

KoBEST: Korean Balanced Evaluation of Significant Tasks

Myeongjun Jang^{1*} Dohyung Kim^{2*} Deuk Sin Kwon² Eric Davis²

¹Department of Computer Science, University of Oxford, UK
myeongjun.jang@cs.ox.ac.uk

²Language Super Intelligence Labs, SK Telecom, South Korea
{bruno, brian.k, eric.davis}@sktair.com

Abstract

A well-formulated benchmark plays a critical role in spurring advancements in the natural language processing (NLP) field, as it allows objective and precise evaluation of diverse models. As modern language models (LMs) have become more elaborate and sophisticated, more difficult benchmarks that require linguistic knowledge and reasoning have been proposed. However, most of these benchmarks only support English, and great effort is necessary to construct benchmarks for other low resource languages. To this end, we propose a new benchmark named **Korean balanced evaluation of significant tasks** (KoBEST), which consists of five Korean-language downstream tasks. Professional Korean linguists designed the tasks that require advanced Korean linguistic knowledge. Moreover, our data is purely annotated by humans and thoroughly reviewed to guarantee high data quality. We also provide baseline models and human performance results. Our dataset is available on the Huggingface¹.

1 Introduction

The NLP field is now facing unprecedented rapid development. A major factor propelling the progress is the existence of unified benchmark datasets like GLUE (Wang et al., 2018), which are designed to assess models’ language understanding capabilities. Such benchmark datasets, enabled modern pre-trained language models (PLMs), such as BERT (Devlin et al., 2019) and GPT-2 and GPT-3 (Radford et al., 2019; Brown et al., 2020), to be assessed in objective and multifaceted manners. The success of GLUE also lead to similar benchmark datasets in a variety of other languages, such as French (Le et al., 2020), Korean (Park et al., 2021), Chinese (Xu et al., 2020) and Indonesian (Wilie et al., 2020).

However, many recent studies reveal that the outstanding performance of PLMs on such benchmark datasets seems plausible but not probable. These studies have found that datasets may contain many spurious artefacts, and the performance of PLMs is enhanced by excessive usage of said artefacts (Habernal et al., 2018; Niven and Kao, 2019; McCoy et al., 2019; Bender and Koller, 2020). Another line of work observed that many PLMs, which showed promising results in GLUE, fall short of expectations for more difficult tasks that require linguistic knowledge (Bhatt et al., 2021) or logical reasoning (Tian et al., 2021). As a result, the importance of well-designed evaluation datasets with higher difficulty-level has been highlighted, and new datasets, such as CHECKLIST (Ribeiro et al., 2020), and LOGICNLI (Tian et al., 2021), have been proposed. Most of them only support specific languages like English, and it requires large efforts to build higher difficulty-level language evaluation suits for other low resource languages, however.

When it comes to the Korean language, two benchmarks are widely used: Korean-NLI & STS (Ham et al., 2020) and KLUE (Park et al., 2021). The former is machine- and human-translated from English natural language inference (NLI) and semantic textual similarity (STS) datasets, which hardly reflect the characteristics of the Korean language. The latter is a Korean version of GLUE benchmark which supports eight tasks, such as NLI, STS, named entity recognition (NER), and relation extraction (RE). Although these tasks are useful for assessing general language ability, it is difficult to ascertain whether a model is able to reason based on more complicated knowledge beyond text form (e.g., passage of time, meaning of text, causality). To this end, we aim to construct a new benchmark dataset in Korean named KoBEST, which consists of five downstream tasks that require advanced knowl-

*These authors contributed equally to this work

¹https://huggingface.co/datasets/skt/kobest_v1

Tasks	Train	Dev	Test	Labels	Text Source
KB-BoolQ	3.7K	700	1.4K	2	Wikipedia
KB-COPA	3.1K	1K	1K	2	N.A
KB-WiC	3.3K	1.3K	1.3K	2	Korean Dictionary
KB-HellaSwag	2K	500	500	4	Wikipedia, YouTube
KB-SentiNeg	3.6K	400	397	2	Product reviews

Table 1: The number of data instances and labels for each downstream task.

edge of Korean. We carefully constructed the data based on the following design principles:

- *Human-driven data annotation*: Our data is purely annotated by humans to prevent incorrect and ambiguous data instances caused by automatic data annotation approach.
- *Leveraging professional linguistic knowledge*: As a result of our collaboration with professional Korean linguists, we are able to collect grammatically correct data with rich vocabulary and expressions.
- *Availability to public*: As a benchmark dataset, it is important to ensure public accessibility. We guarantee that our data is free to use and redistribute.
- *High data quality*: Our data passed thorough reviews driven by both models and humans to deliver high quality data without superficial cues and heuristic artefacts.
- *Avoiding AI ethical issues*: Human review process have been performed to remove toxic content, social biases, and personal information from our data set.

