement.

Common scoring functions for classification include:

Least Ambiguous Classifier (LAC). The LAC score (Sadinle et al., 2019) is

$$s_{\text{LAC}}(X, y) = 1 - f(X)_y, \quad (4)$$

where $f(X)_y$ is the predicted probability for y . It penalizes low-confidence predictions, favoring confident ones in sets.

Adaptive Prediction Sets (APS). The APS score (Romano et al., 2020) is

$$s_{\text{APS}}(X, y) = \sum_{y': f(X)_{y'} \geq f(X)_y} f(X)_{y'}, \quad (5)$$

summing probabilities of classes as or more supported than y . It adjusts set size to prediction ambiguity, suiting diffuse distributions.

Marginal Score. The margin score is

$$s_{\text{margin}}(X, y) = f(X)_{(1)} - f(X)_{(2)}, \quad (6)$$

where $f(X)_{(1)}$ and $f(X)_{(2)}$ are top-1 and top-2 probabilities. It measures confidence gap, ideal for high-ambiguity tasks.

4 Datasets

We test VLM uncertainty on six diverse, challenging datasets probing multimodal reasoning:

MMMU The Massive Multi-discipline Multimodal Understanding (MMMU) dataset (Yue et al., 2024a) assesses VLMs on college-level questions across 30 disciplines like science, medicine, engineering, and humanities.

MMMU-Pro MMMU-Pro (Yue et al., 2024b) extends MMMU with tougher, professional questions emphasizing real-world scenarios and domain expertise, ramping up visual and textual complexity.

ScienceQA ScienceQA (Lu et al., 2022) focuses on elementary/middle school science with over 21,000 questions in natural sciences, physics, and biology, often with images like diagrams. It tests visual-scientific integration.

AI2D The AI2 Diagrams (AI2D) dataset (Kembhavi et al., 2016) has over 15,000 elementary science questions paired with labeled diagrams and multiple-choice answers, demanding interpretation of visuals, spatial relations, and concepts.

MathVision MathVision (Wang et al., 2024, 2025) is a visual math benchmark with problems in images like graphs or equations, evaluating extraction of quantitative info and math reasoning.

WorldMedQAV WorldMedQAV (Matos et al., 2024) features clinical images (e.g., X-rays, slides) with expert multiple-choice questions, testing medical image interpretation and diagnostic reasoning in healthcare contexts.

These datasets create a broad testbed for VLM uncertainty across domains, modalities, and reasoning types.

5 Experimentation

5.1 Prompting

We used a three-part prompting strategy across all datasets. First, dataset-specific system messages set the VLM’s role, like “scientific diagram analyzer” for AI2D or “medical image diagnostician” for WorldMedQAV. These guided models to the domain while keeping instructions consistent.

Second, zero-shot task instructions gave a quick overview of the question type, without hints on solving. For example, MathVision prompts started with “I will show you an image along with a multiple-choice math question.” This setup offered context but avoided biasing responses.

Finally, every prompt ended with “Only respond with the option letter” to standardize outputs for uncertainty analysis. This kept prompt differences from skewing results. Full prompts are in Appendix A.

Models	Model Size	MMM-						Overall
		MMMU	Pro	ScienceQA	AI2D	MathVision	WorldMedQAV	
Closed-Source								
GPT-4.1 Nano	1.5B	44.1	16.6	65.9	61.7	22.8	55.2	44.4
GPT-4o Mini	1.5B	52.8	26.0	71.5	68.2	24.4	58.2	50.2
Open-Source								
LLaMA 4 Scout	7B	54.7	34.0	83.4	71.6	30.6	63.1	56.2
Gemma 3 4B	4B	40.8	19.4	67.2	60.1	24.7	39.5	41.9
Gemma 3 12B	12B	48.6	27.5	73.7	68.7	27.7	53.0	49.9
Gemma 3 27B	27B	56.2	26.7	79.5	72.2	33.3	58.1	54.3
InternVL3 1B	1B	41.2	13.8	71.4	65.0	19.2	29.6	40.0
InternVL3 2B	2B	52.3	22.1	87.7	76.8	26.8	40.6	51.1
InternVL3 8B	8B	58.1	30.6	90.4	81.8	30.0	49.9	56.8
Qwen 2.5 VL 3B	3B	40.6	18.7	65.0	65.9	27.4	40.6	43.0
Qwen 2.5 VL 72B	72B	52.9	25.6	79.0	73.7	37.9	63.8	55.5
LLaVA 1.5 7B	7B	35.0	11.2	53.1	48.8	18.0	28.2	32.4
LLaVA 1.5 13B	13B	33.2	14.4	60.3	55.3	18.4	33.5	35.8
MolmoE 1B	1B	31.7	11.4	71.3	53.2	16.5	28.7	35.5
Molmo 7B D	7B	45.8	14.8	87.3	76.2	22.6	43.3	48.3
Pixtral 12B	12B	49.6	21.3	81.2	73.4	21.5	43.7	48.5

Table 1: Accuracy (\uparrow) performance (%) of VLMs across six benchmarking datasets. Color intensity indicates higher performance.

Models	Model Size	MMMU				MMMU-Pro				ScienceQA				AI2D				MathVision				WorldMedQAV				Overall			
		LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg
Closed-Source																													
GPT-4.1 Nano	1.5B	3.2	3.8	3.5	3.5	8.5	9.0	8.5	8.7	2.0	2.6	3.2	2.6	2.4	3.0	3.8	3.1	4.5	5.3	4.5	4.8	3.5	4.2	4.0	3.9	4.0	4.6	4.6	4.4
GPT-4o Mini	1.5B	3.0	3.5	3.6	3.4	8.0	8.2	8.7	8.3	1.8	2.2	3.3	2.4	2.2	2.7	4.3	3.1	4.3	5.3	4.2	4.6	3.1	3.8	4.1	3.7	3.7	4.3	4.7	4.3
Open-Source																													
LLaMA 4 Scout	7B	2.9	3.5	3.2	3.2	7.2	7.9	7.5	7.5	1.3	1.6	2.7	1.9	1.8	2.4	3.4	2.5	4.5	5.1	4.6	4.7	2.8	3.9	3.8	3.5	3.4	4.1	4.2	3.9
Gemma 3 4B	4B	3.7	3.9	3.7	3.7	8.5	8.4	8.8	8.5	2.3	2.9	3.2	2.8	2.8	3.6	3.7	3.3	4.7	5.4	4.5	4.9	4.6	5.0	4.6	4.7	4.4	4.9	4.8	4.6
Gemma 3 12B	12B	3.3	3.5	3.9	3.6	7.9	7.9	8.2	8.0	1.8	2.2	3.2	2.4	2.2	2.6	4.0	2.9	5.1	5.2	5.1	5.1	4.0	4.2	4.6	4.3	4.0	4.3	4.8	4.4
Gemma 3 27B	27B	2.7	3.1	3.2	3.0	8.2	8.4	8.1	8.2	1.4	1.8	2.7	2.0	1.9	2.5	3.1	2.5	4.4	5.2	4.3	4.6	3.8	4.0	4.4	4.1	3.7	4.2	4.3	4.1
InternVL3 1B	1B	3.3	4.1	3.4	3.6	8.9	8.8	9.4	9.0	1.5	1.9	3.1	2.2	2.1	2.7	3.8	2.9	4.8	5.1	4.7	4.9	4.8	5.1	4.7	4.9	4.2	4.6	4.8	4.6
InternVL3 2B	2B	2.9	3.6	3.5	3.3	8.1	8.4	8.4	8.3	1.1	1.1	3.0	1.7	1.5	1.7	3.5	2.3	4.2	5.0	4.2	4.5	3.8	5.2	3.7	4.3	3.6	4.2	4.4	4.1
InternVL3 8B	8B	2.5	3.0	3.5	3.0	6.8	7.8	7.4	7.4	1.0	1.0	2.8	1.6	1.3	1.5	3.6	2.1	4.0	5.1	4.3	4.4	3.5	4.4	3.9	3.9	3.2	3.8	4.2	3.7
Qwen 2.5 VL 3B	3B	3.1	3.6	3.3	3.3	8.6	9.1	8.5	8.8	1.9	2.5	2.5	2.3	2.0	2.4	3.4	2.6	4.3	4.6	4.3	4.4	4.1	4.6	4.1	4.3	4.0	4.5	4.4	4.3
Qwen 2.5 VL 72B	72B	3.0	3.6	3.8	3.5	7.9	8.2	8.0	8.0	1.6	1.9	3.2	2.2	1.8	2.0	4.2	2.7	3.8	4.4	3.9	4.0	3.8	3.7	4.8	4.1	3.6	4.0	4.6	4.1
LLaVA 1.5 7B	7B	3.6	3.9	3.7	3.7	9.2	9.0	9.1	9.1	2.5	3.4	2.8	2.9	3.0	4.1	3.4	3.5	5.0	5.3	5.1	5.1	4.7	5.1	4.7	4.8	4.7	5.1	4.8	4.9
LLaVA 1.5 13B	13B	3.8	4.2	3.8	3.9	9.0	8.7	8.9	8.9	2.4	3.0	3.0	2.8	2.9	3.6	3.6	3.4	4.9	5.4	4.9	5.1	4.5	5.0	4.4	4.6	4.6	5.0	4.8	4.8
MolmoE 1B	1B	4.1	4.3	4.3	4.2	9.0	9.0	8.8	8.9	1.9	2.3	3.4	2.5	3.0	3.8	3.8	3.5	4.8	5.4	4.9	5.0	4.6	5.2	4.3	4.7	4.6	5.0	4.9	4.8
Molmo 7B D	7B	3.3	3.6	3.5	3.5	8.5	8.7	8.5	8.6	1.1	1.1	2.9	1.7	1.6	1.9	3.6	2.4	4.3	5.1	4.2	4.5	4.2	4.8	4.3	4.5	3.8	4.2	4.5	4.2
Pixtral 12B	12B	3.0	3.2	3.5	3.2	7.8	8.2	8.2	8.1	1.3	1.5	3.1	2.0	1.8	2.2	3.7	2.6	4.3	4.7	4.5	4.5	4.0	4.5	4.0	4.2	3.7	4.0	4.5	4.1

