

# Pretraining and Updates of Domain-Specific LLM: A Case Study in the Japanese Business Domain

Kosuke Takahashi, Takahiro Omi, Kosuke Arima

Stockmark

{kosuke.takahashi, takahiro.omi, kosuke.arima}@stockmark.co.jp

Tatsuya Ishigaki

National Institute of Advanced Industrial Science and Technology

ishigaki.tatsuya@aist.go.jp

## Abstract

The development of Large Language Models (LLMs) in various languages has been advancing, but the combination of non-English languages with domain-specific contexts remains underexplored. This paper presents our findings from training and evaluating a Japanese business domain-specific LLM designed to better understand business-related documents, such as the news on current affairs, technical reports, and patents. Additionally, LLMs in this domain require regular updates to incorporate the most recent knowledge. Therefore, we also report our findings from the first experiments and evaluations involving updates to this LLM using the latest article data, which is an important problem setting that has not been addressed in previous research. From our experiments on a newly created benchmark dataset for question answering in the target domain, we found that (1) our pretrained model improves QA accuracy without losing general knowledge, and (2) a proper mixture of the latest and older texts in the training data for the update is necessary. Our pretrained model and business domain benchmark are publicly available<sup>1</sup> to support further studies.

## 1 Introduction

The development of Large Language Models (LLMs) has seen significant progress across various languages. However, the combination of language-specific and domain-specific contexts remains underexplored. This study focuses on a Japanese LLM tailored for the business domain, addressing the need for models that can understand and process business-related documents such as articles on current affairs, corporate activities, and social issues.

As demonstrated by the examples of business-related questions in Table 1, accurately answering these questions requires specialized knowledge

Category	Question
Current Affairs	<i>Which country joined NATO in April 2023 in response to Russia's invasion of Ukraine?</i>
Corporate Activities	<i>Which Japanese startup has been developing perovskite solar cells since 2022?</i>
Social Issues	<i>What is carbon neutrality?</i>
Trends	<i>What is a dark store?</i>

Table 1: Examples of the business related questions translated from the original Japanese.

about current events, corporate activities, and social issues. Despite the success of LLMs in various language tasks, most models are pretrained on datasets consisting primarily of general English data. Consequently, their performance in other languages, including Japanese, and in domain-specific tasks may remain limited.

Our research aligns with two key directions in LLM development: language-specific models and domain-specific models. Language-specific pretrained models have been developed for several languages (Zhang et al., 2021; dbmdz, 2023; tokyotech llm, 2023), while domain-specific models have excelled in areas such as finance (Scao et al., 2023). Combining these aspects has been shown promising since the traditional advent of domain-specific BERTs (Beltagy et al., 2019; Ishigaki et al., 2023). However, many non-English languages, including Japanese, lack such combinations, especially for recent decoder-only architectures like GPTs (Brown et al., 2020). Addressing this gap requires developing fundamental domain-specific GPT. As a case study, we present the first Japanese LLM with 13 billion parameters specifically for the business domain.

LLMs in the business domain needs to be updated regularly to address the latest business-related queries. For example, LLMs pretrained before 2023 do not know results of the Olympic games in 2024. While the importance of continual

<sup>1</sup><https://huggingface.co/stockmark>

updates, there has been limited research on how we can better update LLMs with the latest knowledge without losing the general knowledge. This paper focuses on a research question about the data used for updates: how can we mix recent texts with general texts to ensure that the LLM incorporates up-to-date information while still retaining general knowledge? To address this gap, we continually pretrain our LLM with recent business documents with various ways of the mixture, ensuring that its knowledge remains current. This approach aims to improve the model’s accuracy in reflecting the latest business trends and information.

Evaluating a domain-specific model presents unique challenges. In addition to using general-domain benchmarks such as lm-evaluation-harness (Gao et al., 2023), we introduce a new business-specific benchmark. This benchmark consists of business questions across three tasks whose inputs are 1) question-only, 2) question with automatically retrieved context, and 3) question with manually retrieved context. Manual evaluations on this benchmark reveal that our pretrained model outperforms existing general-domain Japanese LLMs in accuracy, particularly in the question-only setting. Furthermore, models updated with the latest knowledge show improved accuracy in answering questions about recent events.

Our contributions are threefold: (1) we pretrain the first and largest Japanese business domain-specific LLM; (2) we present the first experiments on LLMs updated with the latest business documents and demonstrate that a proper mixture of recent and older texts in the training data is necessary; (3) we create a new benchmark with questions designed for three distinct tasks; and (4) we demonstrate the effectiveness of our pretrained model in domain-specific tasks. This study establishes a foundational environment for comparing future language- and domain-specific LLMs and provides valuable insights for researchers working with LLMs in other domains and languages. Our model and benchmark questions are publicly available <sup>2</sup>.

## 2 Methods

Our approach consists of three steps: collecting datasets, filtering them, and pretraining an LLM

<sup>2</sup><https://huggingface.co/stockmark/stockmark-13b>  
<https://huggingface.co/datasets/stockmark/business-questions>

from scratch. The trained LLM is later used for updating.

### 2.1 Dataset

To construct our pretraining dataset, we collected 19.8% domain-specific texts and 80.2% general-domain texts in Japanese.

