

KazMMLU: Evaluating Language Models on Kazakh, Russian, and Regional Knowledge of Kazakhstan

Mukhammed Togmanov^{1*} Nurdaulet Mukhituly^{1*} Diana Turmakhan^{1*}
Jonibek Mansurov¹ Maiya Goloburda¹ Akhmed Sakip¹ Zhuohan Xie¹
Yuxia Wang¹ Bekassyl Syzdykov¹ Nurkhan Laiyk¹ Alham Fikri Aji¹
Ekaterina Kochmar¹ Preslav Nakov^{1,2} Fajri Koto¹

¹Mohamed bin Zayed University of Artificial Intelligence

²Institute of Foundation Models

{mukhammed.togmanov,nurdaulet.mukhituly,diana.turmakhan}@mbzuai.ac.ae

Abstract

Despite having a population of twenty million, Kazakhstan’s culture and language remain underrepresented in the field of natural language processing. Although large language models (LLMs) continue to advance worldwide, progress in Kazakh language has been limited, as seen in the scarcity of dedicated models and benchmark evaluations. To address this gap, we introduce KazMMLU, the first MMLU-style dataset specifically designed for Kazakh language. KazMMLU comprises 23,000 questions that cover various educational levels, including STEM, humanities, and social sciences, sourced from authentic educational materials and manually validated by native speakers and educators. The dataset includes 10,969 Kazakh questions and 12,031 Russian questions, reflecting Kazakhstan’s bilingual education system and rich local context. Our evaluation of several state-of-the-art multilingual models (Llama3.1, Qwen-2.5, GPT-4, and DeepSeek V3) demonstrates substantial room for improvement, as even the best-performing models struggle to achieve competitive performance in Kazakh and Russian. These findings highlight significant performance gaps compared to high-resource languages. We hope that our dataset will enable further research and development of Kazakh-centric LLMs.¹

1 Introduction

With a population exceeding twenty million, the Republic of Kazakhstan in Central Asia remains underrepresented in the field of natural language processing (NLP) (Joshi et al., 2020). This gap is highlighted by the limited progress in developing large language models (LLMs) and evaluation benchmarks specifically tailored to the languages and the cultural context of Kazakhstan. Kazakh, a

*These authors contributed equally.

¹KazMMLU can be accessed at <https://huggingface.co/datasets/MBZUAI/KazMMLU>.




Figure 1: Overview of the dataset. This diagram illustrates the distribution of questions by educational level (High School and University) and language (Kazakh and Russian), along with the variety of subjects covered.

Turkic language spoken by more than fourteen million people (around 70% of the population), holds substantial cultural and geopolitical significance in Central Asia. Russian, used by approximately 15% of the population, serves as the country’s second primary language.²

Although Kazakh appears in certain multilingual datasets (Yeshpanov et al., 2024, 2022; Yeshpanov and Varol, 2024), most of these resources rely heavily on translations from English, lacking the cultural richness essential for inclusive LLM development. Previous work has primarily addressed classic NLP tasks such as named entity recognition (Yeshpanov et al., 2022) and sentiment analysis (Yeshpanov and Varol, 2024). Meanwhile, recent developments in LLM research have shifted toward more reasoning-focused evaluation (Meta AI, 2024; OpenAI, 2024), highlighting a clear research

²<https://glottolog.org/>

gap for inclusive NLP in the Kazakh context.

Here, we aim to bridge this gap. In particular, we introduce KazMMLU, a curated dataset of school- and university-level questions from Kazakhstan, available in both Kazakh and Russian. KazMMLU follows the framework of the Massive Multitask Language Understanding (MMLU) dataset (Hendrycks et al., 2021; Koto et al., 2024; Li et al., 2024; Koto et al., 2023), which features multiple-choice questions across various subjects and education levels. MMLU has become a standard benchmark for evaluating LLMs’ reasoning and knowledge capabilities (Meta AI, 2024; Team et al., 2024; Yang et al., 2025). Unlike general MMLU, KazMMLU incorporates Kazakhstan-specific content, including topics on Kazakh history, traditions, and linguistics, while also reflecting the country’s multilingual landscape by providing questions in both Kazakh and Russian.

As shown in Figure 1, the dataset is divided into two categories: High School and University. KazMMLU consists of approximately 48% of the questions in Kazakh and 52% in Russian. The High School section includes questions in both Kazakh and Russian, covering subjects such as Mathematics, Physics, and Kazakh Literature. The University section only features questions in Russian, focusing on professional disciplines such as Law, Economics, and Medicine. This structure aligns with Kazakhstan’s bilingual education system and provides a more representative benchmark for evaluating LLMs in the region. KazMMLU is sourced from authentic educational materials, including national exams, textbooks, and professional certification repositories. Each question is accompanied by metadata, including the subject, level, source, and correct answer key, ensuring transparency and usability for downstream evaluations.

Our contributions can be summarized as follows:

- We present the first **MMLU-style dataset** specifically tailored to the Kazakhstan context, covering diverse subject areas across educational and professional levels. The dataset is made available in both Kazakh and Russian.
- We evaluate various **multilingual LLMs**, including Llama3.1 (Meta AI, 2024), Qwen (Yang et al., 2025), GPT-4o (OpenAI, 2024), BLOOMZ (Muennighoff et al., 2023), mT0 (Muennighoff et al., 2023), and DeepSeek V3 (DeepSeek-AI, 2024), across different model sizes.

- We conduct a thorough analysis of the top-performing open-source models across various dimensions, encompassing (1) individual **subject areas, educational levels, and Kazakhstan-specific topics**, (2) **few-shot inference performance**, (3) **model confidence**, and (4) the **influence of negation** on model performance. This comprehensive evaluation framework allows us to identify the performance gaps and opportunities for improvement in multilingual LLMs when applied to Kazakh and Russian contexts.

