

CLlck: A Benchmark Dataset of Cultural and Linguistic Intelligence in Korean

Eunsu Kim[◊], Juyoung Suk[†], Philhoon Oh[†], Haneul Yoo[◊], James Thorne[†], Alice Oh[◊]

[◊]School of Computing, KAIST, Dajeon, South Korea

[†]GSAI, KAIST, Seoul, South Korea

{kes0317, scottosuk0306, philhoonoh, haneul.yoo, thorne}@kaist.ac.kr , alice.oh@kaist.edu

Abstract

Despite the rapid development of large language models (LLMs) for the Korean language, there remains an obvious lack of benchmark datasets that test the requisite Korean cultural and linguistic knowledge. Because many existing Korean benchmark datasets are derived from the English counterparts through translation, they often overlook the different cultural contexts. For the few benchmark datasets that are sourced from Korean data capturing cultural knowledge, only narrow tasks such as bias and hate speech detection are offered. To address this gap, we introduce a benchmark of Cultural and Linguistic Intelligence in Korean (CLlck), a dataset comprising 1,995 QA pairs. CLlck sources its data from official Korean exams and textbooks, partitioning the questions into eleven categories under the two main categories of language and culture. For each instance in CLlck, we provide fine-grained annotation of which cultural and linguistic knowledge is required to answer the question correctly. Using CLlck, we test 13 language models to assess their performance. Our evaluation uncovers insights into their performances across the categories, as well as the diverse factors affecting their comprehension. CLlck offers the first large-scale comprehensive Korean-centric analysis of LLMs' proficiency in Korean culture and language. CLlck is publicly available at: <https://github.com/rladmstn1714/CLlck>.

Keywords: Evaluation, Benchmark, Korean, Culture

1. Introduction

Recent advancements in Large Language Models (LLMs) have been significant, particularly for a small group of high-resource languages including English. In these languages, LLMs frequently attain or surpass human-level proficiency in numerous Natural Language Processing (NLP) tasks that necessitate comprehension of everyday life and the subtleties of linguistic nuances. However, despite a concerted effort in developing Korean large-scale language models, there remains a significant performance gap for benchmark tasks in the Korean language. For example, KoGPT3-39B underperforms on the Korean HellaSwag task (Jang et al., 2022) by 20%, compared to similar scale English Falcon40B model (Almazrouei et al., 2023) in the original English version of the task (Zellers et al., 2019) even though human annotators can attain the same performance. Instances that contain cultural and linguistic knowledge that deviate from English and other well-represented languages are often incorrectly answered by models.

Current Korean evaluation datasets for LLMs show significant limitations for comprehensive assessment. Existing tasks are either too simple (Ham et al., 2020a; Park et al., 2021) or mainly derived from English benchmarks¹, failing to capture the key aspects of the Korean language or

¹<https://huggingface.co/spaces/upstage/open-ko-llm-leaderboard>

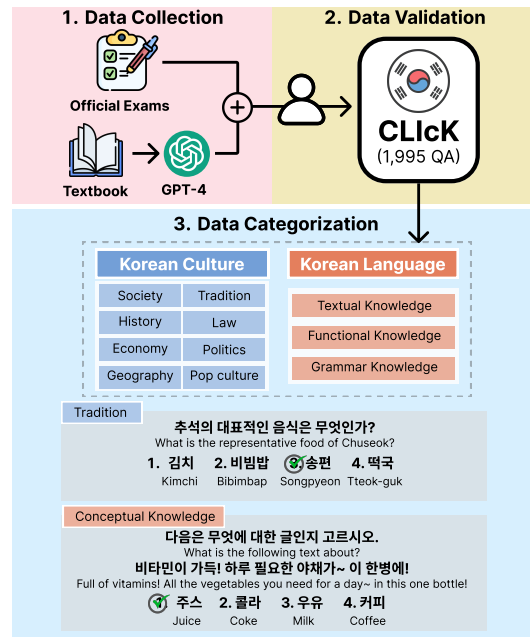


Figure 1: Overview of the CLlck dataset curation and categorization process. Data is sourced from official exams and textbooks and validated by authors. The dataset is categorized into Korean Culture and Korean Language, further divided into 11 sub-categories.

culture. While several Korean-centric datasets have been introduced, these typically target specific tasks such as bias and hate speech detection (Jin et al., 2023; Jeong et al., 2022) and hence

cannot support general LLM evaluation.

To bridge this resource gap, we construct and release CLlck, a culturally-aware evaluation benchmark dataset encompassing 1,995 instances across 11 categories representing facets of the Korean culture, ranging from everyday life to specific subject areas, as well as Korean grammar and linguistics. To ensure high quality, we collect and curate all samples from official examinations and textbooks, which are then rigorously inspected and categorized by four native speakers of Korean.

We subsequently evaluate five families of open-source LLMs with different parameter sizes and two proprietary LLMs with CLlck. The open-source models exhibit especially low accuracies, ranging between 10% and 50%. Meanwhile, proprietary LLMs like GPT-3.5 and Claude-2 outperform these but still perform poorly in some categories. Notably, compared to the general population of Korean test-takers, GPT-3.5 scores in the lowest 11th percentile, which contrasts with its achievement in the top 13th percentile on the English Scholastic Aptitude Test (SAT).

The primary contributions of our work include:

- We construct and publicly release CLlck, a benchmark dataset to evaluate LLMs' cultural and linguistic understanding of Korean.
- We provide a fine-grained categorization of the requisite knowledge to answer each query in the dataset.
- We empirically evaluate 13 model configurations on CLlck, demonstrating the limitations of LLMs and motivating further research on cultural and linguistic benchmarks.

