

KMMMU: Evaluation of Massive Multi-discipline Multimodal Understanding in Korean Language and Context

Nahyun Lee^{1,4,5*} Guijin Son^{2,4,5*} Hyunwoo Ko^{4,5} Chanyoung Kim^{3,5}
Junyoung An² Kyubeen Han⁵ Youngjae Yu² Il-Youp Kwak^{1†}

¹Chung-Ang University ²Seoul National University ³SK A.X ⁴OnlineAI ⁵HAE-RAE

Abstract

Multimodal evaluation is strengthened when benchmarks reflect the languages, institutions, and disciplinary conventions that shape expert visual reasoning. We introduce **KMMMU**, a native Korean benchmark built from Korean-native examination materials. KMMMU contains 3,466 questions spanning nine disciplines and nine visual modality categories. For targeted evaluation on challenging questions and Korean-specific knowledge, we additionally provide KMMMU-Adv and KMMMU-Local, respectively. Experiments show that the strongest evaluated open-weight model reaches 42.05% on the KMMMU, while the best evaluated proprietary model achieves 52.42% on the KMMMU-Adv. Performance varies substantially across disciplines and answer formats, with Korean-specific questions showing gaps of up to 13.43%. Error analysis suggests that inspected failures often involve convention-to-label mapping, few-shot symbolic induction, localized knowledge recall, and domain-specific standards, not reasoning depth alone. KMMMU provides a testbed for native Korean expert-level multimodal understanding beyond English-centric or translation-based evaluation.

¹

1 Introduction

Multimodal Large Language Models (MLLMs) have shown strong performance on a range of vision–language tasks, including visual recognition, document understanding, and multimodal question answering (Alayrac et al., 2022; Li et al., 2023a; Liu et al., 2023; Team and Google, 2023; Bai et al., 2025). As MLLMs are deployed across a wider range of languages and domains, evaluation must account for tasks where linguistic context, visual




Figure 1: Comparison of English (MMMU, MMMU-Pro), Japanese (JMMMU, JMMMU-Pro), and Korean (others) multimodal benchmarks. Each point is positioned by benchmark size (x-axis, log scale) and difficulty proxy (100 – peak public score), with lighter colors indicating more recent releases. Shaded regions mark two common limitations: small size (left) and high peak model performance (bottom).

structure, domain knowledge, official standards, and specialized terminology jointly shape model performance. However, most widely used multimodal benchmarks remain English-centric or are derived from translated datasets (Li et al., 2023b; Yue et al., 2024). Korean multimodal benchmarks have made important progress, but many either target narrower settings or have reported peak scores high enough to leave limited separation among current strong models (Figure 1).

To bridge this gap, we introduce **KMMMU**, a collection of **3,466** questions drawn from Korean assessment sources, spanning nine disciplines, nine visual modality categories, and both *multiple-choice* and *open-form* question formats. To support practitioners with lightweight evaluation, we further provide two subsets: **KMMMU-Local**, which targets Korean-specific localized knowledge, and **KMMMU-Adv**, a challenging subset consisting of questions selected models fail to answer.

Experiments on KMMMU reveal several consistent findings. We evaluate 26 multimodal models spanning open-weight and proprietary systems, and among them the strongest evaluated open-weight model reaches 42.05% on the full set, while the best

*Equal contribution.

†Corresponding author: ikwak2@cau.ac.kr

¹Dataset is available at <https://huggingface.co/datasets/HAERAE-HUB/KMMMU>

Pro (Hong et al., 2025) reports that the gap between translated MMMLU (OpenAI, 2024) and locally authored Korean professional exams is small in medicine but substantially larger in law-related domains, where country-specific knowledge is indispensable. These gaps have so far been examined only in text, however; how Korean characters embedded in images or culturally grounded visual content affect multimodal model performance remains largely unexplored. KMMMU follows the MMMU paradigm in the same sense as JMMMU and CMMMU. The name refers to this evaluation paradigm, namely broad expert-level multimodal understanding measured with examination questions in multiple-choice and open-form formats. KMMMU does not reuse or adapt any MMMU item, and every question comes from a pre-existing Korean examination.

Benchmark	Year	Size	Filtering Method
HellaSwag <i>Commonsense NLI</i>	2019	70K	Adversarially selected hard generated distractors
MMLU-Pro <i>Text-only exams</i>	2024	12K	Remove items solved by > 4 of 8 LLMs
SimpleQA <i>Short-form factuality</i>	2024	4.3K	Collect questions adversarially against GPT-4 responses
MMMU-Pro <i>Multimodal exams</i>	2025	3.5K	Remove items solved by ≥ 3 of 4 text-only LLMs
KMMLU-Pro <i>Text-only Korean exams</i>	2025	2.8K	Filter items solved by ≥ 4 of 7 small LLMs
HLE <i>Expert-level knowledge</i>	2025	2.5K	Frontier-LLM difficulty check before expert review
KMMMU (ours) <i>Multimodal Korean exams</i>	2026	3.4K	Remove items solved by four MLLMs

Table 2: Comparison of benchmark construction strategies that preserve evaluation difficulty through adversarial selection, model-aware filtering, or frontier-model difficulty checks.

3 The KMMMU Benchmark

3.1 Data Collection and Annotation

KMMMU is constructed from Korean-native official examinations and competitions. These sources include the civil service recruitment (PSAT), National Technical Qualifications (NTQ), National Competency Standards exam (NCS), and academic Olympiads (see Appendix A for details). We initially collect approximately 68k raw instances.

We process the collected exam materials into structured multimodal instances using automated extraction followed by source-grounded manual verification. Technical qualification data are collected through web crawling, while other sources are digitized using the *MinerU-2.5* OCR system (Niu et al., 2025). Five native Korean annotators review the items, and to support verification,




Figure 3: **Embedding-based distributional comparison after selection.** Question text embeddings from the original 68k corpus, the *KMMMU*, and the *KMMMU-Adv* are projected using PCA followed by 3D UMAP.

we built a custom interface to compare parsed instances with the original source materials. The verification covered question text, answer choices, LaTeX formulas, image crops, image references, and answer keys (see Appendix B for detailed procedure). Because many source materials are digitized from PDF documents, the benchmark may be less exposed to direct web-scale crawl contamination than datasets collected from public web QA repositories. We therefore expect this process to reduce contamination risk.

3.2 KMMMU Dataset Construction

Starting from approximately 68k manually verified questions, we apply multi-stage filtering to retain challenging instances. We evaluate the candidate pool with four MLLMs in sequence: Phi-3.5-Vision-Instruct (Abdin et al., 2024), InternVL-3.5-38B (Wang et al., 2025), Gemini-2.5-Flash-Lite, and Gemini-2.5-Flash (Comanici et al., 2025). At each stage, questions answered correctly by the corresponding model are removed from the candidate pool (see Appendix B.3 for retention statistics).

This filtering step follows a common practice in recent benchmark construction (Zellers et al., 2019; Wang et al., 2024; Yue et al., 2025; Hong et al., 2025; Wei et al., 2024; Phan et al., 2025). As summarized in Table 2, model-based filtering is critical for producing challenging, high-quality benchmarks that preserve evaluation headroom against frontier models. At the same time, aggressive filtering can distort the original distribution, so it is equally important to ensure that the filtered set retains broad coverage of the source corpus.

Therefore, we additionally assess whether filtering substantially changes the original distribution. Using multilingual-e5-large embeddings projected with PCA ($n = 50$) and 3D UMAP, we




Figure 4: **Discipline-wise visual modality composition of KMMMU.** Stacked bars show the number of questions for each visual modality in each discipline, with total counts shown beneath the labels. Scatter points indicate *Korean-specific* items overlaid on the corresponding discipline-modality segments, and jittered randomly.

find that both *KMMMU* and *KMMMU-Adv* largely preserve the global structure of the original 68k-sample pool, as shown in Figure 3. Dimension-wise KL divergence also remains low for *KMMMU* ($D_{KL} = 0.11\text{--}0.15$), suggesting that filtering increases difficulty without strongly distorting the original corpus structure. Appendix E provides additional density comparisons and KL analyses.

In short, we build *KMMMU*, a collection of **3,466** questions sourced entirely from native Korean materials without translation, the first *MMMU*-style benchmark of its kind for the Korean cultural sphere. Figure 2 shows representative instances across disciplines.

3.3 Taxonomy and Dataset Composition

KMMMU is designed to evaluate expert-level multimodal understanding across diverse domains. Each instance is annotated along four axes: discipline, visual modality, question format, and Korean-specificity (detailed in Section 3.4). Initial labels are assigned by Gemini-2.5-Flash (Comanici et al., 2025) and validated through human audit and consolidation.

Figure 4 presents the discipline-wise distribution of *KMMMU*. Overall, *Engineering* forms the largest portion of the dataset, followed by *Natural Science* and *CS&IT*. Meanwhile, *Diagrams* and *Text, Code & Documents* constitute the most frequent visual modalities (see Appendix D.3 for detailed examples).

3.4 Construction of KMMMU-Adv and KMMMU-Local

To further analyze model limitations, we construct *KMMMU-Adv*, a challenging subset of **627** questions answered incorrectly by all three baseline

models: Gemma-3-27B (Team et al., 2025), Qwen3-VL-235B-Thinking (Bai et al., 2025), and GPT-5-nano (OpenAI, 2025). This subset corresponds to 18% of the full *KMMMU* dataset and provides a baseline-relative stress test for instances that remain difficult. Additional composition details are provided in Appendix Figure 12.

We also provide *KMMMU-Local*, a subset of **300** Korean-specific items designed to assess whether strong multimodal performance transfers to localized institutional, cultural, and terminological knowledge. The Korean-specificity label identifies items whose solutions depend on Korean-localized knowledge, including Korean laws, standards, institutions, geography, history, culture, or official terminology beyond general world knowledge. Following the labeling pipeline described in Section 3.3, we manually audit 300 random instances and verify all items labeled as *KMMMU-Local*. As shown by the overlaid scatter points in Figure 4, *KMMMU-Local* items are concentrated in institutionally grounded disciplines such as *Business & Public* and *Law & Ethics*, where solving questions often requires knowledge of Korean administrative categories, regulations, or official standards. By contrast, many STEM and technical items, even when drawn from Korean exams, can still be solved using general domain knowledge.

4 Experimental setup

4.1 Evaluated Models

We evaluate a diverse set of multimodal models covering both open-weight and proprietary systems. The models are organized according to whether they employ explicit reasoning during inference.

The **Open-weight Non-Reasoning** group includes Gemma-3 (Team et al., 2025) (4B, 12B, 27B), Qwen3-VL (Bai et al., 2025) (2B, 4B, 8B, 32B, 30B-A3B, 235B-A22B), and Llama-4 (Meta, 2025b,a) (Scout and Maverick), along with the Korean models VARCO-VISION-2.0 (Cha et al., 2025) (1.7B, 14B) and HyperCLOVAX-SEED-Vision-3B (NAVER HyperCLOVAX, 2025).

The **Open-Weight Reasoning** group includes Qwen3-VL-Thinking (Bai et al., 2025) (30B-A3B, 32B, 235B-A22B).