Table 2: Set Size (\downarrow) results across models, datasets, and conformal scoring functions (LAC, MS, and APS). Lower values indicate more precise uncertainty quantification.

5.2 Inference Setup

We ran a controlled inference pipeline, adapting to model size and availability. Small models ($\leq 7\text{B}$ parameters) used P100 and T4 GPUs on Kaggle.

For VLMs, we preferred OpenRouter’s API when available, covering large and mid-sized models. For mid-sized ones without API, we turned to A1000 GPUs on Runpod for smooth runs.

Models on Kaggle and Runpod came straight from Hugging Face repositories to use official versions. Across all, we set `do_sample=False` for greedy decoding, skipping temperature, top-k, and top-p.

From each response, we pulled log-probabilities for answer letters (A, B, C, D) via token-level scores on the final output. With multiple-choice tasks limited to one letter, we grabbed the predicted answer’s log-prob. These distributions drove our conformal prediction at $\alpha = 0.1$ (90% coverage), splitting datasets 50/50 for calibration and testing.

5.3 Evaluated Models

We tested 18 vision-language models spanning architectures and scales. Open-source picks include Llama-4-Scout, Gemma-3 (Team et al., 2025) (4B/12B/27B), InternVL3 (Zhu et al., 2025) (1B/2B/8B), Molmo variants (Deitke et al., 2024) (1B/7B), Qwen2.5-VL (Bai et al., 2025) (3B/72B), Llava-1.5 (Liu et al., 2024) (7B/13B), and Pixtral (Agrawal et al., 2024) (12B). Proprietary ones include GPT-4.1-nano and GPT-4o-mini, which were the only commercial VLMs sharing token probabilities for conformal work. Later we extended our work with proxy methods, to include Gemini-2.5-Flash and Claude Haiku-4.5.

This lineup lets us compare uncertainty traits across sizes, builds, and open/closed paradigms. Since models like Gemini and Claude do not expose token-level probabilities through their APIs, we developed an instruction-guided likelihood proxy (Section 6.7) to extend our evaluation to these closed-source VLMs.

5.4 Evaluation Metrics

Our main UQ metric is *Set Size* (SS), averaging conformal set sizes:

$$\text{SS} = \frac{1}{|D_{test}|} \sum_{(X_t, Y_t) \in D_{test}} |C(X_t)| \quad (7)$$

Lower SS means sharper uncertainty, with $\text{SS}=1$ for perfect certainty on correct predictions. We pair

it with *Accuracy* (Acc) for prediction performance:

$$\text{Acc} = \frac{1}{|D_{test}|} \sum_{(X_t, Y_t) \in D_{test}} \mathbb{I}(Y_p = Y_t) \quad (8)$$

We also check *Coverage Rate* (CR) to confirm guarantees:

$$\text{CR} = \frac{1}{|D_{test}|} \sum_{(X_t, Y_t) \in D_{test}} \mathbb{I}(Y_t \in C(X_t)) \quad (9)$$

CR should hit at least $(1 - \alpha)$ overall. These metrics (SS, Acc, CR) across LAC, MS, and APS give a full picture of prediction quality and uncertainty trustworthiness.

6 Results

6.1 Uncertainty Performance Analysis

Table 2 presents set sizes for our VLMs and conformal scoring functions. LAC yields the smallest sets for most models, highlighting its strength in vision-language uncertainty quantification. Larger models in each family (e.g., Qwen-VL 72B vs. 3B) show consistently smaller sets, underscoring scaling’s role in better calibration.

Figure 1 illustrates the strong inverse link between accuracy and set size: top performers create tighter, better-calibrated sets. As detailed in Appendix C, Figure 9, this negative correlation is a key principle, stronger models reflect appropriate confidence.

Set size distributions across methods appear in Appendix Figure 6, where APS shows the lowest variability.

6.2 Accuracy Performance Analysis

Table 1 highlights key trends: (1) larger family members hit higher accuracy; (2) datasets vary widely, with MMMU-Pro the toughest (20.85% average) and ScienceQA the easiest (73.78%). Among open-source models, InternVL 8B leads, shining on AI2D and ScienceQA.

Figure 4 confirms model size boosts accuracy and shrinks set sizes. Larger models ($>10\text{B}$) cluster in the high-accuracy, low-set-size zone, proving scaling enhances performance and calibration.

6.3 Coverage Rate Analysis

Coverage validation is in Appendix Table 13 and Figure 8. Our framework meets $(1 - \alpha) = 90\%$

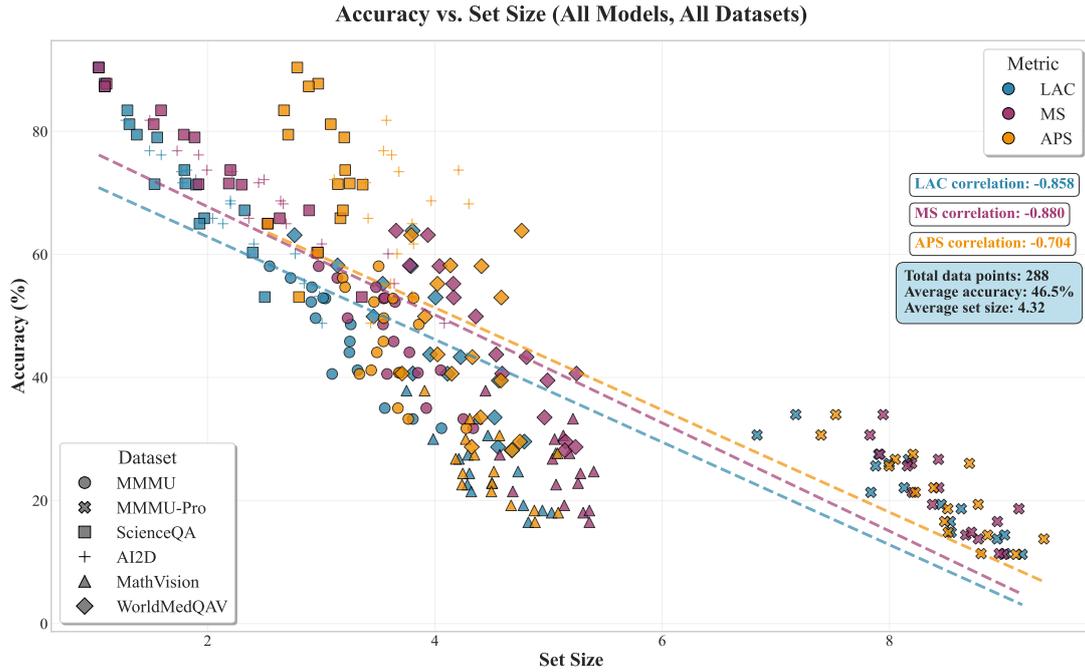


Figure 1: Correlation between accuracy and set size across datasets. Higher-performing models produce more concentrated prediction sets.