2. Related Work

2.1. General English Benchmarks

The General Language Understanding Evaluation (GLUE) and SuperGLUE benchmarks were established to evaluate models in a wide range of tasks such as sentiment analysis, natural language inference, and question-answering (Wang et al., 2018, 2020). HellaSwag (Zellers et al., 2019) and CosmoQA (Huang et al., 2019) further include commonsense reasoning for a more robust evaluation. However, due to the rapid advancement in AI communities, achieving human-level state-of-the-art performances on these datasets has become increasingly common (Martínez-Plumed et al., 2021). To properly evaluate the capabilities of models in the era of LLMs, more challenging benchmarks have been introduced. For example, MMLU (Hendrycks et al., 2021) and

AGIEval (Zhong et al., 2023) consist of questions originally designed for humans, while BIG-bench (Srivastava et al., 2023) aims to cover diverse topics, comprising 204 tasks to address the limitations in LLMs.

2.2. Multilingual and Commonsense

To investigate how LMs can comprehend or generate text in other languages, there have been efforts to construct multilingual datasets. For instance, XGLUE (Liang et al., 2020) encompasses 100 languages that can be employed for both pre-training and evaluating cross-lingual tasks. XTREME (Hu et al., 2020) introduces an evaluation framework for cross-lingual benchmarks, while MEGA (Ahuja et al., 2023) focuses on assessing LLMs, providing 16 NLP datasets ranging from low-resource to high-resource languages. In addition, datasets that primarily focus on specific target languages, such as Chinese, Indian, and African languages, have been introduced (Huang et al., 2023; Dodapaneni et al., 2023; Adebara et al., 2023).

The popularity of commonsense datasets has increased because they reflect a wide array of sociocultural knowledge shared by humans (Liu and Singh, 2004). These datasets incorporate everyday concepts, such as CommonsenseQA (Talmor et al., 2019), scientific knowledge like ARC (Clark et al., 2018), and simple arithmetic reasoning like GSM8K (Cobbe et al., 2021). These datasets can be seen as a representation of general and practical knowledge that aligns with human intentions. Consequently, certain datasets incorporate or employ translated portions from English datasets (Seo et al., 2022), potentially overlooking subtle linguistic or cultural differences that may not be apparent to the audience (Tandon et al., 2018).

Lee et al. 2023a demonstrated that language models fail to capture biases in different languages due to their cultural insensitivity, which can have societal impacts (Tamkin et al., 2021). Furthermore, Ma et al. 2022 emphasized the importance of cultural background and showed that integrating cultural knowledge can improve models performance. These findings illustrate the need for cultural evaluation datasets. However, building a cultural evaluation dataset from scratch is challenging since it entails significant time and resources while relying on translated datasets fails to incorporate cultural knowledge in different languages.

2.3. Korean Datasets

Previous benchmarks for Korean language models focused on specific tasks such as paraphrase detection, natural language inference (NLI), machine reading comprehension (MRC), and hate speech detection (Yang et al., 2019; Ham et al.,

2020b; Lim et al., 2019; Moon et al., 2020). The Korean Language Understanding Evaluation (KLUE) benchmark (Park et al., 2021), similar to GLUE, introduced eight downstream tasks for the Korean language. However, these benchmarks lacked tasks for advanced reasoning, which is inadequate for LLM evaluation. Recently published Open Ko-LLM LeaderBoard¹ aimed to address this issue; however, its closed-source nature and reliance on translations may not fully represent the nuances of the Korean language context.

Contemporary benchmarks prioritize preserving linguistic and cultural nuances of the target language in translation. For instance, Jin et al. (2023) introduced the Korean Bias Benchmark for Question Answering (KoBBQ) based on the original BBQ dataset (Parrish et al., 2022). KoBBQ first classified the translations into 3 distinct categories: SIMPLY-TRANSLATED, where translations are appropriate for the Korean knowledge, SAMPLE-REMOVED where translated sentences are removed due to their lack of relevance to Korean cultural context, and TARGET-MODIFIED where target translations are adjusted to align with Korean culture background. This classification adjusts knowledge differences available in languages and reflects subtle cultural nuances in Korean.

Recent efforts have revolved around creating authentic Korean datasets from scratch. For instance, the Korean Offensive Language Dataset (Jeong et al., 2022) gathers toxic sentences from news articles and YouTube platforms, while the Korean Balanced Evaluation of Significant Tasks (Jang et al., 2022) is entirely annotated by humans for five distinct tasks. The HAE-RAE Benchmark (Son et al., 2023) provides Korean Reading Comprehension datasets sourced from the original Korean Corpus. However, these datasets are constrained to specific tasks and may not be suitable for evaluating diverse topics related to Korean cultural and linguistic knowledge. With the rapid growth of the language model’s capacity, there is a growing need for fine-grained evaluation to assess the cultural and commonsense knowledge within LLMs (Ye et al., 2023).

In this paper, we introduce CLlck comprising 1,995 questions that require a wide range of Korean linguistic and cultural knowledge along with reasoning capacity, which is close to real-world settings. Additionally, it directly sources content from the original Korean Corpus, including six different Korean native exams, ensuring cultural authenticity and facilitating comparisons between models and human-level scores. Lastly, CLlck provides fine-grained evaluations of LLMs on eleven diverse topics, thereby contributing to further research on the assessment of Korean cultural and linguistic knowledge within LLMs.