2 Related Work

Language Models in Kazakh and Russian
Prominent models such as OpenAI’s ChatGPT, Anthropic’s Claude, and Yandex’s Yandex-GPT are designed to handle multiple languages, including Russian and Kazakh, enabling a wide range of applications from translation to content generation (OpenAI, 2024; Anthropic, 2023; Yandex, 2023). Additionally, open-source models like Meta’s Llama series provide multilingual support (Meta AI, 2024). While these models can produce text in Kazakh, they were not specifically trained or fine-tuned for it. In contrast, the Aya model, an open-access multilingual LLM, supports 101 languages, including Kazakh (Üstün et al., 2024).

Kazakh-specific language models remain limited, with most multilingual models providing only partial support. To address this, ISSAI (2024) and Koto et al. (2025) introduced KazakhLLM and Sherkala, both fine-tuned on Kazakh data from Llama. However, their evaluations have largely relied on machine-translated benchmarks.

NLP Benchmark for Kazakhstan Context
Evaluating LLMs across diverse linguistic and cultural contexts is increasingly critical; however, existing benchmarks overlook Kazakhstan. While benchmarks such as XCOPA (Ponti et al., 2020), XGLUE (Liang et al., 2020), and XTREME (Hu et al., 2020) assess cross-lingual performance, they exclude Kazakh, and GlobalMMLU (Singh et al., 2024) lacks Kazakhstan-specific content. Recent efforts like Kardeş-NLU (Senel et al., 2024) explore cross-lingual transfer for Turkic languages including Kazakh, but focus on NLU tasks specifically, without domain-specific reasoning or local context. In contrast, several Kazakh-specific datasets exist, including KazNERD (Yesh-

panov et al., 2022) for named entity recognition, KazQAD (Yeshpanov et al., 2024) for question answering, and KazSANDRA (Yeshpanov and Varol, 2024) for sentiment analysis. However, these datasets focus on narrow tasks and do not assess reasoning, factual recall, or domain-specific knowledge. TUMLU (Isbarov et al., 2025) provides a native-language benchmark for Turkic languages including Kazakh, but lacks Kazakhstan-specific context or domain-specific evaluation. To address these limitations, KazMMLU presents a large-scale, Kazakhstan-specific benchmark covering STEM, humanities, and social sciences. Unlike previous datasets, KazMMLU supports a holistic evaluation of reasoning and domain-specific knowledge, offering a more accurate assessment of multilingual LLM capabilities and advancing AI for low-resource languages.

To further illustrate the differences between KazMMLU and previous benchmarks, we compare it with two existing datasets, SIGTURK (Maxutov et al., 2024) and INCLUDE (Romanou et al., 2024), in Table 1. As shown, KazMMLU is the only dataset that incorporates **real-world educational materials**, **professional subjects**, and **domain-specific reasoning** in both Kazakh and Russian, offering a more **localized** and **comprehensive** evaluation of LLMs. This structured assessment underscores the **importance of country-specific benchmarks** in multilingual NLP research and contributes to bridging the gap in Kazakh language understanding.

3 KazMMLU

In Kazakhstan, the K-12 education system operates in a multilingual setting, with Kazakh as the primary language of instruction in most schools. However, Russian and other languages, such as Uzbek, Uyghur, and Tajik, are also used in specific regions. The curriculum includes core subjects such as mathematics, science, and history, with students required to study both Kazakh and Russian as part of their language education. At the university and professional levels, Russian remains the dominant language of instruction, especially in fields such as law, medicine, economics, and engineering.

KazMMLU spans a broad range of subjects across STEM, humanities, social sciences, and professional studies. As shown in Figure 2, the dataset has strong coverage in STEM fields, with Biol-




Figure 2: Subject-wise distribution of questions in KazMMLU.

ogy, Mathematics, and Physics comprising a significant portion of the questions. Humanities and social sciences are also well represented, particularly Kazakh History, World History, and Law, reflecting their central role in the Kazakh educational system.

In addition to academic subjects, KazMMLU includes questions from professional domains such as Economics, Finance, Jurisprudence, and Medicine, enabling the evaluation of language models in specialized areas. The dataset features content in both Kazakh and Russian, ensuring balanced linguistic representation. All questions are written in the Cyrillic script, which remains the prevailing standard in Kazakhstan’s formal education system. This broad subject and language coverage makes KazMMLU a robust benchmark for evaluating multilingual language models across diverse domains.

To illustrate the question format in KazMMLU, Figure 3 presents a sample multiple-choice question in Kazakh. Answering this question requires an understanding of Kazakhstan’s local context, as it covers topics such as history and geography. This example highlights the dataset’s diversity across subjects and difficulty levels.

3.1 Data Construction

To construct KazMMLU, we adopted a systematic approach inspired by MMLU datasets (Koto et al., 2023, 2024; Li et al., 2024). The dataset comprises questions sourced from national exams, textbooks, and professional certification materials such as iTest.kz, ymnik.kz, oltest.kz and Book - Shyn Kitap. To ensure diversity, we employed three data collection strategies: (1) automated online crawling, (2) manual transcription from scanned books, and (3) manual extraction from online sources. A

Feature	KazMMLU	SIGTURK	INCLUDE
Public Dataset Size	23,000 questions in Kazakh and Russian	3,000 questions exclusively in Kazakh	23,741 total questions, including 500 in Kazakh
Languages Covered	Kazakh, Russian	Kazakh	44 (including Kazakh)
Kazakh-Specific Content	Yes , sourced from local curriculum, national exams	Limited (Kazakh NLP tasks)	Limited
Education Levels	High School, University	Not explicitly structured	General education
Subjects Covered	STEM, Humanities, Social Sciences, Law, Medicine	QA, MT, causal reasoning	Broad general knowledge
Task Type	Bilingual MCQs reflecting real-world knowledge	QA, classification, generative tasks	General MCQs across languages
Model Benchmarking	41 LLMs (GPT-4o, Llama3.1, DeepSeek V3, etc.)	7 models on Kazakh NLP tasks	Multiple LLMs across 44 languages

Table 1: Comparison of **KazMMLU** with SIGTURK and INCLUDE.

