

Korean Canonical Legal Benchmark: Toward Knowledge-Independent Evaluation of LLMs’ Legal Reasoning Capabilities

Hongseok Oh*
University of Seoul
cxv0519@uos.ac.kr

Wonseok Hwang*,†
University of Seoul
wonseok.hwang@uos.ac.kr

Kyoung-Woon On†
LBOX
kyoungwoon.on@lbox.kr

Abstract

We introduce the **Korean Canonical Legal Benchmark (KCL)**, a benchmark designed to assess language models’ legal reasoning capabilities independently of domain-specific knowledge. KCL provides question-level supporting precedents, enabling a more faithful disentanglement of reasoning ability from parameterized knowledge. KCL consists of two components: (1) **KCL-MCQA**, multiple-choice problems of 283 questions with 1,103 aligned precedents, and (2) **KCL-Essay**, open-ended generation problems of 169 questions with 550 aligned precedents and 2,739 instance-level rubrics for automated evaluation. Our systematic evaluation of 30+ models shows large remaining gaps, particularly in KCL-Essay, and that reasoning-specialized models consistently outperform their general-purpose counterparts. We release all resources, including the benchmark dataset and evaluation code, at <https://github.com/lbox-kr/kcl>.

1 Introduction

Large reasoning models (LRMs) trained to reason explicitly in the verbal space have shown superior performance over general large language models (Guo et al., 2025; OpenAI, 2025b). However, their capabilities have been evaluated mainly in formally structured domains such as mathematics, STEM, and code (Cobbe et al., 2021; Rein et al., 2024; Jain et al., 2025). A recent knowledge-independent, rule-following benchmark takes a step toward assessing LRMs’ reasoning independently of background knowledge (Ma et al., 2025).

However, real-world applications require solving problems in vertical domains such as biomedicine, finance, and law, where the types of reasoning are diverse and context dependent. In these domains, it is more challenging to isolate reasoning ability

from the memorized knowledge. Tasks in the legal domain are representative cases in which domain knowledge and logical reasoning are tightly intertwined, requiring complex context understanding and reasoning (Sadowski and Chudziak, 2025). However, the need for country-specific expertise makes legal data less accessible for LLM development. Consequently, prior benchmarks mainly focused on knowledge itself (Chalkidis et al., 2022; Hwang et al., 2022; Fei et al., 2024).

Meanwhile, several recent benchmarks attempt to incorporate legal reasoning more explicitly, but remain limited along a few common dimensions. Guha et al. (2023) define tasks aligned with the Issue, Rule, Application, Conclusion (IRAC) framework and evaluate them using exact match, but are restricted to English-language contexts and do not provide supporting knowledge. Similarly, Fan et al. (2025) construct IRAC-based tasks from English and German law exam questions and adopt a reference-based LLM-as-a-Judge evaluation scheme, yet likewise omit explicit knowledge resources. Kim et al. (2024) introduce reasoning-focused tasks in the Korean legal domain, but these are largely confined to relatively simple yes/no classifications (e.g., causal reasoning, statement consistency, case relevance), and their Korean Bar Exam MCQA does not include related precedents, making it difficult to disentangle reasoning ability from prior knowledge. Li et al. (2024) offer a large-scale and comprehensive evaluation of Chinese legal knowledge and reasoning, but the benchmark is predominantly multiple-choice-based and does not explicitly evaluate statute-/precedent-grounded faithfulness at the instance level. Focusing on retrieval, Zheng et al. (2025) propose legal RAG benchmarks that highlight the reasoning required to formulate retrieval queries, yet their tasks remain limited to multiple-choice and binary questions and are restricted to U.S. legal materials. Finally, while Lee et al. (2025) and Chlapanis et al. (2024) move

* Most work was done while at LBOX.

† Corresponding authors.

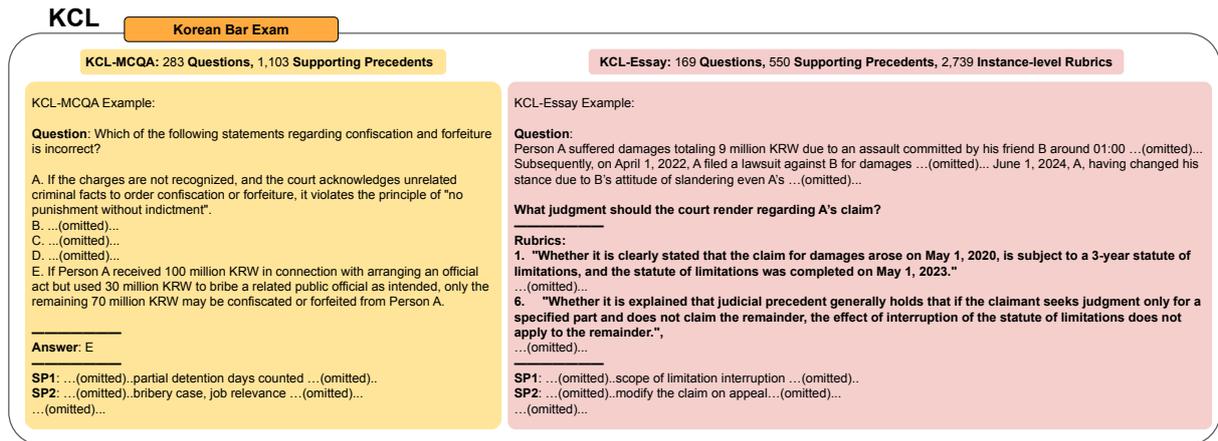


Figure 1: An overview of the KCL benchmark and representative samples translated into English.

beyond simple classification, addressing provision-grounded open-ended QA with fidelity-aware evaluation and legal argument continuation in ECHR cases, respectively, their scope remains constrained: the former focuses on statute-centric reasoning without precedent-based analysis, and the latter formulates legal reasoning as multiple-choice next-statement prediction rather than fully grounded end-to-end reasoning.

Going beyond prior benchmarks, our goal is to enable the evaluation of reasoning capabilities independently of parameterized knowledge, thereby providing more meaningful insights for LLM development. To this end, we introduce the Korean Canonical Legal Benchmark (KCL), a Korean Bar Exam-based benchmark that provides question-level supporting precedents. KCL further distinguishes itself in two ways: it enables the automatic evaluation of open-ended generation with rubric-based LLM-as-a-Judge, and it targets Korean, a low-resource language, thereby broadening the scope of legal LLM evaluation beyond English-centric contexts.