Because of the cost constraints, we report the **Proprietary model** group *only on the KMMMU-Adv*: GPT-5 and GPT-5-mini (OpenAI, 2025), Claude-Opus-4.5 and Claude-Sonnet-4.5 (Anthropic, 2025), Grok-4 and Grok-4.1-Fast (xAI,

Model	Arts & Design	Business & Public	CS & IT	Engineering	General	Law & Ethics	Math & Stats	Natural Sciences	Social Sciences	Overall Acc.
Open-Weight Multilingual Non-Reasoning Models										
Gemma-3-4B-IT	29.96 _{1.58}	26.51 _{1.58}	18.83 _{2.40}	23.19 _{0.66}	19.24 _{1.67}	22.76 _{2.27}	17.03 _{1.81}	15.25 _{1.13}	19.29 _{2.32}	20.49 _{0.41}
Gemma-3-12B-IT	28.27 _{1.58}	25.85 _{2.45}	20.99 _{1.57}	22.95 _{0.53}	14.09 _{2.03}	31.09 _{1.98}	16.74 _{2.94}	18.49 _{1.29}	22.83 _{2.25}	21.59 _{0.55}
Gemma-3-27B-IT	23.63 _{6.72}	31.02 _{2.85}	25.41 _{1.91}	22.48 _{0.33}	17.62 _{0.38}	30.77 _{3.93}	23.35 _{0.95}	21.74 _{1.35}	26.64 _{3.04}	23.78 _{0.88}
Qwen3-VL-2B-IT	24.05 _{1.03}	22.88 _{1.33}	19.55 _{0.52}	13.75 _{0.65}	8.67 _{2.03}	23.08 _{2.08}	8.52 _{1.36}	9.61 _{0.60}	15.75 _{0.85}	14.24 _{0.34}
Qwen3-VL-4B-IT	25.32 _{1.03}	34.43 _{1.38}	26.75 _{1.62}	16.73 _{0.79}	13.28 _{1.38}	24.36 _{2.52}	13.22 _{0.95}	17.68 _{0.70}	21.78 _{0.67}	19.88 _{0.42}
Qwen3-VL-8B-IT	25.74 _{1.58}	37.95 _{1.23}	29.63 _{1.97}	20.86 _{1.64}	13.82 _{1.76}	33.01 _{2.40}	22.61 _{1.16}	22.63 _{0.62}	29.53 _{0.96}	24.56 _{0.77}
Qwen3-VL-30B-A3B-IT	28.27 _{8.35}	46.09 _{2.04}	31.58 _{1.43}	22.38 _{1.04}	13.55 _{2.51}	25.96 _{4.08}	30.54 _{3.00}	32.64 _{0.15}	39.37 _{0.64}	29.45 _{0.98}
Qwen3-VL-32B-IT	27.43 _{1.19}	52.15 _{1.64}	36.93 _{0.29}	28.64 _{0.75}	22.76 _{1.76}	32.69 _{2.08}	43.91 _{1.45}	36.29 _{0.66}	48.95 _{1.86}	35.65 _{0.48}
Qwen3-VL-235B-A22B-IT	25.32 _{2.73}	57.87 _{0.82}	37.45 _{1.79}	31.38 _{0.56}	24.66 _{2.68}	35.58 _{2.36}	53.01 _{1.26}	41.40 _{0.64}	53.81 _{1.45}	39.44 _{0.08}
Llama-4-Scout-17B-16E-IT	26.58 _{1.03}	51.05 _{1.09}	36.01 _{1.54}	30.81 _{1.06}	23.85 _{2.33}	33.65 _{1.36}	34.07 _{0.91}	30.86 _{0.06}	39.63 _{1.65}	33.67 _{0.33}
Llama-4-Maverick-17B-128E-IT	27.43 _{2.60}	55.89 _{2.06}	39.81 _{1.82}	34.09 _{1.04}	20.87 _{2.03}	38.46 _{0.79}	43.61 _{0.95}	39.29 _{0.46}	48.56 _{0.81}	38.95 _{0.47}
Open-Weight Korean Non-Reasoning Models										
HyperCLOVAX-SEED-Vision-3B	22.36 _{2.98}	26.73 _{1.23}	24.28 _{0.15}	17.22 _{0.21}	17.07 _{1.15}	29.81 _{0.79}	13.66 _{0.36}	13.63 _{0.34}	17.59 _{0.49}	18.14 _{0.15}
VARCO-VISION-2.0-1.7B	23.63 _{6.64}	26.40 _{2.16}	29.12 _{2.92}	25.90 _{0.57}	19.24 _{1.01}	24.68 _{2.40}	19.82 _{2.36}	19.38 _{0.66}	20.73 _{2.92}	23.59 _{0.91}
VARCO-VISION-2.0-14B	28.69 _{3.63}	34.76 _{1.75}	33.02 _{0.25}	27.96 _{0.48}	20.05 _{2.03}	33.01 _{1.98}	24.08 _{1.85}	23.52 _{0.30}	27.30 _{1.65}	27.55 _{0.37}
Open-weight Multilingual Reasoning Models										
Qwen3-VL-30B-A3B-Thinking	24.68 _{0.63}	52.31 _{0.83}	35.96 _{2.62}	27.54 _{0.77}	21.54 _{1.22}	34.13 _{2.40}	46.92 _{1.10}	38.20 _{0.12}	44.88 _{0.39}	35.47 _{0.68}
Qwen3-VL-32B-Thinking	25.74 _{3.16}	58.75 _{1.77}	41.77 _{1.24}	28.10 _{0.44}	20.05 _{2.07}	33.97 _{2.52}	49.93 _{0.91}	39.86 _{0.45}	51.18 _{1.47}	37.80 _{0.49}
Qwen3-VL-235B-A22B-Thinking	22.36 _{4.30}	62.38 _{0.47}	40.33 _{1.27}	32.79 _{0.98}	30.08 _{2.89}	32.05 _{2.76}	56.09 _{0.55}	45.30 _{0.80}	57.87 _{0.85}	42.05 _{0.48}

Table 3: **Accuracy (%) on the KMMU by disciplines.** Overall accuracy is averaged across disciplines. Mean accuracy is reported in percentage, with standard deviation shown as a subscript. Best in each model group is shown in **bold**.

2026), Gemini-3-Pro (Google Cloud, 2025b), Gemini-3-Flash (Google Cloud, 2025a), and Mistral-Large-3-675B-IT (Mistral AI, 2025).

4.2 Evaluation Protocols

All evaluations are conducted zero-shot using a shared prompt template, with provider-recommended decoding values throughout (see Appendix F for details).

For scalable scoring across both multiple-choice and free-form questions, we extract and normalize the final answer from each response, and Gemini-2.5-Flash with temperature 0 then decides whether it matches the gold answer. We call this procedure LLM-based answer scoring, since it checks answer-level correctness rather than overall response quality. Section 4.3 validates it against human correctness labels. Each model is evaluated over three trials, and we report mean accuracy and standard deviation.

4.3 Is LLM-based Answer Scoring Reliable?

To validate the LLM scorer’s evaluation capability, we conduct a human alignment study on 600 examples, sampled from six model runs (100 outputs each) and balanced across question formats. Three annotators assign binary labels for model-response correctness and mark whether each response is complete or not (e.g., terminated mid-reasoning or degeneration).

As shown in Figure 5, inter-annotator agreement is consistently high, ranging from 0.91 to 0.99, which indicates that correctness labels are generally well defined. LLM scorer also aligns well with




Figure 5: **Human alignment of the LLM scorer.** Each x-axis label indicates the source model of the sampled responses. We independently sample and re-annotate 100 outputs per model. H-H Agr. denotes inter-annotator agreement between human annotators, and LLM-H Agr. denotes agreement between human annotations and the LLM scorer. The LLM scorer is Gemini-2.5-Flash, the same judge used throughout our evaluation pipeline.

human annotations, achieving agreement between 0.88 and 0.98 across models. Although alignment varies by model, lower LLM–human agreement tends to coincide with lower human–human agreement, suggesting that these cases are better explained by outputs that are difficult for both humans and the LLM judge to interpret than by bias toward a particular model family. Some annotation noise therefore remains inevitable, but we reduce its impact by evaluating each model over three independent runs and reporting mean performance and standard deviation. For more details on judging validation analysis, see Appendix I.2.

5 Results

5.1 Main Results

KMMU remains challenging, even for strong multimodal models. As shown in Table 3, on the

Model	Arts & Design	Business & Public	CS & IT	Engineering	General	Law & Ethics	Math & Stats	Natural Sciences	Social Sciences	Overall Acc.
Claude-Opus-4.5	28.00 _{6.53}	32.38 _{3.56}	25.73 _{2.19}	23.12 _{2.31}	21.93 _{3.28}	24.07 _{6.93}	19.54 _{3.25}	22.94 _{1.22}	35.19 _{4.72}	24.51 _{0.54}
Claude-Sonnet-4.5	16.00 _{0.00}	24.76 _{5.39}	14.62 _{4.60}	18.58 _{2.04}	12.28 _{5.41}	37.04 _{5.24}	17.24 _{5.63}	18.40 _{0.81}	19.44 _{2.27}	18.55 _{1.31}
Gemini-3-Flash	48.00 _{0.00}	42.86 _{7.00}	48.54 _{1.65}	46.52 _{1.32}	21.93 _{3.28}	46.30 _{9.44}	37.93 _{0.00}	46.10 _{1.84}	56.48 _{3.46}	45.14 _{0.98}
Gemini-3-Pro	50.67 _{4.99}	61.90 _{1.35}	57.31 _{0.83}	54.04 _{2.28}	27.19 _{1.24}	46.30 _{5.24}	48.28 _{2.82}	52.60 _{3.23}	58.33 _{2.27}	52.42 _{0.94}
Mistral-Large-3-675B-IT	20.00 _{3.27}	16.19 _{1.35}	15.79 _{3.79}	15.32 _{2.28}	4.39 _{1.24}	16.67 _{4.54}	6.90 _{2.82}	16.23 _{0.53}	21.30 _{3.46}	15.15 _{1.17}
GPT-5	20.00 _{3.27}	26.67 _{3.56}	26.90 _{2.19}	30.35 _{1.40}	24.56 _{3.28}	18.52 _{2.62}	34.48 _{2.82}	33.55 _{2.51}	45.37 _{3.46}	30.57 _{0.40}
GPT-5-Mini	14.67 _{6.80}	25.71 _{2.33}	14.04 _{2.48}	18.58 _{1.40}	17.54 _{4.96}	24.07 _{11.42}	29.89 _{3.25}	24.89 _{1.70}	31.48 _{1.31}	21.32 _{1.04}
Grok-4	28.00 _{0.00}	24.29 _{4.29}	27.19 _{2.63}	26.38 _{1.28}	18.42 _{2.63}	16.67 _{0.00}	34.48 _{0.00}	22.73 _{4.55}	31.94 _{1.39}	25.44 _{0.88}
Grok-4.1-Fast	26.67 _{3.77}	20.95 _{1.35}	17.54 _{1.43}	20.00 _{0.35}	9.65 _{2.48}	7.41 _{2.62}	12.64 _{3.25}	22.73 _{1.59}	25.93 _{3.46}	19.78 _{0.13}

Table 4: **Accuracy (%) on the KMMU-Adv by disciplines.** Overall accuracy is averaged across disciplines. Mean accuracy is reported in percentage, with standard deviation shown as a subscript. The best result is shown in **bold**.

full set, the strongest open-weight model reaches 42.05% overall accuracy, while the best Korean-focused open-weight model, VARCO-VISION-2.0-14B, reaches 27.55%. This gap suggests that Korean-language specialization alone is insufficient for expert-level multimodal reasoning, and that strong performance still depends heavily on overall model capacity.

Model scale consistently improves performance, but explicit test-time reasoning yields smaller and less consistent gains. Within the Qwen3-VL family, larger models generally outperform smaller ones, with especially large gains in disciplines such as *Math & Stats* and *Social Sciences*. This suggests that train-time scaling, including greater model capacity and broader knowledge acquisition, remains important for KMMU. By contrast, reasoning variants show only modest or uneven improvements over their non-reasoning counterparts, indicating that simply increasing test-time deliberation is not sufficient when errors stem from missing knowledge, weak grounding, or incorrect multimodal interpretation.

Performance also varies substantially across disciplines. **While stronger models improve markedly in some areas, *General* and *Arts & Design* remain persistent bottlenecks, with only limited gains even at larger scales.** This pattern suggests that KMMU requires more than surface-level recognition, demanding multimodal grounding, contextual interpretation, and discipline-specific knowledge.