Kazakh History, High School	
1930 жылдары өзге республикалардың жоғарғы оқу орындары мен техникумдарында білім алған қазақстанның жастар саны:	The number of young people from Kazakhstan who studied at universities and technical schools in other republics in the 1930s:
A. 25 мыңдай	A. About 25 thousand
B. 20 мыңдай	B. About 20 thousand
C. 15 мыңдай	C. About 15 thousand
D. 30 мыңдай	D. About 30 thousand
E. 40 мыңдай	E. About 40 thousand
Geography, High School	
Трактор, ауыл шаруашылық машиналар шығаратын қалалары:	Cities producing tractors and agricultural machinery:
A. Теміртау, Павлодар	A. Temirtau, Pavlodar
B. Астана, Теміртау	B. Astana, Temirtau
C. Арқалық, Павлодар	C. Arkalyk, Pavlodar
D. Балқаш, Ақтау	D. Balkash, Aktau
E. Астана, Павлодар	E. Astana, Pavlodar

Figure 3: Examples of a Kazakh History question and a Geography question from KazMMLU. The **left** side shows the original text, and the **right** side provides the English translation for reference. Correct answers are marked in bold.

detailed breakdown of dataset sources is provided in Appendix 6.

For automatic online crawling, we collect question texts, multiple-choice options, correct answer keys, and metadata. For books, authors manually scan materials and apply document processing for machine-readable conversion. Two expert workers fluent in Kazakh and Russian manually extract questions from scanned books and online sources, recording metadata such as source, country, subject, level, and answer key. In total, we compiled 23,000 questions—85% obtained through automated crawling and the remainder through manual extraction.

Only questions with valid answer keys were included, while multimodal ones requiring images or videos were excluded. For context-dependent questions, annotators ensured necessary context was

included. A training workshop clarified guidelines, and weekly check-ins ensured consistency. Annotators were competitively compensated to maintain quality.

3.2 Quality Control

Our quality control process primarily targets the automatically-crawled data, as the other two data collection strategies involve direct human involvement. To ensure accuracy, we recruited two professional annotators, each holding at least a bachelor’s degree and fluent in both Kazakh and Russian. They manually reviewed all questions to verify correctness and completeness. Any question containing errors or missing components (e.g., incomplete contexts or broken answer options) was discarded. Through this extensive human verification, every question included in KazMMLU undergoes manual validation, ensuring a high-quality dataset.

Additionally, we developed scripts to detect duplicates, verify metadata, and eliminate errors like duplicate questions, incorrect answer keys, and formatting issues, enhancing dataset reliability

3.3 Data Statistics

KazMMLU comprises 23,000 multiple-choice questions spanning two educational levels: high school and university. As shown in Figure 1, 48% of the dataset consists of high school-level questions in Kazakh, 36% in Russian, and 16% university-level questions in Russian. Table 2 outlines the subject distribution, covering STEM, Humanities, Social Sciences, and Languages. Notably, the Humanities, Social Sciences, and Language sections contain extensive Kazakhstan-specific knowledge. The dataset maintains a balanced distribution be-

Group	Subjects
Humanities	Culture and Art (U), Kazakh History (H), Kazakh Literature (H), Philosophy and Psychology (U), Russian Literature (H), World History (H)
Language	Kazakh Language (H), Reading Literacy (H), Russian Language (H)
Other	General Education Disciplines (U)
STEM	Biology (H), Chemistry (H), Informatics (H), Math (H), Math Literacy (H), Medicine (U), Physics (H)
Social Science	Accounting and Auditing (U), Economics and Entrepreneurship (U), Education and Training (U), Finance, Credit, Insurance (U), General Education Disciplines (U), Geography (H), Jurisprudence (U), State and Law (U), Management and Marketing (U), Social Science (U)

Table 2: Subject groups covered by KazMMLU. “H” indicates high school subjects, and “U” indicates university subjects.

Group	Question	Answer
High School	78.3	16.6
University	84.4	29.6
Humanities	81.3	19.1
Language	49.3	20.3
Others	82.0	37.1
STEM	83.8	15.1
Social Science	52.9	15.4

Table 3: Average question and answer length (in characters) for each educational group and subject area.

tween Kazakh and Russian, reflecting the bilingual nature of Kazakhstan’s education system.

Table 3 presents the average question and answer lengths across educational levels and subject areas. While the overall question length remains relatively consistent between high school and university levels, answer lengths (in characters) tend to be longer at the university level. Additionally, questions in Humanities and STEM subjects are generally longer compared to those in Social Sciences and Languages.

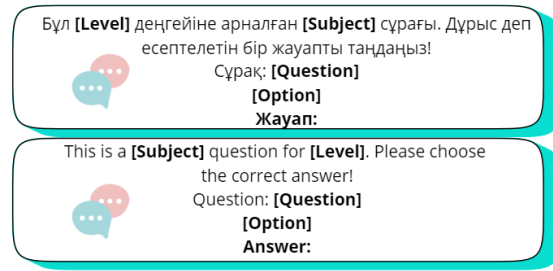


Figure 4: Prompt templates in Kazakh and English.

4 Experiment

4.1 Setup

We evaluate 27 multilingual LLMs of varying sizes in both zero-shot and few-shot settings. Our selection spans a range of architectures, including BLOOM (Scao et al., 2023), Aya (Üstün et al., 2024), Mistral (Jiang et al., 2023), Gemma (Team et al., 2024), Qwen (Yang et al., 2025), GPT-4o (OpenAI, 2024), and DeepSeek V3 (DeepSeek-AI, 2024). To ensure regional relevance, we also include models tailored to specific languages: YandexGPT³ and Vikhr (Nikolich et al., 2024) for Russian, as well as KazakhLLM-8B⁴ and SherKala-Chat-8B (Koto et al., 2025) for Kazakh. These regionally fine-tuned models are widely recognized for their focus on local languages.