2 KCL Dataset Construction

KCL consists of two complementary benchmarks derived from the Korean Bar Exam: KCL-MCQA, a multiple-choice benchmark, and KCL-Essay, an open-ended generation benchmark (Figure 1).

Motivation KCL is designed to evaluate models along two distinct dimensions. The first dimension, knowledge coverage, measures how much legal knowledge a model has internally parameterized, which can be assessed directly from question-answer pairs. The second, evidence-grounded reasoning, captures a model’s ability to apply rea-

soning when relevant legal knowledge is supplied externally.

Disentangling these two axes is particularly difficult in law, where specialized, country-specific expertise is required and high-quality data remains scarce. To address this challenge, KCL pairs each question with supporting precedents extracted directly from expert solutions. These precedents are not answer labels but question-aligned knowledge that provides a foundation for reasoning.

Design KCL evaluates Large Language Models (LLMs) in two settings: a vanilla setting, where models answer questions without extra resources to reveal their knowledge coverage, and a w/ supporting precedents setting, where supporting precedents are provided to assess evidence-grounded reasoning. Comparing the two allows developers to diagnose and improve knowledge acquisition and reasoning independently, leading to more precise evaluation and better-targeted model development.

2.1 KCL-MCQA

We constructed the multiple-choice set by curating questions from the 2024 and 2025 Korean Bar Exams. All answer options in KCL-MCQA are taken verbatim from the official Korean Bar Exam. A prior study (Kim et al., 2024) reported that, even under a RAG setting, using only publicly available datasets of judicial precedents and statutes not directly connected to the task yielded only negligible performance gains (Table 2 in (Kim et al., 2024)). As a result, it was difficult to assess how models would perform if provided with sufficient knowledge prior to reasoning.

To address this, we first curated questions by

filtering expert commentaries that cited at least one precedent. We then retrieved the referenced precedents for each question using a Korean legal document search engine¹, and defined them as supporting precedents. The resulting benchmark, KCL-MCQA, consists of 283 five-choice questions, with an average of 3.9 supporting precedents per question, for a total of 1,103 precedents (Table 1).

2.2 KCL-Essay

To examine LLMs on legal reasoning tasks, we further focused on the case-based questions of the Korean Bar Exam that correspond to the Multi-state Essay Exam (MEE) of the U.S. Uniform Bar Exam. They consist of open-ended generation questions that require candidates to identify legal issues, apply statutes and precedents, and construct well-reasoned conclusions (IRAC).

To build KCL-Essay, similar to KCL-MCQA, we first collected case-based questions and their expert commentaries from the past five years (2021-2025) of the Korean Bar Exam. We then curated 169 questions for which a model’s answer would need to reference at least one precedent. And we retrieved supporting precedents again. On average, we obtained 3.3 precedents per question, resulting in a total of 550 precedents (Table 1).

Unlike KCL-MCQA, KCL-Essay is an open-ended generation task, and we therefore adopted an LLM-as-a-Judge approach for automated evaluation. Following prior studies (Fan et al., 2024; Kim et al., 2025), we constructed instance-level rubrics to provide finer-grained criteria tailored to the context of each question.

To build the rubrics, we first drafted an initial set for a representative problem in each subject (civil, criminal, public) with assistance from a licensed Korean attorney. Using these initial rubrics as a one-shot seed, we prompted Gemini 2.5 Flash to generate rubric sets for all problems with the question text and expert solution. Finally, three licensed Korean attorneys reviewed the 744 generated rubrics, correcting 13 of them (2%). Given this low error rate, we built a total of 2,739 rubrics automatically.

In the original essay-type Korean Bar Exam, each question is assigned points reflecting its difficulty. We preserved these weightings and evenly distributed the total score across its rubrics, yield-

Dataset	KCL-MCQA	KCL-Essay
# Questions	283	169
Mean tokens per question	552	745
Mean # rubrics per question	-	16.2
Mean # supporting precedents	3.9	3.3
Mean tokens per precedent	4,631	3,648
Max tokens per precedent	13,833	749,580

¹ <https://github.com/openai/tiktoken>

Table 1: Statistics of KCL-MCQA and KCL-Essay.¹

ing a total score of 2,905 points for the entire benchmark. We report model scores as percentages of the total in the results below.

3 Experimental Setup

Inference We systematically evaluated 30+ models across the two components, KCL-MCQA and KCL-Essay. Following Fan et al. (2025), we categorized the models into four types: Large Reasoning, Small Reasoning, Large General, and Small General. For open-API models, since detailed information was not disclosed, we classified them as either Small or Large according to whether they were the smaller or larger variant within the same release family. For open-weight models, we use 32B parameters for dense models and 120B for MoE models as thresholds to classify them as Large or Small. Models trained to reason in the verbal space were categorized as Reasoning models. We used vLLM (Kwon et al., 2023) to evaluate open-weight models, and the usage details for open-API models are provided in Appendix C. Beyond well-known models, we additionally evaluate LLMs developed in Korea, including HCX SEED (Team, 2025), Kanana (Team et al., 2025b), A.X (SKT-AI, 2025), and Exaone (Research et al., 2025), along with multilingual-focused models such as Aya Expans (Dang et al., 2024)

Evaluation We evaluated KCL-Essay using Gemini 2.5 Flash as the judge model. In preliminary comparisons (see Figure 2), larger judge models showed slightly higher correlations, but Gemini 2.5 Flash was selected for its favorable cost and latency profile while retaining comparable agreement with human experts. The use of instance-level rubrics yielded a strong correlation between human experts and LLM: ~ 0.9 for questions with total scores below 20 points (80% of questions), ~ 0.7 for those above 20 points, and ~ 0.8 overall (Pearson; Figure 2 in Appendix A.1). Assuming that the original total scores from the bar exam reflect the true difficulty of the questions, these results sug-