For the domain-specific data, we created our original “Curated Business Corpus” and also used patent documents from the Japan Patent Office. These two corpora are designed to provide the business knowledge and technical terminology necessary for the target domain. The Curated Business Corpus was built by curating publicly available web pages published up to September 2023. We identified relevant pages using predefined URLs and a list of cue words, extracting pages that matched the URL patterns or contained at least one of the keywords. The URLs and cue words were selected to cover various aspects of the business domain, including chemistry, materials, biology, engineering, economics, current affairs, and social trends.

To ensure a diverse training dataset, we also included general-domain data from sources such as Wikipedia, CC100, mC4, and Common Crawl. These datasets provide essential general knowledge, which is crucial for handling a wide range of natural language processing tasks. Most other LLMs utilize similar general-domain datasets, though often with different filtering strategies.

### 2.2 Filtering

Filtering is essential for enhancing dataset quality. We implemented a three-step pipeline: language identification, removal of noisy characters, and deduplication.

**Language Identification:** Language identification is crucial for language-specific datasets. We used a two-method pipeline to identify non-Japanese documents: a library-based method and a language characteristics-based method. Initially, we employed the “xlm-roberta-base-language-detection” library <sup>3</sup>. For documents with uncertain results from this tool, we applied the franc library <sup>4</sup>. Any texts not recognized as Japanese were removed from the dataset.

**Noise Character Removal:** Low-quality Japanese texts often lack proper sentence struc-

<sup>3</sup><https://huggingface.co/papluca/xlm-roberta-base-language-detection>

<sup>4</sup><https://github.com/woorm/franc>

ture. Noisy texts may consist only of dates, HTML tags of menu bars, URLs, or lack end-of-sentence punctuation marks, such as “。” (Japanese period). We removed such texts if they were deemed non-sentential to ensure the dataset’s quality and to focus the model on relevant linguistic features. Additionally, because English sentences end with a period (“.”), using punctuation as a clue helps in refining English sentences. Our characteristics-based method, while simple, is effective in distinguishing languages with specific features, such as Thai which does not use punctuation.

**Deduplication:** Finally, we deduplicated the collected documents and sentences to eliminate identical entries. At the document level, we used Python’s built-in hash function to convert documents into hashed values, which allowed us to remove duplicates efficiently. At the sentence level, we counted the frequency of sentences and removed those appearing more than 15 times. Both document-level and sentence-level deduplication were performed using exact matching. This step prevents the model from fixating on repeated data and ensures dataset diversity.

Following these preprocessing steps, our dataset comprised a total of 220 billion tokens, as detailed in Table 2.

### 2.3 Pretraining

We used the Llama2 architecture (Touvron et al., 2023) with 13 billion parameters for our model. Our hyperparameters were aligned with those reported in the Llama2 paper. Instead of further training Meta’s Llama2 weights, our model was pretrained from scratch.

To balance the training data for each epoch, we adopted the weighting strategy used by Llama2, but doubled the amount of Wikipedia data for two main reasons. First, Wikipedia data is relatively clean. Second, it contains a wealth of content relevant to the business domain. Additionally, we doubled the size of our Curated Business Corpus to enhance the integration of domain-specific knowledge. By increasing the proportion of clean, domain-specific data, we aimed to mitigate the impact of noisy data sources like mC4 and Common Crawl.

For infrastructure, we utilized AWS’s Trainium, a hardware accelerator specifically designed for high-performance machine learning computations. We deployed our training scripts on 16 trn1.32xlarge instances, each equipped with Trainium, to create a robust distributed learning

Dataset	Num. of tokens [billion]
Curated business corpus	9.1
Patent	34.8
Wikipedia	1.0
CC100	10.9
mC4	53.2
Common Crawl	112.9

Table 2: The size of preprocessed dataset used for pre-training.

Dataset	Num. of examples	Language
Dolly	15,015	translated Japanese
OASST	88,838	translated Japanese
Alpaca	51,716	translated Japanese
Ichikara	10,329	Human-authored Japanese

Table 3: Candidates of datasets for instruction tuning.

environment. The distributed learning process was managed using the neuronx-nemo-megatron library<sup>5</sup>, which is available on AWS’s custom accelerators. The pretraining phase spanned 30 days to complete one epoch of training data.

### 2.4 Updating the Pretrained Model with Latest Business Documents

In real-world applications, such as domain-specific question answering (QA), LLMs must be updated to incorporate the most recent knowledge. The pretrained model discussed in the previous sections was trained on texts published up until September 2023. Consequently, it may struggle to answer questions about events occurring after October 2023.

To address this limitation, we continued training our pretrained model with the latest business documents published in October and November 2023. However, we are mindful of the issue of “catastrophic forgetting” (French, 1999), where acquiring new knowledge can unintentionally displace previously learned information.

To mitigate catastrophic forgetting, we employ a strategy that blends the latest business documents with randomly selected, non-latest documents from the Curated Business Corpus. This approach is inspired by previous research (Scialom et al., 2022). We introduce a hyperparameter  $r$ , which represents the proportion of instances sampled from the non-latest document set. For instance, if  $r$  is set to 0.3, then 30% of the continual update data comes from the Curated Business Corpus.