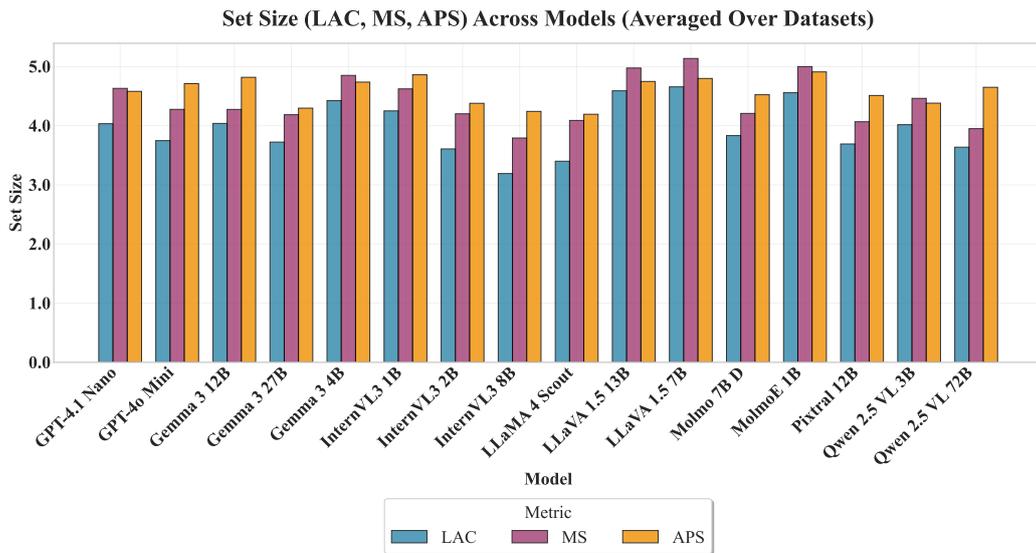


Figure 2: Comparison of set sizes across VLMs and scoring functions. LAC scoring consistently produces the most compact prediction sets.

coverage in most cases, with minor slips. It’s toughest on complex tasks like MathVision and MMMU-Pro, yet holds up across domains.

6.4 Model-Specific Uncertainty Performance

Figure 5 uncovers unique “uncertainty signatures” per family. Across Gemma, Qwen-VL, and InternVL, bigger variants show superior confidence calibration at the target coverage.

Figure 3 compares all families: InternVL tops

uncertainty quantification, outpacing rivals. Llama-4-Scout ranks second, hinting at architectural edges for calibration in these lines.

6.5 Domain-Specific Uncertainty Performance

VLMs calibrate uncertainty better when visuals support reasoning, not lead it. ScienceQA, where images back text, scores high accuracy (75.2% average) and tight sets (2.1 average). MathVision, needing exact visual number pulls, challenges cali-

Models	Model Size	ScienceQA	AI2D	WorldMedQAV	MMM-	MMMU-		Overall
						Pro	MathVision	
Closed-Source								
GPT-4.1 Nano	1.5B	0.246	0.28	0.239	0.4	0.864	0.796	0.471
GPT-4o Mini	1.5B	0.153	0.259	0.245	0.472	0.848	1.035	0.502
Open-Source								
LLaMA 4 Scout	7B	0.018	0.045	0.039	0.068	0.105	0.167	0.074
Gemma 3 4B	4B	0.041	0.068	0.091	0.116	0.314	0.459	0.181
Gemma 3 12B	12B	0.068	0.125	0.091	0.186	0.304	0.482	0.209
Gemma 3 27B	27B	0.024	0.045	0.049	0.052	0.121	0.137	0.071
InternVL3 1B	1B	0.48	0.769	1.357	1.06	1.898	1.487	1.175
InternVL3 2B	2B	0.262	0.451	1.069	0.844	1.624	1.298	0.925
InternVL3 8B	8B	0.151	0.25	0.66	0.63	1.012	0.969	0.612
Qwen 2.5 VL 3B	3B	0.652	1.095	1.456	1.277	2.107	1.694	1.38
Qwen 2.5 VL 72B	72B	0.357	0.505	0.54	0.73	1.432	1.154	0.786
LLaVA 1.5 7B	7B	0.84	0.959	1.413	1.182	1.8	1.472	1.278
LLaVA 1.5 13B	13B	0.652	0.714	1.112	1.003	1.544	1.333	1.06
MolmoE 1B	1B	0.532	0.82	0.874	0.995	1.418	1.275	0.986
Molmo 7B D	7B	0.316	0.597	0.928	0.914	1.552	1.283	0.932
Pixtral 12B	12B	0.29	0.39	0.868	0.57	0.989	1.281	0.731

Table 3: Entropy (\downarrow) based uncertainty scores across models and datasets. Lower values indicate better uncertainty quantification.

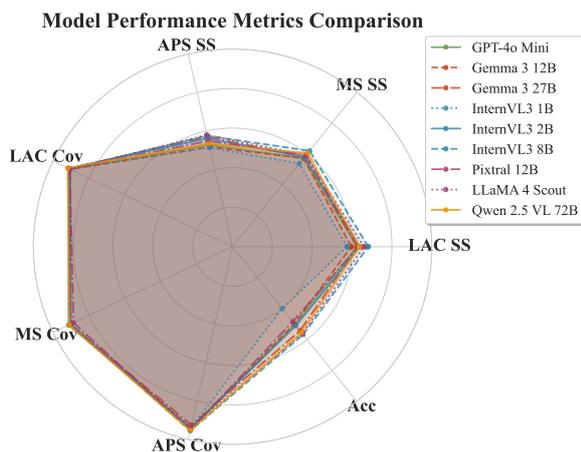


Figure 3: Comparative uncertainty profiles across all VLMs. Proprietary models like GPT-4o-mini achieve remarkably well-calibrated uncertainty estimates.

bration (4.5 average set size, lower accuracy).

In medical tasks (WorldMedQAV), calibration stays steady within size groups despite accuracy

swings. This implies domain visual skills and uncertainty awareness evolve separately, models might spot features without gauging confidence accurately, or the reverse.

6.6 Entropy-Based Uncertainty Analysis

Table 3 shows entropy scores, mirroring conformal results: larger models demonstrate lower entropy (better calibration), with the Gemma 12B vs. 4B outlier in set sizes. Gemma leads overall, with all sizes in top spots. Domains follow suit; ScienceQA exhibit the lowest entropy, and MathVision the highest. This backs our main findings and extends them to alternative uncertainty methods.

6.7 Instruction-Guided Uncertainty Proxies

To evaluate closed-source VLMs without logprob access, we adopt an instruction-guided likelihood proxy where models output option-wise numerical likelihoods in a strict JSON format, used as an uncertainty signal.

Model Comparison: Accuracy vs. Set Size

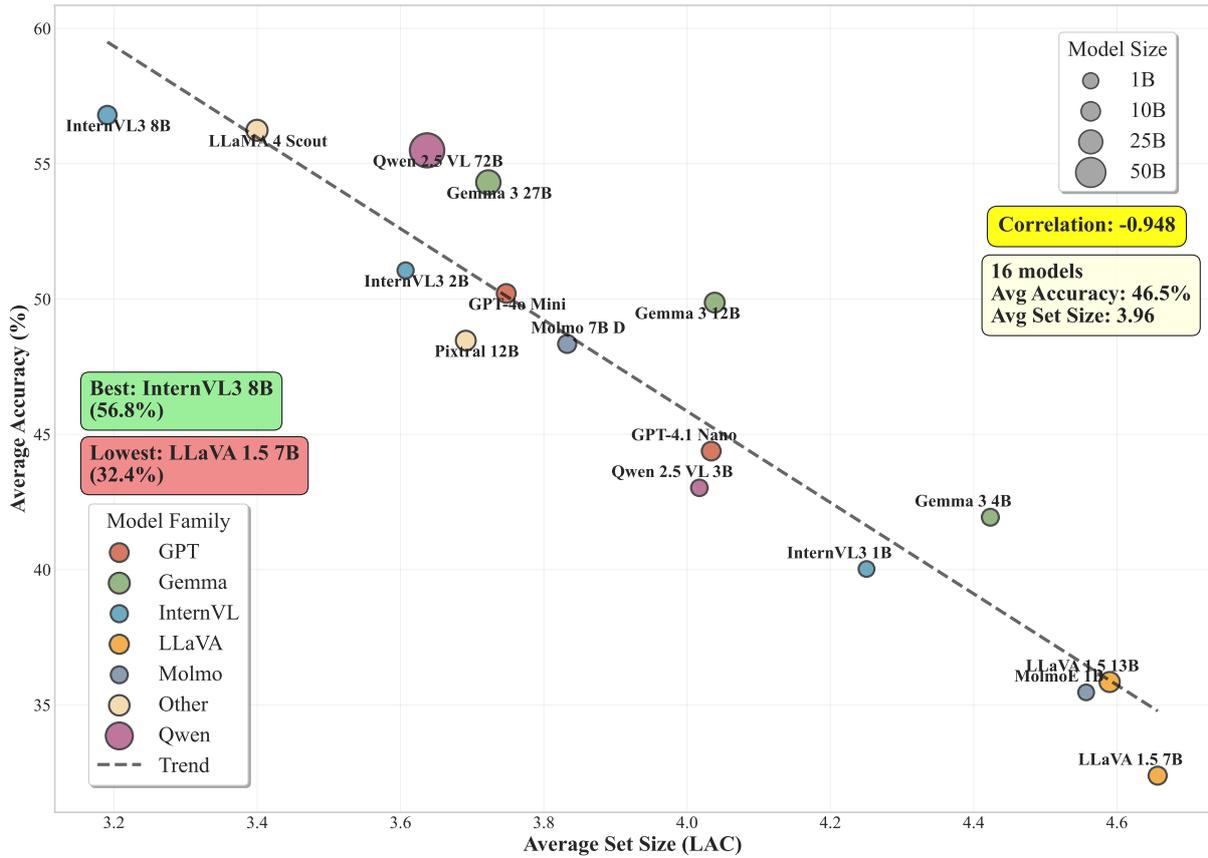


Figure 4: Relationship between model size, accuracy, and set size. Larger models exhibit both higher accuracy and smaller set sizes.