Next, we evaluated widely used Korean PLMs on the KoBEST dataset. Specifically, we conducted fine-tuning, zero-shot, one-shot, and few-shot experiments. The experimental results can serve as a baseline for performance on KoBEST. Participants also provided human performance baselines for all of our tasks. Our results suggest that modern PLMs and a large-size generative language model (GLM) are far from reaching human-level language ability.

2 KoBEST Downstream Tasks

2.1 Overview

The KoBEST benchmark consists of the following five downstream tasks:

1. KoBEST-BoolQ (KB-BoolQ): identify whether a given question is true or false considering a paragraph.
2. KoBEST-COPA (KB-COPA): select an al-

ternative which is a cause/effect of a given premise.

3. KoBEST-WiC (KB-WiC): identify whether the meaning of a target word is the same or different in two given contexts.
4. KoBEST-HellaSwag (KB-HellaSwag): select a correct sentence among four candidates that is likely to appear after a given context.
5. KoBEST-SentiNeg (KB-SentiNeg): predict the polarity of a negated sentence.

The number of training/development/test data points is illustrated in Table 1.

2.2 KoBEST-BoolQ

Data/Task Description We built the KB-BoolQ dataset by referencing boolean questions (BoolQ) task (Clark et al., 2019). A data point consists of a paragraph, question, and label. The task aims to evaluate models’ understanding of the paragraph by asking a true/false question. An example is presented in Table 2.

We extracted paragraphs from Korean Wikipedia². To cover diverse materials, we first choose topics, such as Science/Technology and Art/Culture, by referring to previous works regarding Korea written/spoken language (Seo and Kim, 2005; Seo, 2007). Then, we defined keywords for each topic and selected documents containing enough information regarding the keyword. Next, we extracted paragraphs for each document and generated corresponding questions that could be answered as true/false based on the paragraph.

Guidelines Annotators were instructed to construct the KB-BoolQ dataset following the guidelines described below.

1. Paragraphs should be evenly extracted from various domains and topics to minimise bias.
2. Questions should be answered only with the information presented in the paragraph. We set this guide for two reasons: 1) to exclude the impact of pre-trained commonsense knowledge for decision making and 2) annotators have different viewpoints regarding the boundary of commonsense knowledge, which can cause uneven task difficulty.
3. Questions should be written in clear, unambiguous, easy-to-understand language. A true/false judgement should be obvious from a human perspective.

²<https://ko.wikipedia.org/wiki/>

KB-BoolQ	<p>Paragraph: 구한말, 통영 안뒤산 기슭 간창골에 김봉제 형제가 살았다. 김봉제는 관약국을 경영하며 부를 누렸는데, 선비적 성품을 지녔던 형과 반대로 막냇동생 김봉룡은 성질이 포악했다. 어느 날 봉룡은 아내였던 숙정을 사모하던 나그네를 살해하였고, 숙정은 누명을 벗으려고 비상을 먹고 자살한다. (At the end of the Joseon Dynasty, the Kim Bong-je brothers lived in Ganchanggol at the foot of Andui Mountain in Tongyeong. Kim Bong-je enjoyed wealth while running a government office, but his younger brother Kim Bong-ryong had a violent temper, contrary to his brother who had a scholarly character. One day, Bongryong killed a traveler who adored his wife Sukjeong, and Sukjeong commits suicide to clear his name.)</p> <p>Question: 봉룡은 숙정을 죽였는가? (Did Bongryong kill Sukjeong?) Answer: False</p>
KB-COPA	<p>Premise: 전쟁이 시작되었다. (The war had begun.) Question: 결과 (Effect)</p> <p>Alternative 1: 병사들이 집으로 돌아왔다. (Soldiers returned home.)</p> <p>Alternative 2: 병사들이 전투에 파견되었다. (Soldiers were sent to battle.)</p> <p>Correct Alternative: 2</p>
KB-WiC	<p>Context 1: 망가진 엔진은 수리 가 불가능하다. (It is impossible to repair a broken engine.)</p> <p>Context 2: 이 배는 수리 에 들어간 지 일주일이었다. (The ship has been under repair for a week.)</p> <p>Target Word: 수리 (repair) Answer: True</p>
KB-HellaSwag	<p>Context: 양궁 선수들이 경기장으로 입장한다. 관중들이 함성을 지르고 응원한다. 선수들이 상대팀과 악수하고 자리로 돌아온다. 코치가 전략을 설명하고 화이팅을 외친다. (Archery players enter the stadium. The crowd shouts and cheers. The players shake hands with the opposing team and return to their seats. The coach explains the strategy and shouts "Go for it.")</p> <p>Ending 1: 활이 과녁에 적중했다. (The arrow hits the target.)</p> <p>Ending 2: 선수가 심호흡을 하고 활을 쏜다. (The player takes a deep breath and shoots an arrow.)</p> <p>Ending 3: 선수가 활을 들어 과녁을 조준한다. (The player raises his bow and aims at the target.)</p> <p>Ending 4: 선수들이 각자 자리에 서서 활을 꺼낸다. (The players stand in their own positions and take out their bows.)</p> <p>Correct Ending: 4</p>
KB-SentiNeg	<p>Sentence 1: 뚜껑이 잘 열려요! (The lid opens well!) Label 1: 긍정 (Positive)</p> <p>Sentence 2: 뚜껑이 잘 안열려요! (The lid does not open well!) Label 2: 부정 (Negative)</p>

Table 2: Examples of development set from the KoBEST tasks. The variables of each task are highlighted in bold. Text written in parenthesis is the English translated version of the original data points.