Our experiments in the next sections show how

<sup>5</sup><https://github.com/aws-neuron/neuronx-nemo-megatron>

different values of  $r$  affect the performance of domain-specific QA and analyze the extent of catastrophic forgetting.

### 3 Experiments

To evaluate our LLM, we have created a benchmark, Business Question Benchmark, for business domain question answering. We then compared our model against various Japanese LLMs using this benchmark, as well as common Japanese benchmarks, lm-evaluation-harness.

#### 3.1 Benchmarks

**Business Question Benchmark:** This benchmark consists of 50 questions written in natural language, each paired with relevant articles. The questions cover a range of topics, including recent events, company activities, social issues, and business trends. Each question is associated with web pages retrieved through automatic and manual methods.

The benchmark offers three QA settings; 1) **NoContext-QA:** In this setting, no web pages are provided as context. This allows us to assess the LLMs’ ability to generate answers based solely on their internal knowledge, without external information, 2) **AutoRAG-QA:** for this task, we use a search engine to find the most relevant web pages for each question. The highest-ranked page with available body text is selected as the context. This setting helps evaluate how well LLMs can generate answers considering both their existing knowledge and potentially irrelevant information from the web page, 3) **ManualRAG-QA:** Here, we manually select a web page that contains answers to the question. The RAG tasks are more about comprehension, requiring the model to understand and extract information from a specific page rather than relying solely on its internal knowledge.

For AutoRAG-QA and ManualRAG-QA, the text from each web page is truncated to 1000 characters. The prompts used for these tasks are detailed in Appendix 2.

Responses generated by the LLMs during the evaluation were manually assessed by an NLP researcher. The evaluation involved a binary judgment of responses as either correct or incorrect based on two main criteria: 1) **Content Faithfulness:** The response must accurately answer the question without any factual errors, 2) **Response Appropriateness:** If the question included specific instructions (e.g., “provide only one example”), the

response is considered correct only if it follows these instructions.

The evaluator considered a response correct if it met both criteria. Manual evaluation was preferred over automatic metrics (e.g., BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), BERTScore (Zhang et al., 2020)) because these metrics do not adequately assess factual correctness, which is crucial in the business domain. Additionally, redundant parts of responses, such as repeated sentences, were disregarded, focusing only on unique content.

**lm-evaluation-harness:** The lm-evaluation-harness framework <sup>6</sup> is a well-established tool for evaluating Japanese LLMs. It includes eight Japanese NLP tasks spanning various domains of language comprehension and generation. These tasks include reading comprehension (JSQuAD (Kurihara et al., 2022)), QA (JCommonsenseQA (Kurihara et al., 2022) and JAQKET <sup>7</sup>), Natural language inference (JNLI (Kurihara et al., 2022)), Summarization (XLSum-ja <sup>8</sup>), Co-reference resolution (WinoGrande-ja (Takahashi et al., 2023)), and Math problems (MGSM (Zhang et al., 2023)).

Due to the unavailability of the MARC-ja dataset (Kurihara et al., 2022), we conducted evaluations on the remaining seven tasks. Given that LLM performance can be significantly influenced by the prompts or templates used, we evaluated all available templates for each task and reported the highest score.

#### 3.2 Compared Models and Instruction Tuning

We compare our pretrained model with six existing Japanese LLMs and three multilingual LLMs. The Japanese LLMs are categorized into two groups: models pretrained from scratch and models that have undergone continual pretraining from multilingual models.

Among the models pretrained from scratch, we compare three: 1) **llm-jp-13b**(LLM-jp et al., 2024), which is pretrained on Japanese Wikipedia, mC4, English Wikipedia, The Pile (Gao et al., 2020), and The Stack (Kocetkov et al., 2022); 2) **plamo-13b**<sup>9</sup>, which is pretrained on English texts from RedPa-

<sup>6</sup><https://github.com/Stability-AI/lm-evaluation-harness>

<sup>7</sup><https://www.nlp.ecei.tohoku.ac.jp/projects/jaquet/>

<sup>8</sup>[https://huggingface.co/datasets/mkshing/xlsum\\_ja](https://huggingface.co/datasets/mkshing/xlsum_ja)

<sup>9</sup><https://huggingface.co/pfnet/plamo-13b>

jama and Japanese texts from mC4 and Japanese Wikipedia; and 3) **weblab-10b**<sup>10</sup>, which is pretrained on Japanese texts from mC4 and English texts from The Pile. The first two models have 13 billion parameters, while the last has 10 billion.

We also include models that were continually pretrained from existing multilingual models: 4) **nekomata-14b**(Sawada et al., 2024), which is trained from qwen-14b on Japanese CC100, C4, Wikipedia, Oscar, Rinna curated corpus, and The Pile; 5) **ELYZA-japanese-13b**<sup>11</sup>, which is trained from Llama-2-13b on Japanese texts from Oscar and Wikipedia; and 6) **Swallow-13b**(Fujii et al., 2024), which is trained from Llama-2-13b on Japanese texts from Wikipedia, meticulously cleaned Common Crawl, and English texts from RefinedWeb (Penedo et al., 2023) and The Pile. Models 4) has 14 billion parameters, while the others have 13 billion.