3. CLlck Dataset

CLlck contains 1,995 QA pairs, organized into two main categories and 11 subcategories of multiple-choice QA about Korean facts. Dataset statistics are reported in Tables 1 and 2. The dataset is constructed in three stages: (1) Data Collection (§ 3.1), (2) Data Validation (§ 3.2), and (3) Data Categorization (§ 3.3), summarized in Figure 1.

Category		# of Samples		
		Textbook	Exams	Total
Korean Culture	Society	284	25	309
	Tradition	161	61	222
	History	0	280	280
	Law	51	168	219
	Politics	79	5	84
	Economy	57	2	59
	Geography	39	92	131
	Pop culture	15	26	41
Korean Language	Textual	0	285	285
	Functional	0	133	133
	Grammar	0	232	232
Total		1245	750	1995

Table 1: Statistics of CLlck per categories. ‘From Textbook’ denotes data from the KIIP textbook; ‘From Exams’ refers to data from official exams. ‘Number of samples’ indicates unique QA pairs in the dataset.

Code	Subject	# of Samples
CSAT	Language	226
	Geography	30
TOPIK	Language	237
PSE	Language	14
	History	189
PSAT	Constitution	168
KHB	History	47
Kedu	Language	173
	Culture	161

Table 2: Statistics of CLlck per Exams. Code and Subject refer to the exam’s codes and its utilized subjects.

3.1. Data Collection

To collect data, we employ two approaches: 1) following the AGIEval dataset (Zhong et al., 2023) to select questions from standardized Korean exams, and 2) using GPT-4 to generate new questions based on textbooks. In this process, we utilize six exams and one textbook related to Korean language and cultural knowledge as sources (the

Source	Code	Type	Subject
Korean Immigration and Integration Program	KIIP	Textbook	Culture
College Scholastic Ability Test of Korea	CSAT	Exam	Language Geography
Test of Proficiency in Korean	TOPIK	Exam	Language
National Public Service Examination - Grade 9	PSE	Exam	Language History
Public Service Aptitude Test	PSAT	Exam	Constitution
Korean History Exam-Basic	KHB	Exam	History
Test of Teaching Korean as a Foreign Language	Kedu	Exam	Language Culture

Table 3: Overview of data sources used in the CLICK dataset, detailing each source’s code, type (textbook or exam), and covered subjects related to Korean language and culture.

genres of these resources are summarized in Table 3).

Selecting Questions from Exams We obtain test data from six Korean examinations and receive permission from the relevant institutions. Descriptions for each examination can be found in Appendix A. We use Clova OCR² to extract text from the exams, excluding images and tables.

Generating Questions Using GPT-4 We use this approach to introduce novel cultural questions, which are not covered in the exams. Building on established practices in question generation (Zhou et al., 2017; Kurdi et al., 2019), we feed GPT-4 the full text of each chapter from the KIIP textbook, prompting it to produce multiple-choice questions, their corresponding choices and answers based strictly on the book’s content. Each question was verified before inclusion in the dataset. More detailed information on GPT-4 question generation, including used prompts and procedures, can be found in Appendix B.

3.2. Data Validation

We scrutinize transcribed exams for OCR errors and validate the GPT-4 generated questions according to the following criteria:

1. Questions are solely based on the given text.
2. Information in Questions remains consistent over time.
3. Questions should centrally relate to Korea.
4. Questions should be objective and free from bias.

By checking the first criterion, we ensure that questions originate exclusively from the textbook’s content, eliminating any influence of GPT-4’s internal knowledge. Additionally, the fourth criterion helps

²<https://clova.ai/ocr/>

maintain dataset objectivity by excluding instances connected to subjective beliefs or biases, like gender, politics, or international relations.

Human Annotation Four of the authors, who are Korean native speakers, validated to ensure its relevance and quality. Validation checks are categorized into three types; valid, needs modification, and invalid, with the following process:

1. Initially, three annotators reviewed each sample. If two or more annotators considered a sample invalid, it was discarded. 15.9% of the data was labeled invalid by one annotator. Samples that needed modification were revised by one of the annotators.
2. For the remaining invalid and modified samples, a second round of annotation was conducted by three annotators. After this phase, 3.9% of the initial set still had discrepancies.
3. The four annotators involved in the previous steps discuss the disagreements. Only samples with unanimous agreement between all four annotators are included in the dataset.

Following the validation process, the initial dataset, which consisted of 1,985 samples, was reduced to 1,245 samples, accounting for 62.9% of the original dataset.

3.3. Data Categorization

For each instance, we provide fine-grained annotation of which aspects of Cultural and Linguistic intelligence are required to answer the question. Summary statistics are presented in Table 1.

Cultural Intelligence We adopt eight subcategories based on the KIIP textbook. The primary chapters of the KIIP basic textbook encompass *Society, Culture, Politics, Economy, Law, History, and Geography*. Within the *Culture* chapter, there are subsections on *Tradition* and *Pop Culture*. We label each instance with respect to its cultural knowledge with one of the following labels: *Society, Tradition, Pop Culture, Politics, Economy, Law, History, and Geography*.

Linguistic Intelligence We follow definitions of linguistic knowledge from Bachman and Palmer (1996). Specifically, we annotate instances for *Textual Knowledge*, which concerns organizing utterances into coherent texts with cohesion and rhetorical structures; *Functional Knowledge*, focusing on the communicative roles of language, especially ideational, manipulative, heuristic, and imaginative functions; and *Grammatical Knowledge*, addressing the organization of utterances with an

Type	Model	
API-based LLMs	GPT-3.5-turbo, Claude-2	
Open-source LMs	Multilingual	Korean-specialized
	LLaMA2-chat (7B,13B)	LLaMA2-Ko(7B) ³ KULLM-v2 (Lee et al., 2023b) KoAlpaca ⁴ Polyglot-Ko (Ko et al., 2023) (1.3B,3.8B,5.8B,12.8B)

Table 4: Model Selection for Our Experiment. The numbers in parentheses indicate the parameters for models.

emphasis on vocabulary, syntax, and phonology/graphology. We exclude the *socio-linguistic knowledge* category in our work because it is largely subsumed by our annotations in the Cultural Intelligence category.