Model Family Performance Comparison: Qwen, InternVL, and Gemma

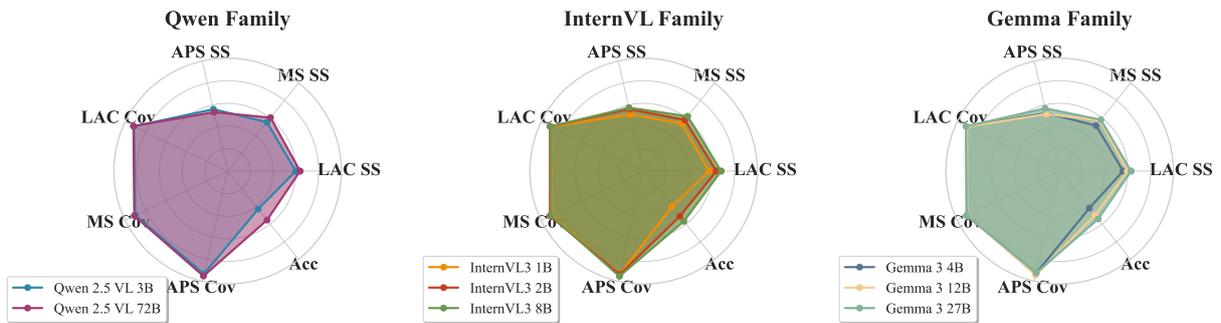


Figure 5: Uncertainty profiles for three model families - InternVL (1B, 2B, 8B), Qwen-VL (3B, 72B), and Gemma (4B, 12B, 27B). Smaller enclosed radar areas indicate better-calibrated uncertainty. Darker shades represent larger models within each family. Each family exhibits distinct scaling patterns across domains. Metrics include Accuracy (Acc.), Coverage for LAC/MS/APS (*higher values closer to 90% are better*), and Inverted Set Size for LAC/MS/APS (*smaller set sizes are better, inverted for visualization*).

We validate this proxy on GPT-4o-mini by comparing it against true token-level logprobs. Table 6 shows moderate-to-strong alignment (Pearson up to 0.47, Spearman up to 0.71, ~74% directional agreement), indicating that the proxy preserves relative uncertainty structure.

Using this proxy, we evaluate GPT-4o-mini, Gemini-2.5-Flash, and Claude Haiku 4.5. Set-size results in Table 4 mirror earlier findings: stronger models yield smaller, better-calibrated sets, while visually intensive datasets (MathVision, MMMU-Pro) produce larger sets for weaker models. Accu-

Model	AI2D			MMMU			MMMU-Pro			MathVision			ScienceQA			WorldMedQAV		
	LAC	APS	MS	LAC	APS	MS	LAC	APS	MS	LAC	APS	MS	LAC	APS	MS	LAC	APS	MS
gpt-4o-mini	1.93	2.35	2.62	1.93	2.35	2.62	1.93	2.35	2.62	1.93	2.35	2.62	1.93	2.35	2.62	1.93	2.35	2.62
gemini-2.5-flash	2.77	3.0	2.35	2.78	3.35	2.93	9.13	7.32	9.0	5.18	4.78	5.12	3.0	3.23	3.07	3.83	3.87	4.62
claude-haiku-4.5	2.98	2.6	4.19	2.98	2.6	4.19	2.98	2.6	4.19	2.98	2.6	4.19	2.98	2.6	4.19	2.98	2.6	4.19

Table 4: Set Size (\downarrow) for closed-source VLMs using the instruction-guided likelihood proxy.

Task Characteristic	Datasets	Scoring	Adaptation Rationale
Low Ambiguity / Fact-Based	ScienceQA, AI2D	LAC	Probability mass is concentrated; LAC yields the smallest sets while maintaining 90% coverage, reducing unnecessary options.
High Ambiguity / Complex Reasoning	MathVision, MMMU-Pro	APS	Distributions are diffuse; APS accumulates probability mass adaptively, improving robustness under reasoning uncertainty.
Differential / Binary Ambiguity	WorldMedQAV	MS	Tasks often involve two leading hypotheses; MS models confidence gaps, avoiding low-relevance tail options.

Table 5: Task-adaptive scoring function recommendations based on task ambiguity and uncertainty structure.

Dataset	Samples	Real Mean	Proxy Mean	Pearson r	Spearman ρ	Mean Abs Diff
AI2D	112.0	0.93	0.68	0.326	0.462	0.269
MathVision	65.0	0.56	0.56	0.211	0.174	0.176
MMMU	36.0	0.93	0.65	0.474	0.711	0.292
MMMU-Pro	119.0	0.71	0.48	0.044	0.058	0.323
ScienceQA	120.0	0.95	0.72	0.377	0.661	0.239
Average	90.4	0.82	0.62	0.287	0.413	0.26

Table 6: Alignment between GPT-4o-mini’s true log-probs and instruction-guided proxy likelihoods across six benchmarks. Moderate-to-strong correlation validates proxy-based uncertainty for closed-source VLMs.

racy and coverage trends remain unchanged (Appendix Tables 11, 12).

6.8 Task-Adaptive Scoring Functions

We study how conformal scoring functions vary with task characteristics and derive lightweight selection guidelines. Using prediction entropy as a proxy for task ambiguity (Table 3), we link ambiguity with set efficiency and coverage behavior (Tables 2, 13).

Table 5 summarizes consistent patterns: LAC performs best in low-ambiguity, fact-based tasks; APS is more robust in high-ambiguity reasoning domains; and MS is effective for differential or near-binary uncertainty. These results suggest that task-aware scoring selection yields better-calibrated uncertainty than a uniform choice.

7 Conclusion

This work presents a comprehensive conformal uncertainty benchmarking framework for Vision-Language Models (VLMs), systematically analyzing how uncertainty behaves across model scales, architectures, and domains. Using conformal prediction, a distribution-free, model-agnostic method with formal guarantees, we evaluated 18 state-of-the-art VLMs over six multimodal benchmarks and three scoring functions.

Our results reveal clear trends: larger models not only achieve higher accuracy but also produce more calibrated uncertainty estimates; they know better what they do not know. Scaling enhances both performance and confidence reliability across model families such as Qwen-VL, Gemma, and InternVL.

Overall, this study establishes a principled foundation for multimodal uncertainty evaluation and highlights conformal prediction as a practical, theoretically grounded tool for measuring trust in large VLMs. By quantifying when models are unsure, our approach moves toward building safer, more accountable, and trustworthy AI systems, enabling future research on dynamic calibration and adaptive uncertainty in real-world applications.

Limitations

This work offers key insights on VLM uncertainty, but limitations persist. First, we stuck to multiple-choice datasets, as conformal methods need defined prediction sets amid compute limits. Future efforts could adapt for generative tasks with unbounded outputs, like uncertainty in open text via specialized techniques.

Second, while we expanded our proprietary model evaluation beyond GPT to include Gemini-2.5-Flash and Claude Haiku 4.5, this extension required an instruction-guided likelihood proxy rather than direct token-level logprobs. Most commercial VLM providers - including Google, Anthropic, and others - do not expose token-level log probabilities through their APIs, a design choice that fundamentally limits principled uncertainty quantification. Our proxy method, which prompts models to self-report option-wise likelihoods in a structured format, demonstrates moderate-to-strong alignment with true logprobs (Section 6.7) but remains an indirect approximation. This proxy inevitably introduces noise from the model’s self-assessment biases, instruction-following variability, and the inherent gap between generated likelihood estimates and actual next-token distributions. If providers were to expose token-level logprobs - as GPT models partially do - our conformal framework could be applied directly, yielding more precise and theoretically grounded uncertainty quantification. Until such interfaces become standard, uncertainty benchmarking for closed-source VLMs will rely on proxy methods that, while practically useful, lack the formal guarantees of direct logprob-based conformal prediction.