Additionally, we compare our models with multilingual LLMs: 7) **gpt-3.5-turbo-0125**, 8) **gpt-4-1106-preview**, selected from the OpenAI API; and 9) **Llama-2-13b-hf** (Touvron et al., 2023)<sup>12</sup>, a 13-billion parameter English LLM capable of generating Japanese text.

LLMs that are not instruction-tuned often struggle to provide accurate answers, highlighting the importance of instruction-tuning. We evaluated several datasets for instruction-tuning in Japanese, including Ichikara (Sekine et al., 2023), Alpaca (Taori et al., 2023; Shimizu, 2023), Dolly (Conover et al., 2023; Kuniyoshi, 2023a), and OASST (Köpf et al., 2023; Kuniyoshi, 2023b). Statistics for these datasets are provided in Table 3. Preliminary experiments shown in Table 4 reveal that the use of Ichikara dataset consistently performs well on question-answering tasks, with scores of 0.78 on JSQuAD, 0.86 on JAQKET, and 0.84 on JCommonsenseQA. Therefore, we chose the Ichikara dataset for instruction-tuning our LLMs. We use Low-Rank Adaptation (LoRA) (Hu et al., 2021) for the instruction tuning.

## 4 Results

### 4.1 Results on Business Question Benchmark:

Table 5 displays the results from our domain-specific benchmark, categorizing the models into

<sup>10</sup><https://huggingface.co/matsuo-lab/weblab-10b>

<sup>11</sup><https://huggingface.co/elyza/ELYZA-japanese-Llama-2-13b>

<sup>12</sup><https://huggingface.co/elyza/ELYZA-japanese-Llama-2-13b>

three groups: Japanese LLMs pretrained from scratch, Japanese LLMs with continual pretraining, and multilingual LLMs.

In the **NoContext-QA** setting, our pretrained model achieves the highest accuracy of 0.90. This is notably higher compared to other Japanese LLMs such as nekomata-14b and Swallow-13b, which both score 0.78. This performance indicates that our model effectively utilizes domain-specific knowledge stored internally without relying on external documents.

For the **ManualRAG-QA** task, our model again performs best among LLMs pretrained from scratch, achieving an accuracy of 0.84. In contrast, the other models in this category score 0.76 (llm-jp-13b), 0.82 (plamo-13b), and 0.52 (weblab-10b). A similar trend is observed in the **AutoRAG-QA** task, where our model scores 0.74, while other models score 0.62 (llm-jp-13b), 0.64 (plamo-13b), and 0.34 (weblab-10b), respectively. In these RAG settings, which require comprehension and the ability to extract information from documents, our models perform better than other full-scratch LLMs. However, Japanese LLMs with continual pretraining, such as ELYZA-japanese-13b and Swallow-13b, achieve higher scores of 0.94 and 0.90 in **ManualRAG-QA**, respectively. These values suggest that continual pretraining works better in comprehension.

When comparing pretraining strategies, continual pretraining outperforms full-scratch models except for our model. For instance, in the **NoContext-QA** task, models with continual pretraining, such as Swallow-13b and nekomata-14b, achieve scores of 0.78, whereas the best pretrained full-scratch baseline, llm-jp-13b, scores 0.34. Despite the overall advantage of continual pretraining, our model performs the best in the **NoContext-QA** task, highlighting the value of domain-specific data to strengthen the internal domain-specific knowledge in LLMs.

For tasks involving context, such as **ManualRAG-QA** and **AutoRAG-QA**, our model achieves the highest scores of 0.84 and 0.74, respectively, among full-scratch models. However, models with continual pretraining outperform in these tasks with scores of 0.94 and 0.82. This may be attributed to the use of Llama2 as a base model for continual pretraining, which benefits from being pretrained on a larger dataset and thus has stronger language comprehension capabilities.

model	JSQuAD	JAQKET	JCom.	XW.	JNLI	MGSM	XLSum	Ave.
Ichikara	<b>0.78</b>	0.85	<b>0.84</b>	0.75	0.49	0.08	0.08	<b>0.550</b>
Alpaca	0.73	0.85	0.81	0.73	<b>0.57</b>	0.07	0.07	0.545
Dolly	0.77	<b>0.87</b>	0.81	0.73	0.53	0.08	0.08	0.547
OASST	0.77	0.85	0.81	0.74	0.25	0.07	0.08	0.510

Table 4: Preliminary experiments: the comparison between datasets for instruction tuning. All base-model is our proposed 13-billion Japanese business domain specific LLM. JCom. and XW. refer to the JCommonsenseQA and XWinograd datasets, respectively. The Ave. column shows the averaged values over all datasets.

Although this study focuses on models pre-trained from scratch, future work could explore the potential performance gains from continually training a strong existing model, such as Llama2, on a domain-specific corpus.

## 4.2 Results on General-domain Benchmark:

Table 6 presents the results on the general-domain benchmark, “lm-evaluation-harness.” Among the models pretrained from scratch, our model demonstrates superior performance in terms of the averaged score i.e., ours achieves 0.55 on average while the best performing Japanese LLMs pretrained from scratch achieve 0.49. In particular, in QA tasks without context, such as JAQKET and JCommonsenseQA, our model achieves the best scores 0.78 and 0.85, respectively.