Specific Procedures Because questions generated from the textbook already align to Cultural Intelligence categories, we do not require further annotation for these instances. Similarly, for exams that focus on a single subject (e.g. *History*), no additional categorization is performed. For the CSAT-Korean exam, problems are categorized into speaking, writing, language, media, literature, and reading. In terms of our defined categories, speaking and writing correspond to *Functional knowledge*, language aligns with *Grammar knowledge*, and both literature and reading are associated with *Textual knowledge*. Other Korean Language exams offer solutions detailing the problem’s category. Based on this information, a single annotator validates the label assignment.

4. Evaluation

We evaluate the capabilities of established language models using our CLiCK dataset. We compare a range of models which have a variety of exposure to the Korean language (detailed in Table 4). For API-based LLMs, we conduct experiments between September and October of 2023.

Prompt Types Our prompts are derived from Jin et al. (2023). Following Izacard et al. (2023), we apply cyclic permutation for each question to mitigate option’s order effects in the prompt to the language model. We report the average over three different wordings of the prompt. Depending on whether the question requires the model to read background information, we prompt the model with context (type 1) or without context (type 2) according to the examples which are provided below.

Type 1: Sample with Context

주어진 맥락을 천천히 읽고, 질문에 대한 적절한 정답을 A, B, C, D 중에 골라 알파벳 하나로 답하십시오.
(Read the given context, and choose the correct answer to the question from options A, B, C, or D. Respond with a single alphabet.)

맥락 (Context): {CONTEXT}
질문 (Question): {QUESTION}
보기 (Options):
A: {A}, B: {B}, C: {C}, D: {D}
정답 (Answer):

Type 2: Sample without Context

주어진 질문을 천천히 읽고, 적절한 정답을 A, B, C, D 중에 골라 알파벳 하나로 답하십시오.
(Read the given Question, and choose the correct answer from options A, B, C, or D. Respond with a single alphabet.)

질문 (Question): {QUESTION}
보기 (Options):
A: {A}, B: {B}, C: {C}, D: {D}
정답 (Answer):

Evaluation Methodology We adopt the evaluation methodology from MMLU, which also aligns with prevalent LLM evaluation frameworks such as EleutherAI lm-harness⁵ and OpenAI Evals⁶. For open-source models, we examine the output probabilities of option ID tokens (A/B/C/D or A/B/C/D/E) concatenated with the option string, selecting the most probable answer as the model prediction. For API-based LLMs (GPT-3.5-turbo and Claude-2), the evaluation involves comparing the generated response with the labeled answer. Here, the decoding temperature is set to 0. Though our prompt directly asks the model to output only the option ID, we notice that these models may at times produce verbose responses. Therefore, we adopt the acceptance criteria from Jin et al. (2023), allowing answers that: i) mention only one alphabet from the given options, ii) exactly match a term provided in the options, iii) include specific expressions clearly intended to convey the answer, such as ‘the answer is -’, or iv) present the answer distinctly as per conditions i) to iii), followed by further explanation. Responses that don’t meet these conditions are considered as out-of-option answers.

Evaluation Metric Our primary evaluation metric is accuracy. We report the average accuracy over the entire dataset. As we prompt the model 3 times and adopt cyclic permutation for each instance, the total number of experiments per instance is $3N$, where N denotes the “number of

³<https://huggingface.co/beomi/llama-2-ko-7b>

⁴<https://github.com/Beomi/KoAlpaca>

⁵<https://github.com/EleutherAI/lm-evaluation-harness>

⁶<https://github.com/openai/evals>

		Polyglot-Ko				KULLM		KoAlpaca		LLaMA-Ko		LLaMA		GPT-3.5	Claude2
		1.3B	3.8B	5.8B	12.8B	5.8B	12.8B	5.8B	12.8B	7B	7B	13B			
Korean Culture	History	26.30	24.71	25.52	24.43	26.48	25.07	26.05	25.84	26.38	30.75	30.73	31.32	35.00	
	Geography	30.18	28.72	29.06	30.12	27.21	28.66	28.53	30.01	33.21	23.10	25.20	45.42	43.30	
	Law	38.44	40.16	40.70	43.44	41.67	41.90	40.67	40.13	40.02	45.13	44.12	55.31	57.09	
	Politics	30.53	32.00	27.74	27.15	26.96	22.68	23.42	28.79	36.03	27.31	26.43	47.75	60.89	
	Society	34.69	34.31	35.95	37.37	35.95	37.37	33.33	36.44	32.10	39.48	40.93	60.48	62.43	
	Tradition	32.48	34.01	34.97	33.96	35.86	34.63	32.80	35.45	33.60	33.88	36.12	50.16	52.10	
	Economy	42.54	42.62	42.25	45.03	43.86	44.08	43.35	43.79	45.32	45.83	46.27	47.59	53.62	
	Pop culture	29.77	33.68	32.64	29.59	33.60	32.76	34.02	32.63	27.20	33.45	36.41	68.61	59.56	
Average	32.71	32.90	33.14	33.40	33.79	33.51	32.33	33.80	33.26	35.44	36.22	49.30	51.72		
Korean Language	Textual	23.44	23.57	23.27	22.96	24.52	24.65	23.07	24.19	26.75	24.73	24.29	53.19	55.86	
	Functional	23.77	21.67	22.64	19.84	20.06	19.38	24.76	20.50	26.31	27.04	30.50	32.62	32.88	
	Grammar	21.87	21.79	23.64	23.05	24.69	25.67	24.03	22.05	23.04	29.32	26.52	38.85	43.95	
	Average	22.88	22.38	23.27	22.24	23.50	23.78	23.87	22.42	25.69	27.17	26.71	42.32	45.39	

Table 5: Accuracy of the models by category. The highest accuracy for each category is in bold. The top-performing open-source models are marked in blue.