¹LBOX, <https://lbox.kr/v2>

KCL-MCQA		
Type	Model Name	score (% , \uparrow) vanilla \rightarrow w/ supporting precedents
Large Reasoning	Gemini 2.5 Pro	66.1 \pm 1.0 \rightarrow 89.3 \pm 0.9
	Claude Opus 4.1*	57.2 \pm 1.0 \rightarrow 83.3 \pm 0.7
	Claude Sonnet 4.5	54.9 \pm 0.7 \rightarrow 85.2 \pm 1.3
	GPT-5	53.9 \pm 1.1 \rightarrow 84.2 \pm 0.2
	Claude 3.7 Sonnet	50.4 \pm 1.0 \rightarrow 82.0 \pm 0.3
	o3	50.2 \pm 1.2 \rightarrow 83.4 \pm 1.8
	Claude Sonnet 4*	45.7 \pm 2.1 \rightarrow 78.6 \pm 1.6
	Qwen3 235B Thinking	41.5 \pm 1.3 \rightarrow 76.9 \pm 1.4
DeepSeek-R1*	40.9 \pm 1.9 \rightarrow 72.4 \pm 1.3	
Average		51.2 \rightarrow 81.7
Small Reasoning	Gemini 2.5 Flash	51.8 \pm 0.6 \rightarrow 83.7 \pm 0.0
	GPT-5 Mini	40.9 \pm 1.3 \rightarrow 76.8 \pm 1.4
	o4-mini	37.7 \pm 1.2 \rightarrow 74.6 \pm 0.8
	Exaone 4.0 32B	36.4 \pm 1.6 \rightarrow 72.3 \pm 2.2
	Qwen3 32B	35.0 \pm 3.2 \rightarrow 65.4 \pm 1.3
	HCX SEED Think 14B*	31.3 \pm 0.8 \rightarrow 57.6 \pm 0.3
	gpt-oss 120B	29.2 \pm 1.6 \rightarrow 64.1 \pm 1.2
	Average	
Large General	A.X 4.0	49.1 \pm 0.5 \rightarrow 71.3 \pm 0.4
	GPT-4.1	47.9 \pm 1.0 \rightarrow 72.0 \pm 0.4
	Claude 3.5 Sonnet*	45.1 \pm 0.3 \rightarrow 70.8 \pm 0.9
	Llama 4 Maverick	43.3 \pm 0.4 \rightarrow 72.3 \pm 1.0
	GPT-4o	43.3 \pm 2.0 \rightarrow 63.8 \pm 1.0
	Qwen3 235B Instruct	40.9 \pm 1.3 \rightarrow 74.2 \pm 0.8
	Kimi K2 Instruct	39.5 \pm 1.2 \rightarrow 71.6 \pm 1.4
	DeepSeek-V3	37.3 \pm 0.7 \rightarrow 69.3 \pm 0.3
Average		43.3 \rightarrow 70.7
Small General	Gemma 3 27B	35.1 \pm 0.3 \rightarrow 53.6 \pm 0.4
	GPT-4o-mini	29.9 \pm 1.5 \rightarrow 46.4 \pm 0.9
	GPT-4.1-mini	29.4 \pm 0.7 \rightarrow 60.5 \pm 1.2
	Kanana 1.5 8B	29.4 \pm 0.8 \rightarrow 47.5 \pm 1.4
	Aya Expans 32B	28.4 \pm 0.2 \rightarrow 42.3 \pm 0.4
	Claude 3.5 Haiku*	28.0 \pm 1.3 \rightarrow 54.5 \pm 1.4
	Phi-4-mini	22.0 \pm 0.8 \rightarrow 30.6 \pm 0.4
	Average	
Random		20

Table 2: Evaluation results on KCL-MCQA, averaged over three independent runs. Precedents exceeding the maximum context length were truncated (indicated by *). "Average" denotes the average performance within each group.

gest that while LLM-as-a-Judge with instance-level rubrics is generally effective, caution is needed when evaluating complex examples. Further details are explained in the Appendix.

4 Analysis

4.1 The Impact of Legal Knowledge on Problem Solving Performance

We evaluated various LLMs on the vanilla setting first, where supporting precedents are not provided, to examine differences in parameterized knowledge across models.

As shown in Table 2, performance is strongly associated with model size. The average performance followed the order of Large Reasoning, Large General, Small Reasoning, and Small General. This accords with findings from prior work that larger models tend to memorize more during training (Tirumala et al., 2022). Among Large

KCL-Essay		
Type	Model Name	score (% , \uparrow) vanilla \rightarrow w/ supporting precedents
Large Reasoning	Gemini 2.5 Pro	60.5 \pm 0.5 \rightarrow 74.9 \pm 0.5
	GPT-5	57.5 \pm 1.3 \rightarrow 72.8 \pm 0.3
	o3	51.2 \pm 0.6 \rightarrow 69.4 \pm 0.8
	Claude Sonnet 4.5	47.6 \pm 1.5 \rightarrow 68.5 \pm 0.4
	Claude Opus 4.1*	45.5 \pm 0.2 \rightarrow 60.3 \pm 0.5
	Claude 3.7 Sonnet	43.6 \pm 0.7 \rightarrow 63.6 \pm 0.8
	Qwen3 235B Thinking	39.2 \pm 1.1 \rightarrow 63.9 \pm 0.1
	Claude Sonnet 4*	35.5 \pm 1.2 \rightarrow 54.9 \pm 0.2
DeepSeek-R1*	35.0 \pm 0.5 \rightarrow 53.1 \pm 0.2	
Average		46.2 \rightarrow 64.6
Small Reasoning	Gemini 2.5 Flash	53.7 \pm 1.1 \rightarrow 71.1 \pm 0.1
	GPT-5 Mini	40.3 \pm 0.5 \rightarrow 65.5 \pm 0.5
	Exaone 4.0 32B	32.6 \pm 0.4 \rightarrow 47.5 \pm 1.1
	gpt-oss 120B	31.6 \pm 0.5 \rightarrow 56.0 \pm 0.3
	o4-mini	31.3 \pm 0.9 \rightarrow 54.5 \pm 0.3
	Qwen3 32B	29.2 \pm 0.3 \rightarrow 55.4 \pm 1.2
	HCX SEED Think 14B*	26.2 \pm 0.2 \rightarrow 42.8 \pm 0.7
	Average	
Large General	GPT-4.1	42.4 \pm 0.8 \rightarrow 65.1 \pm 0.6
	Qwen3 235B Instruct	38.9 \pm 0.9 \rightarrow 63.0 \pm 0.5
	A.X 4.0	34.7 \pm 0.6 \rightarrow 52.4 \pm 0.9
	DeepSeek-V3	34.3 \pm 0.6 \rightarrow 57.6 \pm 1.5
	Claude 3.5 Sonnet*	33.7 \pm 0.3 \rightarrow 50.3 \pm 0.7
	Kimi K2 Instruct	33.3 \pm 0.3 \rightarrow 55.7 \pm 0.6
	Llama 4 Maverick	25.6 \pm 0.3 \rightarrow 44.3 \pm 0.8
	GPT-4o	24.3 \pm 0.4 \rightarrow 42.0 \pm 0.1
Average		33.4 \rightarrow 53.8
Small General	GPT-4.1-mini	30.9 \pm 0.7 \rightarrow 56.1 \pm 1.0
	Gemma 3 27B	23.6 \pm 0.4 \rightarrow 44.0 \pm 0.5
	Kanana 1.5 8B	21.8 \pm 0.8 \rightarrow 49.1 \pm 0.4
	Claude 3.5 Haiku*	21.2 \pm 0.6 \rightarrow 39.9 \pm 0.1
	Aya Expans 32B	20.2 \pm 0.3 \rightarrow 37.2 \pm 0.1
	GPT-4o-mini	16.1 \pm 0.3 \rightarrow 32.6 \pm 0.4
	Phi-4-mini	6.0 \pm 0.3 \rightarrow 18.9 \pm 0.1
	Average	