## 4.3 Analysis of Updated Models

We compare the updated model with the original pretrained model. The left part of Table 7 shows the validation loss values on two datasets: the 8,192 documents that were used for the validation portion during pretraining (Pretrain Data) and the latest business documents used for updating the model (Latest Data). If our LLM were able to learn the latest knowledge without losing general knowledge, it would achieve low validation loss on both Pretrain Data and Latest Data. The loss value for the updated model with  $r = 0.0$ , indicating that it was trained solely on the latest documents, is 2.05 on Latest Data. This represents an improvement over the original pretrained model’s loss of 2.25, suggesting that the updated model acquired the latest knowledge better. However, the loss on the Pretrain Data increased from 2.11 to 2.19 compared to the original model, which suggests the worse fit to the non-latest documents.

Increasing the value of  $r$ , which incorporates more pretraining data along with the latest documents, mitigates the increase in loss on the pretraining data. For example, with  $r$  set to 0.3, the loss is 2.12 on Pretrain Data, which is nearly equivalent

to the original model’s loss of 2.11, while the loss on the latest documents remains stable at 2.04.

The right part of Table 7 shows the accuracy of our compared updated LLMs in question answering about latest news. We created a set of 10 questions (LatestQ) about recent business topics from October to November 2023. We selected topics that had shown a notable increase in search engine access compared to September 2023. As result, all the updated models achieve a higher accuracy (0.90) compared to the pretrained model (0.30), indicating that a Japanese- and business-specific model can acquire new information through continual learning. Unlike the results regarding loss, we did not observe significant differences in accuracy among models with different values of  $r$  on the Business Question Benchmark (Non-LatestQ) where we obtained high values e.g., 0.90 (Pretrained model and Updated model( $r = 0.3$ )) or 0.92 (Updated model ( $r = 0.0, 0.1$ )). We conclude that using only the latest articles can lead to degradation in both loss and accuracy and incorporating around 10% of non-latest articles proves effective, while including 30% is excessive.

## 4.4 Examples of Outputs

Table 8 shows the examples of the outputs from our pretrained model and compared LLMs. When we input a question “Which two banks failed in March 2023?”, our pretrained model correctly generates two banks: the Silicon Valley Bank and the Signature Bank. Other three Japanese LLMs’ outputs contain hallucinations, which are indicated in bold in the table. For example, Wells Fargo, Bank of America, and Citigroup did not fail in March 2023. We observed a similar tendency for the other questions.

## 5 Related Work

Recent major LLMs are multilingual models such as Llama2 (Touvron et al., 2023), OpenAI’s GPTs (OpenAI), and BLOOM (Scao et al., 2023). The datasets used for such major models often in-

model	NoContext-QA	ManualRAG-QA	AutoRAG-QA
Japanese LLMs trained from scratch			
- Our model	<b>0.90</b>	<b>0.84</b>	<b>0.74</b>
- 1) llm-jp-13b	0.34	0.76	0.62
- 2) plamo-13b	0.34	0.82	0.64
- 3) weblab-10b	0.26	0.52	0.34
Japanese LLMs with continual pretraining			
- 4) nekomata-14b	<b>0.78</b>	0.74	0.76
- 5) ELYZA-japanese-13b	0.32	<b>0.94</b>	0.70
- 6) Swallow-13b	<b>0.78</b>	0.90	<b>0.82</b>
Multilingual LLMs			
- 7) gpt-3.5-turbo-0125	0.54	0.62	0.34
- 8) gpt-4-1106-preview	<b>0.78</b>	<b>0.94</b>	<b>0.86</b>
- 9) Llama-2-13b-hf	0.24	0.84	0.64

Table 5: Results on business-specific benchmark. Each score stands for the accuracy. The values obtained from the best performing models in each category is shown in bold.

model	JSQuAD	JAQKET	JCom.	XW.	JNLI	MGSM	XLSum	Ave.
Japanese LLMs Pretrained from Scratch								
- Ours	<b>0.78</b>	<b>0.85</b>	<b>0.84</b>	<b>0.75</b>	<b>0.49</b>	<b>0.08</b>	0.08	<b>0.55</b>
- 1) weblab-10b	0.72	0.43	0.65	0.67	0.30	0.02	0.05	0.41
- 2) plamo-13b	0.68	0.69	0.64	0.68	0.41	0.02	<b>0.10</b>	0.46
- 3) llm-jp-13b	0.69	0.76	0.79	0.70	0.37	0.02	<b>0.10</b>	0.49
Japanese LLMs with Continual Pretraining								
- 4) nekomata-14b	<b>0.87</b>	0.88	<b>0.94</b>	<b>0.80</b>	<b>0.65</b>	<b>0.36</b>	<b>0.22</b>	<b>0.68</b>
- 5) ELYZA-13b	0.79	0.75	0.87	0.78	0.51	0.10	0.19	0.57
- 6) Swallow-13b	0.86	<b>0.91</b>	0.91	0.72	0.52	0.18	0.20	0.61
Multilingual LLMs pretrained from Scratch								
- 9) Llama-2-13b-hf	0.81	0.75	0.82	0.63	0.47	0.12	0.21	0.54

Table 6: Results on lm-evaluation-harness. JCom. and XW. refer to the JCommonsenseQA and XWinograd datasets, respectively. The Ave. column shows the averaged values over all datasets. OpenAI’s GPTs cannot be used for this experiments because the model parameters are not public.

clude a high percentage of English data. Therefore, performances in languages other than English have been still underexplored. Studies targeting non-English languages, such as Chinese (Zhang et al., 2021), German (dbmdz, 2023), and Japanese (tokyotech llm, 2023), increase the proportion of non-English texts used in pretraining. Whereas, we propose an approach to build a large size of Japanese specific corpus by filtering out other languages with language identification libraries and a noise character detection, and we pretrain Japanese-specific LLM.