Factor	F	p
Model Scale	0.33	.57
Korean corpus	0.30	.59

Table 6: Results from the ANOVA Test considering two factors. The columns ‘ F ’ and ‘ p ’ represent the F -value and p -value, respectively. ‘Korean corpus’ refers to the supplemental Korean dataset used during training.

options”. Instance accuracy is computed as:

$$p_{\text{accuracy}} = \frac{\text{count}(\text{correct answers})}{3N} \quad (1)$$

4.1. Experimental Results

We present comprehensive results of evaluating 13 LLMs with CLiCk in Table 5. We find that open-source models with fewer than 13B parameters generally exhibit a low accuracy in the range of 10-50%. While Claude-2 and GPT-3.5 surpass those smaller models in most categories, their performance remains similar in *History*, *Economy*, and *Functional Knowledge*. Claude-2 outperforms GPT-3.5 in all categories except *Geography* and *Pop Culture*.

Model Scale We use the ANOVA test on Polyglot-Ko, KULLM, KoAlpaca, LLaMA models to assess sensitivity to number of parameters. Test statistics are reported in Table 6. We find that model scale does not have a statistically significant impact on accuracy ($F_{1,43} = 0.33, p = .57$).

Korean Corpus Scale While the KULLM and KoAlpaca models are fine-tuned with additional Korean corpora based on the Polyglot-Ko models, and the LLaMa-Ko model is a Korean-specialized version of LLaMA, our findings in

Table 6 suggest that additional training on Korean datasets doesn’t significantly enhance their comprehension of Korean culture and language ($F_{1,54} = 0.30, p = .59$). In fact, the accuracy is even lower than that of their base models. It’s notable that llama-2 chat, despite deriving only 0.06% of its training data from Korean sources, outperforms most Korean-specialized models by a small margin.

Our results reveal that despite their extensive pretraining on vast datasets and large scales, GPT-3.5-turbo and Claude-2 perform similarly to open-source models in specific categories like *History*, *Economy*, and *Functional knowledge*. The results and analysis in this section give insights that simply amassing more data and enlarging the model size (currently common practices in LM) may not be the optimal solution for enhancing the cultural intelligence of Language Models in non-English languages.

5. Discussion

5.1. Difficulty Analysis

Overall Difficulty To measure the challenging problems within our dataset, we first define the difficulty by the number of problems with an accuracy below the random selection threshold ($1/N$), where N denotes the number of choices belonging to questions. If the accuracy of the models is below this threshold, we mark it as a challenging problem for a model and vice versa. As depicted in Figure 2, open-source models, such as Polyglot-Ko, KULLM, KoAlpaca, LLaMa2-Ko, and LLaMa2-chat, have difficulty with over 60% of our dataset, with a shared difficulty spanning 35.5%. LLMs like GPT-3.5 and Claude-2 face challenges with over 30%, sharing 12.6% of problems. Only in 0.6% of the data were none of the




Figure 2: Portions of challenging samples encountered by models of varying sizes. The sky-blue bar represents the shared portion of KULLM, KoAlpaca, and LLaMa2-chat, while the gray bar corresponds to GPT-3.5 and Claude-2.

models able to find the correct solution. This emphasizes the CLICK dataset’s inherent complexity.

Qualitative Study We analyze the tendencies of the models by examining all samples they perform well on and those they don’t. We observe no discernible trends that are specific to each category, and observe a wide range of performance across samples derived from the same category and source, even though they have similar levels of difficulty. For example, Problem 1 and Problem 2 are questions from the Korean society category, derived from KIIP textbook, and both inquire about Korean-specific matters.

Problem 1

질문 (Question): {한국 정부의 주택 정책이 아닌 것은? (What is not the Korean government’s housing policy?)}
 A: {국민임대주택 (National rental housing)}
 B: {공공임대주택 (Public rental housing)}
 C: {보금자리주택 (Bogumjari house/Nest house)}
 D: {단기경매주택 (Short auction house)}
 정답 (Answer): {단기경매주택 (Short auction house)}

Problem 2

질문 (Question): {남편이 아내의 오빠를 어떻게 부르는가? (How does a husband call his wife’s older brother?)}
 A: {처형 (Cheohyeong/Sister-in-law)}
 B: {처제 (Cheoje/Sister-in-law)}
 C: {처남 (Cheonam/Brother-in-law)}
 D: {형님 (Hyeongnim/Older brother-in-law)}
 정답 (Answer): {형님 (Hyeongnim/Older brother-in-law)}

Problem 2 asks about the everyday life of *Korean society*, whereas problem 1 seeks more professional information, making problem 1 appear more challenging. However, all 13 models achieved 100% accuracy on Problem 1, yet none correctly answered Problem 2. Therefore, 1) there is less alignment between the model’s perceived difficulty and those by humans, and 2) such cultural and linguistic intelligence contexts are challenging for the model to comprehend.