A similar pattern appears on *KMMU-Adv* (Table 4). Gemini-3-Pro achieves the best overall accuracy at 52.42%, followed by Gemini-3-Flash at 45.14%, while the remaining models perform substantially worse. Discipline-level variation also remains strong: *General* is again one of the weakest areas, with Gemini-3-Pro reaching only 27.19%, far below its scores in other disciplines. Taken together, these results show that **KMMU-Adv**

sharply exposes residual weaknesses in reasoning, multimodal understanding, and discipline-specific interpretation.

Model	Non Korean-spec	Korean -spec	Raw gap	Controlled gap
Open-Weight Multilingual Models				
Qwen3-VL-8B-IT	24.33	27.11	+2.78	+4.01
Qwen3-VL-30B-A3B-IT	30.16	22.00	-8.16	-5.97
Qwen3-VL-30B-A3B-Thinking	36.32	26.67	-9.65	-7.99
Qwen3-VL-32B-IT	36.56	26.22	-10.33	-8.09
Qwen3-VL-32B-Thinking	38.70	28.33	-10.37	-10.00
Qwen3-VL-235B-A22B-IT	40.62	27.11	-13.51	-13.43
Qwen3-VL-235B-A22B-Thinking	43.18	30.22	-12.96	-11.17
Llama-4-Maverick-17B-IT	39.96	28.44	-11.51	-12.35
Llama-4-Scout-17B-IT	34.27	27.44	-6.83	-7.97
Open-Weight Korean Models				
HyperCLOVAX-SEED-Vision-3B	17.57	24.22	+6.66	+5.15
VARCO-VISION-2.0-14B	27.70	26.11	-1.59	-4.18

Table 5: **Performance on KMMU-Local subset.** Raw gap is the accuracy difference between Korean-specific and non-Korean-specific questions; controlled gap is the discipline-controlled difference. Negative values indicate worse performance on Korean-specific questions. The largest positive gap is green shaded, and the largest negative gap is red shaded.

5.2 What makes KMMU-Local hard?

We investigate whether the difficulty of *KMMU-Local* stems from Korean linguistic processing or from the localized knowledge the items embed. Because *KMMU-Local* items are unevenly distributed across disciplines, particularly in *Business & Public* and *Law & Ethics*, we report both raw gaps and discipline-controlled gaps in Table 5.

Multilingual scale does not close the gap. Most multilingual open-weight models show lower accuracy on *KMMU-Local* even after controlling for discipline composition. The gap is especially pronounced among the strongest models: Qwen3-VL-235B-A22B-IT, Qwen3-VL-235B-A22B-Thinking, and Llama-4-Maverick-17B-IT show controlled gaps of -13.43 , -11.17 , and -12.35 percentage points, respectively. For Qwen3-VL-235B-A22B-IT, the 95% CI of this gap is $[-19.27, -7.44]$, and the deficit is concentrated in *Business & Public* and *Law & Ethics* (Appendix H.2).

Model	Ko Acc.	En Acc.	Δ	Ko/En-only	p
GPT-5	42.5	29.0	+13.5	33/6	1.43×10^{-5}
Gemini-3-Flash	58.0	40.0	+18.0	38/2	1.49×10^{-9}

Table 6: **Cross-lingual ablation on 200 Korean-specific paired items.** We translate the textual question and answer options into English. Ko/En-only denotes paired discordant cases answered correctly only in Korean or only in English. p is from an exact McNemar test.

Korean-focused training is not sufficient either. HyperCLOVAX-SEED-Vision-3B shows a positive controlled gap of +5.15 points, but VARCO-VISION-2.0-14B still shows a negative gap of -4.18 points. Being a Korean-focused model thus helps in some cases but does not reliably resolve the difficulty of Korean-specific institutional and terminological knowledge.

The deficit is localized knowledge, not Korean-language processing. To isolate the role of the input language, we conduct a controlled cross-lingual ablation on 200 Korean-specific paired items, translating the textual question and answer options into English while keeping the visual input, answer, and question type fixed (Table 6). Both GPT-5 and Gemini-3-Flash perform *worse* on the English translations: GPT-5 drops from 42.5% to 29.0% and Gemini-3-Flash from 58.0% to 40.0%, with paired discordance statistically significant in both cases and substantially more Korean-only correct cases than English-only correct cases. Translating the question text strips localized terminology, institutional references, and the grounding between Korean context and visual evidence; the difficulty therefore reflects the localized knowledge that current models do not yet encode, regardless of language coverage or Korean focus. See Appendix H.1 for the full paired-outcome breakdown.

6 Error Analysis

In this section, we conduct a targeted manual error analysis over selected model generations. The full inspection protocol is provided in Appendix J.

Across inspected cases, errors split into three patterns: incomplete answer composition (Section 6.1), knowledge recall and category matching failures (Section 6.2), and weak rule induction on symbolic problems (Section 6.3).

6.1 Post-perceptual Effects of Reasoning

Different failure patterns across quantitative domains. Although reasoning improves overall performance in some quantitative disciplines, espe-

cially *Math & Stats* (43.91→49.93), its remaining failures follow different patterns across domains. To examine this, we sampled 25 items each from *Math & Stats*, *Engineering*, and *Natural Sciences* among questions answered correctly by Qwen3-VL-32B-IT but incorrectly by Qwen3-VL-32B-Thinking, as these disciplines show contrasting reasoning effects. The clearest pattern appears in *Math & Stats*. In our inspected sample, 72% (18/25) of these reversals were not caused by obviously worse intermediate reasoning, but by *answer finalization failure*. The thinking model often developed a partially correct or plausible solution path, but stopped before producing a fully resolved final answer. By contrast, reversals in *Engineering* and *Natural Sciences* more often reflected incorrect problem framing than incomplete finalization. In these cases, the thinking model sometimes appears to map partial visual or textual cues onto a familiar *device type*, *curve pattern*, *control category*, or *physical scenario* too early, and then elaborate that interpretation into a coherent but incorrect solution (see Appendix J.2 for detailed examples).

Reasoning gains in answer composition tasks.

Among cases where Qwen3-VL-32B-Thinking succeeds and Qwen3-VL-32B-IT fails, the clearest gains appear on questions that require *answer composition*. This is especially visible for open-ended questions that ask for multiple requested outputs and for multiple-choice questions with multiple correct answers. In such cases, the *Thinking* variant does a notably better job in formatting their responses for questions requiring multiple outputs, while the *Instruct* variant often misses to do so, even after solving correctly. The non-reasoning model is more likely to provide a subset of the requested components whereas the reasoning model is more likely to preserve the required answer structure and return all necessary components.

This tendency is also reflected in the aggregate pattern across question types (Figure 7). Reasoning gains are large for *multiple-answer* and *open-form* questions. These findings suggest that explicit reasoning helps most with constraint tracking, structured decomposition, and complete answer assembly, so its benefits appear more in *output completeness* than in knowledge recovery, helping explain why the gains in the main results are uneven.

Across visual modalities, the two variants show broadly similar performance, and inspected differences rarely come from one variant clearly reading




Figure 6: Example of a Korean-specific regulatory category mismatch. Qwen3-VL-235B-A22B-IT reads the table correctly, but maps *small vehicle* to the wrong category and applies the wrong standard. This is a failure of institutional knowledge recall and lexical category matching, not OCR.




Figure 7: Reasoning gain by discipline and question type in Qwen3-VL-32B (IT vs. Thinking).

or missing the image while the other does not (Appendix J.2, Figure 14).

6.2 Knowledge recall and category matching failure in Korean-specific questions

Interestingly, although Qwen3-VL-235B-A22B-IT substantially outperforms HyperCLOVAX-SEED-Vision-3B on the full benchmark (39.44 vs. 18.14), the gap narrows on Korean-specific items. In a 300-item comparison, the two models achieve relatively similar performance, scoring 83/300 and 72/300, respectively (Figure 5). This reduced separation suggests that general reasoning ability provides limited advantage on Korean-specific items, many of which depend on regulation-specific knowledge or fine-grained administrative distinctions.

Figure 6 is a representative failure case concerning regulation-specific terminology. Korean law distinguishes *소형차* (small vehicle), defined by an engine displacement of 1000 to 1600 cc, from *승용차* (passenger vehicle), which refers to cars with up to 15 seats. However, Qwen3-VL-235B-A22B-IT appears to collapse both terms into the same English expression during intermediate reasoning,

producing an incorrect answer. Similar patterns are reported by Son et al. (2025b), who show that multilingual models often translate inputs into a preferred language, introducing noise and reducing task performance. Thus, the main bottleneck is not visual perception but the mapping between Korean legal terms and their institution-specific categories.

6.3 Disciplinary Bottlenecks

Among the subject groups, *Arts & Design* and *General* remain consistently difficult across models, suggesting bottlenecks that are not readily resolved by either scale or reasoning mode. The inspected failures point to different sources of difficulty in the two categories.

In *General*, many failures occur in items sourced from the KLO (Korea Linguistic Olympiad) exam. These questions often combine several skills in a single problem, including linguistic puzzle solving, Korean orthography, semantic change, dictionary ordering, and symbol-to-sound or symbol-to-word rule induction from a small set of examples. In many failures, models identify some local patterns but fail to recover the full rule system, leading to plausible but unsupported answers. This points to a limitation in few-shot symbolic induction.

In *Arts & Design*, many failures instead involve exact technical labels for specialized visual conventions. As shown in Appendix J.3, some models recognize the relevant visual cues but choose a nearby standardized term rather than the correct one. This is especially common when distinguishing related expert categories such as *local* versus *partial projection*, *cutting line* versus *revolved-section line*, or similar notation symbols. These cases suggest

that the bottleneck is not basic visual recognition, but precise mapping from visual conventions to domain-specific terms. Taken together, these results suggest two complementary directions for improvement. For *Arts & Design*, models may need broader exposure to specialized materials, and for *General*, models may need training examples that require rule induction, linguistic reasoning, and knowledge retrieval within the single problem.

7 Conclusion

We introduce **KMMMU**, a native Korean benchmark for expert-level multimodal understanding in culturally and institutionally grounded settings. Across 3,466 carefully verified questions, KMMMU shows that current MLLMs remain far from robust on Korean real-world assessment materials. Our findings suggest that key failures arise less from insufficient reasoning depth than from weak convention-to-label mapping, few-shot symbolic induction, localized knowledge recall, and familiarity with domain-specific standards and terminology. These bottlenecks help explain the limited and uneven gains from reasoning, the persistent difficulty of Korean-specific content, and strong disciplinary variation in performance. We hope KMMMU to serve as a rigorous benchmark for evaluating expert-level Korean multimodal understanding and a practical testbed for developing more culturally grounded and institutionally aware MLLMs.

Limitations

Scope and factor separation. KMMMU targets expert-level multimodal understanding in native Korean exam-style settings. Its results may not generalize to everyday perception, interactive image use, or non-exam multimodal tasks. KMMMU-Local is a diagnostic subset for localized institutional, cultural, legal, and terminological knowledge, but Korean language processing, domain difficulty, visual complexity, and local knowledge often co-occur. Thus, our controlled analyses help characterize these factors but do not fully separate their independent effects.

Filtering and subset bias. Model-based filtering is used as a difficulty-selection signal, while item validity is supported separately by Korean examination sources, answer keys, OCR checks, image-text verification, and human review. Because the final set retains items missed by multiple models, KMMMU is not a random sample of Korean exam ques-

tions. KMMMU-Adv is also model-dependent and should be interpreted as a baseline-relative stress test, not a model-neutral measure of intrinsic difficulty.

Annotation, evaluation, and contamination.

KMMMU uses LLM-assisted metadata labels and LLM-based answer scoring for mixed-format answers. Although we audit labels, manually verify Korean-specific positives, and validate judge outputs against human annotations, residual taxonomy and grading noise may remain. Our source selection and prefix-completion analysis reduce contamination concerns, but they cannot rule out exposure to source documents or near-duplicates because model training data are not fully transparent. Finally, results reflect the evaluated models and zero-shot protocol; newer models, prompts, decoding settings, retrieval, or tool use may change scores and rankings.

Acknowledgments

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Ministry of Science and ICT (RS-2026-25477127).

