

GreekMMLU: A Native-Sourced Multitask Benchmark for Evaluating Language Models in Greek

Yang Zhang^{1*}, Mersin Konomi^{1*}, Christos Xypolopoulos^{1,3},
Konstantinos Divriotis², Konstantinos Skianis⁴, Giannis Nikolentzos⁵,
Giorgos Stamou³, Guokan Shang^{2†}, Michalis Vazirgiannis^{1,2†}

¹LIX, Ecole Polytechnique, ²MBZUAI, ³National Technical University of Athens,
⁴University of Ioannina, ⁵University of Peloponnese

†Correspondence: guokan.shang@mbzuai.ac.ae, mvazirg@lix.polytechnique.fr

Abstract

Large Language Models (LLMs) are commonly trained on multilingual corpora that include Greek, yet reliable evaluation benchmarks for Greek—particularly those based on authentic, native-sourced content—remain limited. Existing datasets are often machine-translated from English, failing to capture Greek linguistic and cultural characteristics. We introduce GreekMMLU¹, a native-sourced benchmark for massive multitask language understanding in Greek, comprising 21,805 multiple-choice questions across 45 subject areas, organized under a newly defined subject taxonomy and annotated with educational difficulty levels spanning primary to professional examinations. All questions are sourced or authored in Greek from academic, professional, and governmental exams. We publicly release 16,857 samples and reserve 4,948 samples for a private leaderboard² to enable robust and contamination-resistant evaluation. Evaluations of over 80 open- and closed-source LLMs reveal substantial performance gaps between frontier and open-weight models, as well as between Greek-adapted models and general multilingual ones. Finally, we provide a systematic analysis of factors influencing performance—including model scale, adaptation, and prompting—and derive insights for improving LLM capabilities in Greek.

1 Introduction

Large language models have achieved strong performance across a wide range of natural language understanding and reasoning tasks, largely driven by large-scale training on multilingual corpora (Brown et al., 2020; Grattafiori et al., 2024; Üstün et al., 2024). Contemporary models are designed to support dozens or even hundreds of languages, a choice driven by the need to maximize data scale

*These authors contributed equally.

¹<https://github.com/mersinkonomi/GreekMMLU>

²<https://hf.co/spaces/yangzhang33/GreekMMLU-Leaderboard>

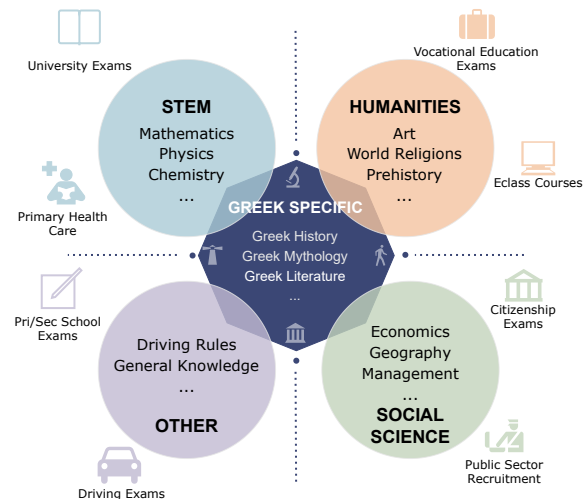


Figure 1: GreekMMLU task overview.

and leverage cross-lingual transfer to improve general reasoning capabilities.

Furthermore, recent research has increasingly emphasized extending LLM capabilities to lower-resource languages (Voukoutis et al., 2024; Rousis et al., 2025; Martins et al., 2025; Shang et al., 2025a,b). In practice, while Greek is often present in the long tail of web-scale training corpora (Grattafiori et al., 2024), it is rarely prioritized as a core language, resulting in significantly lower representation compared to major European languages. Despite this incidental exposure, relatively little work systematically reports LLM performance on Greek, largely due to the absence of large-scale, native-language evaluation benchmarks. Existing evaluations are commonly based on machine-translated datasets originally designed for English (Xuan et al., 2025; Voukoutis et al., 2024), which fail to capture the linguistic and cultural characteristics of authentic Greek language use.

A widely adopted benchmark for assessing broad knowledge and reasoning is the Massive Multi-

task Language Understanding (MMLU) benchmark (Hendrycks et al., 2021), which evaluates models across diverse subjects including STEM, social science and humanities. While MMLU has become a cornerstone of LLM evaluation, it is inherently grounded in English and the educational and cultural context of the United States. Extending MMLU to other languages through translation introduces well-known limitations, including translationese, semantic drift, altered difficulty calibration, and the preservation of source-language cultural priors (Singh et al., 2025; Artetxe et al., 2020; Li et al., 2024). As a result, translated benchmarks conflate native language understanding with cross-lingual transfer and often provide a misleading assessment of model capabilities.