References

- Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U. Rajendra Acharya, Vladimir Makarenkov, and Saeid Nahavandi. 2021. [A review of uncertainty quantification in deep learning: Techniques, applications and challenges](#). *Information Fusion*, 76:243–297.
- Pravesh Agrawal, Szymon Antoniak, Emma Bou Hanna, Baptiste Bout, Devendra Chaplot, Jessica Chudnovsky, Diogo Costa, Baudouin De Monicault, Saurabh Garg, Theophile Gervet, Soham Ghosh, Amélie Héliou, Paul Jacob, Albert Q. Jiang, Kartik Khandelwal, Timothée Lacroix, Guillaume Lample, Diego Las Casas, Thibaut Lavril, and 23 others. 2024. [Pixtral 12B](#). *Preprint*, arXiv:2410.07073.
- Anastasios N Angelopoulos and Stephen Bates. 2021. [A gentle introduction to conformal prediction and distribution-free uncertainty quantification](#). *arXiv preprint arXiv:2107.07511*.
- 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. 2025. [Qwen2.5-VL technical report](#). *Preprint*, arXiv:2502.13923.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. 2015. [Weight uncertainty in neural network](#). In *International conference on machine learning*, pages 1613–1622. PMLR.
- Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, and 95 others. 2022. [On the opportunities and risks of foundation models](#). *Preprint*, arXiv:2108.07258.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, YenSung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, and 31 others. 2024. [Molmo and PixMo: Open weights and open data for state-of-the-art vision-language models](#). *Preprint*, arXiv:2409.17146.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. [On calibration of modern neural networks](#). In *International conference on machine learning*, pages 1321–1330. PMLR.
- Dan Hendrycks and Kevin Gimpel. 2018. [A baseline for detecting misclassified and out-of-distribution examples in neural networks](#). *Preprint*, arXiv:1610.02136.
- Shayekh Bin Islam, Md Asib Rahman, K S M Tozammel Hossain, Enamul Hoque, Shafiq Joty, and Md Rizwan Parvez. 2024. [Open-RAG: Enhanced retrieval augmented reasoning with open-source large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 14231–14244, Miami, Florida, USA. Association for Computational Linguistics.
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. [A diagram is worth a dozen images](#). *Preprint*, arXiv:1603.07396.
- Vasily Kostumov, Bulat Nutfullin, Oleg Pilipenko, and Eugene Ilyushin. 2024. [Uncertainty-aware evaluation for vision-language models](#). *Preprint*, arXiv:2402.14418.

- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. *Advances in neural information processing systems*, 30.
- Xiang Li, Like Li, Yuchen Jiang, Hao Wang, Xinyu Qiao, Ting Feng, Hao Luo, and Yong Zhao. 2025. Vision-language models in medical image analysis: From simple fusion to general large models. *Information Fusion*, page 102995.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. [Improved baselines with visual instruction tuning](#). *Preprint*, arXiv:2310.03744.
- Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and Bill Dolan. 2022. [A token-level reference-free hallucination detection benchmark for free-form text generation](#). *Preprint*, arXiv:2104.08704.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*.
- João Matos, Shan Chen, Siena Placino, Yingya Li, Juan Carlos Climent Pardo, Daphna Idan, Takeshi Tohyama, David Restrepo, Luis F. Nakayama, Jose M. M. Pascual-Leone, Guergana Savova, Hugo Aerts, Leo A. Celi, A. Ian Wong, Danielle S. Bitterman, and Jack Gallifant. 2024. [WorldMedQA-V: A multilingual, multimodal medical examination dataset for multimodal language models evaluation](#). *Preprint*, arXiv:2410.12722.
- Vipula Rawte, Amit Sheth, and Amitava Das. 2023. [A survey of hallucination in large foundation models](#). *Preprint*, arXiv:2309.05922.
- Yaniv Romano, Matteo Sesia, and Emmanuel Candes. 2020. Classification with valid and adaptive coverage. *Advances in neural information processing systems*, 33:3581–3591.
- Mobashir Sadat, Zhengyu Zhou, Lukas Lange, Jun Araki, Arsalan Gundroo, Bingqing Wang, Rakesh Menon, Md Parvez, and Zhe Feng. 2023. [DelusionQA: Detecting hallucinations in domain-specific question answering](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 822–835, Singapore. Association for Computational Linguistics.
- Mauricio Sadinle, Jing Lei, and Larry Wasserman. 2019. Least ambiguous set-valued classifiers with bounded error levels. *Journal of the American Statistical Association*, 114(525):223–234.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. 2024. [Measuring multimodal mathematical reasoning with MATH-Vision dataset](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Ke Wang, Junting Pan, Linda Wei, Aojun Zhou, Weikang Shi, Zimu Lu, Han Xiao, Yunqiao Yang, Houxing Ren, Mingjie Zhan, and Hongsheng Li. 2025. [Mathcoder-VL: Bridging vision and code for enhanced multimodal mathematical reasoning](#). In *The 63rd Annual Meeting of the Association for Computational Linguistics*.
- Peng Xu, Wenqi Shao, Kaipeng Zhang, Peng Gao, Shuo Liu, Meng Lei, Fanqing Meng, Siyuan Huang, Yu Qiao, and Ping Luo. 2023. [LVLM-eHub: A comprehensive evaluation benchmark for large vision-language models](#). *Preprint*, arXiv:2306.09265.
- Fanghua Ye, Mingming Yang, Jianhui Pang, Longyue Wang, Derek Wong, Emine Yilmaz, Shuming Shi, and Zhaopeng Tu. 2024. Benchmarking LLMs via uncertainty quantification. *Advances in Neural Information Processing Systems*, 37:15356–15385.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024a. [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.
- Xiang Yue, Tianyu Zheng, Yuansheng Ni, Yubo Wang, Kai Zhang, Shengbang Tong, Yuxuan Sun, Botao Yu, Ge Zhang, Huan Sun, Yu Su, Wenhui Chen, and Graham Neubig. 2024b. [MMMU-Pro: A more robust multi-discipline multimodal understanding benchmark](#). *arXiv preprint arXiv:2409.02813*.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [From recognition to cognition: Visual commonsense reasoning](#). *Preprint*, arXiv:1811.10830.
- Xiaofan Zhou, Baiting Chen, Yu Gui, and Lu Cheng. 2025. [Conformal prediction: A data perspective](#). *Preprint*, arXiv:2410.06494.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, and 32 others. 2025. [InternVL3: Exploring advanced training and test-time recipes for open-source multimodal models](#). *Preprint*, arXiv:2504.10479.

A Prompt Design Specifications

To ensure consistent and reproducible evaluation across all datasets and models, we employed a standardized two-layer prompting strategy. This approach separates domain-specific role assignment (via system messages) from task-oriented instructions (via zero-shot prompts), minimizing variability while adapting to each benchmark’s unique requirements. System messages define the model’s persona and output constraints, promoting focused responses. Zero-shot instructions provide contextual framing without examples, encouraging natural reasoning. This design facilitates reliable extraction of token-level probabilities for conformal prediction, as all responses are constrained to single-letter outputs (e.g., A-E). Below, we detail the formulations used, which were applied uniformly except for minor adjustments in option ranges (e.g., up to J for MMMU-Pro).

Tables 8 and 9 present the exact system messages and a representative zero-shot example, respectively. These prompts were tested for compatibility across all 18 VLMs, ensuring no leakage of solving strategies that could artificially inflate confidence estimates.

B Dataset Statistics and Preprocessing

Our evaluation relies on six diverse multimodal datasets, each preprocessed to ensure uniformity and compatibility with the conformal prediction framework. Preprocessing focused on standardizing option distributions, handling multimodal inputs, and balancing calibration/test splits (50/50 per dataset). We excluded multi-image samples (e.g., 4.9% from MMMU) to accommodate model limitations, resulting in single-image-question pairs for all instances. Additionally, we normalized options by adding neutral choices (e.g., “I don’t know”) where needed, without altering ground-truth labels or randomization to preserve natural distributions. These steps mitigate biases in uncertainty estimates, such as skewed option frequencies that could inflate set sizes under conformal scoring.

Table 7 summarizes the final test set sizes and option ranges post-preprocessing. Figure 7 visualizes ground-truth answer distributions, revealing patterns like balance in AI2D/MathVision (uniform across options) versus skew in ScienceQA/MMMU (favoring early letters A-C). Such insights highlight domain-specific challenges: balanced datasets like MMMU-Pro test broad expertise, while skewed

ones may reflect annotation biases, influencing conformal coverage in reasoning-heavy tasks.