References

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, and 1 others. 2024. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, and 1 others. 2022. Flamingo: a visual language model for few-shot learning. In *NeurIPS*.
- Anthropic. 2025. [Introducing Claude Opus 4.5](#). Anthropic Newsroom. Published: 2025-11-24. Accessed: 2026-01-05.
- 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. 2025. [Qwen3-vl technical report](#). *Preprint*, arXiv:2511.21631.
- Young-rok Cha, Jeongho Ju, SunYoung Park, Jong-Hyeon Lee, Younghyun Yu, and Youngjune Kim. 2025. Varco-vision-2.0 technical report. *arXiv preprint arXiv:2509.10105*.
- Byungjin Choi, Seongsu Bae, Sunjun Kweon, and Edward Choi. 2026a. Kormedmcqa-v: A multimodal benchmark for evaluating vision-language models

- on the korean medical licensing examination. *arXiv preprint arXiv:2602.13650*.
- Dasol Choi, Guijin Son, Hanwool Lee, Minhyuk Kim, Hyunwoo Ko, Teabin Lim, Ahn Eungyeol, Jungwhan Kim, Seunghyeok Hong, and Youngsook Song. 2026b. What users leave unsaid: Under-specified queries limit vision-language models. *arXiv preprint arXiv:2601.06165*.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Google Cloud. 2025a. [Gemini 3 Flash \(Preview\) | Generative AI on Vertex AI](#). Google Cloud Documentation. Release date: 2025-12-17. Accessed: 2026-01-05.
- Google Cloud. 2025b. [Gemini 3 Pro \(Preview\) | Generative AI on Vertex AI](#). Google Cloud Documentation. Release date: 2025-11-18. Accessed: 2026-01-05.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Seokhee Hong, Sunkyoung Kim, Guijin Son, Soyeon Kim, Yeonjung Hong, and Jinsik Lee. 2025. From kmmlu-redux to kmmlu-pro: A professional korean benchmark suite for llm evaluation. *arXiv preprint arXiv:2507.08924*.
- Taebaek Hwang, Minseo Kim, Gisang Lee, Seonuk Kim, and Hyunjun Eun. 2025. Kreta: A benchmark for korean reading and reasoning in text-rich vqa attuned to diverse visual contexts. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 33409–33420.
- Jeongho Ju, Daeyoung Kim, SunYoung Park, and Youngjune Kim. 2024. [Varco-vision: Expanding frontiers in korean vision-language models](#). *Preprint*, arXiv:2411.19103.
- Yoonshik Kim and Jaeyoon Jung. 2025. Koffvqa: An objectively evaluated free-form vqa benchmark for large vision-language models in the korean language. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 575–585.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew Peters, Ashish Sabharwal, and Yejin Choi. 2020. Adversarial filters of dataset biases. In *International conference on machine learning*, pages 1078–1088. Pmlr.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *NeurIPS*.
- Meta. 2025a. [meta-llama/Llama-4-Maverick-17B-128E-Instruct](#). Hugging Face model card and weights. Model release date: 2025-04-05. Accessed: 2026-01-05.
- Meta. 2025b. [meta-llama/Llama-4-Scout-17B-16E](#). Hugging Face model card and weights. Model release date: 2025-04-05. Accessed: 2026-01-05.
- Mistral AI. 2025. [Mistral Large 3 \(v25.12\) | Mistral Docs](#). Mistral Documentation. Dated: 2025-12-02. Accessed: 2026-01-05.
- NAVER HyperCLOVAX. 2025. [HyperCLOVAX-SEED-Vision-Instruct-3B](#). Hugging Face model card and weights. Model release date: 2025-04-24 (as stated in repository license). Accessed: 2026-01-05.
- Junbo Niu, Zheng Liu, Zhuangcheng Gu, Bin Wang, Linke Ouyang, Zhiyuan Zhao, Tao Chu, Tianyao He, Fan Wu, Qintong Zhang, Zhenjiang Jin, Guang Liang, Rui Zhang, Wenzheng Zhang, Yuan Qu, Zhifei Ren, Yuefeng Sun, Yuanhong Zheng, Dongsheng Ma, and 42 others. 2025. [Mineru2.5: A decoupled vision-language model for efficient high-resolution document parsing](#). *Preprint*, arXiv:2509.22186.
- Shota Onohara, Atsuyuki Miyai, Yuki Imajuku, Kazuki Egashira, Jeonghun Baek, Xiang Yue, Graham Neubig, and Kiyoharu Aizawa. 2025. Jmmmu: A japanese massive multi-discipline multimodal understanding benchmark for culture-aware evaluation. 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 932–950.
- OpenAI. 2024. Multilingual massive multitask language understanding (mmmlu). <https://huggingface.co/datasets/openai/MMMLU>.
- OpenAI. 2025. [Introducing GPT-5](#). OpenAI. Published: 2025-08-07. Accessed: 2026-01-05.
- ChaeHun Park, Yujin Baek, Jaeseok Kim, Yu-Jung Heo, Du-Seong Chang, and Jaegul Choo. 2025. Evaluating visual and cultural interpretation: The k-viscuit benchmark with human-vlm collaboration. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21960–21974.
- Sanghee Park and Geewook Kim. 2025. Evaluating multimodal generative ai with korean educational standards. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 671–688.

- Long Phan, Alice Gatti, Ziwen Han, Nathaniel Li, Josephina Hu, Hugh Zhang, Chen Bo Calvin Zhang, Mohamed Shaaban, John Ling, Sean Shi, and 1 others. 2025. Humanity’s last exam. *arXiv preprint arXiv:2501.14249*.
- Guijin Son, Hanwool Lee, Sungdong Kim, Seungone Kim, Niklas Muennighoff, Taekyoon Choi, Cheonbok Park, Kang Min Yoo, and Stella Biderman. 2025a. Kmmlu: Measuring massive multitask language understanding in korean. 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 4076–4104.
- Guijin Son, Donghun Yang, Hitesh Laxmichand Patel, Amit Agarwal, Hyunwoo Ko, Chanuk Lim, Srikant Panda, Minhyuk Kim, Nikunj Drolia, Dasol Choi, and 1 others. 2025b. Pushing on multilingual reasoning models with language-mixed chain-of-thought. *arXiv preprint arXiv:2510.04230*.
- Gemini Team and Google. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivi re, and 1 others. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, and 1 others. 2025. Internvl3.5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, and 1 others. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290.
- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. Measuring short-form factuality in large language models. *arXiv preprint arXiv:2411.04368*.
- xAI. 2026. [Models and Pricing \(xAI API Documentation\)](#). xAI Developer Docs. Accessed: 2026-01-05.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567.
- Xiang Yue, Tianyu Zheng, Yuansheng Ni, Yubo Wang, Kai Zhang, Shengbang Tong, Yuxuan Sun, Botao Yu, Ge Zhang, Huan Sun, and 1 others. 2025. Mmmu-pro: A more robust multi-discipline multimodal understanding benchmark. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15134–15186.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. Swag: A large-scale adversarial dataset for grounded commonsense inference. In *Proceedings of the 2018 conference on empirical methods in natural language processing*, pages 93–104.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 4791–4800.
- Ge Zhang, Xinrun Du, Bei Chen, Yiming Liang, Tongxu Luo, Tianyu Zheng, Kang Zhu, Yuyang Cheng, Chunpu Xu, Shuyue Guo, and 1 others. 2024. Cmmu: A chinese massive multi-discipline multimodal understanding benchmark. *arXiv preprint arXiv:2401.11944*.
- Wenxuan Zhang, Mahani Aljunied, Chang Gao, Yew Ken Chia, and Lidong Bing. 2023. M3exam: A multilingual, multimodal, multilevel benchmark for examining large language models. *Advances in Neural Information Processing Systems*, 36:5484–5505.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2024. Agieval: A human-centric benchmark for evaluating foundation models. In *Findings of the association for computational linguistics: NAACL 2024*, pages 2299–2314.

A Data Sources and Collection Scope

KMMMU is collected from four high stakes sources in South Korea. We summarize the collection scope for each source below.

A.1 PSAT

We annotate ten years of past examinations from civil service recruitment tracks. The PSAT includes Language Logic, Data Interpretation, and Situational Judgment sections that assess logical reasoning and information integration.

A.2 National Technical Qualifications

We collect fifteen years of questions from 252 distinct certification exams, including Information Processing Engineer, Electric Engineer, and Fire Safety Manager. These exams cover a wide range of technical domains across industrial and engineering fields.

A.3 Olympiads

To incorporate academically challenging reasoning problems, we gather ten years of Olympiad questions spanning middle school, high school, and university levels. The collected problems focus primarily on mathematics and science.

A.4 NCS

We include three years of National Competency Standards examinations covering all ten competency areas, such as Communication, Numeracy, and Problem Solving. These exams are used in recruitment for public sector organizations.

B Annotation and Quality Control Details

The construction of KMMMU uses a rigorous pipeline that combines automated processing with human verification to ensure high data fidelity.

B.1 Human Verification Interface

We utilized a custom built annotation tool to verify and correct the output of the OCR pipeline. Raw data digitized by MinerU-2.5 (Niu et al., 2025) often contained artifacts and formula errors. Figure 8 shows the interface where five Korean annotators reviewed the parsed content against the original PDF source. Annotators were instructed to

- Correct LaTeX formatting for mathematical formulas

- Verify that image references in the text matched the cropped images
- Discard questions where essential visual information was illegible or missing.

All five annotators are native Korean speakers with at least a bachelor’s degree and prior experience in annotation or dataset curation. They are also familiar with AI-related workflows, which helped them reliably identify OCR artifacts, formula corruption, and image–text mismatches during verification.

B.2 Automatic Labeling and Taxonomy Consolidation

We annotate several auxiliary attributes to support analysis and stratified reporting, including discipline, visual modality type, question format and a Korean-specific flag. All taxonomy labels are assigned using Gemini-2.5-Flash. For each labeling job, the model is given the question text and its associated image, and outputs the most appropriate label.

We use an open labeling step that does not constrain predictions to a fixed label set. This reduces forced assignments when an instance does not cleanly match a predefined taxonomy. All label types are generated independently.

Manual audit and consolidation We conduct a manual audit by randomly sampling around 300 instances and reviewing the assigned labels. Based on the audited outputs, we consolidate the discipline taxonomy through human curation into 45 sub-discipline categories and 9 macro discipline categories.

Verification of Korean-specific cases Because false positives can inflate localization analyses, we manually verify all instances labeled as Korean-specific. We confirm that each positive case requires Korean-specific knowledge or context rather than general world knowledge expressed in Korean.

B.3 Adversarial Filtering Protocol

To construct a challenging diagnostic benchmark, we apply a multi-stage adversarial filtering pipeline (Zellers et al., 2018; Le Bras et al., 2020) as a difficulty-oriented selection mechanism. Model-based filtering is used only to select challenging items; item validity is established separately through source-grounded answer keys and human verification. Starting from a cleaned and




Figure 8: Annotation tool interface used for OCR verification. The tool displays the original PDF page on the left and the parsed text and images on the right, allowing annotators to correct OCR errors and validate image cropping in real time.

manually verified pool of 68,214 questions, we apply the following procedure.

1. **Data cleaning and de-duplication.** We remove samples with invalid image links and de-duplicate near-duplicate questions across exam years using image and text similarity checks.
2. **Model-based adversarial filtering.** We sequentially filter the remaining candidate pool using four multimodal models: Phi-3.5-Vision-Instruct (Abdin et al., 2024), InternVL-3.5-38B (Wang et al., 2025), Gemini-2.5-Flash-Lite, and Gemini-2.5-Flash (Comanici et al., 2025). Each model is evaluated in a zero-shot setting, and questions answered correctly at each stage are removed from the candidate pool.
3. **Final retention.** Questions that remain unsolved after all four filtering stages are retained in the final benchmark.

Table 7 reports the number of items retained after each filtering stage. The resulting KMMM benchmark contains 3,466 curated questions.