Table 3: Evaluation results on KCL-Essay. The experimental settings are identical to those in Table 2.

Reasoning models, Gemini 2.5 Pro performs highest overall, followed by Claude and GPT-5 effectively showing their parameterized legal knowledge in a low-resource language such as Korean.

Next, in the w/ supporting precedents setting, where corresponding precedents were provided as in-context input, scores increased substantially compared to the vanilla setting across all models. Interestingly, the average performance gap between the Large General and Small Reasoning types narrows to just 0.1%p. This suggests that while Small Reasoning models have lower knowledge coverage than Large General models, they may nevertheless demonstrate superior evidence-grounded reasoning. Also note that many models surpass the human passing threshold (68.7% (Seo, 2025)) underscoring that high-quality legal data remains scarce.

4.2 On the Importance of Reasoning for KCL-Essay

Unlike KCL-MCQA, in KCL-Essay, the model exhibits average performance in the order of Large

Reasoning, Small Reasoning, Large General, and Small General in both vanilla and w/ supporting precedents settings. Furthermore, GPT-5 and o3 outperform the Claude model family across all settings. These results suggest that KCL-Essay requires deeper reasoning ability.

Nevertheless, even in the w/ supporting precedents setting, the absolute scores remained relatively low, unlike in KCL-MCQA. Even the best-performing model, Gemini 2.5 Pro, still exhibited a remaining gap of about 25%p, underscoring the need for further advances in legal reasoning.

4.3 On the Performance of Open-Weight Models in KCL

We analyze the performance of open-weight models. As shown in Table 2, for KCL-MCQA, it is noteworthy that in the Large Reasoning and Small Reasoning types, the best-performing models across all settings were open-API models. By contrast, in the Large General and Small General types, the best models in the vanilla setting were the open-weight models A.X 4.0 (SKT-AI, 2025) and Gemma 3 27B (Team et al., 2025a), respectively, although they fell behind open-API models in the w/ supporting precedents setting. This suggests that open-weight models may achieve competitive performance in terms of knowledge coverage within certain vertical domains, while still showing limitations in deeper reasoning.

For KCL-Essay, as shown in Table 3, the top-performing model in every type was consistently an open-API model. Although open-weight LLMs have demonstrated strength in STEM and code-related benchmarks (Guo et al., 2025), they continue to face clear challenges in specialized domains requiring expert-level reasoning, where API-based models hold a substantial advantage. These results highlight the value of KCL as a reliable reference for advancing the open-weight model ecosystem.

4.4 Effects of Language Specialization

Given that our benchmark data is based on Korean, a low-resource language, we also conducted a closer analysis of models specialized in Korean. In KCL-MCQA, as shown in Table 2, HCX SEED Think 14B achieved higher performance in the vanilla setting than the recent open-weight LLM gpt-oss 120B. Similarly, as noted in Section 4.3, A.X 4.0 surpassed all other open-API and open-

weight models in the Large General type. However, both models exhibited limited improvement in the w/ supporting precedents setting and were eventually outperformed by other models.

In KCL-Essay, interestingly, Exaone 4.0 32B and A.X 4.0, both trained specifically for Korean, even outperformed gpt-oss 120B and Kimi K2 Instruct in the vanilla setting, but underperformed relative to them in the w/ supporting precedents setting. Although these models benefit from Korean-specific knowledge in terms of knowledge coverage, their reasoning ability still lags behind, underscoring the need for further advances in reasoning-focused training.

In contrast to KCL-MCQA, in KCL-Essay the HCX SEED Think 14B (Team, 2025), despite being specifically trained for Korean, recorded the lowest score among models in the Small Reasoning category across all settings. This supports our earlier finding in Section 4.2 that KCL-Essay places a stronger emphasis on expert-level reasoning rather than knowledge coverage alone.

Taken together, this analysis underscores the importance of KCL as a benchmark in the broader context of developing national AI strategies.

5 Conclusion

We propose KCL for knowledge-independent legal reasoning: KCL-MCQA, with 283 questions and 1,103 aligned precedents, and KCL-Essay, with 169 problems, 550 precedents, and 2,739 rubrics. With supporting precedents, many frontier LLMs surpass 80% on KCL-MCQA; without them, errors largely reflect missing knowledge. KCL-Essay remains difficult: the best score is 74.9% even with precedents. KCL’s supporting precedents and instance-level rubrics enable controlled, reproducible evaluation. Overall, KCL is designed to foster models that *reason*, rather than merely *recall* knowledge, by disentangling evidence-grounded reasoning from parameterized knowledge when evaluating model capabilities.

Limitation

KCL focuses exclusively on Korean legal tasks, limiting direct applicability to other legal systems and languages. For KCL-Essay, we rely on LLM-as-a-Judge (Gemini 2.5 Flash) for automated scoring, which, despite validated correlation with human expert evaluations, may not capture all nuances of legal reasoning quality. The w/

supporting precedents setting assumes perfect knowledge retrieval and does not fully isolate the retrieval-reasoning hybrid challenge of identifying which parts of provided precedents are relevant. Future work on KCL-Essay could benefit from aligning the evaluation rubric with the IRAC framework, which would allow disentangling and assessing reasoning at each stage more explicitly.

Another limitation is that our benchmark evaluates model performance on questions that are already annotated with correct precedents in expert solutions. In real-world deployments, legal AI systems are typically implemented as retrieval-augmented generation (RAG) pipelines. Because our study does not model an end-to-end pipeline that performs retrieval and reasoning jointly, we do not account for noise, omissions, or misranking introduced during retrieval. Nevertheless, the benchmark explicitly provides question-aligned precedents, which enables an additional and complementary use case: the annotated precedents serve as gold references, allowing retrieval components to be evaluated.