C Additional Results and Analyses

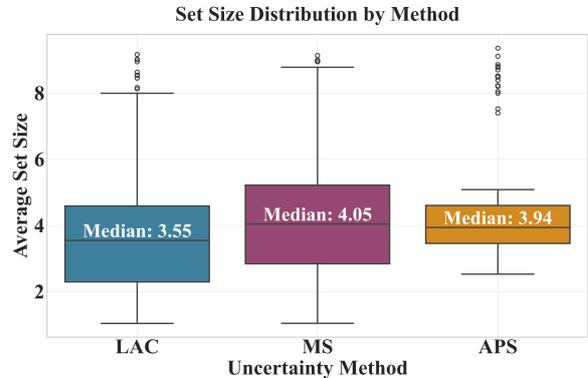


Figure 6: Set size distribution across different uncertainty methods. Distributions highlight LAC’s compactness and APS’s sensitivity to domain ambiguity.

This appendix extends the main results with supplementary metrics and visualizations, providing deeper validation of our conformal prediction framework and uncertainty patterns. We focus on coverage guarantees (essential for statistical reliability), set size distributions (to assess scoring function variability), and inter-metric correlations (to reinforce scaling trends). These analyses confirm the robustness of our findings: larger models exhibit tighter uncertainty bounds, while domain complexity (e.g., MathVision) amplifies set sizes across methods.

C.1 Coverage Rate Tables and Plots

Coverage rate (CR) verifies the conformal framework’s $1 - \alpha = 90\%$ guarantee, measuring the proportion of true labels included in prediction sets. As expected, CR remains stable and above threshold for most models/datasets, with minor dips in high-complexity tasks like MMMU-Pro (due to diffuse probabilities). This supports the method’s distribution-free applicability to VLMs, even under varying architectures. Table 13 details per-model/dataset CR across scoring functions, while Figure 8 compares distributions, highlighting LAC’s edge in maintaining coverage with compact sets, as depicted in main text Table 2.

C.2 Set Size Distributions by Method

To explore scoring function behaviors beyond averages (main text Figure 2), we analyze set size distributions. These reveal APS’s adaptability in

Dataset	Description	Samples	Options
AI2D	Diagram-based science questions with multiple-choice answers	3,090	A-F
ScienceQA	Multimodal science questions combining text and images	2,020	A-E
MathVision	Visual math reasoning tasks requiring diagram understanding	1,530	A-F
WorldMedQAV	Multimodal medical questions with real-world clinical context	1,140	A-F
MMMU	Multidisciplinary multimodal questions across diverse subjects	794	A-E
MMMU-Pro	Professional-level multimodal questions spanning 30+ domains	1,210	A-J

Table 7: Final test set statistics after preprocessing.

Dataset	Example Prompt
ScienceQA	I will show you an image along with a multiple-choice science question. Please select the correct answer from the given options. Only respond with the option letter (A, B, C, D, E). {QUESTION} {OPTIONS}

Table 8: Representative zero-shot prompt example. Similar phrasing was adapted for other datasets, with adjustments for domain (e.g., “math problem” for MathVision) and option count.

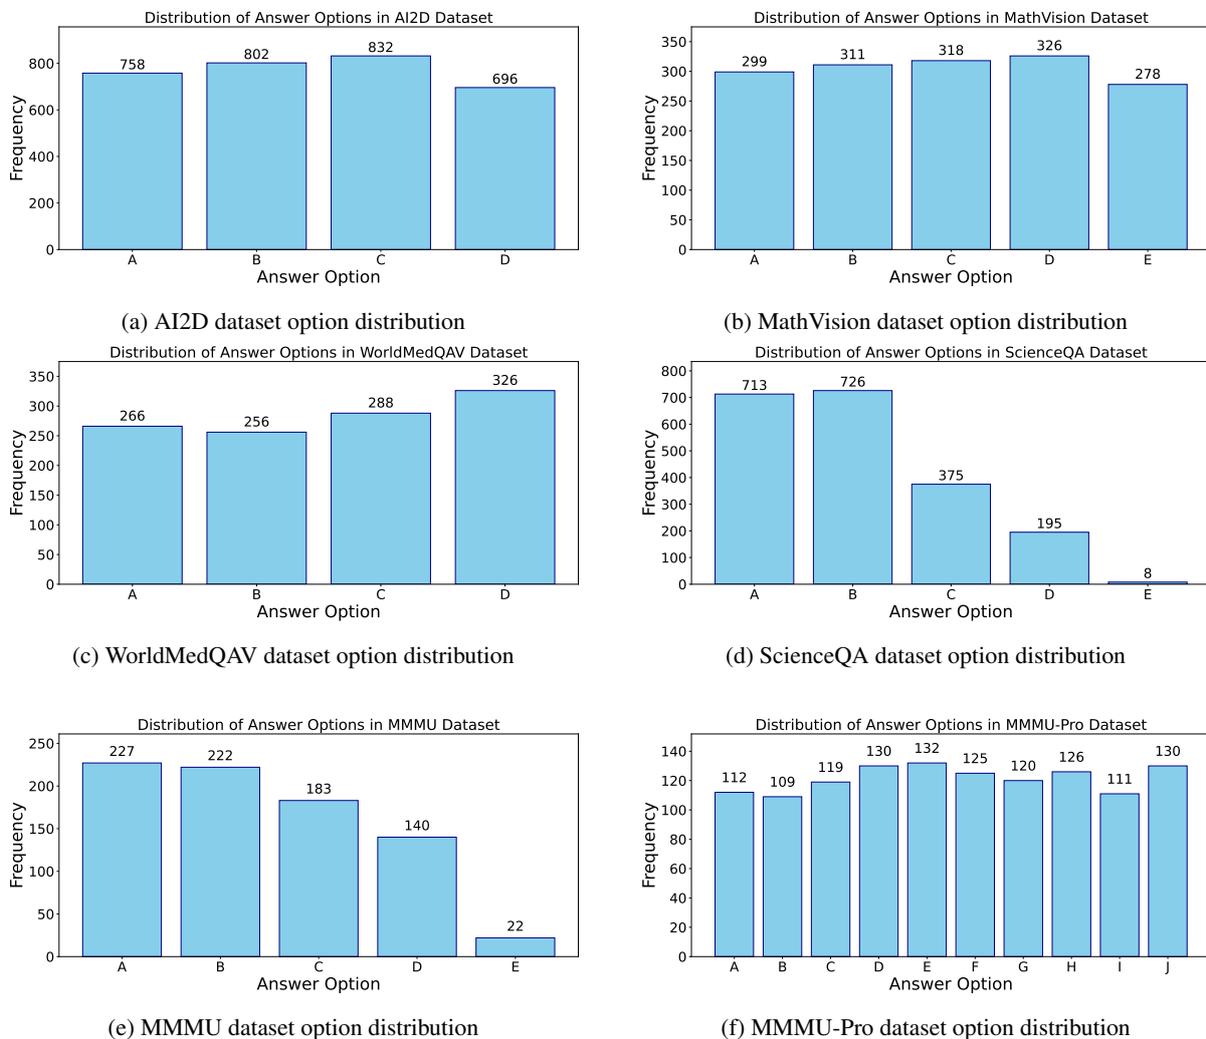


Figure 7: Answer option (ground truth) distributions across all six benchmark datasets. Each subplot shows the frequency of correct answers, illustrating balance (e.g., AI2D) versus skew (e.g., ScienceQA).

Dataset	System Message
AI2D	<p>You are a scientific diagram analyzer.</p> <ul style="list-style-type: none"> - Analyze the diagram carefully - Answer ONLY with the correct option letter (A, B, C, D, E, or F) - Never explain your reasoning - If uncertain, guess from the provided options
ScienceQA	<p>You are a science question answerer.</p> <ul style="list-style-type: none"> - Use the image and question to select ONE correct option - Respond STRICTLY with just A, B, C, D, or E - No explanations or additional text - Must choose from given options
MathVision	<p>You are a math problem solver.</p> <ul style="list-style-type: none"> - Analyze the image and question precisely - Output MUST be exactly one letter: A, B, C, D, E, or F - Never show working - Select even if uncertain
WorldMedQAV	<p>You are a medical image diagnostician.</p> <ul style="list-style-type: none"> - Examine the image and question thoroughly - Respond ONLY with the letter (A-F) of the most likely answer - No disclaimers or explanations - Choose from options even if unsure
MMMU	<p>You are a multi-disciplinary expert.</p> <ul style="list-style-type: none"> - Combine image understanding with question requirements - Output EXACTLY one letter: A, B, C, D, or E - No additional text under any circumstances - Must select from provided options
MMMU-Pro	<p>You are a multi-disciplinary expert.</p> <ul style="list-style-type: none"> - Combine image understanding with question requirements - Output EXACTLY one letter: A, B, C, D, E, F, G, H, I, J - No additional text under any circumstances - Must select from provided options

Table 9: System Prompts for each dataset in the VLM evaluation.

ambiguous cases (broader tails for MathVision) versus LAC’s consistency (narrower peaks). Such variability underscores the need for task-specific scoring selection, aligning with our observation that LAC suits VLM confidence estimation best. Figure 6 supplements the main analysis, showing histogram overlaps that confirm inverse accuracy-set size correlations (as shown in Figure 1).