Figure 3: Box-and-whisker plot of uncertainty score of challenging samples across the models. ‘x’ marks the mean, and the horizontal bar represents the median. For Polyglot and LLaMa, consistency rises with increasing model scale. Polyglot 1.3B has a median near 1, while LLaMa’s median is closer to 0.

Why Do Models Struggle? As mentioned in § 4, we conducted a total of $3N$ experiments for each problem sample, utilizing three different prompts and through cyclic permutation. The model’s uncertainty on each sample, calculated using normalized Shannon entropy (Shannon, 1948), is defined as follows:

$$\text{Uncertainty score} = -\frac{1}{\log N} \sum_{i \in \text{options}} p_i \log p_i$$

$$\text{where } p_i = \frac{\text{count}(i)}{3N} \quad (2)$$

and i represents each option.

The score is normalized between 0 and 1, considering varying numbers of total options. A score nearer to 0 implies high consistency, whereas one near 1 indicates almost random selection. For each model’s challenging problem, as defined at the beginning of § 5.1, we calculate an uncertainty to analyze the reasons behind their low accuracy on our dataset.

The plot in Figure 3 illustrates the uncertainty scores of various models. Smaller models (polyglot 1.3B, 3.8B) tend to randomly select answers without consistency, and as the size increases (12.8B, 5.8B), they tend to choose wrong answers consistently. In other words, models are not getting our dataset wrong because they are unsure of the answer, but they choose specific wrong answers at a high rate. Upon analyzing GPT-3.5 and Claude-2, we observe an ambiguous pattern, as their mean and median uncertainty scores both approximate 0.5.

5.2. Comparisons to Human Level

We compare model accuracy to human performance on exams. Since our dataset does not encompass the actual score distribution for the problems, a simple score conversion, the ratio of correctly answered questions, is applied to facilitate a comparative analysis. We utilize exam statistics from CSAT Korean, TOPIK and Kedu for Korean Culture to assess the models' performances against actual exam takers.

CSAT Korean An average is taken of the statistics spanning from 2017 to 2020, a total of five years. The CSAT is divided into tiers from 1 (best) to 9 (worst). The `Polyglot-Ko` (12.8B), `KULLM` (5.8B, 12.8B), and `KoAlpaca` (12B) models performed at the 9th level, corresponding to the lowest 4% of Korean high school seniors (3rd-year students). The remaining models are at the 8th grade level, representing the lowest 11%.

TOPIK TOPIK is evaluated on an absolute scale, with levels ranging from 1 to 6, with higher being better. `Claude-2` achieved a level 6, implying it can perform language functions required for specialized research or professional tasks relatively accurately and fluently. While it doesn't reach the proficiency of a native speaker, it doesn't face difficulties in functional performance or expressing meanings. `GPT-3.5` achieved a level 5, indicating it can appropriately differentiate language use depending on formal and informal, as well as spoken and written contexts. The other models score too low and fall outside the measurement range.

Kedu for Korean Culture We analyze data over a five-year, from 2014 to 2018. All model results fall below the average of participants, which stood at 49.9% correct answers. `Claude-2` and `GPT-3.5` are lower than the average by about 10%, scoring 39.1% and 37.0%, respectively. Meanwhile, other models lagged by more than 20% compared to the average.

6. Conclusion

We introduced CLlck, A Benchmark Dataset of Cultural and Linguistic Intelligence in Korean. CLlck emerges as a uniquely Korean-centric dataset sourced from Korean examinations and textbooks. The dataset is categorized into two main category and 11 sub-categories, which enable fine-grained and Korean-centric evaluation. Through our analyses and experiments, we observed that five open-source models struggle with over 60% of the data. Proprietary LLMs outperform these models yet still require further im-

provement. Interestingly, simply scaling up the model or fine-tuning it with additional Korean corpora doesn't guarantee enhanced Korean linguistic and cultural knowledge of models. This implies that models find it challenging to understand non-English linguistic and cultural intelligence, highlighting the need for more tailored methods in further research.

Ethics Statement

This work presents CLlck dataset, a free and open evaluation benchmark of cultural and linguistic intelligence in Korean. Our dataset contains multiple-choice QAs with four or five choices. The data sources are publicly available exams and published articles from the Korean government, and we receive official permission from the related agencies. We scrutinize all samples in the dataset to ensure that no personally identifiable information or sensitive content is included. Furthermore, we have made our best effort to deliver Korean culture correctly by 1) generating samples from reliable grounds and 2) iteratively validating samples.

Acknowledgements

This project was funded by Institute of Information communications Technology Planning Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2022-0-00184, Development and Study of AI Technologies to Inexpensively Conform to Evolving Policy on Ethics).

Bibliographical References

- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Mérouane Debbah, Étienne Goffinet, Daniel Hesslow, Julien Launay, Quentin Malaric, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The falcon series of open language models](#).
- L.F. Bachman and A.S. Palmer. 1996. *Language Testing in Practice: Designing and Developing Useful Language Tests*. Language applied linguistic. OUP Oxford.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2023. [Atlas: Few-shot learning with retrieval augmented language models](#). *Journal of Machine Learning Research*, 24(251):1–43.