2.3 KoBEST-COPA

Data/Task Description We referenced choice of plausible alternatives (COPA) (Roemmele et al., 2011) to construct the KB-COPA dataset. The data has four variables: premise, two alternatives, and a question that asks a model to decide the cause or effect of the premise from the two alternatives. An example is available in Table 2.

Guidelines We provided the following guidelines to annotators for generating data instances.

1. The alternatives should belong to a similar area, e.g., states and actions. This rule is introduced to preclude systems from making decisions based on situational difference, not the meaning of alternatives.
2. The alternatives should contain a keyword related to that of premise. For instance, in the example presented in Table 2, the keyword of the premise is “전쟁 (war)”, and both alternatives contain the related same keyword “병사 (soldier)”. We introduce this guideline to increase the task’s difficulty by making the alternatives belong to the same category.
3. All the premises and alternatives should be written concisely so that the content can

be understood intuitively. Therefore, using proper noun, slang, and redundant expressions should be avoided.

4. All sentences should be written in the past tense. In the Korean language, simple present can cause confusion because it has indication of tense and sometimes can imply present progressive. On the other hand, the past tense is morphologically clear and is able to convey meaning without confusion.
5. All sentences must include a subject. Although the subject is frequently omitted in Korean, it is difficult to infer the cause or effect without a subject because all the premises and alternatives are quite short. Therefore, even though such sentences are slightly unnatural in Korean, we guide annotators to insert a subject.

2.4 KoBEST-WiC

Data/Task Description KB-WiC is a task that determines whether a word has the same connotation in different contexts. We referenced words in context (WiC) (Pilehvar and Camacho-Collados, 2019) when building the dataset. An instance is composed of a target homonym and two different

contexts that contain the target word. Table 2 provides an example for the KB-WiC task. Unlike the original WiC dataset that includes various word forms, we only used words with the same form, so as to focus more on recognising a word’s meaning without the distraction of various forms.

Guidelines To construct the KB-WiC dataset, we instructed annotators to follow these guidelines.

1. A target word should be listed in the National Institute of the Korean Language Basic Korean Dictionary³ or Korea University Korean Dictionary (Hong and Kim, 2009). We exclude words not registered in the dictionaries because they can cause ambiguous criteria for determining an answer. This is despite the fact that they are generally used in daily life.
2. For generating a data point where an answer is *False*, only a **homonym** should be used as a target word because a **polysemy** makes the task considerably more challenging, even for native speakers.
3. The part of speech (PoS) tag of a target word should be a noun, pronoun, numeral, or dependent noun⁴. We introduce this guideline because the four PoS tags have a fixed form and distinct meaning in Korean.
4. The contexts should be extracted from example sentences in the dictionaries to make it possible to clearly understand the sense of a target word only using the given context.

2.5 KoBEST-HellaSwag

Data/Task Description This task evaluates whether a system can utilize passage of time and order to complete the last sentence in a series of sentences. We referenced the HellaSwag dataset (Zellers et al., 2019) to build our version but modified the task to consider specific characteristics of the Korean language.

The original HellaSwag benchmark was designed to ascertain whether a LM can generate a plausible ending sentence given a relevant subject and context. In Korean, however, subjects are typically omitted. As a result, if the ending sentence is generated from the subject, the sentence becomes awkward and barely resembles a plausible Korean

³<https://stdict.korean.go.kr/main/main.do>

⁴A noun that cannot be used without the help of other words in Korean.

sentence. Evaluating such unnatural sentences is not in line with the purpose of KoBEST, so we modify the task to predict the most plausible final sentence among four candidates. An example instance is available in Table 2.

Guidelines We instruct annotators to build the data based on the following guidelines.

1. The annotators should generate or modify free-text descriptions of YouTube videos and Wikipedia documents that progress with the passage of time.
2. At least three sentences should be included in the context. A system should have as much context as possible to generate a plausible ending sentence.
3. All the candidate-ending sentences should be thematically related to the context. The answer should only be able to be found by analysing the passage of time among the sentences, not via the topic or keywords. This guideline is introduced to prevent low task difficulty by generating alternative endings that simply contradict the correct ending.