These limitations are particularly salient for Greek. Modern Greek exhibits rich morphology, flexible word order, and complex negation, all of which are difficult to evaluate through translated data. Furthermore, many knowledge-intensive domains—such as history, law, civics, and professional certification—are tightly coupled to national curricula, legal frameworks, and cultural context. Evaluating such domains using translated English benchmarks obscures important gaps in localized knowledge and overestimates real-world utility.

To address these issues, we introduce GreekMMLU, the first large-scale, native-sourced benchmark for evaluating massive multitask language understanding in Greek. GreekMMLU consists of 21,805 Multiple-Choice Questions (MCQ) across 45 manually defined subject areas spanning STEM, Humanities, Social Sciences, and Other domains, drawn from authentic academic, professional, and governmental examinations. All materials were collected exclusively from sources explicitly released under open-access licenses or educational reuse terms, ensuring ethical data use and legal compliance. The benchmark covers a wide range of educational and difficulty levels, including primary school, secondary school, university, and professional level, and includes eight Greek Specific subject areas requiring localized cultural knowledge. All questions are originally sourced or authored in Greek, preserving linguistic nuance and culturally grounded content.

Our contributions are summarized as follows:

- We introduce **GreekMMLU**, the first large-scale, fully native-sourced benchmark. It comprises 21,805 multiple-choice questions, organized under a **carefully defined subject taxonomy** with 45

subjects and **systematically annotated with educational difficulty levels** ranging from primary education to professional examinations.

- We conduct a large-scale evaluation of **80+ open- and closed-source LLMs** on GreekMMLU, revealing clear and consistent trends: closed-source frontier models substantially outperform open-weight alternatives, model scale strongly correlates with performance and general-purpose multilingual models exhibit pronounced weaknesses on Greek-specific and culturally grounded subject areas.

- We provide an in-depth analysis of factors influencing performance on native Greek language understanding, examining the effects of model scale, instruction tuning, prompting strategies, subject domains, and educational levels. Our findings show that Greek-adapted training leads to significant and systematic gains, particularly on tasks requiring culturally grounded knowledge.

- We release and maintain an official leaderboard for GreekMMLU, with separate public and private subsets to support fair, contamination-resistant evaluation.

2 Related Work

Language Models in Greek Greek remains underrepresented in large-scale language modeling, as most multilingual LLMs include it only implicitly and without explicit tokenizer design or data balancing, resulting in weaker performance compared to high-resource languages (Chowdhery et al., 2023; Touvron et al., 2023). While models such as Qwen (Yang et al., 2025) and Gemma (Team et al., 2025) can perform Greek tasks, they do not explicitly report Greek training data and performance. Earlier multilingual models like BLOOMZ and mT0 (Muennighoff et al., 2023) explicitly included Greek, but only at a very limited scale (around 0.03%). More recent efforts have addressed this gap through targeted pretraining. EuroLLM (Martins et al., 2025) increased coverage of European languages, including Greek. Voukoutis et al. (2024) introduced *Meltemi*, the first openly released Greek-centric LLM, trained with large-scale Greek corpora and a Greek-aware tokenizer, achieving substantial gains on Greek benchmarks. Building on this work, Roussis et al. (2025) proposed *LLaMA-Krikri*, further expanding Greek coverage through increased Greek data. Beyond general-purpose models, domain-specific efforts such as *Plutus-8B* demonstrate the benefits

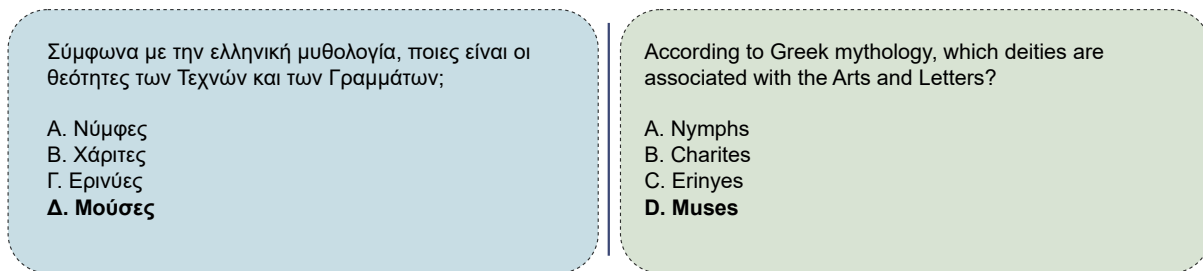


Figure 2: Example of a mythology question from GreekMMLU. **Left** shows the content structure derived from native sources and **right** is the English translation. The bold options represent the correct answer keys.