For the evaluation purpose, we use two distinct prompt configurations to examine the effect of prompt language: (1) a Kazakh prompt with English (Latin script) alphabetic output and (2) an English prompt with English alphabetic output, as illustrated in Figure 4. For placeholders such as [Subject] and [Level], we use Kazakh in the Kazakh prompt and translate them into English for the English prompt. However, the question and answer choices remain in their original language (Kazakh or Russian). A complete example prompt is provided in Appendix B (Figure 10).

Following prior studies (Koto et al., 2023; Li et al., 2024), we adopt different answer selection methods based on model accessibility. For open-weight models, we apply the *next-token prediction* approach, computing probabilities for each multiple-choice option (A, B, C, D, or E) and selecting the one with the highest probability. This method is well-suited for autoregressive models that perform token-wise scoring. For closed-weight models (e.g., GPT-4o, Yandex-GPT, and

³yandex.cloud/en/services/yandexgpt

⁴huggingface.co/issai/LLama-3.1-KazLLM-1.0-8B

DeepSeek V3), we use a *free-text generation* strategy, prompting the model to generate a textual response, from which the predicted answer is extracted via string matching. This approach is necessary due to the lack of direct token probability outputs in closed-source APIs. Regarding inference settings, we used the default configurations provided by their respective APIs or platforms. For evaluation, we use accuracy as the primary metric, following prior studies (Koto et al., 2023; Li et al., 2024).

4.2 Results and Analysis

First, we observe that LLMs achieve higher accuracy when prompted in English, as shown in Table 4 and Table 7. To provide clearer insights into model performance, we focus on English-prompted results in the main body of the paper.

Results across all models Table 4 presents the average accuracy for each subject area across 27 models on the KazMMLU using English prompts. The performance analysis reveals several notable patterns. GPT-4o and DeepSeek V3 emerge as the top performers, achieving remarkably similar average scores of 76.6% and 76.9% respectively, outperforming other models. Gemma-2-27B instruction-tuned model achieved the highest open-source score (57.4%), followed by Llama3.1-70B (56.2%) and its base variant (55.7%). Interestingly, the impact of instruction tuning varies across model families - while Gemma-2-27B-IT showed a slight improvement over its base model (+1.7%), Llama3.1-70B-instruct performed worse than its base variant (-7.9%). However, for smaller models, instruction tuning appears more beneficial, as seen in Llama3.1-8B-instruct outperforming its base model by 4.9% (44.6% vs 39.7%).

Consistently across all models, the Language category seems to be the most challenging, with scores lower than other categories.

Few-Shot Performance As shown in Figure 5, our few-shot results show a consistent improvement across all models as the number of shots increases, with Qwen-2.5-7B and Mistral-7B-v0.3 benefiting the most. English prompts consistently outperform Kazakh prompts in 1, 2, and 3-shot settings, though this trend does not hold in 0-shot, where Kazakh sometimes performs better. Instruction-tuned models also improve, though the gains are smaller, with Qwen-2.5-7B-Instruct (English prompt) increasing from 47.8% (0-shot) to




Figure 5: The few-shot accuracy (%) of LLMs on KazMMLU, averaged across all tasks, comparing base models and instruction-tuned models using Kazakh (dotted lines) and English (solid lines) prompts.

58.9% (3-shot). The largest accuracy jumps occur between 0-shot and 1-shot, indicating that even a single in-context example significantly enhances model understanding. Overall, these results highlight the robustness of few-shot learning across diverse model architectures and prompt settings.

Kazakh vs Russian Performance In Figure 6, we compared model performance across two languages: Kazakh and Russian. To ensure a fair comparison, we only considered the High School level subjects and excluded the Professional & University level tasks because they are not available in Kazakh. The results indicate that GPT-4o achieves the highest accuracy in Kazakh, scoring 76.90%, while DeepSeek performs best in Russian with an accuracy of 81.4%. Llama3.1-70B and Gemma-2-27B-IT show lower but comparable results, with a slight advantage in Russian over Kazakh. Overall, models tend to perform slightly better in Russian than in Kazakh, which could be due to differences in training data availability, language complexity, or tokenization differences.

Results Across Education Level Results in Figure 7 indicate that GPT-4o performs approximately the same across both education levels. Similarly, DeepSeek V3 maintains a balanced performance, but with a slight preference towards High School. In contrast, the open-source models Llama3.1-70B and Gemma-2-27B-IT show considerably lower accuracy and a pronounced gap between educa-

Model	STEM	Social Science	Humanities	Language	Other	Average
Mistral-7B-Instruct-v0.3	37.6	47.4	43.3	30.0	40.6	41.0
Mistral-7B-v0.3	32.6	38.1	36.5	27.4	34.6	34.6
Vikhr-Nemo-12B-Instruct-R-21-09-24	40.7	50.5	48.5	32.7	44.0	44.5
aya-23-35B	35.7	39.0	35.5	29.4	33.2	35.9
aya-23-8B	30.8	32.6	32.1	25.8	26.5	31.0
Bloom-1.1B	23.6	20.7	22.1	21.6	20.8	22.1
Bloomz-1.7B	24.0	22.4	23.5	23.0	25.8	23.4
Bloomz-3B	24.4	23.9	22.8	22.1	25.2	23.7
Bloomz-7B	23.8	24.1	23.4	23.9	22.5	23.8
Gemma-2-27B	55.0	60.3	60.6	37.4	47.3	55.7
Gemma-2-27B-IT	57.3	60.5	63.2	39.1	48.3	57.4
Gemma-2-9B	51.4	57.8	55.5	36.2	42.6	52.3
Gemma-2-9B-IT	50.8	53.0	48.3	35.8	44.3	49.1
KazakhLLM-8B	40.3	45.6	44.0	31.3	38.6	41.7
Llama3.1-70B	58.0	59.1	57.4	41.8	49.3	56.2
Llama3.1-70B-instruct	51.2	48.1	51.3	33.9	45.6	48.3
Llama3.1-8B	37.9	44.6	41.7	28.9	38.3	39.7
Llama3.1-8B-instruct	43.0	47.8	50.6	30.4	39.3	44.6
mt0-large	24.6	23.7	23.9	22.9	23.5	24.0
mt0-xl	28.2	37.0	37.7	26.5	31.9	32.8
mt0-xxl	31.3	40.5	33.0	29.5	37.9	34.4
qwen-2.5-7B	45.2	43.1	44.1	29.4	39.9	42.5
qwen-2.5-7B-instruct	47.5	51.0	51.4	33.9	42.3	47.8
Sherkala-Chat-8B	43.1	49.8	50.0	34.9	38.3	45.6
GPT-4o	70.0	81.9	83.3	73.4	62.1	76.6
DeepSeek V3	77.6	81.3	78.9	61.2	65.1	76.9
YandexGPT	54.8	70.6	63.7	42.6	57.0	60.2