Acknowledgments

We thank Sooyung Park, a licensed lawyer in South Korea, for valuable advice on handling legal data. Hongseok Oh and Wonseok Hwang were supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government(MSIT) (RS-2025-23524855). Hongseok Oh was also supported by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (NO. RS-2022-NR068754)

References

- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J. Hewett, Mojan Javaheripi, Piero Kauffmann, James R. Lee, Yin Tat Lee, Yuanzhi Li, Weishung Liu, Caio C. T. Mendes, Anh Nguyen, Eric Price, Gustavo de Rosa, Olli Saarikivi, and 8 others. 2024. [Phi-4 technical report](#). *Preprint*, arXiv:2412.08905.
- Moonshot AI. 2025. [Kimi-k2 instruct license \(modified mit\)](#). Accessed: 2025-10-02.
- Anthropic. 2023. [Anthropic on bedrock - commercial terms of service](#). Accessed: 2025-10-02.
- Anthropic. 2025a. [Claude 3.7 sonnet and claude code](#). Accessed: 2025-10-02.
- Anthropic. 2025b. [System card addendum: Claude opus 4.1](#). Accessed: 2025-10-02.
- Anthropic. 2025c. [System card: Claude opus 4 and claude sonnet 4](#). Accessed: 2025-10-02.
- Anthropic. 2025d. [System card: Claude sonnet 4.5](#). Accessed: 2025-10-02.
- Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Katz, and Nikolaos Aletras. 2022. [LexGLUE: A benchmark dataset for legal language understanding in English](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4310–4330, Dublin, Ireland. Association for Computational Linguistics.
- Odysseas S. Chlapanis, Dimitrios Galanis, and Ion Androutsopoulos. 2024. [LAR-ECHR: A new legal argument reasoning task and dataset for cases of the European court of human rights](#). In *Proceedings of the Natural Legal Language Processing Workshop 2024*, pages 267–279, Miami, FL, USA. Association for Computational Linguistics.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- NAVER Corporation and NAVER Cloud Corporation. 2025. [Hyperclova x seed 14b think model license agreement](#).
- John Dang, Shivalika Singh, Daniel D’souza, Arash Ahmadian, Alejandro Salamanca, Madeline Smith, Aidan Peppin, Sungjin Hong, Manoj Govindassamy, Terrence Zhao, Sandra Kublik, Meor Amer, Viraat Aryabumi, Jon Ander Campos, Yi-Chern Tan, Tom Kocmi, Florian Strub, Nathan Grinsztajn, Yannis Flet-Berliac, and 26 others. 2024. [Aya expand: Combining research breakthroughs for a new multilingual frontier](#). *Preprint*, arXiv:2412.04261.
- Google DeepMind. 2025a. [Gemini 2.5 flash and 2.5 flash image model card](#). Accessed: 2025-10-02.
- Google DeepMind. 2025b. [Gemini 2.5 pro model card](#). Accessed: 2025-10-02.
- DeepSeek. 2023. [License-model](#). Accessed: 2025-10-02.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.

- Yu Fan, Jingwei Ni, Jakob Merane, Etienne Salimbeni, Yang Tian, Yoan Hermstrüwer, Yinya Huang, Mubashara Akhtar, Florian Geering, Oliver Dreyer, Daniel Brunner, Markus Leippold, Mrinmaya Sachan, Alexander Stremitzer, Christoph Engel, Elliott Ash, and Joel Niklaus. 2025. [Lexam: Benchmarking legal reasoning on 340 law exams](#). *Preprint*, arXiv:2505.12864.
- Zhiyuan Fan, Weinong Wang, Xing W, and Debing Zhang. 2024. [SedarEval: Automated evaluation using self-adaptive rubrics](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 16916–16930, Miami, Florida, USA. Association for Computational Linguistics.
- Zhiwei Fei, Xiaoyu Shen, Dawei Zhu, Fengzhe Zhou, Zhuo Han, Alan Huang, Songyang Zhang, Kai Chen, Zhixin Yin, Zongwen Shen, Jidong Ge, and Vincent Ng. 2024. [LawBench: Benchmarking legal knowledge of large language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7933–7962, Miami, Florida, USA. Association for Computational Linguistics.
- Google. 2025a. [Gemma terms of use](#). Last updated: September 12, 2025; Accessed: 2025-10-02.
- Google. 2025b. [Google apis terms of service](#). Accessed: 2025-10-02.
- Neel Guha, Julian Nyarko, Daniel E. Ho, Christopher Re, Adam Chilton, Aditya Narayana, Alex Chohlas-Wood, Austin Peters, Brandon Waldon, Daniel Rockmore, Diego Zambrano, Dmitry Talisman, Enam Hoque, Faiz Surani, Frank Fagan, Galit Sarfaty, Gregory M. Dickinson, Haggai Porat, Jason Hegland, and 21 others. 2023. [Legalbench: A collaboratively built benchmark for measuring legal reasoning in large language models](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z F Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature*, 645(8081):633–638.
- Wonseok Hwang, Dongjun Lee, Kyoungyeon Cho, Hanuhl Lee, and Minjoon Seo. 2022. [A multi-task benchmark for korean legal language understanding and judgement prediction](#). In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2025. [Livecodebench: Holistic and contamination free evaluation of large language models for code](#). In *The Thirteenth International Conference on Learning Representations*.
- Seungone Kim, Juyoung Suk, Ji Yong Cho, Shayne Longpre, Chaceun Kim, Dongkeun Yoon, Guijin Son, Yejin Cho, Sheikh Shafayat, Jinheon Baek, Sue Hyun Park, Hyeonbin Hwang, Jinkyung Jo, Hyowon Cho, Haebin Shin, Seongyun Lee, Hanseok Oh, Noah Lee, Namgyu Ho, and 13 others. 2025. [The BiGGen bench: A principled benchmark for fine-grained evaluation of language models with language models](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5877–5919, Albuquerque, New Mexico. Association for Computational Linguistics.
- Yeeun Kim, Youngrok Choi, Eunkyung Choi, JinHwan Choi, Hai Jin Park, and Wonseok Hwang. 2024. [Developing a pragmatic benchmark for assessing Korean legal language understanding in large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 5573–5595, Miami, Florida, USA. Association for Computational Linguistics.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP '23*, page 611–626, New York, NY, USA. Association for Computing Machinery.
- Jihyung Lee, Daehui Kim, Seonjeong Hwang, Hyounghun Kim, and Gary Lee. 2025. [KoBLEX: Open legal question answering with multi-hop reasoning](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 4019–4053, Suzhou, China. Association for Computational Linguistics.
- Haitao Li, You Chen, Qingyao Ai, Yueyue Wu, Ruizhe Zhang, and Yiqun LIU. 2024. [Lexeval: A comprehensive chinese legal benchmark for evaluating large language models](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Kaijing Ma, Xeron Du, Yunran Wang, Haoran Zhang, ZhoufutuWen, Xingwei Qu, Jian Yang, Jiaheng Liu, minghao liu, Xiang Yue, Wenhao Huang, and Ge Zhang. 2025. [KOR-bench: Benchmarking language models on knowledge-orthogonal reasoning tasks](#). In *The Thirteenth International Conference on Learning Representations*.
- Meta. 2025a. [Llama 4 community license agreement](#). Accessed: 2025-10-02.
- AI Meta. 2025b. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation. <https://ai.meta.com/blog/llama-4-multimodal-intelligence/>, checked on, 4(7):2025.