2.6 KoBEST-SentiNeg

Data/Task Description Many studies have revealed that PLMs lack understanding of negation expressions (Hossain et al., 2020; Ettinger, 2020; Kassner and Schütze, 2020; Hosseini et al., 2021; Jang et al., 2022). Inspired by the *Negation* capability test of Ribeiro et al. (2020), we designed a similar but enhanced task by utilizing negation to create sentences opposite in meaning. Specifically, we created a two-class sentiment analysis task by generating product reviews based on real product reviews available on e-commerce websites.⁵ We then used the *training* and *development* sets to train a sentiment classification model. Next, we extracted candidates from training data where the polarity switched when transformed into a sentence with the opposite meaning. Finally, we converted each candidate to a sentence with its opposite meaning and reversed the label. The modified candidate is then added to the final test set. We used the following three methods to generate the sentences with opposite meanings.

1. **Adding/removing negation expressions:** We add or remove Korean negation expressions (e.g., “안”, “못”, “지 않다”).

⁵The real product reviews are only used as references and our data is newly generated by human annotators.

2. **Antonym replacement:** A word is replaced with its antonym.
3. **Using both method 1 and 2 or idiom:** Both methods described above are used. If a sentence includes an idiom, we replace it with its corresponding opposite meaning idiom.

Guidelines We instructed the annotators to comply with the following guidelines to generate data points for the KB-SentiNeg task.

1. The sentence should not include the name of specific brands or products. This guideline is meant to avoid any possible legal issues.
2. To generate a new sentence resembling a real product review, typos and spacing errors that frequently occur in Korean spoken language should be included occasionally.

2.7 Evaluation Metrics

All of our downstream tasks have discrete labels. Therefore, we use the $F1$ score as a base criterion to evaluate models' performance.

3 Design Principles

In this section, we provide detailed descriptions about how we attempted to achieve the design principles illustrated in Section 1.

3.1 Human-driven data annotation

Automatic data generation using meta-information, such as a review score and news article category (e.g., Naver Sentiment Movie Corpus⁶, is a widely used approach to rapidly collect a large amount of labelled data. While it is an efficient approach, there exists a high risk of the dataset containing incorrect and ambiguous data points. Such noisy data points are a major issue for evaluation datasets because they can lead to spurious performance increases (and or degradations) in the performance of LMs.

Our data is created purely by human annotations to produce the highest quality dataset with the lowest amount of incorrect and ambiguous data points possible. We hired four annotators who are Korean native speakers and major in *Korean Language Education* or *Korean Language in Literature*. Also, our Korean linguists trained the annotators before the data annotation process to avoid generating possible grammatical errors and unethical expressions.

⁶<https://github.com/e9t/nsmc>

Example #1

P: 날씨가 추워졌다. (The weather has become colder.)

Q: 원인 (Cause)

A1: 겨울이 되었다. (Winter has come.)

A2: 여름이 되었다. (Summer has come.)

Example #2

P: 겨울이 되었다. (Winter has come.)

Q: 결과 (Effect)

A1: 날씨가 추워졌다. (The weather has become colder.)

A2: 날씨가 더워졌다. (The weather became hot.)

Table 3: Examples of the KB-COPA data where the causality is interlocked. **P**, **Q**, **A1** and **A2** denote a premise, question, and alternatives, respectively. The two sentences highlighted in red and blue colours are swapped by changing the question from *Cause* to *Effect*.

3.2 Leveraging linguistic knowledge

Korean benchmark datasets, translated from English datasets (e.g., Kornli and Korsts (Ham et al., 2020)), might include incorrect translations and grammatical errors, particularly if they are machine-translated. Moreover, since the original examples come from English, such benchmark datasets are unlikely to assess properly assess Korean-specific knowledge or language intricacies.

Relying on our in-house Linguistic team, allowed us to mitigate and resolve issues with automatically generated datasets. First, our linguists trained the annotators to generate natural and grammatical Korean sentences and performed thorough reviews of the data. Thanks to their efforts, we have created a highly grammatical and natural. Secondly, the linguists designed tasks and data generation processes that considered the Korean language's characteristics. This is illustrated in guidelines for each task in Section 2. Such guidelines and processes enabled us to create an accurate Korean evaluation dataset with expressive vocabulary and colloquial usage.

3.3 Availability to public

Our data is free from copyright issues. All sentences and answers, except for the paragraphs in KB-BoolQ task, are generated by our annotators from scratch by referencing publicly available sources. Also, the paragraphs in KB-BoolQ were extracted from Wikipedia, which is under the Creative Commons Attribution-ShareAlike license. Therefore, researchers are free to use, modify and redistribute the KoBEST dataset.




Figure 1: The overall process of model-driven review for removing artefacts in training data.

3.4 High data quality

As a benchmark suite, accomplishing high data quality is important to accurately evaluate various LMs. We achieve this purpose through two review phases: *human-driven* and *model-driven* reviews.

Human-driven review After collecting data, our linguists reviewed all data instances and found two major issues in the KB-COPA and KB-HellaSwag datasets.

For the KB-COPA task, we observed many cases with high correlation between data instances. Examples corresponding to this case are presented in Table 3. We conjecture this is because it was easier for the annotators to collect data by simply swapping the premise and answer along with changing the question. We removed or modified such instances because they are near duplicates and harm data diversity.