The language-specific datasets for instruction tuning have been released. For Japanese, Ichikara, Alpaca, Dolly, and OASST are common; however, the performances of models trained on these datasets have not been compared in depth. Our preliminary experiment is the first to compare these datasets, which can be considered as an important contribution for the Japanese LLM commu-

nity. Also, our experiments can provide insights for researchers who focus on other non-English languages.

Domain-specific models are pretrained in two different ways; training from scratch or continual training. Representative examples of the former started for the encoder-only models such as SciBERT (Beltagy et al., 2019) in the science domain and later decoder-only large language models follow, e.g., BloombergGPT (Wu et al., 2023). The continually learned domain-specific LMs include BioBERT (Lee et al., 2019), exBERT (Wang et al., 2020). Scratch approaches are often used for settings where we can get sufficient data. Our setting is categorized into this setting, thus, we use the scratch approach.

The combination of the two directions, i.e., language- and domain-specific settings, is underexplored in Japanese. Japanese domain-specific language model does not exist, except for the pre-

	Loss ↓		Accuracy ↑	
	Pretrain Data	Latest Data	Non-LatestQ.	LatestQ.
Pretrained model	<b>2.11</b>	2.25	0.90	0.30
Updated model ( $r = 0.0$ )	2.19	<b>2.05</b>	<b>0.92</b>	<b>0.90</b>
Updated model ( $r = 0.1$ )	2.14	<b>2.05</b>	<b>0.92</b>	<b>0.90</b>
Updated model ( $r = 0.3$ )	2.12	<b>2.04</b>	0.90	<b>0.90</b>

Table 7: Comparisons of validation losses on the pretrain data and the latest data, and accuracies on the questions about the latest business and non-latest business knowledge (Latest Q. and Non-latest Q.).

Question : Which two American banks went bankrupt in March 2023?	
Model	Generated Text
Ours	The two U.S. banks that failed in March 2023 were Silicon Valley Bank and Signature Bank. (2023年3月に経営破綻した米国の2つの銀行は、シリコンバレー銀行とシグネチャー銀行である。)
llm-jp-13b-v1.0	The two U.S. banks that failed in March 2023 were <b>Wells Fargo and Bank of America</b> . (2023年3月に経営破綻したアメリカの2つの銀行は、ウェルズ・ファーコートとバンク・オブ・アメリカである。)
plamo-13b	The two banks that failed in March 2023 were <b>Bank of America and Citigroup</b> . (2023年3月に経営破綻した2つの銀行は、バンク・オブ・アメリカとシティ・グループである。)
weblab-10b	The two U.S. banks that failed in March 2023 were <b>Bank of America and Wells Fargo</b> . (2023年3月に経営破綻したアメリカの2つの銀行は、Bank of America と Wells Fargo である。)
ChatGPT (gpt-3.5-turbo-0125)	The two American banks that went bankrupt in March 2023 are <b>Regional Asahi Bank and Centry National Bank</b> . (2023年3月に経営破綻したアメリカの2つの銀行は、リージョナル朝陽銀行とセントリーナショナル銀行です。)

Table 8: Translated examples of output generated from different LLMs. The sentences in the buckets are the original outputs in Japanese. Hallucinated terms are presented in bold.

trained BERT for the material science domain (Ishigaki et al., 2023). Our target is decoder-only architecture with many more parameters for the business domain, which has high demand in the industry but is less studied. Our pretrained model is the first Llama-based domain-specific model for Japanese.

Continual pretraining is a promising direction if we have only a small dataset for pretraining. Our experiments for updating the pretrained model with the latest news are categorized in this setting. ‘‘Catastrophic forgetting’’ is a major problem in this case (Ling et al., 2023). Scialom et al. (2022) suggest that the mixture of two types of data can mitigate this problem, thus, we used the technique to mix the latest documents and older ones.

## 6 Conclusion

This paper presented the first Japanese business domain-specific LLM. We pretrain the LLM with 13 billion parameters from scratch and the model is released to be publicly available. We also update the model parameters by the latest articles and confirmed performance gains. For updates,

we conclude that using only the latest articles can lead to performance degradation but incorporating around 10% of non-latest articles proves effective. Comprehensive evaluations demonstrated that our LLM achieves the best accuracy score without retrieval on the domain-specific benchmark we newly released. These results provide valuable insights for researchers working on other domains and languages. For future work, we will explore comparing different ways to train language- and domain-specific LLMs e.g., full-stratch v.s. continual pretraining.