- Hyunwoong Ko, Kichang Yang, Minhø Ryu, Taekyoon Choi, Seungmu Yang, Jiwung Hyun, Sungho Park, and Kyubyong Park. 2023. [A technical report for polyglot-ko: Open-source large-scale korean language models](#).
- Ghader Kurdi, Jared Leo, Bijan Parsia, Uli Sattler, and Salam Al-Emari. 2019. [A systematic review of automatic question generation for educational purposes](#). *International Journal of Artificial Intelligence in Education*, 30:121–204.
- Nayeon Lee, Chani Jung, and Alice Oh. 2023a. [Hate speech classifiers are culturally insensitive](#). In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 35–46, Dubrovnik, Croatia. Association for Computational Linguistics.
- SeungJun Lee, Taemin Lee, Jeongwoo Lee, Yoona Jang, and Heuseok Lim. 2023b. [Kullm: Learning to construct korean instruction-following large language models](#). In *Annual Conference on Human and Language Technology*, pages 196–202. Human and Language Technology.
- Hugo Liu and Push Singh. 2004. [Conceptnet—a practical commonsense reasoning tool-kit](#). *BT technology journal*, 22.
- Weicheng Ma, Samiha Datta, Lili Wang, and Soroush Vosoughi. 2022. [EnCBP: A new benchmark dataset for finer-grained cultural background prediction in English](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2811–2823, Dublin, Ireland. Association for Computational Linguistics.
- Fernando Martínez-Plumed, Pablo Barredo, Seán Ó hÉigeartaigh, and José Hernández-Orallo. 2021. [Research community dynamics behind popular ai benchmarks](#). *Nature Machine Intelligence*, 3(7):581–589.
- Claude Elwood Shannon. 1948. [A mathematical theory of communication](#). *The Bell System Technical Journal*, 27:379–423.
- Alex Tamkin, Miles Brundage, Jack Clark, and Deep Ganguli. 2021. [Understanding the capabilities, limitations, and societal impact of large language models](#).
- Niket Tandon, Aparna S. Varde, and Gerard de Melo. 2018. [Commonsense knowledge in machine intelligence](#). *SIGMOD Rec.*, 46(4):49–52.
- Seonghyeon Ye, Doyoung Kim, Sungdong Kim, Hyeonbin Hwang, Seungone Kim, Yongrae Jo, James Thorne, Juho Kim, and Minjoon Seo. 2023. [Flask: Fine-grained language model evaluation based on alignment skill sets](#).
- Qingyu Zhou, Nan Yang, Furu Wei, Chuanqi Tan, Hangbo Bao, and M. Zhou. 2017. [Neural question generation from text: A preliminary study](#). *ArXiv*, abs/1704.01792.

Language Resource References

- Ife Adebara, AbdelRahim Elmadany, Muhammad Abdul-Mageed, and Alcides Alcoba Inciarte. 2023. [Serengeti: Massively multilingual language models for africa](#).
- Kabir Ahuja, Harshita Diddee, Rishav Hada, Millicent Ochieng, Krithika Ramesh, Prachi Jain, Akshay Nambi, Tanuja Ganu, Sameer Segal, Maxamed Axmed, Kalika Bali, and Sunayana Sitaram. 2023. [Mega: Multilingual evaluation of generative ai](#).
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have solved question answering? try arc, the ai2 reasoning challenge](#).
- 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](#). *arXiv preprint arXiv:2110.14168*.
- Sumanth Doddapaneni, Rahul Aralikkatte, Gowtham Ramesh, Shreya Goyal, Mitesh M. Khapra, Anoop Kunchukuttan, and Pratyush Kumar. 2023. [Towards leaving no indic language behind: Building monolingual corpora, benchmark and models for indic languages](#).
- Jiyeon Ham, Yo Joong Choe, Kyubyong Park, Ilji Choi, and Hyungjoon Soh. 2020a. [KorNLI and KorSTS: New benchmark datasets for Korean natural language understanding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 422–430, Online. Association for Computational Linguistics.
- Jiyeon Ham, Yo Joong Choe, Kyubyong Park, Ilji Choi, and Hyungjoon Soh. 2020b. [Kornli and korsts: New benchmark datasets for korean natural language understanding](#).
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob

- Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations (ICLR)*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. [Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization](#).
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. [Cosmos QA: Machine reading comprehension with contextual commonsense reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2391–2401, Hong Kong, China. Association for Computational Linguistics.
- Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu, Maosong Sun, and Junxian He. 2023. [C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models](#).
- Myeongjun Jang, Dohyung Kim, Deuk Sin Kwon, and Eric Davis. 2022. [KoBEST: Korean balanced evaluation of significant tasks](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3697–3708, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Younghoon Jeong, Juhyun Oh, Jongwon Lee, Jaimeen Ahn, Jihyung Moon, Sungjoon Park, and Alice Oh. 2022. [KOLD: Korean offensive language dataset](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10818–10833, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jiho Jin, Jiseon Kim, Nayeon Lee, Haneul Yoo, Alice Oh, and Hwaran Lee. 2023. [Kobbq: Korean bias benchmark for question answering](#).
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Seungyoung Lim, Myungji Kim, and Jooyoul Lee. 2019. [Korquad1.0: Korean qa dataset for machine reading comprehension](#).
- Jihyung Moon, Won Ik Cho, and Junbum Lee. 2020. [BEEP! Korean corpus of online news comments for toxic speech detection](#). In *Proceedings of the Eighth International Workshop on Natural Language Processing for Social Media*, pages 25–31, Online. Association for Computational Linguistics.
- Sungjoon Park, Jihyung Moon, Sungdong Kim, Won Ik Cho, Ji Yoon Han, Jangwon Park, Chisung Song, Junseong Kim, Youngsook Song, Taehwan Oh, Jooheon Lee, Juhyun Oh, Sungwon Lyu, Younghoon Jeong, Inkwon Lee, Sangwoo Seo, Dongjun Lee, Hyunwoo Kim, Myeonghwa Lee, Seongbo Jang, Seungwon Do, Sunkyoung Kim, Kyungtae Lim, Jongwon Lee, Kyumin Park, Jamin Shin, Seonghyun Kim, Lucy Park, Lucy Park, Alice Oh, Jung-Woo Ha (NAVER AI Lab), and Kyunghyun Cho. 2021. [Klue: Korean language understanding evaluation](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, volume 1. Curran.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel Bowman. 2022. [BBQ: A hand-built bias benchmark for question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2086–2105, Dublin, Ireland. Association for Computational Linguistics.
- Jaehyung Seo, Seunghoon Lee, Chanjun Park, Yoonna Jang, Hyeonseok Moon, Sugyeong Eo, Seonmin Koo, and Heuseok Lim. 2022. [A dog is passing over the jet? a text-generation dataset for Korean commonsense reasoning and evaluation](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2233–2249, Seattle, United States. Association for Computational Linguistics.
- Guijin Son, Hanwool Lee, Suwan Kim, Huiseo Kim, Jaecheol Lee, Je Won Yeom, Jihyu Jung, Jung Woo Kim, and Songseong Kim. 2023. [Hae-rae bench: Evaluation of korean knowledge in language models](#).
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, and et al. 2023. [Beyond the imitation game: Quantifying and extrapolating the capabilities of language models](#).
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [CommonsenseQA](#).