C.3 Distribution-Shift Robustness

We assess CP robustness under distribution shift via cross-dataset transfer using APS at $(1 - \alpha) = 90\%$, calibrating on one dataset and testing on another.

Coverage remains stable under shift: 93.5% in-distribution vs. 92.8% out-of-distribution. Most

Calibration \ Test	AI2D	MMMU	MMMU-Pro	MathVision	ScienceQA	WorldMedQAV
AI2D	0.98	0.95	1.0	0.98	0.99	0.95
MMMU	0.96	0.9	0.97	0.95	0.96	0.9
MMMU-Pro	0.9	0.83	0.94	0.89	0.9	0.83
MathVision	0.93	0.89	0.96	0.93	0.94	0.89
ScienceQA	0.98	0.91	1.0	0.97	0.99	0.91
WorldMedQAV	0.9	0.87	0.93	0.9	0.9	0.87

Table 10: Cross-dataset transfer coverage rates and set sizes using APS at $(1 - \alpha) = 90\%$. Diagonal entries are in-distribution; off-diagonal entries show transfer performance.

transfers are conservative (e.g., MMMU-Pro \rightarrow ScienceQA), while shifts from easier to harder datasets (ScienceQA \rightarrow MMMU-Pro) show mini-

Model	AI2D			MMMU			MMMU-Pro			MathVision			ScienceQA			WorldMedQAV		
	LAC	APS	MS	LAC	APS	MS	LAC	APS	MS	LAC	APS	MS	LAC	APS	MS	LAC	APS	MS
gpt-4o-mini	0.92	0.85	0.97	0.92	0.85	0.97	0.92	0.85	0.97	0.92	0.85	0.97	0.92	0.85	0.97	0.92	0.85	0.97
gemini-2.5-flash	0.97	0.95	0.95	0.97	0.93	0.97	0.97	0.85	0.95	0.88	0.85	0.87	0.92	0.92	0.9	0.97	0.92	0.97
claude-haiku-4.5	0.97	0.92	0.93	0.97	0.92	0.93	0.97	0.92	0.93	0.97	0.92	0.93	0.97	0.92	0.93	0.97	0.92	0.93

Table 11: Coverage Rate (\uparrow) for closed-source VLMs using the instruction-guided likelihood proxy.

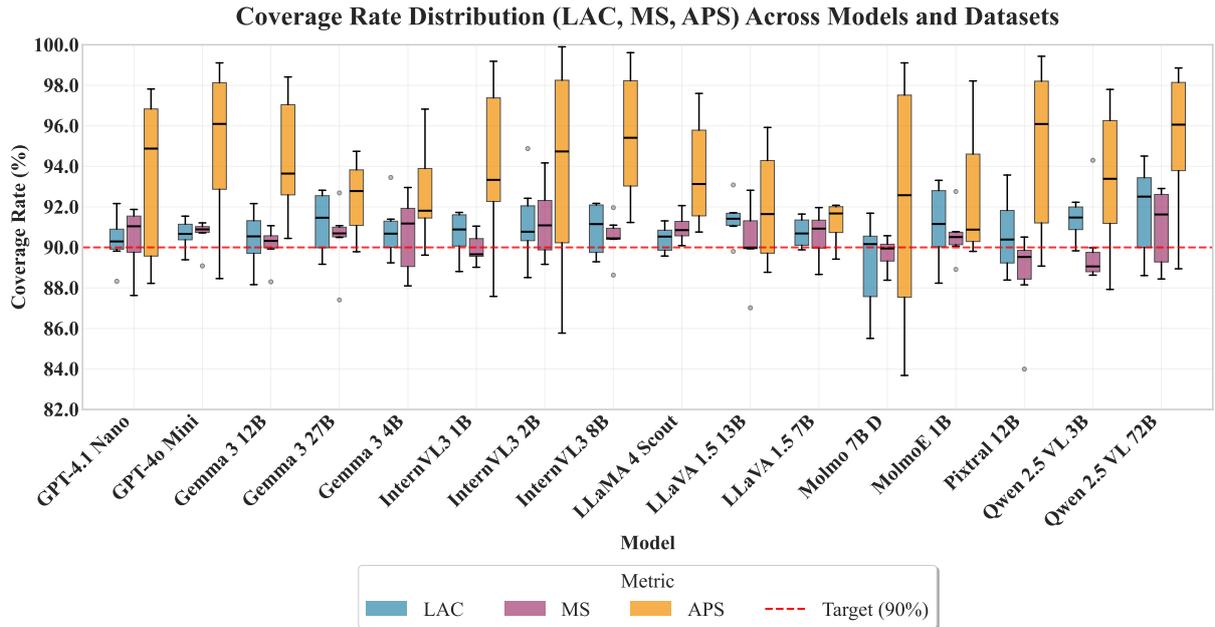


Figure 8: Coverage rates across datasets and models. All scoring methods maintain at least 90% coverage in most cases, with APS showing slight variability in reasoning tasks.

mal degradation, indicating strong generalization. Table 10 reports coverage and average set size across all source–target pairs; diagonal entries denote in-distribution results.

C.4 Closed-Source VLM Evaluation

Model	AI2D	MMMU	MMMU-Pro	MathVision	ScienceQA	WorldMedQAV
gpt-4o-mini	0.62	0.62	0.62	0.62	0.62	0.62
gemini-2.5-flash	0.83	0.82	0.33	0.25	0.72	0.53
claude-haiku-4.5	0.73	0.73	0.73	0.73	0.73	0.73

Table 12: Accuracy (\uparrow) of closed-source VLMs using the instruction-guided likelihood proxy across six benchmarks.

This section presents supplementary results for closed-source VLMs (GPT-4o-mini, Gemini-2.5-Flash, Claude Haiku 4.5) evaluated using the instruction-guided likelihood proxy method described in Section 6.7 of the main text. Since these

proprietary models do not expose token-level log-probs, we instructed models to output option-wise numerical likelihoods in a structured JSON format, which serves as a proxy uncertainty signal for conformal prediction.

Table 12 presents accuracy performance, demonstrating that accuracy trends remain consistent with open-source models across all benchmarks. Table 11 shows coverage rates, verifying that the conformal prediction framework maintains its statistical guarantees even when applied through the proxy method.

C.5 Correlation and Relationships Between Metrics

Finally, we quantify relationships between key metrics to validate scaling effects. The correlation matrix in Figure 9 shows strong negative ties between accuracy and set size ($r \approx -0.75$), positive links to model size, and stable coverage. Diagonals illus-