of Greek-adapted training in specialized settings (Peng et al., 2025). Overall, these works highlight the importance of explicit Greek-focused training for robust Greek language understanding.

LLM Evaluation and Multilingual Benchmarks

General knowledge and reasoning in LLMs are commonly evaluated using multitask benchmarks such as MMLU (Hendrycks et al., 2021), along with datasets like ARC (Clark et al., 2018) and HellaSwag (Zellers et al., 2019). While effective for tracking architectural and scaling progress, benchmarks like MMLU are predominantly English-centric, limiting their ability to assess culturally grounded language understanding. Multilingual extensions often rely on machine translation, including multilingual variants of MMLU (Xuan et al., 2025), but prior work has shown that translated benchmarks suffer from translationese, semantic drift, altered difficulty calibration, and inherited cultural priors, leading to validity concerns (Singh et al., 2025; Artetxe et al., 2020; Li et al., 2024). To address these issues, native-sourced benchmarks such as CMMLU for Chinese (Li et al., 2024), ArabicMMLU for Arabic (Koto et al., 2024), and TurkishMMLU for Turkish (Yüksel et al., 2024) have been proposed, highlighting the need for locally grounded multilingual evaluation.

Greek LLM Evaluation Evaluation of Greek language capabilities has traditionally relied on multilingual benchmarks such as XNLI (Conneau et al., 2018), XTREME (Hu et al., 2020), XQuAD (Artetxe et al., 2020), MASSIVE (FitzGerald et al., 2023), and FLORES-101 (Goyal et al., 2022), where Greek is included as a target language to assess cross-lingual generalization. More recent Greek-centric efforts, such as Meltemi (Voukoutis et al., 2024) and KriKri (Roussis et al., 2025), largely resorted to machine-translated versions of English benchmarks like HellaSwag and MMLU.

Several Greek-specific benchmarks enable a more targeted evaluation. GreekSUM (Evdaimon et al., 2024) introduced the first large-scale abstractive summarization dataset for Greek news. Earlier resources such as eNER (Bartziokas et al., 2020), support named entity recognition in Greek, and recent domain-specific benchmarks like Plutus-ben (Peng et al., 2025) and GreekBarBench (Chlapanis et al., 2025) extend evaluation to financial and legal reasoning tasks. However, these efforts are typically task-specific and limited in scale, reflecting a need for dedicated benchmarks that better capture the linguistic and cultural specific challenges in Greek.

3 The GreekMMLU Dataset

3.1 Overview

GreekMMLU is a **native-sourced benchmark** for evaluating massive multitask language understanding in Greek, composed exclusively of original Greek questions drawn from real-world educational and professional assessments. All questions follow an MCQ format as illustrated in Figure 2, with a variable number of answer options (2–4) and exactly one correct answer, reflecting the diversity of Greek national examinations. Each question is annotated with a difficulty level corresponding to its educational context.

A major effort in creating GreekMMLU involved the systematic structuring of heterogeneous raw exam material. We designed a **custom subject taxonomy** and **carefully assigned each task to an educational difficulty level**, enabling consistent analysis across domains and degrees of specialization. The resulting benchmark spans 45 subject areas, organized into four high-level categories—**STEM, Humanities, Social Sciences, and Other** as shown in Figure 1, covering educational levels from **primary school to university and professional examinations**.

Group	Subjects
Humanities	Art (S, U, Pr), Greek History (P, S, Pr), Greek Literature (S, U), Greek Mythology (P, S, U), Law (S, Pr), Prehistory (P), World History (P, S), World Religions (S, U)
STEM	Agriculture (U, Pr), Biology (P, S), Chemistry (P, U), Civil Engineering (Pr), Clinical Knowledge (Pr), Computer Networks & Security (U), Computer Science (U, Pr), Electrical Engineering (U, Pr), Mathematics (P, U), Medicine (U, Pr), Physics (P, U, Pr)
Social Sciences	Economics (U, Pr), Education (U, Pr), Geography (P, S), Government and Politics (P, S), Greek Traditions (S, Pr), Management (U, Pr), Modern Greek Language (P, S), Accounting (Pr)
Other	Driving Rules (NA), General Knowledge (S, Pr), Maritime Safety and Rescue Operations (Pr)

Table 1: Subject areas in GreekMMLU. “P”, “S”, “U”, “Pr”, and “NA” indicate availability in primary school, secondary school, university, professional, and not available categories, respectively.