Table 4: Performance of different models on the Kazakh MMLU benchmark across different subject categories using **English** prompt. “Average” means the average across all subject areas in KazakhMMLU.




Figure 6: LLM Performance across different languages (only at the high school level).




Figure 7: LLM Performance across different education levels.

tion levels. Llama3.1-70B achieves 57.7% in High School but drops to 53.2% in Professional & University. These results suggest that while proprietary models generalize well across different subject complexities, open-source models struggle more with specialized university-level knowledge.

Negation Sensitivity Analysis Table 5 presents the accuracy of Llama3.1-70B, Gemma-2-27B-IT, and DeepSeek V3 on *negation-sensitive subjects*, comparing performance with and without negation. The results indicate that DeepSeek V3 consistently outperforms both Llama3.1-70B and Gemma-2-27B-IT, demonstrating greater resilience to negation-based reasoning challenges.

To systematically analyze the impact of negation, we employed a negation phrase filtering method inspired by ArabicMMLU (Koto et al., 2024). Specifically, we identified questions containing common negation phrases in Kazakh: жоқ (no), емес (is not), болмайды (not allowed), жарамайды (prohibited), невозможно (impossible), не (not), нельзя (forbidden). After applying this filtering, we obtained a total of 2,554 negation-related questions. To validate our filtering accuracy, we randomly sampled 100 questions and manually inspected them. The detection accuracy exceeded 92%, confirming the reliability of our filtering process.

For Llama3.1-70B and Gemma-2-27B-IT, accu-

Subject	Model	W/o Negation	W/ Negation
<i>Jurisprudence (University)</i>	Llama3.1-70B	56.2	55.2
	Gemma-2-27B-IT	55.2	56.5
	DeepSeek V3	78.1	76.4
<i>Law (High School)</i>	Llama3.1-70B	60.8	59.0
	Gemma-2-27B-IT	59.0	58.1
	DeepSeek V3	79.5	78.1
<i>Reading Literacy (High School)</i>	Llama3.1-70B	57.1	50.0
	Gemma-2-27B-IT	100.0	87.5
	DeepSeek V3	85.7	83.5
<i>Philosophy and Psychology (University)</i>	Llama3.1-70B	55.9	56.6
	Gemma-2-27B-IT	56.6	55.9
	DeepSeek V3	83.2	81.9

Table 5: Model accuracy on negation-sensitive questions across various subjects. Bold values indicate higher accuracy in each category.




Figure 8: Confidence vs. Accuracy for different models in a zero-shot setting. **Confidence (%)** denotes the average probability scores in percentage.

accuracy generally decreases on questions containing negation, suggesting a negative impact on reasoning capabilities. Notably, Llama3.1-70B exhibits a larger drop in accuracy, particularly in *Reading Literacy*, where performance declines from 57.1% to 50.0%. Meanwhile, Llama3.1-70B demonstrates greater robustness in certain cases, showing less fluctuation in performance across negation and non-negation settings.

Model Confidence We analyze whether the evaluated models, including Gemma-2.9B-IT, Qwen-2.5-7B-Instruct, and Llama3.1-70B, are well-calibrated when answering KazMMLU questions by comparing the probability of the correct answers with the actual accuracy for each subject and level combination. The answer probability is obtained through softmax normalization across the available candidate answers. In Figure 8, we observe a strong correlation between model confidence and

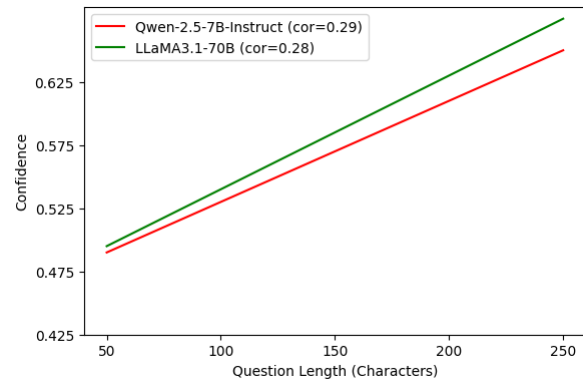


Figure 9: Correlation between model confidence and question length across different models.

accuracy, with correlation scores exceeding 0.9 across the evaluated models. This indicates that models with higher confidence generally produce more accurate predictions.

Additionally, we investigate the correlation between model confidence and question length in Figure 12. The results show that question length has minimal impact on model confidence, as evidenced by the weak correlation scores across all evaluated models. Qwen-2.5-7B-Instruct ($r = 0.29$) and Llama3.1-70B ($r = 0.28$) show a mild positive correlation, suggesting that **longer questions slightly increase model confidence**. However, the effect is weak overall, indicating that confidence calibration remains relatively stable across different question lengths.