C Korean-Specific Context

To provide a concrete illustration of KMMM, Figure 9 presents a *Korean-Specific* instance from

Stage	Filtering model	Input	Removed	Remaining
Initial	—	—	—	68,214
1	Phi-3.5-Vision-Instruct	68,214	18,309	49,905
2	InternVL-3.5-38B	49,905	8,390	41,515
3	Gemini-2.5-Flash-Lite	41,515	18,368	23,147
4	Gemini-2.5-Flash	23,147	19,681	3,466

Table 7: **Stage-wise retention during adversarial filtering.** At each stage, questions answered correctly by the corresponding filtering model are removed from the candidate pool.

the benchmark. Unlike standard multimodal benchmarks, which often emphasize culturally invariant knowledge such as Physics or Mathematics, KMMM includes a dedicated subset of questions that require localized knowledge grounded in Korean institutional and legal contexts. In this example, the input consists of an image containing regulation text and a corresponding question, and the model must interpret the visual text referring to the “extraction area slope criteria” in the specific context of South Korea’s *Mountainous Districts Management Act* to identify the correct legal standard (Option 3). This example shows that solving such questions requires not only optical character recognition, but also grounded knowledge of Korean administrative law.

D Detailed Dataset Statistics

In this section, we provide a granular breakdown of the dataset composition. Beyond the overview (Ta-

Korean Specific Question	
[Original Image] 채취지역의 평균 경사도는 (①) 미하이며마 하 고, 채취등을 완료한 후 절개사면의 기울기(비탈면 의 높이에 대한 수평거리의 비율을 말한다)는 건축 용 석재인 경우에는 (②)기준에 적합할 것. 다만, 채취등을 함으로써 절개사면 없이 평탄지로 될 수 있는 경우에는 그러하지 아니하다.	[Original Question] 문제: 산지관리법령산 대통령령이 정하는 토석채취허가기준 중 산지의 경사도에 관한 기준이다. ()안에 알맞은 것은? 선택지: (1) ① 45도, ② 1: 0.5 이하 (2) ① 45도, ② 1: 0.4 이하 (3) ① 35도, ② 1: 0.5 이하 (4) ① 35도, ② 1: 0.4 이하
[Image Context Translation] Image Context: The average slope of the extraction area must be ① or less. Regarding the slope of the cut surface after extraction... it must comply with the criteria of ②	[English Translation] Question: These are the criteria for soil and stone extraction permits determined by the Presidential Decree under the <i>Mountainous Districts Management Act</i>. What fits in the parentheses ()? Options: (1) ① 45° , ② 1:0.5 or less (2) ① 45° , ② 1:0.4 or less (3) ① 35° , ② 1:0.5 or less (4) ① 35° , ② 1:0.4 or less

Figure 9: **Data Card for a Korean-Specific Question.** The figure aggregates the raw inputs and their translations. **[Original Image]** The original visual input containing a text-rich regulation box. **[Original Question]** The original question text in Korean. **[Translation]** English translations for both the visual context and the question. Correctly answering this question requires retrieving specific legal provisions regarding slope limits for soil extraction permits in South Korea, demonstrating the benchmark’s focus on localized expert knowledge.

ble 8), we report (i) the distribution of fine-grained discipline categories (Table 9), (ii) the question format distribution (Table 10), and (iii) the examples of visual modalities (sub-visual modalities, Figure 10)

Statistic	Count
Total Questions	3,466
KMMMU-Adv	627
Discipline Categories	9
Sub-discipline Categories	45
Visual Modality Types	9
Question Types	5
Questions with in-image texts	2,383 (68.75%)
Korean-specific questions	300 (8.65%)

Table 8: **Dataset distribution outline.** We report counts and percentages for key attributes such as in-image text and Korean-specific content.

D.1 Discipline Category Distribution

Table 9 details the frequency of questions across 45 fine-grained Discipline categories. The distribution reflects the emphasis on STEM (Science, Technology, Engineering, and Mathematics) fields, with *Physics*, *Civil Engineering*, and *Mechanical*

Engineering constituting the largest portions. This heavy tail in engineering disciplines ensures that KMMMU serves as a robust benchmark for technical domain expertise.

D.2 Question Format Distribution

Because KMMMU contains both multiple-choice and free-form items, the answer format affects evaluation difficulty and failure modes. Table 10 reports the distribution of question formats in the benchmark.

D.3 Visual Modality Taxonomy

KMMMU includes a wide range of fine-grained visual types, including circuit, mechanical, and structural diagrams, document-style text images, tables, mathematical figures, charts, maps, symbols, and photographs. For analysis, we consolidate these fine-grained types into 9 macro-level visual modality categories. Technical diagrams constitute a particularly large portion of the dataset, reflecting KMMMU’s emphasis on professional and schematic visual reasoning. Figure 10 presents representative examples of these fine-grained visual types before

Subject	Cnt	Subject	Cnt	Subject	Cnt
Physics	474	General Knowledge & Interdisciplinary	67	Economics	33
Civil & Structural Engineering	332	Industrial & Systems Engineering	64	Earth & Geological Sciences	33
Mechanical Engineering	223	Architecture & Urban Studies	57	Library, Archival & Information Science	30
Electrical & Electronics Engineering	221	Human Resources & Organizational Studies	56	Transportation & Logistics	29
Computer Science	164	Linguistics & Language Studies	54	Agriculture & Life Sciences	27
Statistics & Probability	149	Biology	43	Software & Programming	27
Mathematics	121	Safety, Risk & Reliability Engineering	42	Astronomy & Space Science	26
Business Administration & Management	104	Data Science & Analytics	41	Psychology	22
Environmental Science & Engineering	99	Finance & Accounting	40	Cognitive & Behavioral Sciences	20
Geography & Spatial Studies	95	Sociology & Social Sciences	36	Biomedical & Health Sciences	20
Manufacturing & Production Engineering	90	Chemical Engineering	76	Communication & Media Studies	17
Public Administration & Policy	89	Chemistry	75	Education & Pedagogy	16
Design & Visual Arts	88	Law & Legal Studies	74	Marketing & Consumer Studies	15
Chemical Engineering	76	Materials Science & Metallurgy	69	Ethics & Philosophy	12
Chemistry	75	Information Technology & Systems	69	Artificial Intelligence & Machine Learning	12

Table 9: Distribution of sub-discipline categories in KMMMU.

Question Type	Count	%
Multiple Choice (Single Answer)	2,831	81.68
Multiple Choice (Multiple Answers)	346	9.98
Open Format (Numerical / Calculation)	207	5.98
Open Format (Text)	67	1.93
Essay (Descriptive)	15	0.43
Total	3466	100

Table 10: Question type distribution.

consolidation.

D.4 Question Type Taxonomy

Table 11 summarizes the distribution of answer formats within each macro subject. This table clarifies which subjects are dominated by multiple choice items versus numerical or descriptive responses. It also provides context for interpreting task-wise performance, since answer format affects both evaluation difficulty and failure modes.

E Additional Distributional Analysis

Figure 11 provides per-dimension density comparisons for the original 68k corpus, the KMMMU *Full set*, and the KMMMU-Adv in the 3D UMAP space. Across all three dimensions, the filtered subsets broadly preserve the major density peaks and overall multimodal structure of the original distribution. The KMMMU remains especially close to the original corpus, while the KMMMU-Adv shows a somewhat larger shift in parts of the latent space.

To quantify these differences, we compute the Kullback–Leibler (KL) divergence between the original distribution and each filtered subset along each UMAP dimension. For the *Full set*, the divergence remains low across all three dimensions ($D_{KL} = 0.1184, 0.1459, \text{ and } 0.1437$ for Dimensions 1–3, respectively). The KMMMU-Adv shows similarly low divergence on Dimensions 1 and 2 (0.1081 and 0.1699), but a larger deviation on Di-




Figure 10: Representative examples of fine-grained visual types in KMMMU. Before consolidation into the final macro-level visual modality categories, the dataset included diverse fine-grained visual types, such as specialized engineering diagrams, document-style text images, and South Korean geographic maps.

mension 3 (0.3747). Overall, these results are consistent with the main-text UMAP visualization: adversarial filtering increases difficulty while largely preserving the broader distributional structure of the original corpus.

F Inference Settings

F.1 Decoding Settings

Table 12 summarizes the main decoding settings used in our experiments. For open-weight models, we follow the Hugging Face generation configuration or the official model-card recommendation when available. If no model-specific setting is specified, we use the vLLM default sampling configuration. For proprietary models, we use the default

Discipline	Multiple Choice (Single Answer)	Multiple Choice (Multiple Answers)	Open Format (Numerical)	Open Format (Text)	Essay (Descriptive)
Engineering	1182	25	20	2	0
Natural Sciences	450	180	154	27	11
CS & IT	311	11	2	0	0
Business & Public	263	37	2	2	0
Social Sciences	193	55	2	4	0
Math & Stats	161	36	27	3	0
Law & Ethics	103	1	0	0	0
General/Interdisciplinary	89	0	0	30	4
Arts & Design	79	0	0	0	0

Table 11: Dataset attributes by disciplines and answer type.




Figure 11: Per-dimension density comparison after adversarial filtering. Kernel density estimates over the three UMAP dimensions for the original 68k corpus, the *KMMU*, and the *KMMU-Adv*. The filtered subsets broadly retain the major density peaks and multimodal trends of the original distribution, although the *KMMU-Adv* shows a somewhat larger deviation in Dimension 3.

Model family	Temperature	Top-p	Top-k
Gemma-3	1.0	0.95	64
Qwen3-VL-Instruct	0.7	0.8	20
Qwen3-VL-Thinking	1.0	0.95	20
Llama-4	0.6	0.9	—
HyperCLOVAX-SEED-Vision-3B	1.0	1.0	50
VARCO-VISION-2.0	1.0	1.0	—

Table 12: Decoding settings used for model evaluation. Dashes indicate parameters that were not explicitly specified or not exposed by the corresponding provider interface.

Model	API identifier	Access	Released	Accessed
GPT-5	gpt-5	OpenAI API	2025-08-07	2025-12-29
GPT-5-mini	gpt-5-mini	OpenAI API	2025-08-07	2025-12-31
Claude-Opus-4.5	anthropic/claude-opus-4.5	OpenRouter	2025-11-24	2026-01-01
Claude-Sonnet-4.5	anthropic/claude-sonnet-4.5	OpenRouter	2025-09-29	2025-12-30
Gemini-3-Pro	gemini-3-pro-preview	Gemini API	2025-11-18	2025-12-31
Gemini-3-Flash	gemini-3-flash-preview	Gemini API	2025-12-17	2025-12-31
Grok-4	x-ai/grok-4	OpenRouter	2025-07-09	2025-12-31
Grok-4.1-Fast	x-ai/grok-4.1-fast	OpenRouter	2025-11-19	2025-12-30
Mistral-Large-3	mistralai/mistral-large-2512	OpenRouter	2025-12-02	2025-12-30

Table 13: API identifiers and access information for proprietary models.

API settings provided by each vendor.

F.2 Proprietary Model Provenance

All evaluated proprietary models are publicly available models released in 2025. We access the Gemini family through the Google Gemini Developer API, the GPT family through the OpenAI API, and all other proprietary models through OpenRouter, using the default decoding settings exposed by each endpoint. Table 13 reports the exact API identifiers, release dates, and access dates.

F.3 Prompt Set

We use the same prompts for all evaluated models. The inference prompt asks the model to place its final answer in `\boxed{ }`, and the scoring prompt asks Gemini-2.5-Flash to decide whether the extracted answer matches the gold answer. An empty

extracted answer is scored FALSE.

Zero-shot Inference Prompt

Solve the given question.
You should refer to question sentences and an image.
You can print your reasoning and solving process.
You may answer in any language that helps you solve the problem.
At the end, you MUST present the final answer enclosed in `\boxed{ }`.

The final line MUST strictly follow this format:
`\boxed{your_final_answer}`

Do not put anything outside the `\boxed{ }` in the final line.

LLM-based Answer-Scoring Prompt

You will be given a question, the gold answer, and a model response.
Your task is to judge whether the model response correctly answers the question.