Beyond general academic content, GreekMMLU includes a dedicated subset of **Greek-specific tasks** that require explicitly Greek linguistic and cultural knowledge, such as Greek History, Greek Literature, Greek Mythology, Greek Traditions, and the Modern Greek Language. These tasks are tightly coupled to the local cultural context and cannot be reliably assessed through translated benchmarks, highlighting the importance of native-sourced evaluation.

The dataset is divided into a **public, open-source subset** released for research use and a **private subset** reserved for a leaderboard to support more robust and contamination-resistant evaluation. Representative examples of the dataset are provided in Appendix A.

3.2 Data Collection and Curation

We conducted an exhaustive survey of publicly available Greek testing platforms to compile a corpus of questions and answers. Our sourcing strategy targeted authoritative bodies, identifying a wide spectrum of standardized examinations ranging from primary education to professional licensure. We systematically crawled and ingested data from these repositories, developing custom extraction pipelines to handle heterogeneous file formats—including structured web interfaces, PDF

archives, and DOCX documents. These formats typically reside outside the scope of standard web-crawling pipelines (e.g., Common Crawl) used for LLM pre-training, thereby minimizing the risk of data contamination. This process yielded a raw corpus of diverse subject matter, ensuring the benchmark captures the breadth of the Greek educational and professional curriculum.

Given the predominance of PDF documents in the raw corpus, we utilized *PyMuPDF4LLM*³ for structure-aware text extraction. For legacy scanned files, we applied Tesseract OCR⁴. To mitigate extraction artifacts—particularly in malformed Greek text and mathematical notations—we implemented an LLM-assisted correction phase utilizing Claude 3.5 Sonnet and Haiku. The models were prompted to preserve the original semantic intent verbatim and identify toxic content, ensuring that no generative question creation occurred during the process.

Following this automated restoration, the dataset underwent strict Unicode canonicalization and punctuation standardization (e.g., correcting the Greek question mark ‘;’). Finally, the curated dataset—after filtering and masking any personal identifying information and meaningless ids—was validated by a dedicated team of **five native Greek-speaking experts** holding graduate-level academic qualifications. This team manually reviewed the question–answer pairs to verify linguistic fidelity and filter out remaining processing artifacts, ensuring the benchmark adheres to the rigorous standards of authentic Greek examinations.

3.3 Quality Control

In constructing GreekMMLU, the dataset was derived predominantly from official sources, whose data quality was assumed to be reliable based on their authoritative provenance. Materials from unofficial or secondary sources constituted 26.2% of the corpus (approximately 4,400 samples) and were therefore subjected to comprehensive human review. Each of these samples was manually inspected by our experts, resulting in the identification and correction of errors in 3.8% of the subset. To further ensure minimal residual noise introduced during OCR and parsing stages, we conducted an additional round of expert human validation on a randomly sampled set of 8,000 examples drawn across all sources. This multi-stage verification pipeline underscores the substantial human effort

³<https://github.com/pymupdf/pymupdf4llm>

⁴<https://github.com/tesseract-ocr/tesseract>

invested in ensuring the overall accuracy and robustness of GreekMMLU.

For the final dataset, we performed a rigorous human evaluation. Three Greek-speaking graduate- and professor-level experts randomly selected 5% of the samples from each individual task and manually verified both the question stems and the corresponding ground-truth answers. This evaluation focused on identifying residual OCR errors, semantic drift, or incorrect labeling. Based on this process, we estimated the overall noise level in the dataset—the proportion of QA samples deemed low quality or incorrect—to be approximately 2%, with all identified issues corrected prior to release.

Group	# Questions	# Chars	
		Question	Answer
STEM	6787	94.5	26.5
Humanities	2751	93.8	42.5
Social Sciences	4753	306.4	21.4
Other	2341	86.4	51.9
Primary School	4393	67.2	19.9
Secondary School	1566	747.2	19.7
Professional	8103	106.5	33.0
University	912	90.7	42.2
NA	1658	88.7	57.6

Table 2: Average question and answer length (in characters) for each education group and subject area in GreekMMLU.