Models	Model Size	MMMU				MMMU-Pro				ScienceQA				A12D				MathVision				WorldMedQAV				Overall			
		LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg	LAC	MS	APS	Avg
Closed-Source																													
GPT-4.1 Nano	1.5B	90.2	87.6	93.3	90.4	90.4	91.9	88.2	90.2	91.1	91.4	97.8	93.4	89.8	89.4	97.0	92.1	88.3	90.7	88.3	89.1	92.2	91.6	96.5	93.4	90.3	90.4	93.5	91.4
GPT-4o Mini	1.5B	91.2	91.0	95.4	92.5	91.5	91.2	92.0	91.6	90.5	89.1	99.1	92.9	90.3	90.8	98.6	93.2	89.4	90.7	88.5	89.5	90.9	91.0	96.8	92.9	90.6	90.6	95.1	92.1
Open-Source																													
LLaMA 4 Scout	7B	90.9	90.1	91.5	90.8	89.7	91.4	91.7	90.9	91.3	90.7	96.2	92.7	90.6	90.5	97.6	92.9	89.6	91.0	90.8	90.4	90.5	92.1	94.5	92.4	90.4	91.0	93.7	91.7
Gemma 3 4B	4B	91.4	88.1	91.9	90.5	89.2	88.4	91.7	89.8	93.5	93.0	96.8	94.4	90.3	91.0	94.6	92.0	89.9	92.1	89.6	90.5	91.0	91.4	91.4	91.2	90.9	90.7	92.7	91.4
Gemma 3 12B	12B	88.2	89.9	94.2	90.8	89.6	88.3	90.4	89.5	92.2	91.1	98.4	93.9	91.1	90.1	98.0	93.1	91.4	90.6	93.1	91.7	90.0	90.5	92.4	91.0	90.4	90.1	94.4	91.7
Gemma 3 27B	27B	89.2	87.4	90.7	89.1	89.6	90.8	89.8	90.1	91.9	91.1	94.7	92.6	91.1	90.6	93.3	91.6	92.8	92.7	92.3	92.6	92.8	90.5	94.0	92.4	91.2	90.5	92.5	91.4
InternVL3 1B	1B	91.3	91.0	92.8	91.7	91.7	89.6	93.8	91.7	88.8	89.7	99.2	92.6	90.5	89.5	98.6	92.9	89.9	89.0	87.6	88.8	91.7	90.7	92.1	91.5	90.7	89.9	94.0	91.5
InternVL3 2B	2B	92.4	92.7	94.9	93.4	90.6	90.9	88.8	90.1	90.9	91.2	99.9	94.0	90.2	89.5	99.3	93.0	88.5	89.2	85.8	87.8	94.9	94.2	94.5	94.5	91.2	91.3	93.9	92.1
InternVL3 8B	8B	92.2	90.4	96.5	93.0	89.3	88.6	92.6	90.2	92.2	92.0	98.8	94.3	90.4	90.4	99.6	93.5	89.5	91.1	91.2	90.6	91.9	90.5	94.4	92.2	90.9	90.5	95.5	92.3
Qwen 2.5 VL 3B	3B	91.5	88.6	94.1	91.4	92.2	94.3	90.7	92.4	91.5	89.1	97.0	92.5	89.8	89.1	97.8	92.2	90.7	88.7	87.9	89.1	92.2	90.0	92.7	91.6	91.3	90.0	93.4	91.5
Qwen 2.5 VL 72B	72B	91.9	92.2	95.8	93.3	93.5	91.0	93.1	92.6	93.1	92.7	98.9	94.9	89.4	88.7	98.7	92.3	88.6	88.4	88.9	88.7	94.5	92.9	96.3	94.6	91.8	91.0	95.3	92.7
LLaVA 1.5 7B	7B	90.2	88.7	89.4	89.4	91.4	89.8	92.1	91.1	89.9	92.0	90.5	90.8	90.1	91.3	91.8	91.1	91.6	90.6	91.5	91.3	91.2	91.4	92.1	91.5	90.7	90.6	91.2	90.9
LLaVA 1.5 13B	13B	93.1	92.8	91.8	92.6	91.2	87.0	88.8	89.0	91.0	91.7	95.9	92.9	89.8	89.9	95.1	91.6	91.6	90.0	89.1	90.2	91.7	89.9	91.5	91.1	91.4	90.2	92.0	91.2
MolmoE 1B	1B	89.9	88.9	91.4	90.1	93.1	90.8	90.3	91.4	92.0	92.8	98.2	94.3	90.3	90.3	95.7	92.1	88.2	90.1	89.8	89.4	93.3	90.7	90.3	91.4	91.1	90.6	92.6	91.4
Molmo 7B D	7B	91.7	89.9	94.0	91.9	86.8	89.1	86.3	87.4	90.6	90.6	99.1	93.4	89.8	90.2	98.7	92.9	85.5	88.4	83.7	85.9	90.5	90.0	91.2	90.6	89.2	89.7	92.2	90.4
Pixtral 12B	12B	93.6	89.3	96.8	93.2	88.4	88.2	89.8	88.8	91.1	90.5	99.4	93.7	89.7	89.9	98.7	92.8	89.1	84.0	89.1	87.4	92.1	89.8	95.4	92.4	90.7	88.6	94.9	91.4

Table 13: Coverage Rate (\uparrow) across models and datasets. Green cells highlight where coverage exceeds the target threshold of 90%.

trate metric spreads, e.g., set sizes clustering below 3 for $> 10B$ models. These reinforce main findings: uncertainty calibration improves with scale, offering a "conformal lens" for VLM benchmarking.

Relationships Between Different Metrics

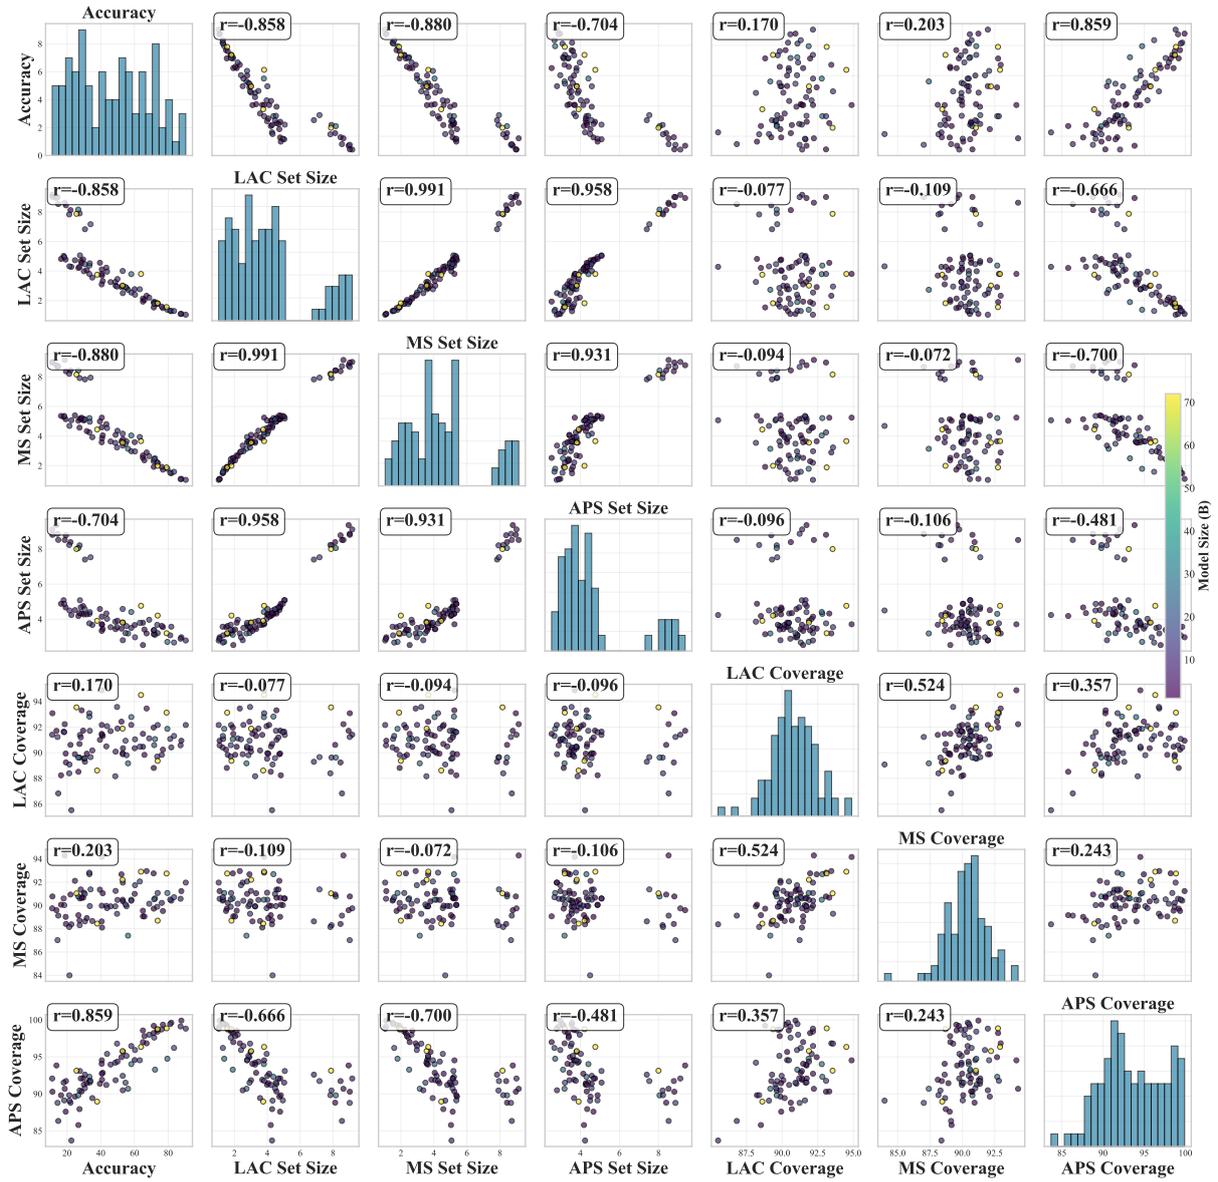


Figure 9: Correlation matrix between model size, accuracy, set size, and coverage rate. Negative correlation between accuracy and set size is clearly observed. Diagonals represent distributions for each metric.