5 Conclusion and Future Work

We introduced KazMMLU, the first large-scale multi-task language understanding dataset designed to evaluate real-world knowledge in Kazakhstan’s

bilingual setting. Through experiments with over 23K multiple-choice questions spanning various subjects and education levels, we observed that models perform much better in Russian than in Kazakh, with proprietary models such as GPT-4o and DeepSeek V3 achieving the highest accuracy.

KazMMLU provides a **bilingual** (Kazakh and Russian) evaluation framework tailored to Kazakhstan’s educational and professional landscape, distinguishing itself from previous multilingual benchmarks. Unlike SIGTURK, which focused on Kazakh NLP tasks, and INCLUDE, which lacks Kazakhstan-specific content, KazMMLU enables a localized and comprehensive assessment of LLMs in Kazakhstan.

Future research directions include extending KazMMLU to multimodal evaluations, improving reasoning-based question assessments, and mitigating biases in data sources. We hope that this benchmark will encourage further development of high-performance LLMs for Kazakh and other low-resource languages.

Limitations

While we are confident that our benchmark will significantly advance the development of Kazakh LLMs, it is important to acknowledge certain limitations that need to be addressed in future research. We outline these limitations as follows:

Limited Modality KazMMLU is focused solely on text-based assessment, and the exploration of multimodal questions (including those involving images, audio, or other media types) has been excluded. Future work could explore the integration of multimodal content to better reflect real-world applications, such as vision-language tasks, speech recognition, and interactive assessments.

Lack of Explicit Reasoning Evaluation While KazMMLU provides a broad and representative set of multiple-choice questions, it does not explicitly evaluate reasoning processes beyond answer selection. Investigating how models approach complex reasoning, justification, and open-ended question answering would be a valuable direction for further improvement.

Static Evaluation Limitation KazMMLU primarily evaluates static model performance on pre-defined questions, which may not fully capture how models generalize to dynamic, real-world language use. Exploring benchmarks that assess interactive

and adaptive reasoning, as well as domain adaptation, could enhance our understanding of model capabilities in evolving contexts.

By addressing these limitations, future research can further refine the evaluation of Kazakh LLMs, ensuring more robust, fair, and practically useful language models for Kazakhstan and beyond.

Ethics and Broader Impact

We adhered to the internal policies of web resources while scraping data and included only publicly available information verified by authorities.

All human subjects in our study provided informed consent, were fully aware of the study’s objectives, and had the right to withdraw at any time. They were also appropriately compensated as part of their job.

References

- Anthropic. 2023. *Claude*. Accessed: 2025-02-10.
- DeepSeek-AI. 2024. *Deepseek llm: Advancing open large language models with high-quality pretraining and instruction tuning*. Available at: <https://huggingface.co/deepseek-ai>.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *International Conference on Learning Representations*.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. *XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation*. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 4411–4421. PMLR.
- Jafar Isbarov, Arofat Akhundjanova, Mammad Hajili, Kavsar Huseynova, Dmitry Gaynullin, Anar Rzayev, Osman Tursun, Ilshat Saetov, Rinat Kharisov, Saule Belginova, et al. 2025. *Tumlu: A unified and native language understanding benchmark for turkic languages*. *arXiv preprint arXiv:2502.11020*.
- ISSAI. 2024. *Llama-3.1-kazllm-1.0-8b*. Accessed: 2025-02-10.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L elio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth ee Lacroix, and William El Sayed. 2023. *Mistral 7b*. *Preprint*, arXiv:2310.06825.

- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Fajri Koto, Nurul Aisyah, Haonan Li, and Timothy Baldwin. 2023. [Large language models only pass primary school exams in Indonesia: A comprehensive test on IndoMMLU](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12359–12374, Singapore. Association for Computational Linguistics.
- Fajri Koto, Rituraj Joshi, Nurdaulet Mukhituly, Yuxia Wang, Zhuohan Xie, Rahul Pal, Daniil Orel, Parvez Mullah, Diana Turmakhan, Maiya Goloburda, Mohammed Kamran, Samujjwal Ghosh, Bokang Jia, Jonibek Mansurov, Mukhammed Togmanov, Debopriyo Banerjee, Nurkhan Laiyk, Akhmed Sakip, Xudong Han, Ekaterina Kochmar, Alham Fikri Aji, Aaryamonvikram Singh, Alok Anil Jadhav, Sathesh Katipomu, Samta Kamboj, Monojit Choudhury, Gurpreet Gosal, Gokul Ramakrishnan, Biswajit Mishra, Sarath Chandran, Avraham Sheinin, Natalia Vasileva, Neha Sengupta, Larry Murray, and Preslav Nakov. 2025. [Llama-3.1-sherkala-8b-chat: An open large language model for kazakh](#). *Preprint*, arXiv:2503.01493.
- Fajri Koto, Haonan Li, Sara Shatnawi, Jad Doughman, Abdelrahman Sadallah, Aisha Alraeesi, Khalid Al-mubarak, Zaid Alyafeai, Neha Sengupta, Shady Shehata, Nizar Habash, Preslav Nakov, and Timothy Baldwin. 2024. [ArabicMMLU: Assessing massive multitask language understanding in Arabic](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5622–5640, Bangkok, Thailand. Association for Computational Linguistics.
- Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. 2024. [CMMLU: Measuring massive multitask language understanding in Chinese](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 11260–11285, Bangkok, Thailand. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Akylbek Maxutov, Ayan Myrzakhmet, and Pavel Braslavski. 2024. [Do LLMs speak Kazakh? a pilot evaluation of seven models](#). In *Proceedings of the First Workshop on Natural Language Processing for Turkic Languages (SIGTURK 2024)*, pages 81–91, Bangkok, Thailand and Online. Association for Computational Linguistics.
- Meta AI. 2024. [Introducing Meta Llama 3: The Most Capable Openly Available LLM to Date](#). Accessed: 2024-11-29.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, et al. 2023. Crosslingual generalization through multitask finetuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111.
- Aleksandr Nikolich, Konstantin Korolev, Sergei Bratchikov, Igor Kiselev, and Artem Shelmanov. 2024. [Vikhr: Constructing a state-of-the-art bilingual open-source instruction-following large language model for russian](#). In *Proceedings of the Fourth Workshop on Multilingual Representation Learning (MRL 2024)*, pages 189–199.
- OpenAI. 2024. [GPT-4o](#). Accessed: 2024-11-29.
- Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korhonen. 2020. [XCOPA: A multilingual dataset for causal common-sense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2362–2376, Online. Association for Computational Linguistics.
- Angelika Romanou, Negar Foroutan, Anna Sotnikova, Zeming Chen, Sree Harsha Nelaturu, Shivalika Singh, Rishabh Maheshwary, Micol Altomare, Mohamed A. Haggag, Imanol Schlag, et al. 2024. [INCLUDE: Evaluating multilingual language understanding with regional knowledge](#). *Preprint*, arXiv:2411.19799.
- Teven Le Scao, Angela Fan, Christopher Akiki, Elie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, 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