3.4 Statistics and Analysis

GreekMMLU consists of **21,805** multiple-choice questions across **45 subjects**, organized into **four supercategories** (STEM, Humanities, Social Sciences, and Other) and annotated with **four difficulty levels** corresponding to *Primary*, *Secondary*, *University*, and *Professional* education. We split about 23% from each subject to create two subsets. The **public subset contains 16,857 questions** and is used for all experiments reported in this work, while the **private subset contains 4,948 questions** and is reserved for a future evaluation leaderboard. A detailed statistical distribution can be found in Table 2 and Appendix B.

As shown in Table 2, most questions are drawn from professional examinations, followed by primary, secondary, and university-level questions, with an additional NA category of 1.7K questions not aligned to a specific educational tier. Average question and answer length generally increases with educational level; secondary school items form a notable exception, with longer prompts but

relatively short answers, reflecting the inclusion of citizenship and civic education tasks that combine context-rich descriptions with concise answer options. Across subject areas, the dataset is broadly balanced, with STEM, humanities, and social sciences each comprising 2.7K–6.8K questions; social science questions are longer on average due to text-based, situational prompts, while STEM questions tend to be more concise and formula-driven.

4 Experiments and Results

4.1 Experimental Setup

We integrate GreekMMLU into the lm-evaluation-harness framework (Gao et al., 2024) to ensure a standardized and reproducible evaluation environment. All models are evaluated under both zero-shot and five-shot settings. We show the detailed setup in Appendix C.

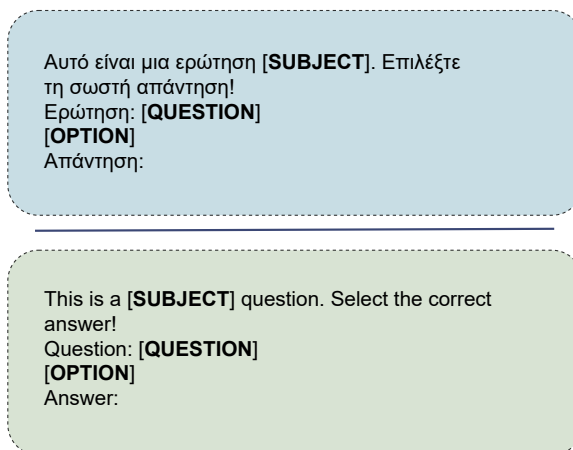


Figure 3: Prompt templates in Greek and English.

Evaluation Strategy. Evaluation depends on access to model outputs. For open-weights models, we use a rank-based multiple-choice setup, selecting the answer with the highest token log-likelihood. For closed-source API models, we use free-form generation, prompting models to output the answer key directly and extracting the predicted Greek label with regular expressions (Li et al., 2024).

Prompting Protocol. Models are evaluated under both zero-shot and five-shot prompting using prompts written entirely in Greek. Each prompt includes a subject-specific instruction, the question stem, and the labeled answer options, following Greek examination conventions with answer labels A, B, Γ, and Δ. In the five-shot setting, five repre-

sentative examples from the task’s development set are prepended to the prompt. The prompt templates are illustrated in Figure 3.

Model Selection. We conduct a comprehensive analysis of diverse model families on the GreekMMLU benchmark. Our study covers three categories. We first establish broad capabilities using **General-Purpose LLMs**, including the Qwen (Yang et al., 2025), Llama (Grattafiori et al., 2024), and Gemma-3 (Team et al., 2025), GLM-4 (GLM et al., 2024) families. We compare these against **Greek-Adapted LLMs**—such as Meltemi (Voukoutis et al., 2024), Krikri (Roussis et al., 2025), and Plutus (Peng et al., 2025)—which are selected to quantify the benefits of native-language specialization. Additionally, we evaluate **Multilingual European LLMs**, represented by the EuroLLM family (Martins et al., 2025) alongside multilingual models like Aya-101 (Üstün et al., 2024), BLOOMZ (Muennighoff et al., 2023), and mT0 (Muennighoff et al., 2023). Finally, we assessed frontier closed-source models like ChatGPT and Gemini, while a **Random Baseline** is included to establish a theoretical lower bound.

4.2 Main Results

The overall zero-shot performance on part of our evaluated models is summarized in Table 3. We show all the evaluated models with both zero-shot and five-shot settings in Appendix D. In this paper, we only report results on the public subset, we provide analysis for the private results in Appendix E.