Return ONLY one of the following tokens on the LAST line.

TRUE
FALSE

Do not explain.
Do not add any other text.




Figure 12: **Discipline-wise visual modality composition of KMMMU-Adv subset.** Stacked bars show the number of questions for each visual modality in each discipline, with total counts shown beneath the labels. Scatter points indicate Korean-specific (KMMMU-Local) items overlaid on the corresponding discipline-modality segments. The KMMMU-Adv is concentrated in *Engineering* and *Natural Sciences*, similar to Full set.

Split	Korean Specific	Not Korean Specific	Ratio (%)
KMMMU	300	3166	8.65
KMMMU-Adv	76	551	12.12

Table 14: Number of Korean-specific questions each in Full set and KMMMU-Adv.

G Analysis of the KMMMU-Adv

G.1 Distributional Characteristics of the KMMMU-Adv

We analyze the structural composition of the KMMMU-Adv to better understand the types of instances that contribute to systematic model failures.

We first examine the prevalence of Korean-specific items (Table 14). Korean-specific questions account for 8.65% of the KMMMU set (300/3,466), but 12.12% of the KMMMU-Adv (76/627). This increase suggests that localized Korean content is somewhat overrepresented among harder examples.

However, because the majority of KMMMU-Adv subset questions are still not Korean-specific, localization alone does not fully explain the difficulty of the subset.

G.2 Model Performance on KMMMU-Adv

Table 15 reports KMMMU-Adv results for open-weight models only. Unlike the main KMMMU-Adv table, which reports proprietary models evaluated directly on KMMMU-Adv, this appendix table is obtained by restricting open-weight models’ full-set evaluation outputs to the retained KMMMU-Adv instances. Performance drops substantially relative to their full-set results across nearly all models, indicating that KMMMU-Adv is meaningfully more difficult for the evaluated open-weight models.

Even the strongest model remains below 20% overall accuracy, with VARCO-VISION-2.0-14B achieving the highest overall score at 19.56%.

Performance also varies considerably across disciplines. For example, *Law & Ethics* and *Arts & Design* remain difficult for most models, while *Engineering* and *Natural Sciences* still show modest separation among stronger systems. At the same time, reasoning models do not exhibit a consistent advantage over non-reasoning models on this subset. This pattern suggests that many KMMMU-Adv subset failures arise not simply from insufficient chain-of-thought depth, but from more persistent limitations in knowledge, grounding, visual interpretation, and answer execution.

H Additional Analyses of Korean-specific Items

H.1 Paired Outcomes of the Cross-lingual Ablation

Table 16 provides a detailed paired-outcome breakdown of the controlled cross-lingual ablation discussed in Section 5.2. All 200 items in this analysis are Korean-specific, after removing non-Korean-specific questions from the cross-lingual subset. The English condition translates only the textual question and answer options, while preserving the original visual input, answer, and task structure.

Across both models, accuracy is lower in the English condition than in the original Korean condition. GPT-5 decreases from 42.5% to 29.0%, and Gemini-3-Flash decreases from 58.0% to 40.0%. The paired outcomes show the same direction. For GPT-5, 33 items are answered correctly only in Ko-

Model	Arts & Design	Business & Public	CS & IT	Engineering	General	Law & Ethics	Math & Stats	Natural Sciences	Social Sciences	Overall Acc.
Open-weight Multilingual Non-Reasoning Models										
Gemma-3-4B-IT	22.67 _{1.89}	21.90 _{1.35}	9.94 _{4.14}	19.01 _{2.95}	14.04 _{1.24}	18.52 _{6.93}	10.34 _{2.82}	10.82 _{1.33}	24.07 _{5.24}	16.06 _{1.43}
Gemma-3-12B-IT	14.67 _{3.77}	8.57 _{4.04}	11.70 _{4.60}	15.32 _{1.84}	6.14 _{3.28}	22.22 _{0.00}	5.75 _{4.30}	8.87 _{1.70}	12.04 _{7.29}	12.01 _{0.72}
Gemma-3-27B-IT	16.44 _{8.46}	11.63 _{10.45}	15.11 _{7.09}	14.04 _{1.78}	14.04 _{12.59}	16.20 _{15.06}	9.51 _{8.57}	12.61 _{6.48}	16.24 _{6.60}	13.72 _{5.88}
Llama-4-Scout-17B-16E-IT	13.33 _{3.77}	9.52 _{1.35}	9.36 _{2.19}	17.73 _{0.53}	10.53 _{2.15}	20.37 _{2.62}	6.90 _{2.82}	11.26 _{2.72}	19.44 _{4.54}	13.98 _{0.53}
Llama-4-Maverick-17B-128E-IT	17.33 _{3.77}	13.33 _{1.35}	20.47 _{1.65}	16.74 _{1.06}	10.53 _{2.15}	22.22 _{4.54}	17.24 _{4.88}	14.29 _{1.91}	21.30 _{3.46}	16.37 _{1.08}
Qwen3-VL-2B-IT	25.33 _{3.77}	9.52 _{4.86}	12.87 _{2.98}	11.49 _{1.84}	3.51 _{2.48}	12.96 _{2.62}	4.60 _{3.25}	6.28 _{0.31}	6.48 _{1.31}	9.73 _{1.14}
Qwen3-VL-4B-IT	18.67 _{1.89}	10.48 _{3.56}	11.70 _{0.83}	12.34 _{0.35}	3.51 _{1.24}	3.70 _{2.62}	3.45 _{2.82}	4.98 _{0.61}	8.33 _{2.27}	9.20 _{0.08}
Qwen3-VL-8B-IT	20.00 _{3.27}	14.29 _{2.33}	16.37 _{3.60}	14.04 _{0.35}	5.26 _{0.00}	5.56 _{0.00}	6.90 _{2.82}	8.01 _{0.81}	14.81 _{2.62}	11.96 _{0.13}
Qwen3-VL-30B-A3B-IT	21.33 _{8.22}	14.29 _{4.04}	6.43 _{2.19}	13.62 _{0.92}	4.39 _{2.48}	9.26 _{6.93}	8.05 _{1.63}	10.17 _{0.61}	16.67 _{6.00}	11.70 _{1.21}
Qwen3-VL-32B-IT	14.67 _{4.99}	12.38 _{4.86}	14.62 _{3.60}	13.90 _{1.45}	7.02 _{3.28}	11.11 _{4.54}	9.20 _{1.63}	11.04 _{1.40}	21.30 _{3.46}	12.92 _{0.78}
Qwen3-VL-235B-A22B-IT	9.33 _{4.99}	12.38 _{2.69}	13.45 _{1.65}	17.02 _{0.92}	7.89 _{4.30}	7.41 _{5.24}	13.79 _{2.82}	14.94 _{1.06}	24.07 _{2.62}	15.05 _{0.87}
Open-weight Korean Non-Reasoning Models										
HyperCLOVAX-SEED-Vision-3B	9.33 _{1.89}	11.43 _{0.00}	16.37 _{2.19}	15.46 _{0.53}	4.39 _{1.24}	11.11 _{0.00}	8.05 _{1.63}	8.66 _{0.61}	14.81 _{2.62}	12.23 _{0.27}
VARCO-VISION-2.0-1.7B	17.33 _{1.89}	21.90 _{4.86}	19.88 _{0.83}	20.43 _{1.81}	7.02 _{1.24}	27.78 _{12.00}	12.64 _{4.30}	11.47 _{0.61}	16.67 _{3.93}	16.96 _{1.54}
VARCO-VISION-2.0-14B	20.00 _{3.27}	23.81 _{4.86}	22.81 _{7.58}	24.40 _{1.12}	3.51 _{1.24}	24.07 _{6.93}	13.79 _{4.88}	15.15 _{2.14}	16.67 _{0.00}	19.56 _{0.72}
Open-weight Multilingual Reasoning Models										
Qwen3-VL-30B-A3B-Thinking	18.00 _{2.00}	11.43 _{0.00}	4.39 _{2.63}	12.55 _{0.21}	6.58 _{1.32}	16.67 _{0.00}	15.52 _{1.72}	11.69 _{0.65}	15.28 _{4.17}	11.80 _{0.48}
Qwen3-VL-32B-Thinking	22.67 _{1.89}	6.67 _{4.86}	16.37 _{1.65}	12.34 _{1.04}	6.14 _{1.24}	16.67 _{4.54}	14.94 _{4.30}	11.47 _{3.84}	13.89 _{6.00}	12.55 _{0.53}
Qwen3-VL-235B-A22B-Thinking	4.00 _{3.27}	6.67 _{5.87}	3.51 _{2.48}	7.80 _{5.54}	4.39 _{6.20}	7.41 _{5.24}	6.90 _{7.45}	9.31 _{6.59}	13.89 _{10.39}	7.66 _{5.44}

Table 15: **Accuracy (%) on the *KMMU-Adv* subset by disciplines in Open-Weight models.** This table reports results recomputed by restricting full-set evaluation outputs to the adversarially filtered *Adv* subset. Overall accuracy is averaged across all disciplines. Mean accuracy is reported in percentage, with standard deviation shown as a subscript. The best result for each discipline and overall accuracy is shown in bold.

Model	n	Ko Acc.	En Acc.	Δ	Both Correct	Ko-only	En-only	Both Wrong
GPT-5	200	42.5	29.0	+13.5	52	33	6	109
Gemini-3-Flash	200	58.0	40.0	+18.0	78	38	2	82

Table 16: **Detailed paired outcomes for the cross-lingual ablation on Korean-specific items.** We evaluate 200 Korean-specific paired items in Korean and English. The English condition translates the textual question and answer options while keeping the other conditions are fixed. Accuracies are percentages. Δ denotes Korean accuracy minus English accuracy. Both Correct and Both Wrong indicate paired cases where the model is correct or incorrect in both language conditions. Ko-only and En-only indicate paired discordant cases where the model answers correctly only in Korean or only in English, respectively. Exact McNemar tests over the paired predictions show significant paired differences for both GPT-5 ($p = 1.43 \times 10^{-5}$) and Gemini-3-Flash ($p = 1.49 \times 10^{-9}$).

rean, compared to 6 items answered correctly only in English. For Gemini-3-Flash, 38 items are answered correctly only in Korean, compared to only 2 items answered correctly only in English. Exact McNemar tests over the paired predictions indicate that these differences are statistically significant for both models.

These results suggest that translating Korean-specific questions into English does not remove the main source of difficulty. Instead, translation may disrupt localized legal, administrative, technical, or institutional terminology, as well as the alignment between the translated question and Korean visual evidence.

H.2 Discipline-level Korean-specific Gaps

Table 17 decomposes the Korean-specific gap of Qwen3-VL-235B-A22B-IT by discipline. The gap is Korean-specific accuracy minus non-Korean-specific accuracy, and negative values indicate lower accuracy on Korean-specific items. The discipline-standardized gap is the average of within-discipline gaps weighted by the discipline distri-

bution of the full benchmark. We estimate 95% confidence intervals with 100,000 bootstrap resamples over the 3,466 items with three complete trials. Both Korean-specific and non-Korean-specific items are presented in Korean under the same zero-shot protocol with the same English instruction prompt, so this gap is computed independently of the cross-lingual ablation in Appendix H.1.

The deficit is concentrated in Business & Public and Law & Ethics. The confidence intervals for the other sufficiently represented disciplines except Natural Sciences and CS & IT include zero, so the aggregate gap is not uniform across disciplines. The Arts & Design estimate is descriptive because it contains only one Korean-specific item, and we report no confidence interval for it. The Math & Stats estimate should also be read with caution because it contains only six Korean-specific items.




Figure 13: **Annotation interface for manual validation of LLM scorer outputs.** For each sample, annotators review the question, image, gold answer, model response, and parsed answer, and record parsing consistency, correctness judgments, metadata consistency, and optional comments.