- OpenAI, :, Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K. Arora, Yu Bai, Bowen Baker, Haiming Bao, Boaz Barak, Ally Bennett, Tyler Bertao, Nivedita Brett, Eugene Brevdo, Greg Brockman, Sebastien Bubeck, and 108 others. 2025. [gpt-oss-120b and gpt-oss-20b model card](#). *Preprint*, arXiv:2508.10925.
- OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, and 401 others. 2024. [Gpt-4o system card](#). *Preprint*, arXiv:2410.21276.
- OpenAI. 2024a. [Gpt-5 system card](#). Accessed: 2025-10-02.
- OpenAI. 2024b. [Openai terms of use](#). Accessed: 2025-10-02.
- OpenAI. 2025a. [Introducing GPT-4.1 in the api](#). Accessed: 2025-10-02.
- OpenAI. 2025b. [Openai o3 and o4-mini system card](#). System card, version published April 16, 2025.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2024. [GPQA: A graduate-level google-proof q&a benchmark](#). In *First Conference on Language Modeling*.
- LG AI Research, :, Kyunghoon Bae, Eunbi Choi, Kibong Choi, Stanley Jungkyu Choi, Yemuk Choi, Kyubeen Han, Seokhee Hong, Junwon Hwang, Taewan Hwang, Joonwon Jang, Hyojin Jeon, Kijeong Jeon, Gerrard Jeongwon Jo, Hyunjik Jo, Jiyeon Jung, Euisoon Kim, Hyosang Kim, and 23 others. 2025. [Exaone 4.0: Unified large language models integrating non-reasoning and reasoning modes](#). *Preprint*, arXiv:2507.11407.
- Albert Sadowski and Jarosław A. Chudziak. 2025. [On verifiable legal reasoning: A multi-agent framework with formalized knowledge representations](#). *Preprint*, arXiv:2509.00710.
- Hayeon Seo. 2025. [슈퍼로이어, 국내 변호사시험 정답률 74%로 합격점 돌파](#). Access Date: 2025-09-27.
- SKT-AI. 2025. [A.x-4.0 readme \(english\)](#). Accessed: 2025-10-02.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025a. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.
- Kanana LLM Team, Yunju Bak, Hojin Lee, Minho Ryu, Jiyeon Ham, Seungjae Jung, Daniel Wontae Nam, Taegyeong Eo, Donghun Lee, Doohae Jung, Boseop Kim, Nayeon Kim, Jaesun Park, Hyunho Kim, Hyunwoong Ko, Changmin Lee, Kyoung-Woon On, Seulye Baeg, Junrae Cho, and 10 others. 2025b. [Kanana: Compute-efficient bilingual language models](#). *Preprint*, arXiv:2502.18934.
- Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijie Chen, Yanru Chen, Yuankun Chen, Yutian Chen, Zhuofu Chen, Jialei Cui, Hao Ding, Mengnan Dong, Angang Du, Chenzhuang Du, Dikang Du, Yulun Du, Yu Fan, and 150 others. 2025c. [Kimi k2: Open agentic intelligence](#). *Preprint*, arXiv:2507.20534.
- NAVER Cloud HyperCLOVA X Team. 2025. [Hyperclova x think technical report](#). *Preprint*, arXiv:2506.22403.
- Kushal Tirumala, Aram H. Markosyan, Luke Zettlemoyer, and Armen Aghajanyan. 2022. Memorization without overfitting: analyzing the training dynamics of large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Lucia Zheng, Neel Guha, Javokhir Arifov, Sarah Zhang, Michal Skreta, Christopher D. Manning, Peter Henderson, and Daniel E. Ho. 2025. [A reasoning-focused legal retrieval benchmark](#). In *Proceedings of the 2025 Symposium on Computer Science and Law, CSLAW '25*, page 169–193, New York, NY, USA. Association for Computing Machinery.

A Evaluation

A.1 Reliability of the LLM-as-a-Judge in KCL-Essay

As a preliminary consistency check, we first scored the expert solution against the finalized rubrics. When evaluating the expert solution as the ground truth against the finalized rubric set, we obtained a mean score of $99.2 \pm 0.1\%$ over three independent runs, confirming near-perfect alignment between the rubric and the ground truth answers.

To further assess reliability, we compared the evaluation results of Gemini with those of human experts. We randomly sampled 641 question-prediction-rubric triplets, corresponding to 41 question-model pairs from the full results in Table 3. Two independent groups of three licensed Korean attorneys each scored all samples independently. As shown in Figure 2, we obtained an overall Pearson correlation of about 0.8 (top of the Figure 2) between Gemini 2.5 Flash and human evaluators.

We further clustered the questions into two groups according to their original bar exam scores, which reflect difficulty levels. First, to determine the cutoff for the top 20%, we examined the score distribution of whole KCL-Essay questions, and found that a score of 20 served as the threshold (Fig. 3 in Appendix). Using this threshold, we classified the 41 question-model pairs with human evaluations into two categories: those with scores up to 20 were labeled as Easy, and those above 20 as Hard. A clear difference emerged: for easier questions (bottom 80%), both the inter-human correlations 0.93 and the human-LLM correlations (0.88) were high. For more difficult questions, however, the correlations dropped substantially: the Gemini-human Pearson correlation was around 0.7, while the inter-human correlation also decreased to about 0.96. These findings suggest that for harder questions, the use of LLM-as-a-Judge requires more careful consideration.