In the KB-HellaSwag task, we found several critical cases where predicting a correct final sentence is quite difficult due to the omission of detailed delineations in the context. This issue occurs because the source that the annotators referenced for generating contexts lacks detailed information occasionally. For such ambiguous instances, we appended additional clues to the context to allow inferring the meaning between the context and the final sentence.

Model-driven review Artefacts existing in training data can lead a model to learn spurious inductive biases, resulting in distorted evaluation results (Gururangan et al., 2018; McCoy et al., 2019; Hossain et al., 2020). Therefore, we conduct a model-driven review process to find and remove such unwanted artefacts. The overall process is illustrated in Figure 1. First, we trained an ad-hoc model with the initial dataset for each task. Next, we generated predictions for the development and test datasets and analysed the results to ascertain whether specific words/patterns/phrases were strongly correlated with labels. Finally, our linguists analyzed the issues and updated the datasets accordingly. We repeated the whole process up

to three times for each dataset. Through this process, we observed serious artefacts, especially in the KB-WiC and KB-COPA datasets. More than 70% of questions containing number-related representations had *False* as a label. Also, the label distributions of data instances containing specific phrases (e.g., "덥다/춥다 (hot/cold)" or 했다/하지 않았다 (did/did not)) were highly skewed towards the *False* label. All such artefacts were successfully removed and modified by our linguists.

3.5 Avoiding AI ethical issues

Social biases embedded in training data can lead to unethical behaviour of language models (Nangia et al., 2020). To mitigate such issues, we made efforts to remove unethical expressions, such as toxic content (e.g., insults, slang, sexual harassment) and social bias (e.g., gender, race, religion). Our linguists clearly instructed the annotators to avoid unethical expressions when generating sentences and extracting paragraphs from Wikipedia for the KB-BoolQ task. Also, linguists reviewed the data for potential ethical issues, as described in Section 3.4

4 Experiments

In this section, we provide model and human performance results.

4.1 Fine-tuning Experiments

4.1.1 Experimental Design

Model Candidates We used the following four pre-trained Korean language models to benchmark our KoBEST dataset:

- **Encoder models:** KoBERT⁷, KoElectra⁸
- **Decoder models:** KoGPT3-1.2B
- **Encoder-Decoder models:** KoBART⁹

For KoBART, we applied the text-to-text multi-task training technique to fine-tune the model. We

⁷<https://huggingface.co/monologg/kobert>

⁸<https://huggingface.co/monologg/koelectra-base-v2-discriminator>

⁹<https://huggingface.co/hyunwoongko/kobart>

Model	KB-BoolQ	KB-COPA	KB-WiC	KB-HellaSwag	KB-SentiNeg	Average
KoBERT (FT)	62.9±3.0	74.6±0.8	77.3±0.8	74.4±0.4	86.8±2.0	75.2
KoElectra (FT)	75.1±1.0	81.5±0.4	79.7±1.8	74.7±0.8	91.9±1.1	80.6
KoGPT3-1.2B (FT)	73.5±1.6	79.3±0.6	68.4±2.2	73.8±1.0	89.5±3.3	77.0
KoBART (FT)	60.6±2.9	56.9±3.2	60.4±4.9	51.4±1.3	88.6±1.2	63.6
KoGPT3-39B ($k=0$)	33.1	76.8	34.7	59.8	57.7	52.8
KoGPT3-39B ($k=1$)	50.2	78.3	51.8	60.2	74.2	66.3
KoGPT3-39B ($k=10$)	46.9	80.9	52.2	58.7	91.6	70.6
Human	95.1	98.1	96.6	92.4	99.0	96.2

Table 4: The test F1 scores of Korean LMs on the KoBEST downstream tasks. The first and second blocks are the experimental results of fine-tuning and zero/few-shot learning, respectively. For the fine-tuning experiments, we repeat each experiment five times and report the average and standard deviation. The best values for each task are written in bold. k refers to the number of few-shot samples.

	BoolQ	COPA	WiC	HellaSwag	SentiNeg	Model	KB-BoolQ	KB-WiC	KB-SentiNeg
b-size	8	16	16	8	16	Single-task	66.5 ±3.1	68.3 ±3.0	87.7±1.4
s-len	256	128	256	256	128	Multi-task	60.6±2.9	60.4±4.9	88.6 ±1.2
lr	$5e^{-6}$	$5e^{-6}$	$1e^{-5}$	$2e^{-5}$	$1e^{-5}$				

Table 5: Batch-size (b-size), maximum input length (s-len), and learning rates (lr) used for the KoBEST benchmark experiments.

Model	\mathcal{F}_{val}	\mathcal{F}_{test}	Δ
KoBERT	99.1	86.8	12.3
KoElectra	99.4	91.9	7.5
KoGPT3-1B	99.8	89.5	10.3
KoBART	98.7	88.6	10.1

Table 6: The average performance gap between validation and test sets of KB-SentiNeg task.

transformed inputs of each text to a free-text input by referencing the transformation formats used in the T5 model (Raffel et al., 2020).