Model performance on GreekMMLU spans a wide spectrum, ranging from near-random accuracy for weak models to strong results for state-of-the-art models. Small and lightly adapted models have performance near the random baseline, while mid-sized models (3-20B) typically achieve moderate performance in the 55–70% range. A clear separation emerges at the top end, where closed-source frontier models substantially outperform all open-weight alternatives. For example, *Gemini 3 Flash* reaches an average accuracy of 93.16%, while *GPT-5.2* and *GPT-4o* achieve 87.75% and 86.81%, respectively, consistently excelling in all subjects. In contrast, the strongest open-weight models—such as *Llama-3.3-70B-Instruct* and *Qwen2.5-72B-Instruct*—peak at 79.56% and 79.70% average accuracy, leaving a substantial gap to the best models. This persistent performance margin may reflect the

advantages of closed-source models. We hypothesize that these advantages could stem from factors such as broader exposure to Greek data during training, as well as large-scale optimization and advanced alignment.




Figure 4: Models’ 0-shot performance on different task levels.

Instruction Tuning Effects. Across nearly all model families, instruction-tuned variants consistently outperform their corresponding base models at similar parameter scales. This improvement is observed across subject categories and is particularly pronounced on Greek-specific tasks. Instruction tuning appears to improve alignment with the multiple-choice evaluation format, Greek prompt structure, and answer selection conventions used in GreekMMLU.

Greek- and European-Centric Models. Models explicitly adapted to Greek or European languages consistently achieve higher accuracy than general multilingual baselines of comparable size. Greek-centric models like LLaMA-Krikri-8B demonstrate notable gains, especially on Greek-specific subjects. Similarly, European-centric models outperform globally trained multilingual models, indicating that regional linguistic focus contributes meaningfully to Greek language understanding.

Performance Across Educational Levels. Figure 4 shows that accuracy is generally higher on primary and secondary school questions and lower on university and professional-level tasks. This pattern appears across model families and scales, including both open-weight and closed-source models. Performance on N/A-level questions typically lies between primary/secondary and univer-

Model	STEM	Humanities	Social Sci.	Other	Average	Greek-specific
General-Purpose LLMs						
GPT-5.2 [†] ‡	86.05	88.27	90.29	85.96	87.75	92.92
GPT-4o [†] ‡	84.54	88.68	89.36	85.42	86.81	93.11
Gemini 3 Flash [†] ‡	92.82	92.88	94.16	91.84	93.16	95.44
Qwen2.5-7B	51.92	56.37	59.34	53.31	55.16	61.07
Qwen2.5-7B-Instruct [†]	59.88	58.33	61.98	58.89	60.25	64.02
Qwen2.5-14B	64.11	64.95	66.04	60.23	64.39	67.95
Qwen2.5-14B-Instruct [†]	65.20	66.41	69.78	62.92	66.61	73.06
Qwen2.5-32B	72.74	71.29	75.26	70.68	73.14	77.40
Qwen2.5-32B-Instruct [†]	72.08	71.47	76.61	69.69	73.22	80.03
Qwen2.5-72B	78.31	78.78	81.35	76.75	79.20	83.91
Qwen2.5-72B-Instruct [†]	78.90	79.14	81.92	76.95	79.70	84.67
Qwen3-30B	70.72	69.06	74.95	59.63	70.56	77.49
Qwen3-30B-Instruct [†]	79.31	74.81	79.33	76.56	78.39	81.80
Llama-2-7b-hf	36.23	35.97	33.80	35.54	35.30	32.84
Llama-2-7b-chat-hf [†]	36.63	34.92	34.26	33.85	35.27	33.61
Llama-3.1-8B	51.08	57.42	54.78	50.62	53.10	54.78
Llama-3.1-8B-Instruct [†]	56.58	62.85	62.56	57.84	59.56	64.75
Llama-3.1-70B	72.42	79.05	78.90	74.46	75.71	82.46
Llama-3.2-1B	37.37	36.51	34.86	36.83	36.35	33.66
Llama-3.2-1B-Instruct [†]	38.01	37.33	36.87	35.94	37.29	35.46
Llama-3.2-3B	41.82	41.67	43.81	44.75	42.82	41.97
Llama-3.2-3B-Instruct [†]	43.77	43.04	45.62	45.64	44.52	46.15
Llama-3.3-70B-Instruct [†]	77.03	82.20	82.92	76.80	79.65	86.94
Gemma-3-4B-pt	50.77	52.67	53.83	52.17	52.21	55.00
Gemma-3-4B-it [†]	59.79	60.43	65.95	62.32	62.24	68.58
Gemma-3-12B-pt	73.30	74.44	78.28	72.27	74.99	80.71
Gemma-3-12B-it [†]	72.95	75.13	79.28	72.62	75.31	82.21
Gemma-3-27B-pt	77.81	79.83	80.52	76.95	78.88	83.33
Gemma-3-27B-it [†]	78.03	79.87	82.19	75.96	79.41	85.33
Aya-101 [†]	52.62	53.75	62.16	57.05	56.73	59.86
BLOOMZ-7b1 [†]	34.27	32.63	30.03	32.30	32.40	28.80
mT0-xxl [†]	52.72	53.13	61.33	36.92	56.57	56.91
GLM-4-9b	63.19	63.81	68.22	63.41	64.98	68.52
GLM-4-9b-chat [†]	61.44	64.26	68.42	65.85	64.68	69.29
Greek and European LLMs						
Llama-Krikri-8B-Base	59.83	68.92	65.56	62.92	63.33	68.72
Llama-Krikri-8B-Instruct [†]	62.62	70.29	70.57	64.11	66.47	74.73
Meltemi-7B-v1.5	52.13	55.23	53.74	52.36	53.11	56.28
Meltemi-7B-Instruct-v1.5 [†]	57.18	63.99	64.31	61.03	60.93	66.42
Plutus-8B-instruct [†]	61.65	69.74	69.73	64.01	65.71	73.96
EuroLLM-1.7B	33.21	34.28	30.54	32.11	32.33	29.86
EuroLLM-1.7B-Instruct [†]	29.47	30.76	28.76	31.76	29.68	30.16
EuroLLM-9B	62.52	71.11	69.46	65.21	66.30	73.74
EuroLLM-9B-Instruct [†]	64.42	73.16	72.24	66.85	68.48	76.86
EuroLLM-22B	66.94	75.35	73.49	68.44	70.43	77.46
EuroLLM-22B-Instruct-2512 [†]	69.78	75.31	74.66	70.08	72.18	78.99
Random Baseline	32.33	28.77	31.86	32.62	30.42	31.59