B LICENSE

B.1 Model

GPT-5, GPT-5 Mini (OpenAI, 2024a) and GPT 4.1 (OpenAI, 2025a), o3 (OpenAI, 2025b), GPT-4o and GPT-4o-mini (OpenAI et al., 2024) were used in compliance with the OpenAI Terms of Use (OpenAI, 2024b). Gemini 2.5 Pro (DeepMind, 2025b), Gemini 2.5 Flash (DeepMind, 2025a) were

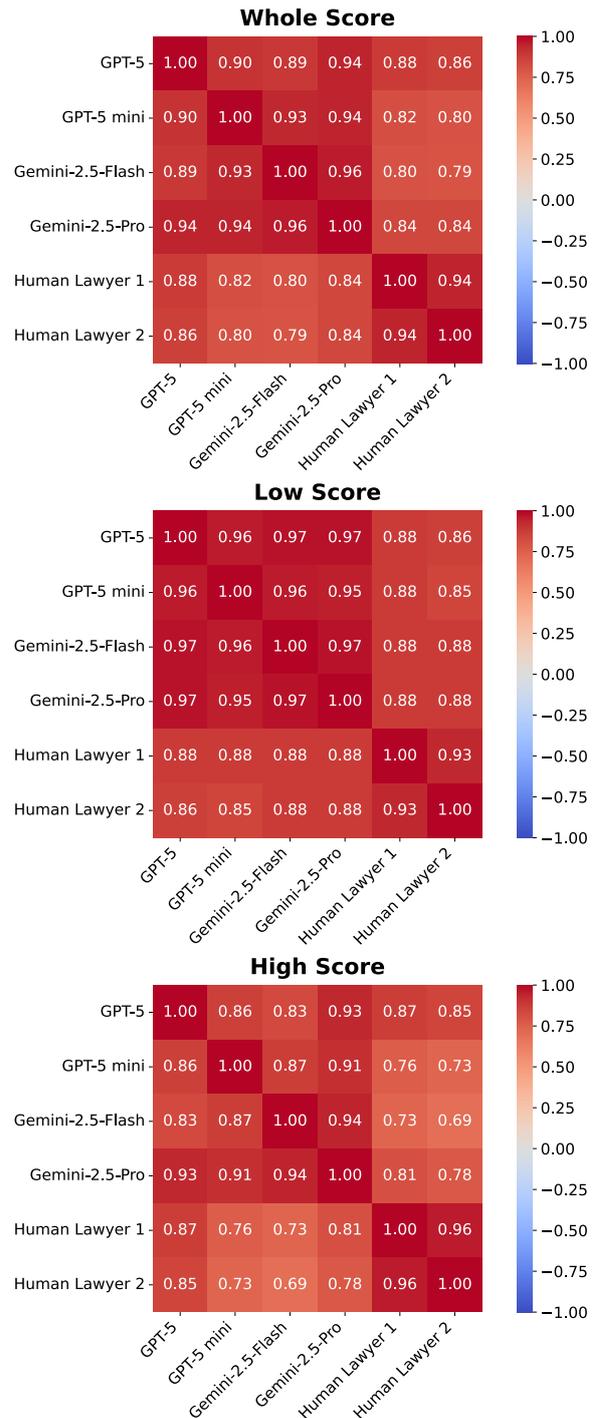


Figure 2: Pearson correlation heatmaps between LLMs and human expert evaluators. Whole Score represents the correlation over all question-prediction pairs derived from 641 triplets, while Low Score corresponds to the correlation for easy questions (bottom 80%), and High Score corresponds to the correlation for difficult questions (top 20%).

used in compliance with the Google APIs Terms of Service (Google, 2025b). Claude Opus 4.1 (Anthropic, 2025b), Claude Sonnet 4.5 (Anthropic, 2025d), Claude Sonnet 4 (Anthropic, 2025c),

Claude 3.7 Sonnet (Anthropic, 2025a) were used in compliance with the Anthropic on Bedrock - Commercial Terms of Service (Anthropic, 2023). DeepSeek-V3 (DeepSeek-AI et al., 2025) is available under DEEPSEEK LICENSE (DeepSeek, 2023). DeepSeek-R1 (Guo et al., 2025) and Phi-4-mini (Abdin et al., 2024) are available under MIT License. Qwen3 235B Thinking, Qwen3 235B, Qwen3 32B Instruct (Yang et al., 2025), gpt-oss 120B (OpenAI et al., 2025), and Kanana 1.5 8B (Team et al., 2025b) are available under Apache License Version 2.0. Exaone 4.0 32B (Research et al., 2025) is available under EXAONE AI Model License Agreement 1.2 - NC. A.X 4.0 (SKT-AI, 2025) is under Qwen LICENSE. Kimi K2 Instruct (Team et al., 2025c) is available under Modified MIT LICENSE (AI, 2025). Llama 4 Maverick (Meta, 2025b) is available under LLAMA 4 COMMUNITY LICENSE (Meta, 2025a). Gemma 3 27B (Team et al., 2025a) is available under the Gemma Terms of Use (Google, 2025a) HCX SEED Think 14B (Team, 2025) is available under the HyperCLOVA X SEED 14B Think Model License Agreement (Corporation and Corporation, 2025). Aya Expanse 32B (Dang et al., 2024) is available under the Creative Commons Attribution 4.0 License (CC BY 4.0).

We used the models in accordance with their respective licenses.

B.2 Data

The Korean Bar Exam is produced by the Ministry of Justice of Korea and distributed under Type 1 of the Korean Open Government License (KOGL), which permits both commercial and non-commercial use provided that the source is credited. Court judgments in Korea are classified as works not subject to copyright protection under Article 7(3) of the Korean Copyright Act and are therefore freely available for use.

B.3 KCL

We release our KCL data and evaluation code under the CC BY-NC 4.0 license.