Training Details We used AdamW optimiser (Loshchilov and Hutter, 2017) for training with a linear learning rate scheduler decaying from $1e^{-2}$. We trained all models for 10 epochs, and used the early stopping method during the training. Different batch sizes and learning rates were used across the tasks, and detailed training hyperparameters are presented in Table 5.

4.1.2 Results and Discussion

The fine-tuning results can be seen in the first block of Table 4. Overall, KoElectra showed the best results, followed by KoGPT-1B, suggesting that model size is not necessarily a requisite for better performance.

Do models understand the opposite in meaning? Table 6 shows the validation and test performance of fine-tuned models on the SentiNeg task. Interestingly, all models show a large performance gap between the validation and test perfor-

Table 7: The performance of KoBART models trained with single- and multi-task manners.

mance in the SentiNeg task. The results suggest that PLMs are vulnerable to a simple negation and antonym-replacement perturbation, even though the data points all originated from the training data. Our results are aligned with previous studies on English data that showed PLMs are incapable of understanding negation expressions (Hossain et al., 2020; Kassner and Schütze, 2020; Etinger, 2020; Hosseini et al., 2021), suggesting that the issue stems from the PLM, not from language itself.

Text-to-Text Multi-task training is not always beneficial Unlike the other three models, KoBART was fine-tuned in a multi-task fashion. However, contrary to the common belief that multi-task training is beneficial in improving performance on benchmark suites (e.g., GLUE (Liu et al., 2019)), in our case, multi-task training produced the worst performance by a large margin. We conducted additional single-task classification experiments on KoBART by introducing a classifier layer. The multiple-choice tasks (i.e., KB-COPA and KB-HellaSwag) are not included in this experiment, as the structure of the BART model is not suitable for the multiple-choice tasks. The results are presented in Table 7. The results show that the single-task model performs better than multi-task models on the KB-BoolQ and KB-WiC tasks by a large margin. We also ascertained that the performance gap was statistically significant ($p < 0.05$) on the two tasks, while there was no significant difference on the KB-SentiNeg task. We conjecture that a leading cause is a misalign-

ment between tasks. All the downstream tasks in KoBEST are independent of each other. However, in the GLUE benchmark, for instance, the sub-tasks are well aligned, containing multiple datasets that share a common objective, e.g., NLI and STS, and it is well studied that the misalignment between task data can cause poor results (Wu et al., 2020).

4.2 Zero/Few Shot Experiments

4.2.1 Experimental Design

The advent of extremely large size GLMs like GPT3 (Brown et al., 2020) has allowed in-context learning (providing the model with a few or no samples) to apply the model to downstream tasks. To this end, we conducted zero-, one- and few-shot experiments by using a Korean GPT3 model trained by Language Superintelligence Labs with 39 billion parameters and 132 billion tokens. We then referenced the work of EleutherAI¹⁰ to design prompts for our zero, one, and few-shot experiments. Several prompt examples are available in Table 8 in the appendix. For multiple-choice problems like **KB-COPA** and **KB-HellaSwag**, we selected the candidate having the lowest perplexity as the prediction.

4.2.2 Results and Discussion

Fine-tuned models are still best. The results, presented in the second block of Table 4, reveal that the fine-tuned models, apart from KoBART, outperform in-context learning methods. Our results are aligned with the work of Brown et al. (2020) that showed GPT3 performance based on few-shot learning is behind that of fine-tuned SOTA in many tasks, including all downstream tasks in SuperGLUE (Wang et al., 2019). Although it is interesting that a large GLM can achieve decent performance with only a few training examples, results suggest that we should be judicious using large GLM in practical applications; especially when considering performance compared to excessively high training costs (i.e., time and resources).

Increasing k is not always beneficial. Few-shot learning approaches with more examples increase performance in general, but merely increasing k does not always lead to better performance. Specifically, in the case of KoBEST, the performance is slightly worse on the **KB-BoolQ** and

KB-HellaSwag tasks. We believe that the length of the input document is a leading cause of this phenomenon. Since the model’s max input length plays a critical role in deciding the maximum number of examples (n) in the prompt (Yang et al., 2021), the available n decreases as the length of prompts increases. However, as we can see in Table 2, the data points of the **KB-BoolQ** and **KB-HellaSwag** tasks have longer inputs than the other tasks. As a result, the prompts for these tasks become very long and likely to exceed the model’s maximum input length. This would result in a sliced prompt that may lack key information the model needs to make a successful prediction.

4.3 Human performance

We asked volunteers to evaluate the dataset to provide human-level performance metrics for KoBEST. Specifically, 10 native Korean evaluators evaluated 100 randomly sampled examples for each downstream task. The results are summarised in the last row of Table 4. The human evaluators outperformed all the PLMs by a large margin, suggesting that modern PLMs need further improvements to achieve human-level language ability.

5 Conclusion

A well-designed benchmark dataset is crucial for an objective and precise evaluation of LMs. Following the GLUE benchmark, more challenging benchmarks have been proposed as modern LMs become more elaborate and sophisticated. However, most of these benchmarks only support English or originate from English (e.g., translation), which hardly captures important characteristics of a specific language.