Table 3: Overall zero-shot performance of different models on the GreekMMLU benchmark. Accuracy (%) is reported. Models marked with [†] are instruction-tuned, and ‡ are closed-source models.

sity/professional levels. Overall, the reduced accuracy at higher levels reflects the increased complexity, specialized knowledge of university and professional examination questions.

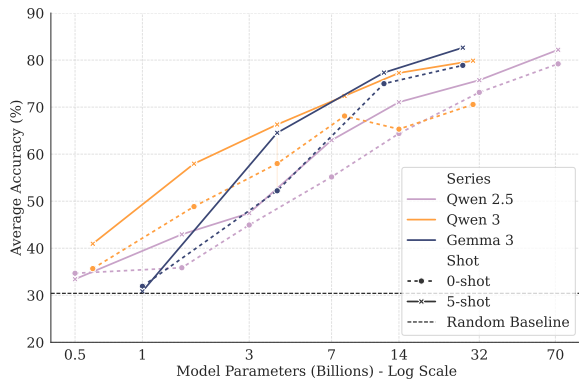


Figure 5: Scaling behavior of average accuracy with respect to model size (in billions of parameters) under zero-shot and five-shot prompting.

4.3 Analysis

Model Scale Effects. Figure 5 shows a clear relationship between model size and performance on GreekMMLU. Very small models (below approximately 2B parameters) generally cannot solve these tasks, with accuracy close to the random baseline across all model families. As model size increases, performance improves consistently. Larger models achieve higher accuracy, reflecting better answer selection and more effective handling of longer and more complex question contexts in Greek. This trend is observed across all evaluated model series, with steady gains as parameters increase, indicating that larger models are better equipped to handle the linguistic and knowledge demands of the benchmark.

Subject-Level and Cultural Differences. Model performance varies across subject domains, with most models achieving higher accuracy in humanities and social sciences than in STEM, reflecting the greater reasoning and technical demands of STEM questions in GreekMMLU. Greek-specific subjects (e.g., history, traditions, mythology) are consistently more challenging than globally shared domains, particularly for general-purpose multilingual models. In contrast, Greek- and European-centric models of comparable size perform better on these culturally grounded tasks, indicating stronger coverage of localized knowledge. Additional analysis is provided in Appendix G.