C Details of Model Usage

In this section, we clarify the versions and configurations of the models referenced in the main text. When using Gemini 2.5 Pro and Gemini 2.5 Flash, we utilized Vertex AI’s gemini-2.5-pro and gemini-2.5-flash, respectively, with dynamic thinking en-

abled. The same configuration was used when employing Gemini 2.5 Flash as an LLM-as-a-Judge. The GPT-5 mentioned in the main text refers to gpt-5-2025-08-07, GPT-5 Mini refers to gpt-5-mini-2025-08-07, o4-mini refers to o4-mini-2025-04-16, and o3 refers to o3-2025-04-16, GPT-4o refers to gpt-4o-2024-08-06, all accessed through the OpenAI API, with reasoning effort set to medium. GPT-4.1 and GPT-4.1-mini were also accessed through the OpenAI API, specifically gpt-4.1-2025-04-14 and gpt-4.1-mini-2025-04-14, respectively. Claude Opus 4.1, Claude Sonnet 4.5, Claude Sonnet 4, and Claude 3.7 Sonnet were accessed via Amazon Bedrock using us.anthropic.claude-opus-4-1-20250805-v1:0, us.anthropic.claude-sonnet-4-5-20250929-v1:0, us.anthropic.claude-sonnet-4-20250514-v1:0, and us.anthropic.claude-3-7-sonnet-20250219-v1:0, us.anthropic.claude-3-5-sonnet-20240620-v1:0, us.anthropic.claude-3-5-haiku-20241022-v1:0 respectively, with the thinking budget set to 8k when using reasoning mode. DeepSeek-R1 and Llama 4 Maverick were also accessed via Amazon Bedrock using us.deepseek.r1-v1:0 and us.meta.llama4-maverick-17b-instruct-v1:0, respectively. All other open-weight models were run using vLLM, and GPU usage amounted to approximately 100 H100 hours in total on the cloud.

D Labeling Detail

D.1 Rubric Extraction Prompt

Table 4 shows the prompt used for rubric extraction in English (translated). We utilized Gemini 2.5 Flash to extract the rubrics, and expert review was completed as described in the main text.

D.2 Annotator Instructions

We collaborated with licensed attorneys using Google Sheets, and provided separate instructions for rubric review and model output evaluation, as depicted in Table 5 and Table 6 in English (translated).

D.3 Annotator Cost

In this section, we discuss the cost associated with expert labeling used in the construction of our benchmark. We engaged licensed Korean lawyers under consulting contracts to conduct rubric reviews and output evaluations. A total of six attorneys participated, and they were compensated at a rate of more than \$200 per hour.

E Additional Stats on KCL

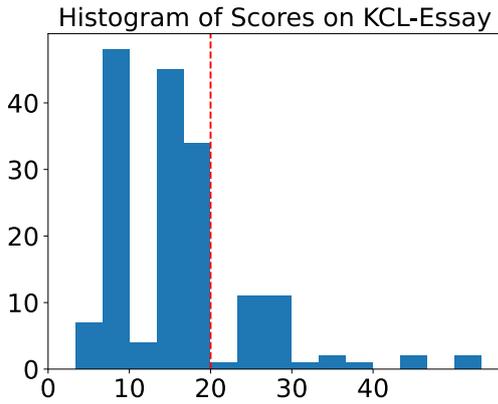


Figure 3: Histogram of score distribution across all 169 questions in KCL-Essay. The cutoff for the top 20% is observed at a score of 20.

As shown in Figure 3, the distribution of question scores in KCL-Essay is highly skewed, with most questions concentrated below a score of 20. To distinguish between relatively easy and more challenging questions in our correlation analysis, we set the top 20% of questions by score (cutoff 20) as the hard subset. This stratification provides a principled way to examine whether the reliability of LLM-as-a-Judge differs depending on question difficulty, rather than relying on an arbitrary threshold.

Rubric Extract Prompt

Please create [rubrics] for grading the [submitted answer sheet] based on the given [model answer] from the Korean Bar Exam commentary.

The generated [rubrics] must allow checking the following points in the [submitted answer sheet]: 1. Whether all key issues are addressed (Issue). 2. Whether the relevant legal principles, major statutes, or conclusions from key precedents are cited (Rule). 3. Whether the legal principles, statutes, and major precedents are properly applied to the given case (Application) and explained logically. 4. Whether the same conclusion (Conclusion) as in the [model answer] is reached.

Please avoid creating grading criteria that only check for formal compliance without substantive content. - **Bad rubric example**: ...(omitted)... - **Good rubric example**: ...(omitted)...

Also, when citing key precedents, do not include the exact case number in the rubric. However, it is important to require mentioning the relevant content. - **Bad rubric example**: ...(omitted)... - **Good rubric example**: ...(omitted)...

Please refer to the following transformation example when writing:

[Example Question]: ...(omitted)...
 [Example Rubric]: ...(omitted)...
 [Question]: ...(omitted)... [GT Answer]: ...(omitted)...

Table 4: Rubric Extraction Prompt

Rubric Evaluation Instruction

Based on the review sheet, please evaluate the appropriateness of the evaluation rubric generated by AI, propose additional rubric if necessary, and provide an overall opinion on their quality.

1. Assessment of Individual Evaluation Criteria (O/X)
 - The sheet [XX Evaluation Criteria Review] contains, for each question from the Korean Bar Exam case questions (Public Law, Criminal Law, Civil Law), the following: “Question Number,” “Question,” “Model Answer (citing Union and Rainbow textbooks),” “Referenced Cases (cases cited in the model answer),” and the evaluation criteria generated by AI.
 - Please determine whether each evaluation rubric is concrete and clear. If so, mark with an ****O****; if not, provide a reason. If the criterion itself is inappropriate, also provide the reason.
2. Proposal of Additional Evaluation Rubric - While reviewing the evaluation criteria in the “Evaluation Criteria Review” sheet, please check if any important issues are missing for each question. If something is missing, please add it in the “Additional Evaluation Criteria” column.
3. Overall Opinion on the Quality of Evaluation Criteria - In the “Overall Opinion” sheet, please provide a general qualitative assessment of how good the evaluation criteria generated by judge AI are. A qualitative review is required rather than numerical scoring.

Table 5: Rubric Evaluation Instruction

Model Output Evaluation Instruction

Each sheet contains the following columns: question (the exam problem), pred (the AI’s answer), and rubric (the evaluation criteria). The problems are drawn from actual Korean bar exam case questions. The AI answers are provided by various models (both strong and weak). The Rubric is designed to resemble an official scoring guideline. The evaluator’s task is to read the problem and the AI’s answer, then decide whether the AI’s answer meets the Rubric. The evaluation should be binary (0 = does not satisfy, 1 = satisfies). The Google Sheet contains two identical sheets. Each sheet must be completed by a different attorney. In other words, an attorney who evaluates Sheet 1 should not participate in evaluating Sheet 2.

Table 6: Model Output Evaluation Instruction