To this end, we propose a new Korean benchmark suite named KoBEST, which consists of five challenging downstream tasks. To overcome the disadvantages of the previous Korean benchmarks, we focused on 1) employing Korean-specific knowledge, 2) achieving high data quality and 3) removing superficial cues. To achieve these goals, we worked with professional Korean linguists and collected data manually and not automatically. We also conducted human- and model-driven review processes to eliminate superficial cues from our dataset. Moreover, we were extra cautious to avoid using unethical expressions.

Finally, we evaluated various PLMs on our new benchmark and provide baseline model and hu-

¹⁰<https://github.com/EleutherAI/lm-evaluation-harness>

man performance metrics. Our experimental results show that current LMs need further improvements to attain human-level language ability. We hope our new benchmark can contribute to advancements in the Korean NLP field.

Acknowledgements

We would like to express our great appreciation to Sunwoo Lee, Seo-jin Lee, and Seokyoung Hong for their valuable assistance and feedback.

References

- Emily M. Bender and Alexander Koller. 2020. [Climbing towards NLU: On meaning, form, and understanding in the age of data](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Shaily Bhatt, Rahul Jain, Sandipan Dandapat, and Sunayana Sitaram. 2021. [A case study of efficacy and challenges in practical human-in-loop evaluation of NLP systems using checklist](#). In *Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval)*, pages 120–130, Online. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Christopher Clark, Kenton Lee, Ming-Wei Chang, Tom Kwiatkowski, Michael Collins, and Kristina Toutanova. 2019. [BoolQ: Exploring the surprising difficulty of natural yes/no questions](#). 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 2924–2936, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Allyson Ettinger. 2020. [What bert is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R Bowman, and Noah A Smith. 2018. [Annotation artifacts in natural language inference data](#). In *NAACL-HLT (2)*.
- Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. 2018. [The argument reasoning comprehension task: Identification and reconstruction of implicit warrants](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1930–1940, New Orleans, Louisiana. Association for Computational Linguistics.
- Jiyeon Ham, Yo Joong Choe, Kyubyong Park, Ilji Choi, and Hyungjoon Soh. 2020. [Kornli and korsts: New benchmark datasets for korean natural language understanding](#). *arXiv preprint arXiv:2004.03289*.
- I.S. Hong and H.G. Kim. 2009. [Korea university korean dictionary](#). *Research Institute of Korean Studies, Korea University, Korea*.
- Md Mosharaf Hossain, Venelin Kovatchev, Pranoy Dutta, Tiffany Kao, Elizabeth Wei, and Eduardo Blanco. 2020. [An analysis of natural language inference benchmarks through the lens of negation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9106–9118, Online. Association for Computational Linguistics.
- Arian Hosseini, Siva Reddy, Dzmitry Bahdanau, R Devon Hjelm, Alessandro Sordani, and Aaron Courville. 2021. [Understanding by understanding not: Modeling negation in language models](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1301–1312, Online. Association for Computational Linguistics.
- Myeongjun Jang, Frank Mtumbuka, and Thomas Lukasiewicz. 2022. [Beyond distributional hypothesis: Let language models learn meaning-text correspondence](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2030–2042, Seattle, United States. Association for Computational Linguistics.
- Nora Kassner and Hinrich Schütze. 2020. [Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7811–7818, Online. Association for Computational Linguistics.

- Hang Le, Loïc Vial, Jibril Frej, Vincent Segonne, Maximin Coavoux, Benjamin Lecouteux, Alexandre Al-lauzen, Benoit Crabbé, Laurent Besacier, and Didier Schwab. 2020. [FlauBERT: Unsupervised language model pre-training for French](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2479–2490, Marseille, France. European Language Resources Association.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. [Multi-task deep neural networks for natural language understanding](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496, Florence, Italy. Association for Computational Linguistics.
- I. Loshchilov and F. Hutter. 2017. Fixing weight decay regularization in adam. *ArXiv*, abs/1711.05101.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. [CrowS-pairs: A challenge dataset for measuring social biases in masked language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. Association for Computational Linguistics.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Sungjoon Park, Jihyung Moon, Sungdong Kim, Won Ik Cho, Jiyoung Han, Jangwon Park, Chisung Song, Junseong Kim, Yongsook Song, Taehwan Oh, et al. 2021. [KLUE: Korean language understanding evaluation](#). *arXiv preprint arXiv:2105.09680*.
- Mohammad Taher Pilehvar and Jose Camacho-Collados. 2019. [WiC: the word-in-context dataset for evaluating context-sensitive meaning representations](#). 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 1267–1273, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21:1–67.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. Choice of plausible alternatives: An evaluation of commonsense causal reasoning. In *2011 AAAI Spring Symposium Series*.
- Sang Kyu Seo. 2007. The 21st century sejong project special corpus construction (1998 2007). In *Annual Conference on Human and Language Technology*, pages 317–322.
- Sang Kyu Seo and Hyeong Jeong Kim. 2005. On some practical conditions for developing spoken language corpus. *Language Facts and Perspectives*, 16:5–29.
- Jidong Tian, Yitian Li, Wenqing Chen, Liqiang Xiao, Hao He, and Yaohui Jin. 2021. [Diagnosing the first-order logical reasoning ability through LogicNLI](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3738–3747, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. [SuperGLUE: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- 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.
- Bryan Wilie, Karissa Vincentio, Genta Indra Winata, Samuel Cahyawijaya, Xiaohong Li, Zhi Yuan Lim, Sidik Soleman, Rahmad Mahendra, Pascale Fung, Syafri Bahar, and Ayu Purwarianti. 2020. [IndoNLU: Benchmark and resources for evaluating Indonesian natural language understanding](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 843–857, Suzhou, China. Association for Computational Linguistics.