A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2020. [Superglue: A stickier benchmark for general-purpose language understanding systems](#).

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. [PAWS-X: A cross-lingual adversarial dataset for paraphrase identification](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3687–3692, Hong Kong, China. Association for Computational Linguistics.

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, Florence, Italy. Association for Computational Linguistics.

Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2023. [Agieval: A human-centric benchmark for evaluating foundation models](#).

Appendix

A. Exam Description

College Scholastic Ability Test of Korea (CSAT)

⁷ CSAT, endorsed by Korean universities, assesses scholastic aptitude based on the Korean high school curriculum and is administered annually by the Korean Ministry of Education.

Test of Proficiency in Korean (TOPIK) ⁸

TOPIK measures and evaluates the proficiency of learners of Korean as a second language.

National Public Service Examination - Grade 9 (PSE) ⁹

PSE evaluates the knowledge and abilities of individuals who wish to work in the public sector. We use the history section and the language section which assesses proficiency in grammar, vocabulary, and reading comprehension.

Public Service Aptitude Test (PSAT) ⁹

PSAT assesses the aptitudes essential for performing public duties at a higher standard than the PSE. Our study focuses on the Korean Constitution section of the PSAT subjects.

Korean History Exam (KHE) ¹⁰

KHE measures the historical literacy of Korean citizens.

Test of Teaching Korean as a Foreign Language (Kedu) ¹¹

Kedu certifies individuals aspiring to teach Korean to overseas Koreans or foreigners. It covers both Korean language and culture.

The Korean Immigration and Integration Program (KIIP) ¹²

KIIP assists foreigners in integrating into Korean society. We use the basic level KIIP textbook and generate QA pairs using

B. Question Generation Using GPT-4

We employed the GPT-4 language model for the generation of multiple-choice questions, focusing on content extracted from the Korean Integration and Immigration Program (KIIP) textbooks. The process involved several key steps to ensure the generation of high-quality, relevant educational questions.

Text extraction We extracted text from the KIIP textbooks to serve as the basis for our question generation. To manage this extensive textual data effectively, the extracted text was split into smaller, manageable chunks. This splitting was accomplished using a *RecursiveCharacterTextSplitter* from Langchain. The parameters for this splitter were chosen based on initial experiments to balance between maintaining textual coherence and ensuring manageable chunk sizes for processing.

Prompt The core of our question generation process involved prompting GPT-4 with a specific structure to ensure that the questions generated were diverse, relevant, and adhered to a specific format. The prompt used for GPT-4 was as follows:

Original Korean Prompt

```
다음 제시문을 읽고 4지 선다 문제 10개를 만들어줘.
문제를 만들 때 내용은 겹치지 않게 하고 형식은 다음을 포함한
json 형식으로 만들어줘. question_id는 {current_cnt}부터 시작해서
1부터 증가해줘.
"cite": 제시문에서 문제를 만드는 데 사용한 문장
"question_id": {문제 번호}
"question": {문제}
"choices": {보기}
"answer": {답}
"제시문": {content}
```

English Translated Prompt

```
Read the following passage and create 10 multiple-choice
questions based on it.
Ensure that the content of each question does not overlap and format
them in JSON format including the following elements.
The question_id should start from {current_cnt} and increase by 1 there-
after.
"cite": The sentence from the passage used to create the question
"question_id": {Question number}
"question": {Question}
"choices": {Options}
"answer": {Answer}
"Passage": {content}
```

This prompt was designed to instruct GPT-4 to produce a set of 10 multiple-choice questions for each text chunk, each with a unique question identifier ("question_id"). The format of the prompt ensured that the questions did not overlap in content and were presented in a structured JSON format. This format included a citation from the text ("cite"), the question itself ("question"), multiple choices ("choices"), and the correct answer ("answer"). This citation was not just for referencing purposes but also served a vital role in the validation process. It allowed human reviewers to quickly ascertain the validity of the generated question and to confirm that the provided answer was indeed correct according to the textbook content.

Verification As a final step in the process, any instances generated by GPT-4 that did not conform to the specified format were identified and removed.

⁷<https://www.suneung.re.kr/>

⁸<https://www.topik.go.kr/>

⁹<https://www.gosi.kr/>

¹⁰<https://www.historyexam.go.kr>

¹¹<https://www.q-net.or.kr>

¹²<https://www.immigration.go.kr>