Discipline	n_{KS}	n_{Non}	Acc-KS	Acc-Non	Gap	95% CI
Business & Public	76	227	23.2	69.5	-46.2	[-55.4, -36.6]
Law & Ethics	82	22	27.6	65.2	-37.5	[-46.7, -28.3]
Arts & Design [†]	1	78	0.0	25.6	-25.6	-
Natural Sciences	31	792	17.2	42.4	-25.1	[-34.6, -14.1]
CS & IT	16	308	14.6	38.6	-24.1	[-38.4, -6.6]
General	28	95	21.4	25.6	-4.2	[-17.9, 10.5]
Math & Stats [†]	6	221	50.0	53.1	-3.1	[-17.0, 10.8]
Engineering	44	1,185	31.1	31.4	-0.3	[-12.2, 11.9]
Social Sciences	16	238	66.7	52.9	+13.7	[-7.0, 32.8]
Raw gap	300	3,166	-	-	-13.51	[-17.94, -9.02]
Discipline-standardized gap	300	3,166	-	-	-13.43	[-19.27, -7.44]

Table 17: **Discipline-level Korean-specific gaps of Qwen3-VL-235B-A22B-IT.** KS and Non denote Korean-specific and non-Korean-specific items. Accuracy and gaps are in percentage points. [†] denotes disciplines with fewer than ten Korean-specific items.

I Reliability of LLM-based Answer Scoring

I.1 Annotation Protocol

To validate the reliability of our evaluation pipeline, we conducted a manual annotation study using a custom annotation interface (Figure 13). Three annotators independently reviewed each sample with access to the question, associated image, gold answer, model response, and parsed answer.

For each sample, annotators first evaluated whether the parsed answer faithfully reflected the answer expressed in the original model response, labeling it as *match*, *mismatch*, or *no_answer*. Here, *mismatch* indicates that the parser failed to preserve the intended answer, while *no_answer* indicates that the model response itself did not contain a

complete answer.

Annotators then assessed correctness with respect to the gold answer in two ways: once based on the full model response and once based on the parsed answer, each labeled as *correct*, *incorrect*, or *no_answer*. This design allowed us to distinguish parsing failures from genuine model errors. In addition, annotators verified the consistency of the recorded question type and image type using *match*, *mismatch*, or *unsure*, and could provide free-form comments for ambiguous cases.

I.2 Human Alignment Results

The alignment study contains 600 examples drawn from six model runs, with 100 outputs sampled from each run and balanced across question formats. Table 18 reports pairwise human agreement and judge-human alignment under both parsed-answer-based and full-response-based evaluation. Human agreement is consistently high, with pairwise agreement ranging from 0.91 to 0.99, indicating that the annotation task is generally well defined. Agreement is slightly lower for some reasoning-heavy outputs, where responses are longer and often contain partially correct intermediate reasoning without a clearly finalized answer.

Overall, parsed-answer-based judging aligns more closely with human labels than full-response judging. Averaged across the six runs, parsed-

Model	Human agr.	No-answer rate	Overall judge-human agr.		No-answer acc.		Answered acc.	
			Parsed	Response	Parsed	Response	Parsed	Response
Qwen3-2B-IT	0.99	0.51	0.980	0.980	0.96	0.96	1.000	1.000
Qwen3-30B-IT	0.97	0.18	0.959	0.928	0.94	0.72	0.962	0.975
Qwen3-32B-Thinking	0.93	0.31	0.925	0.495	1.00	0.20	0.889	0.734
Qwen3-30B-Thinking	0.91	0.22	0.880	0.326	0.95	0.15	0.729	0.347
VARCO2-14B	0.97	0.04	0.948	0.929	1.00	1.00	0.946	0.926
GPT-5-Mini	0.95	0.01	0.937	0.892	1.00	1.00	0.936	0.891
Average	0.95	0.21	0.921	0.742	0.98	0.67	0.910	0.796

Table 18: **Human agreement and judge-human alignment across six model runs.** “Human agr.” denotes pairwise human agreement. “Parsed” and “Response” report judge-human alignment under parsed-answer-based and full-response-based evaluation, respectively. The no-answer rate denotes the proportion of outputs without a complete final answer. The no-answer and answered columns further break down judge-human accuracy by response completeness.

answer-based judging achieves 0.921 overall agreement, compared with 0.742 agreement for full-response judging. The gap is especially large for reasoning models, where full responses can contain ambiguous or incomplete reasoning traces even when the parsed final answer is easier to evaluate consistently.

Table 18 also breaks down judge-human accuracy by response completeness. On *no_answer* cases, parsed-answer-based judging reaches 0.98 accuracy on average, whereas full-response judging drops to 0.67. On answered cases, parsed-answer-based judging also remains stronger, achieving 0.910 accuracy compared with 0.796 for full-response judging. These results indicate that the benefit of parsed-answer-based evaluation is not limited to incomplete outputs; it also improves alignment when the model produces a final answer.

Manual inspection of disagreement cases suggests that many residual mismatches arise from answer-formatting and completion issues. In multiple-choice questions, some models produce the content of the correct option rather than its explicit index, which can cause an otherwise correct response to be judged as incorrect. More broadly, disagreement is concentrated in cases where the response contains extended or partially correct reasoning but fails to end with a clearly finalized answer. Taken together, these results support our use of parsed-answer-based judging as the primary evaluation protocol, especially for long or reasoning-heavy model outputs.

J Error Analysis Details

J.1 Error Inspection Methodology

To investigate the mechanisms underlying the patterns in Tables 3–5, we conducted targeted manual

error inspection over three focused subsets.

First, for paired reasoning comparison, we examined reversal cases between Qwen3-VL-32B-IT and Qwen3-VL-32B-Thinking. To analyze domain-specific reasoning effects, we sampled 25 items each from *Math & Stats*, *Engineering*, and *Natural Sciences* among questions answered correctly by Qwen3-VL-32B-IT but incorrectly by Qwen3-VL-32B-Thinking, yielding 75 inspected reversals in total.

Second, for Korean-specific failures, we analyzed incorrect outputs from a 300-item comparison set between Qwen3-VL-235B-A22B-IT and HyperCLOVAX-SEED-Vision-3B. We randomly sampled 25 incorrect cases from each model for qualitative inspection, focusing on recurring patterns of localized knowledge failure, regulatory category mismatch, and terminology grounding errors.

Third, to characterize persistent disciplinary bottlenecks, we additionally inspected representative failure cases from *Arts & Design* and *General*, focusing on Qwen3-VL-235B-A22B-IT with reference to corresponding Qwen3-VL-235B-A22B-Thinking outputs where relevant.

Each inspected case was reviewed by two authors, who examined the image, question, model output, and ground-truth answer. Disagreements were resolved through discussion.

J.2 Additional Qualitative Examples for Post-perceptual Reasoning Effects

Figure 14 reports performance by visual modality for Qwen3-VL-32B-IT and Qwen3-VL-32B-Thinking. Across most modality categories, the two variants remain broadly similar, with no consistent pattern indicating that explicit reasoning systematically improves raw visual evidence ex-




Figure 14: **Accuracy by visual modality for Qwen3-VL-32B-IT and Qwen3-VL-32B-Thinking.** Performance remains broadly similar across visual modality categories, suggesting that explicit reasoning does not systematically change raw visual evidence extraction. The main differences appear to arise after evidence extraction, such as in task framing, constraint tracking, and answer finalization.

traction.

To complement the aggregate modality comparison, we include a representative reversal case from *Natural Sciences* where Qwen3-VL-32B-IT answers correctly but Qwen3-VL-32B-Thinking fails in Figure 15. This example illustrates the broader pattern discussed in the main text: the difference does not arise from one variant clearly seeing the image while the other does not, but from how the extracted evidence is framed and translated into a final judgment.

In this case, the *Thinking* variant does not fail because it misses the basic visual structure or the relevant physical relation. Instead, it becomes anchored on an overly rigid conceptual rule and evaluates the option through that internal schema rather than the condition stated in the question itself. This qualitatively matches the pattern in our reversal inspection for *Natural Sciences*, where errors often stem from premature commitment to an incorrect problem frame rather than from missing visual evidence.

Figure 16 shows a representative *Engineering* reversal in which both variants initially misread the same FT diagram, but differ in whether they recover from that early structural error. The underlined spans mark the decision points that anchor each model’s subsequent reasoning. For Qwen3-VL-32B-IT, the underlined text shows an initial misinterpretation followed by an explicit self-correction once the model recognizes that the top gate is an AND gate, allowing it to recover the correct answer. By contrast, Qwen3-VL-32B-Thinking remains committed to the mistaken assumption that the top event is connected to an OR gate. That early structural error then propagates through the entire derivation,

leading the model to produce a fully consistent but fundamentally incorrect recovery analysis. This example matches the broader pattern we observed in *Engineering*: failures often arise from incorrect diagram-level structure interpretation, after which the model elaborates a coherent solution under the wrong logical frame.

J.3 Additional Qualitative Examples for Disciplinary Bottlenecks

For disciplinary bottlenecks, we use Qwen3-VL-235B-A22B-IT as a consistent reference point for representative qualitative examples. In addition inspected cases, we also examined corresponding outputs from Qwen3-VL-235B-A22B-Thinking, and observed qualitatively similar failure patterns. The examples in this appendix illustrate recurring errors in exact convention-to-label mapping in *Arts & Design* and few-shot symbolic induction or terminology grounding in *General*.

Example 1: Expert category mismatch in *Arts & Design*. Figure 17 shows a representative *Arts & Design* failure in which the model captures the coarse spatial organization of an architectural plan but fails to assign the correct standardized category. Rather than identifying the plan type required by the question, the model overinterprets the stacked horizontal layout as evidence for an orthogonal or cross-shaped structure, and then reinforces this mistaken frame with a plausible but incorrect villa association. This is not a low-level perception failure: the model recognizes salient geometric structure, but fails at precise convention-to-label mapping among closely related expert categories. The case therefore illustrates a recurring bottleneck in *Arts & Design*, where errors arise not from missing the visual content altogether, but from overconfident misclassification of specialized visual conventions into the wrong technical label.

Example 2: Few-shot symbolic induction failure in *General*. As shown in Figure 18, this item requires the model to infer a latent mapping between diagram configurations and Shanghaiese expressions from a small set of paired examples, and then apply that rule to unseen cases. The model produces a long, locally plausible analysis, but fails to recover the full underlying correspondence system. Instead, it partially matches surface patterns and then drifts into self-invented regularities, yielding answers that are structurally plausible but incorrect. This example illustrates a recurring *General* bot-

[Question - English Translation]
The following figure shows the energy band structure of a certain solid. Select all correct statements about energy band structures and semiconductors from the options below.

ㄱ. In an insulator, the valence band is completely filled with electrons, and the band gap is larger than that of a semiconductor.

ㄴ. Even if a voltage greater than the band gap is applied, no current flows in an insulator.

ㄷ. Even if a band gap exists, the material is metallic if the Fermi level lies inside the conduction band.

ㄹ. When an electron transitions from the conduction band to the valence band, the frequency of the emitted light increases as the band gap becomes larger.

[Image]

[Gold Answer] ㄱ, ㄷ, ㄹ

[Qwen3-VL-32B-IT Response - English Translation]
... (truncated)

ㄴ. Even if a voltage greater than the band gap is applied, no current flows in an insulator.

✗ Incorrect.

This statement is wrong. Even if the applied voltage is greater than the band gap, current generally does not flow in an insulator, but if a very large voltage (e.g., the breakdown voltage) is applied, dielectric breakdown can occur and current may flow. In other words, the absolute statement that "no current flows even if a larger voltage is applied" is incorrect. Therefore, this statement is inaccurate.

ㄷ. Even if a band gap exists, the material is metallic if the Fermi level lies inside the conduction band.

✓ Correct.

A metal is characterized by a partially filled conduction band or by overlap between the valence and conduction bands, so that there is no band gap or only a very small one. However, if the Fermi level lies inside the conduction band, electrons can move easily, so the material has very high conductivity and behaves like a metal. This distinguishes it from semiconductors and insulators. Therefore, this statement is correct.