Sen Wu, Hongyang R Zhang, and Christopher Ré. 2020. [Understanding and improving information transfer in multi-task learning](#). In *International Conference on Learning Representations*.

Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Ye Chen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. [CLUE: A Chinese language understanding evaluation benchmark](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4762–4772, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. 2021. An empirical study of GPT-3 for few-shot knowledge-based VQA. *arXiv preprint arXiv:2109.05014*.

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.

A Appendix

KB-BoolQ	<p>Inputs: Paragraph: 구한말, 통영 안뒤산 기슭 간창골에 김봉제 형제가 살았다. 김봉제는 관악국을 경영하며 부를 누렸는데, 선비적 성품을 지녔던 형과 반대로 막냇동생 김봉룡은 성질이 포악했다. 어느 날 봉룡은 아내였던 숙정을 사모하던 나그네를 살해하였고, 숙정은 누명을 벗으려고 비상을 먹고 자살한다.</p> <p>Question: 봉룡은 숙정을 죽였는가? Answer: False</p> <p>Prompt Design: Answer: “예” if Answer is True else “아니오” Format: “{Paragraph} 질문: {Question} 답변: {Answer}.”</p> <p>Example: 구한말, 통영 안뒤산 기슭 간창골에 김봉제 형제가 살았다. 김봉제는 관악국을 경영하며 부를 누렸는데, 선비적 성품을 지녔던 형과 반대로 막냇동생 김봉룡은 성질이 포악했다. 어느 날 봉룡은 아내였던 숙정을 사모하던 나그네를 살해하였고, 숙정은 누명을 벗으려고 비상을 먹고 자살한다. 질문: 봉룡은 숙정을 죽였는가? 답변: 아니요.</p>
KB-COPA	<p>Inputs: Premise: 전쟁이 시작되었다. Question: 결과 Answer Alternative: 병사들이 전투에 파견되었다.</p> <p>Prompt Design: Connector: “왜냐하면” if Question is “원인” else “그래서” Format: “{Premise} {Connector} {Answer Alternative}”</p> <p>Example: 전쟁이 시작되었다. 그래서 (“왜냐하면” if question is 원인) 병사들이 전투에 파견되었다.</p>
KB-WIC	<p>Inputs: Context 1: 망가진 엔진은 수리가 불가능하다. Context 2: 이 배는 수리에 들어간 지 일주일이었다. Target Word: 수리 Answer: True</p> <p>Prompt Design: Answer: “예” if Answer is True else “아니오” Format: “문장1: {Context1} 문장2: {Context2} 두 문장에서 {Target Word}가 같은 뜻으로 쓰였나? {Answer}”</p> <p>Example: 문장1: 망가진 엔진은 수리가 불가능하다. 문장2: 이 배는 수리에 들어간 지 일주일이었다. 두 문장에서 수리가 같은 뜻으로 쓰였나? 예.</p>
KB-HellaSwag	<p>Inputs: Context: 양궁 선수들이 경기장으로 입장한다. 관중들이 함성을 지르고 응원한다. 선수들이 상대팀과 악수하고 자리로 돌아온다. 코치가 전략을 설명하고 화이팅을 외친다. Correct Ending: 선수들이 각자 자리에 서서 활을 꺼낸다.</p> <p>Prompt Design: 문장: {Context} {Correct Ending}</p> <p>Example: 문장: 양궁 선수들이 경기장으로 입장한다. 관중들이 함성을 지르고 응원한다. 선수들이 상대팀과 악수하고 자리로 돌아온다. 코치가 전략을 설명하고 화이팅을 외친다. 선수들이 각자 자리에 서서 활을 꺼낸다.</p>
KB-SentiNeg	<p>Inputs: Sentence: 뚜껑이 잘 안열려요! Answer: 부정</p> <p>Prompt Design: Format: “문장: {Sentence} 긍부정: {Answer}”</p> <p>Example: 문장: 뚜껑이 잘 안열려요! 긍부정: 부정</p>

Table 8: Prompt designs for each task used in in-context learning experiments. Example data points presented in Table 2 are used.