## 7 Ethics Statement

The proposed model is still in the early stages of research and development, and its output has not yet been adjusted to align with safety considerations. However, adjustments will be made in the future to ensure that the model takes safety into account.



## References

- Satanjeev Banerjee and Alon Lavie. 2005. **METEOR: An automatic metric for MT evaluation with improved correlation with human judgments**. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. **SciBERT: A pretrained language model for scientific text**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP2019)*, pages 3615–3620, Hong Kong, China. 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.
- Mike Conover, Matt Hayes, Ankit Mathur, Jianwei Xie, Jun Wan, Sam Shah, Ali Ghodsi, Patrick Wendell, Matei Zaharia, and Reynold Xin. 2023. **Free dolly: Introducing the world’s first truly open instruction-tuned llm**.
- dbmdz. 2023. **German gpt-2 model**. Online.
- Robert M. French. 1999. **Catastrophic forgetting in connectionist networks**. *Trends in Cognitive Sciences*, 3(4):128–135.
- Kazuki Fujii, Taishi Nakamura, Mengsay Loem, Hiroki Iida, Masanari Ohi, Kakeru Hattori, Hirai Shota, Sakae Mizuki, Rio Yokota, and Naoaki Okazaki. 2024. **Continual pre-training for cross-lingual LLM adaptation: Enhancing japanese language capabilities**. In *First Conference on Language Modeling*.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, Shawn Presser, and Connor Leahy. 2020. **The pile: An 800gb dataset of diverse text for language modeling**. *Preprint*, arXiv:2101.00027.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2023. **A framework for few-shot language model evaluation**.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. **Lora: Low-rank adaptation of large language models**.
- Tatsuya Ishigaki, Yui Uehara, Goran Topić, and Hiroya Takamura. 2023. **Pretraining language-and domain-specific bert on automatically translated text**. In *Proceedings of the 14th International Conference on Recent Advances in Natural Language Processing*, pages 548–555.
- Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis, Yacine Jernite, Margaret Mitchell, Sean Hughes, Thomas Wolf, Dzmitry Bahdanau, Leandro von Werra, and Harm de Vries. 2022. **The stack: 3 tb of permissively licensed source code**. *Preprint*, arXiv:2211.15533.
- Shohei Kuniyoshi. 2023a. **databricks-dolly-15k-ja**.
- Shohei Kuniyoshi. 2023b. **oasst1-89k-ja**.
- Kentaro Kurihara, Daisuke Kawahara, and Tomohide Shibata. 2022. **JGLUE: Japanese general language understanding evaluation**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2957–2966, Marseille, France. European Language Resources Association.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi-Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. 2023. **Openassistant conversations – democratizing large language model alignment**. *Preprint*, arXiv:2304.07327.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2019. **Biobert: a pre-trained biomedical language representation model for biomedical text mining**. *Bioinformatics*.
- Chen Ling, Xujiang Zhao, Jiaying Lu, Chengyuan Deng, Can Zheng, Junxiang Wang, Tanmoy Chowdhury, Yun Li, Hejie Cui, Xuchao Zhang, Tianjiao Zhao, Amit Panalkar, Wei Cheng, Haoyu Wang, Yanchi Liu, Zhengzhang Chen, Haifeng Chen, Chris White, Quanquan Gu, Jian Pei, and Liang Zhao. 2023. **Domain specialization as the key to make large language models disruptive: A comprehensive survey**. *Preprint*, arXiv:2305.18703.
- LLM-jp, :, Akiko Aizawa, Eiji Aramaki, Bowen Chen, Fei Cheng, Hiroyuki Deguchi, Rintaro Enomoto, Kazuki Fujii, Kensuke Fukumoto, Takuya Fukushima, Namgi Han, Yuto Harada, Chikara

Hashimoto, Tatsuya Hiraoka, Shohei Hisada, Sosuke Hosokawa, Lu Jie, Keisuke Kamata, Teruhito Kanazawa, Hiroki Kanezashi, Hiroshi Kataoka, Satoru Katsumata, Daisuke Kawahara, Seiya Kawano, Atsushi Keyaki, Keisuke Kiryu, Hirokazu Kiyomaru, Takashi Kodama, Takahiro Kubo, Yohei Kuga, Ryoma Kumon, Shuhei Kurita, Sadao Kurohashi, Conglong Li, Taiki Maekawa, Hiroshi Matsuda, Yusuke Miyao, Kentaro Mizuki, Sakae Mizuki, Yugo Murawaki, Ryo Nakamura, Taishi Nakamura, Kouta Nakayama, Tomoka Nakazato, Takuro Niitsuma, Jiro Nishitoba, Yusuke Oda, Hayato Ogawa, Takumi Okamoto, Naoaki Okazaki, Yohei Oseki, Shintaro Ozaki, Koki Ryu, Rafal Rzepka, Keisuke Sakaguchi, Shota Sasaki, Satoshi Sekine, Kohei Suda, Saku Sugawara, Issa Sugiura, Hiroaki Sugiyama, Hisami Suzuki, Jun Suzuki, Toyotaro Suzumura, Kensuke Tachibana, Yu Takagi, Kyosuke Takami, Koichi Takeda, Masashi Takeshita, Masahiro Tanaka, Kenjiro Taura, Arseny Tolmachev, Nobuhiro Ueda, Zhen Wan, Shuntaro Yada, Sakiko Yahata, Yuya Yamamoto, Yusuke Yamauchi, Hitomi Yanaka, Rio Yokota, and Koichiro Yoshino. 2024. [Llm-jp: A cross-organizational project for the research and development of fully open japanese llms](#). *Preprint*, arXiv:2407.03963.