Abdulummin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu, Jonathan Chang, Jörg Frohberg, Joseph To-bing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tan-guy, 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 Espe-jel, 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 Has-san Muhammad, Shanya Sharma, Shayne Long-pre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vas-silina Nikoulina, Veronika Laippala, Violette Lep-ercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Tal-lat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajy-oti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Sai-ful Bari, Maged S. Al-shaibani, Matteo Manica, Ni-hal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Ur-mish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin 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, San-chit Gandhi, Shaden Smith, Stéphane Requena, Suraj Patil, Tim Dettmers, Ahmed Baruwa, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névéol, Charles Lover-ing, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bog-danov, 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, Ruo Chen 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, Al-ice Rueda, Amanda Pestana, Amir Feizpour, Am-mar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Unldreaj, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini, Bahareh Behroozi, Benjamin Ajibade, Bharat Sax-ena, Carlos Muñoz Ferrandis, Daniel McDuff, Dan-ish 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 Bhat-tacharya, Irene Solaiman, Irina Sedenko, Isar Ne-jadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim El-badri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Ra-jani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Al-izadeh, 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, An-ima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourrier, Daniel León Perifán, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Flor-ian 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 Sivara-man, 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 Mihalj-cic, 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, Si-mon Ott, Sinee Sang-aaroonsiri, Srishti Kumar, Ste-fan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kosa, 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. *arXiv preprint arXiv:2211.05100*.

Lütfi Kerem Senel, Benedikt Ebing, Konul Baghirova, Hinrich Schuetze, and Goran Glavaš. 2024. **Kardeş-NLU: Transfer to low-resource languages with the help of a high-resource cousin – a benchmark and evaluation for Turkic languages.** In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1672–1688, St. Julian’s, Malta. Association for Computational Linguistics.

Shivalika Singh, Angelika Romanou, Clémentine Four-rier, David I Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio,

- Wei Qi Leong, Yosephine Susanto, et al. 2024. Global mmlu: Understanding and addressing cultural and linguistic biases in multilingual evaluation. *arXiv preprint arXiv:2412.03304*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshov, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonnell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iverson, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Cogan, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya, Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024. **Gemma 2: Improving open language models at a practical size**. *Preprint*, arXiv:2408.00118.
- Ahmet Üstün, Viraat Aryabumi, Zheng Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. **Aya model: An instruction finetuned open-access multilingual language model**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15894–15939, Bangkok, Thailand. Association for Computational Linguistics.
- Yandex. 2023. **Yandexgpt**. Accessed: 2025-02-10.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, Keming Lu, Keqin Bao, Kexin Yang, Le Yu, Mei Li, Mingfeng Xue, Pei Zhang, Qin Zhu, Rui Men, Runji Lin, Tianhao Li, Tianyi Tang, Tingyu Xia, Xingzhang Ren, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yu Wan, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zihan Qiu. 2025. **Qwen2.5 technical report**. *Preprint*, arXiv:2412.15115.
- Rustem Yeshpanov, Pavel Efimov, Leonid Boytsov, Ardak Shalkarbayuli, and Pavel Braslavski. 2024. **KazQAD: Kazakh open-domain question answering dataset**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 9645–9656.
- Rustem Yeshpanov, Yerbolat Khassanov, and Huseyin Atakan Varol. 2022. **KazNERD: Kazakh named entity recognition dataset**. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 417–426, Marseille, France. European Language Resources Association.
- Rustem Yeshpanov and Huseyin Atakan Varol. 2024. **KazSAnDRA: Kazakh sentiment analysis dataset of reviews and attitudes**. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 9657–9667, Torino, Italia. ELRA and ICCL.

A Appendix

A.1 Additional Prompt Examples

Figure 10 illustrates the Kazakh and English prompts used in our evaluation. The formatting emphasizes key aspects such as the subject (e.g., **Law**) and the educational level (e.g., **High School**). The structure maintains consistency in question presentation, ensuring uniformity across both languages.

Бұл **Орта Мектеп** деңгейіне арналған **Заң** сұрағы. Дұрыс деп есептелетін бір жауапты таңдаңыз!

Сұрақ: Қазақстан Республикасы Конституциясының V бөлімінің атауы

- A. Соттар мен сот төрелігі
- B. Парламент
- C. Президент
- D. Үкімет

Жауап:

This is a **Law** question for **High School** Level. Please choose the correct answer!

Question: What is the title of Part V of the Constitution of the Republic of Kazakhstan?

- A. Courts and justice
- B. Parliament
- C. President
- D. Government

Answer:

Figure 10: Examples of Kazakh and English prompts. The placeholders are dynamically replaced based on the question context.