... (truncated)

Thus, as the band gap becomes larger, the emitted energy increases, and the frequency also increases.

Therefore, the correct statements are: ㄱ, ㄷ, ㄹ

Final answer: $\boxed{\text{ㄱ, ㄷ, ㄹ}}$

[Qwen3-VL-32B-Thinking Response - English Translation]
Okay, let's tackle this problem. So, the question is about the energy band structure of a solid, specifically a semiconductor. The image shows two bands: the conduction band (전도띠) on top and the valence band (원자가띠) at the bottom, with a gap labeled "띠틈" between them.

... (truncated)

Option ㄱ: "절연체는 원자가띠에 전자가 완전히 채워져있고 띠틈이 반도체 보다 크다." (Insulators have the valence band completely filled with electrons, and the band gap is larger than in semiconductors.) This seems right. Insulators have a large band gap (usually more than 3 eV), while semiconductors have a smaller one (like 1-2 eV). The valence band is full in both insulators and semiconductors at absolute zero, but the gap size is the key. So ㄱ is correct.

... (truncated)

Option ㄷ: "띠틈이 있더라도 페르미 준위가 전도띠 내부에 위치하면 금속이다." (Even if there is a band gap, if the Fermi level is inside the conduction band, it's a metal.)

Wait, metals don't have a band gap. In metals, the conduction band and valence band overlap, so there's no gap. If there's a band gap (as in semiconductors or insulators), then the Fermi level can't be inside the conduction band unless the gap is zero. So if there is a band gap, the Fermi level can't be inside the conduction band. So this statement is incorrect. Because if there's a band gap, then the conduction band is empty (at 0K), so the Fermi level would be in the gap. So ㄷ is wrong.

... (truncated)

ㄹ: If there's a band gap, and the Fermi level is in the conduction band, then it's a metal. But metals don't have a band gap. So if there's a band gap, it's not a metal. So ㄹ is wrong.

... (truncated)

So the correct answers are $\boxed{\text{ㄱ, ㄹ}}$

Figure 15: **Rigid conceptual framing in a *Natural Sciences* reversal.** Qwen3-VL-32B-IT correctly applies the relevant conductivity criterion, whereas Qwen3-VL-32B-Thinking overcommits to an overly rigid band-gap-based schema and rejects the crucial statement about the Fermi level in the conduction band.

[Question - English Translation]
In the system analyzed by the fault tree (FT) shown in the figure, all components corresponding to the basic events are currently failed.

If the repair team restores Component 1 through Component 5 in order, at the completion of which component repair does the system become operational?

(1) Component 2 (2) Component 3 (3) Component 4 (4) Component 5

[Gold Answer] Option 2

[Qwen3-VL-32B-IT Response - English Translation]
The given FT has the following structure:

- The top event T (system failure) occurs through an AND gate over the three lower events.
- Each lower event is composed of an OR gate.

So, T occurs only when all three OR-gate outputs occur.

... (truncated)

The system becomes operational only if T does not occur.

... (truncated)

At first, the model says:
If even one of the three OR-gate outputs becomes normal, T does not occur and the system becomes normal.

Therefore, when Component 1 is repaired, the system becomes operational.

Then it revises its interpretation:
Ah! Because the top gate is an AND gate, T does not occur only when all three OR-gate outputs are normal.

****T=OR1∧OR2∧OR3****

So, for the system to be operational, OR1, OR2, and OR3 must all be normal.

It then tracks the repair sequence:

- After Component 1: OR1 = normal, OR2 = normal, OR3 = failed
→ T=normal∧normal∧failed=failed
- After Component 2: OR1 = normal, OR2 = normal, OR3 = failed
→ failed
- After Component 3: OR1 = normal, OR2 = normal, OR3 = normal
→ system operational

Final answer: $\boxed{(2) \text{ Component 3}}$

[Qwen3-VL-32B-Thinking Response - English Translation]
... (truncated)

The top event T is connected to an OR gate, and each of the three sub-gates is also an OR gate.

So, the system fails if any of the three sub-gates is true. Therefore, the system is operational only when all three sub-gates are false.

... (truncated)

The conditions are:

- First sub-gate (X1,X2,X3)(X1, X2, X3)(X1,X2,X3): false only if X1 = 0, X2 = 0, X3 = 0
- Second sub-gate (X1,X4)(X1, X4)(X1,X4): false only if X1 = 0, X4 = 0
- Third sub-gate (X3,X5)(X3, X5)(X3,X5): false only if X3 = 0, X5 = 0

... (truncated)

So, the system is operational only when all five components are restored. It then tracks the repair sequence:

- After repairing X3: First sub-gate = 0, second = 1, third = 1
→ T = 1, system failed
- After repairing X4: First = 0, second = 0, third = 1
→ T = 1, system failed
- After repairing X5: All sub-gates = 0
→ T = 0, system operational

... (truncated)

The top event T is connected to an OR gate. So T occurs if any of the three sub-gates occur. For T not to occur, all three sub-gates must not occur.

... (truncated)

Final answer: $\boxed{(4) \text{ Component 5}}$

Figure 16: **Structural misinterpretation in an *Engineering* reversal (Fault Tree analysis).** Qwen3-VL-32B-IT partially corrects an early gate-level misinterpretation, whereas Qwen3-VL-32B-Thinking persists with an incorrect top-level gate reading and derives the wrong recovery point.




Figure 17: **Exact architectural category misclassification in Arts & Design.** The model recognizes the overall structure of the plan, but fails to map it to the correct standardized architectural category. Instead, it overcommits to an orthogonal-plan interpretation and supports it with a plausible but incorrect villa association.




Figure 18: **Few-shot symbolic induction failure in General.** A Shanghainese-language item requiring the model to infer a latent mapping from a small set of diagram-expression pairs and apply it to new cases. Although Qwen3-VL-235B-A22B-IT produces a detailed step-by-step analysis, it fails to recover the full correspondence system and instead relies on partial surface analogies, leading to a plausible but incorrect answer.




Figure 19: **Exact standards and rule-criterion misapplication in *General*.** An item testing the official Romanization rule for hyphen use. The model gives a broadly reasonable explanation of the rule, but selects the wrong option because it applies an approximate plausibility-based criterion rather than the exact condition required by the formal standard.

tleneck in KMMMU: some items require few-shot rule induction from sparse symbolic evidence, not just fluent explanation or broad world knowledge.

Example 3: Exact standards and rule-criterion misapplication in *General*. Figure 19 shows a case that depends on precise application of an official Romanization rule, rather than broad linguistic plausibility alone. The model produces a reasonable explanation and discusses related principles, but selects the wrong option because it applies an approximate criterion instead of the exact standard required by the question. This pattern appears in multiple *General* items that test official terminology, orthographic conventions, or certification-style definitions: the model often gives a broadly sensible account, but misses the precise condition that determines correctness.

K Ablation Study

K.1 Evaluation of Image Dependency

Model	Original	Text-only
Gemini-3-Flash	45.15 _{0.98}	20.43 _{0.48}
GPT-5-Mini	21.32 _{0.33}	9.75 _{0.85}

Table 19: **Ablation results for image dependency.** We compare the average accuracy and standard deviation (Acc_{std}) of models on the original multimodal dataset versus the text-only baseline.

To verify that KMMMU functions as a genuinely multimodal benchmark, we measure how strongly performance depends on access to visual information. We conduct a text-only ablation in which mod-

els receive the textual question and answer options, but the associated image is removed.

Table 19 shows substantial performance drops for both Gemini-3-Flash and GPT-5-Mini under this setting. For Gemini-3-Flash, accuracy declines from 45.15% to 20.43%, a drop of 24.72 percentage points. For GPT-5-Mini, accuracy declines from 21.32% to 9.75%, a drop of 11.57 percentage points. These results indicate that many KMMMU questions cannot be solved reliably from text alone.

We also manually inspected the 60 cases answered correctly by GPT-5-Mini in the text-only setting. A non-trivial subset of visually accompanied items remains solvable without direct image access, but these cases often do not reflect genuine visual understanding. Instead, they typically fall into several recurring patterns: (i) the textual prompt already specifies most of the decisive constraints, making the image largely auxiliary; (ii) the answer can be inferred from strong domain priors or option elimination rather than visual grounding; (iii) the option structure, numerical form, or canonical diagram schema enables answer reconstruction without actual image reading; and (iv) in some quantitative science items, the core reasoning is already determined by symbolic conditions in the text, with the image serving mainly as contextual support.

Overall, these findings support the multimodal validity of KMMMU while also clarifying that visual accompaniment and strict image-essentiality are not identical. Although removing images causes large performance drops, some items remain text-

solvable because they contain sufficient textual, structural, or prior-driven cues to permit correct answering without direct image use.

K.2 Data Contamination Analysis

To probe possible memorization, we run a prefix-completion test in which models receive the first 35% of a question together with its image and are asked to generate the remaining continuation. We restrict this analysis to questions longer than 150 tokens, since shorter exam-style items often begin with generic instructions that provide too little question-specific content for a meaningful reconstruction test.

We evaluate three frontier models under two settings. In the first, no additional metadata is provided. In the second, we provide the exam name and year as a potential memorization trigger.

We evaluate generated continuations using Gemini-2.5-Flash as a judge. For each continuation, the judge assigns (i) an exactness rating from 0 to 100 based on overlap with the reference continuation and preservation of key details, and (ii) categorical labels indicating refusal or hallucination. Thus, exactness reflects judge-rated reconstruction fidelity, whereas refusal and hallucination characterize distinct failure behaviors rather than the same measurement scale.

As shown in Table 20, GPT-5-Mini and Gemini-3-Pro exhibit very high refusal rates across all four sources, typically accompanied by near-zero exactness ratings. This suggests that these models usually do not attempt faithful continuation under the prefix-completion setup.

Gemini-3-Flash attempts continuation more often, yielding lower refusal rates and somewhat higher exactness ratings than the other two models. However, its exactness remains low overall, and the hinted setting does not produce a consistent increase across sources. Moreover, hallucination rates remain high, indicating that many attempted continuations are low-fidelity generations rather than faithful reconstructions.

Taken together, these results do not provide strong evidence that benchmark performance is driven by simple memorization of question continuations. If contamination were a major driver under this setup, we would expect more consistently faithful reconstruction and clearer improvement when exam metadata is provided as a hint. Instead, the dominant pattern is either refusal or low-fidelity continuation.

Source	Model	Exactness _{NH}	Exactness _H	Refusal _{NH}	Refusal _H	Hallucination _{NH}	Hallucination _H
PSAT	Gemini-3-Flash	14.20	13.14	17.53	24.70	76.10	68.13
	Gemini-3-Pro	0.34	0.46	95.60	94.78	3.20	3.21
	GPT-5-Mini	0.00	0.00	99.60	100.00	0.00	0.00
NCS	Gemini-3-Flash	10.05	6.54	29.08	50.20	50.20	35.57
	Gemini-3-Pro	0.18	0.04	99.21	100.00	0.40	0.00
	GPT-5-Mini	0.65	0.45	95.65	100.00	1.19	1.19
NTQ	Gemini-3-Flash	15.70	18.58	19.66	15.61	61.54	76.84
	Gemini-3-Pro	0.33	2.03	93.86	79.15	4.44	16.92
	GPT-5-Mini	0.45	0.25	97.96	99.83	0.51	0.17
Olympiads	Gemini-3-Flash	10.48	10.91	23.12	37.56	55.58	49.48
	Gemini-3-Pro	0.13	0.38	97.92	97.41	1.56	2.33
	GPT-5-Mini	0.52	0.18	94.56	98.96	2.85	1.04

Table 20: **Prefix-completion analysis for potential data contamination.** Models are given the first 35% of each question together with the associated image, and asked to generate the remaining continuation. *Exactness* denotes a judge-assigned 0–100 faithfulness rating with respect to the reference continuation. *Refusal* and *Hallucination* denote judge-labeled failure modes, reported as percentages, under the No-hint (*NH*) and Hinted (*H*) settings.