OpenAI. [Openai api](#). Online.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, pages 311–318. Association for Computational Linguistics.

Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. [The refinedweb dataset for falcon llm: Outperforming curated corpora with web data, and web data only](#). *Preprint*, arXiv:2306.01116.

Kei Sawada, Tianyu Zhao, Makoto Shing, Kentaro Mitsui, Akio Kaga, Yukiya Hono, Toshiaki Wakatsuki, and Koh Mitsuda. 2024. [Release of pre-trained models for the Japanese language](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13898–13905. <https://arxiv.org/abs/2404.01657>.

Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gellé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron

Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jörg Froberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, María Grandury, Mario Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto Luis López, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, Shayne Longpre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vasilina Nikoulina, Veronika Laippala, Violette Lepereq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesh Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Saiful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zhengxin Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, Sanjit Gandhi, Shaden Smith, Stéphane Requena, Suraj Patil, Tim Dettmers, Ahmed Baruwah, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névél, Charles Lovering, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, Jan-Christoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Na-

- joung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, Shachar Mirkin, Shani Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Uldreaj, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Ajibade, Bharat Saxena, Carlos Muñoz Ferrandis, Daniel McDuff, Danish Contractor, David Lansky, Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani, Fatima Mirza, Frankline Onon-iwu, Habib Rezanejad, Hessie Jones, Indrani Bhattacharya, Irene Solaiman, Irina Sedenko, Isar Nadjdholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Rajani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Alizadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguier, Thanh Le, Tobi Oyebade, Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourrier, Daniel León Perinán, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrmann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthik Rangasai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A Castillo, Marianna Nezhurina, Mario Sängler, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel De Wolf, Mina Mihaljic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chandrasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sinee Sang-aroonsiri, Srishti Kumar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2023. [Bloom: A 176b-parameter open-access multilingual language model](#). *Preprint*, arXiv:2211.05100.
- Thomas Scialom, Tuhin Chakrabarty, and Smaranda Muresan. 2022. [Fine-tuned language models are continual learners](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6107–6122, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Satoshi Sekine, Maya Ando, Hisami Suzuki, Daisuke Kawahara, Naoya Inoue, and Kentaro Inui. 2023. [Ichikara : Japanese instruction dataset for llms](#).
- Ryo Shimizu. 2023. [alpaca\\_ja](#).
- Keigo Takahashi, Teruaki Oka, and Mamoru Komachi. 2023. [Effectiveness of pre-trained language models for the japanese winograd schema challenge](#). *Journal of Advanced Computational Intelligence and Intelligent Informatics*, 27(3):511–521.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. [Stanford alpaca: An instruction-following llama model](#). [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- tokyotech llm. 2023. [Swallow](#). Online.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yunying Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *Preprint*, arXiv:2307.09288.
- Zihan Wang, Karthikeyan K, Stephen Mayhew, and Dan Roth. 2020. [Extending multilingual BERT to low-resource languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2649–2656, Online. Association for Computational Linguistics.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David Rosenberg, and Gideon Mann. 2023. [Bloomberggpt: A large language model for finance](#). *Preprint*, arXiv:2303.17564.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations*,

ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net.

Zhengyan Zhang, Xu Han, Hao Zhou, Pei Ke, Yuxian Gu, Deming Ye, Yujia Qin, Yusheng Su, Haozhe Ji, Jian Guan, et al. 2021. Cpm: A large-scale generative chinese pre-trained language model. *AI Open*, 2:93–99.

Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2023. Multi-modal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.

## A Appendix

### A.1 Templates used in Experiments

Table 9 shows the versions of templates used in our experiments.

### A.2 Prompts used in Experiments

Please answer the question briefly. Question: {question} ### Output:
--

Figure 1: Translated prompt for NoContext-QA. In the experiment, Japanese prompt was used.

Please answer to the given question. If the answer to the question is included in the article text, please use the answer from the text. If the article does not contain the answer please state that "the article does not contain the answer" and answer the question using your knowledge. Question: {question} Article Text: {the first 1k characters} ### Output:
---

Figure 2: Translated prompt for ManualRAG-QA and AutoRAG-QA. In the experiment, Japanese prompt was used.

Task	Task Version	Prompt Version	Number of Few-shot
JSQuAD	1.1	0.1, 0.2, 0.3	2
JAQKET	v2-0.2	0.1, 0.2, 0.3	1
JCommonsenseQA	1.1	0.1, 0.2.1, 0.3	3
JWinograd	ja	-	0
JNLI	1.3	0.2, 0.3	3
MGSM	1.0	0.0, 0.3	5
XLSum	ja-1.0	0.0, 0.3	1

Table 9: Settings of the template of the lm-evaluation-harness. We experimented every model in Table 6 with the templates listed here. The scores reported in Table 6 are the highest among the variants of the templates.