B Appendix B:

B.1 Additional Multiple-Choice Question Example

In addition to the Kazakh-language question provided in the main document, we present an example of a Russian-language multiple-choice question (Figure 11). This figure highlights a range of subjects from social sciences to STEM disciplines, demonstrating the dataset’s diversity.

B.2 Breakdown of Dataset Sources

Table 6 provides a categorized breakdown of dataset sources used in KazMLLU, covering national exams, professional certification tests, and textbooks.

B.3 Evaluation Results using Kazakh Prompt

Table 7 presents the experimental results using Kazakh prompts across various models.

Culture and Art, University	
Когда в Казахстане появились театры, концертные залы и филармонии: A. Конец 20 века B. 19 век C. 18 век D. Первая половина 20 века E. 21 век	When the theater, concert hall and philharmonic appeared in Kazakhstan: A. The end of the 20th century B. 19th century C. 18th century D. First half of the 20th century E. 21th century
Social Science, University	
По Конституции РК не должны состоять в партиях A. Члены Парламента B. Работники Аппарата Президента C. Военнослужащие, работники органов национальной безопасности, правоохранительные органы, судьи D. Студенты E. Государственные чиновники	According to the Constitution of the Republic of Kazakhstan, there should be parties A. Member of Parliament B. Employees of the Presidential Executive Office C. Military personnel, employees of national security agencies, law enforcement agencies, judges D. Students E. Government officials

Figure 11: Example of a Russian-language multiple-choice question from KazMLLU. The **left** side shows the original text, while the **right** side provides the English translation for illustrative purposes. The bold options represent the correct answer keys.

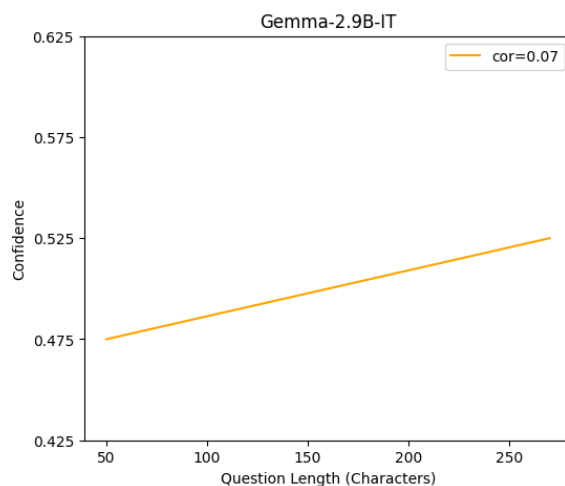


Figure 12: Correlation between model confidence and question length.

Main Source	Language	Level	Subjects
itest.kz	Kazakh	High School	Biology, Chemistry, Geography, Informatics, Kazakh History, Kazakh Literature, Law, Math, Math Literacy, Physics, Reading Literacy, Russian Language, Russian Literature, World History
oltest.kz	Russian	University and Professional	Accounting and Auditing, Biology, Culture and Art, Economics and Entrepreneurship, Education and Training, Finance, General Education Disciplines, Jurisprudence, Management and Marketing, Medicine, Philosophy and Psychology, Social Science
ymnik.kz	Russian	High School	Biology, Geography, Kazakh History, Kazakh Language, World History
Book - Shyn Kitap	Kazakh	High School	Biology, Geography, Kazakh History, Kazakh Language, World History

Table 6: Breakdown of dataset sources in KazMMLU, categorized by domain and subject area. The alternating row colors improve readability.

Model	STEM	Social Science	Humanities	Language	Other	Average
Mistral-7B-Instruct-v0.3	34.9	43.3	38.3	27.9	31.9	37.3
Mistral-7B-v0.3	29.5	29.3	27.7	24.6	25.2	28.4
Vikhr-Nemo-12B-Instruct-R-21-09-24	40.3	52.7	52.7	32.8	41.9	45.9
aya-23-35B	30.5	34.1	33.0	25.8	26.5	31.5
aya-23-8B	27.2	26.9	31.2	23.2	22.8	27.4
Bloom-1.1B	23.4	19.4	20.3	22.9	19.1	21.4
Bloomz-1.7B	22.8	20.6	23.1	23.3	20.1	22.2
Bloomz-3B	23.2	20.6	23.6	21.9	17.8	22.2
Bloomz-7B	22.4	20.0	21.3	21.1	21.5	21.3
Gemma-2-27B	57.4	63.6	64.2	41.1	46.6	58.7
Gemma-2-27B-it	59.3	62.6	63.3	40.9	46.3	58.8
Gemma-2-9B	47.9	54.0	54.0	33.7	39.3	49.3
Gemma-2-9B-it	50.4	51.6	48.8	35.2	41.9	48.5
KazakhLLM-8B	36.9	39.4	47.0	28.4	33.9	38.9
Llama3.1-70B	58.8	64.3	64.5	44.4	50.3	59.9
Llama3.1-70B-instruct	56.7	56.2	59.8	39.5	50.3	55.2
Llama3.1-8B	34.3	39.5	38.6	28.0	31.5	36.0
Llama3.1-8B-instruct	37.7	40.2	42.6	28.0	31.5	38.3
mt0-large	24.4	24.6	22.0	22.9	24.8	23.8
mt0-xl	29.1	38.9	36.9	26.8	34.6	33.6
mt0-xxl	29.9	38.4	35.5	29.4	36.2	33.8
Sherkala-Chat-8B	36.6	39.5	43.4	30.0	34.6	38.2
qwen-2.5-7B	46.7	50.2	44.9	30.5	42.3	45.5
qwen-2.5-7B-instruct	48.0	52.3	56.6	33.2	43.3	49.4
GPT4-o	70.5	82.1	83.2	72.7	62.1	76.7
DeepSeek V3	45.4	48.0	47.8	37.4	33.3	45.5

Table 7: Performance of different models on the Kazakh MMLU benchmark across different subject categories using **Kazakh** prompt. "Average" means the average across all subject areas in KazakhMMLU.