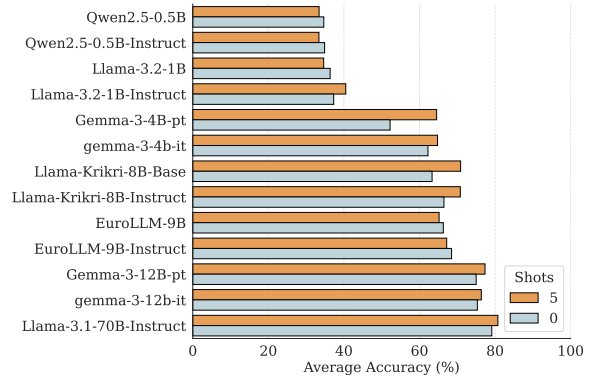


Figure 6: Comparison of average accuracy across models under zero-shot and five-shot prompting.

Zero-Shot vs. Five-Shot Performance. Figure 6 shows that five-shot prompting does not improve performance for models less than 2B, which remain close to the random baseline. In contrast, larger models consistently benefit from additional in-context examples. The improvement is most pronounced for Gemma and Llama families, where five-shot prompting yields noticeable accuracy boosts. Overall, the effectiveness of five-shot prompting is positive and can provide meaningful gains for mid- and large-scale ones.




Figure 7: Calibration behavior across models on GreekMMLU.

Calibration Analysis Our analysis of subject-level calibration on 5-shot results reveal notable differences across model families (Figure 7). The Greek-specialized Llama-Krikri-8B shows strong alignment between confidence and accuracy ($r = 0.93$), while generic multilingual models such as Qwen-2.5-7B are less well calibrated ($r = 0.74$). Earlier-generation models (e.g., Llama-2-7B) exhibit pronounced miscalibration ($r = 0.13$). Although larger models (e.g., Llama-3.1-70B) reduce calibration errors, language-specific training re-

mains the primary driver of reliable calibration. We provide extended calibration analysis in Appendix F.

Correlation Between Question Length and Model Confidence We further experiment with the correlation between the question length and model prediction confidence in Appendix H across different model families and sizes. We found that the model confidence has little or no correlation with question length.

5 Conclusion

We introduced GreekMMLU, the first large-scale, native-sourced benchmark for evaluating massive multitask language understanding in Greek. GreekMMLU comprises 21,805 multiple-choice questions across 45 subjects and multiple educational levels, enabling linguistically and culturally grounded evaluation beyond machine-translated benchmarks. Our comprehensive evaluation highlights substantial performance gaps across model families and demonstrates the benefits of instruction tuning and Greek-adapted training, particularly on Greek-specific domains. We hope GreekMMLU encourages future models to place greater emphasis on native Greek capability and supports the development of more authentic, culturally grounded Greek language models.

Limitations

GreekMMLU is limited to multiple-choice questions drawn from formal educational and professional settings, enabling standardized evaluation but not capturing open-ended generation, interactive reasoning, or informal language use. Although all content is natively sourced in Greek, the benchmark primarily reflects standard Modern Greek as used in official curricula, with limited coverage of regional dialects and colloquial registers, and consists exclusively of text-based inputs without multimodal content. Finally, as with other large-scale benchmarks based on naturally occurring real-world data, potential overlap with the training corpora of some models cannot be entirely ruled out.

6 Acknowledgments

This work was partially supported by the ANR/HELAS chair (ANR-CHIA-0020-01), led by M. Vazirgiannis, which funded members of the authoring group. We also extend our gratitude to C.

Stathopoulos for sharing physics teaching Q&A materials, and to I. Evdaimon and M. Lioudakis for their essential contributions to our initial data collection and processing steps.

References

- Mikel Artetxe, Sebastian Ruder, and Dani Yogatama. 2020. [On the cross-lingual transferability of monolingual representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4623–4637, Online. Association for Computational Linguistics.
- Nikos Bartziokas, Thanassis Mavropoulos, and Constantine Kotropoulos. 2020. [Datasets and Performance Metrics for Greek Named Entity Recognition](#). In *11th Hellenic Conference on Artificial Intelligence (SETN 2020)*, SETN 2020, pages 160–167, New York, NY, USA. Association for Computing Machinery.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Odysseas S Chlapanis, Dimitrios Galanis, Nikolaos Aletras, and Ion Androutsopoulos. 2025. Greekbarbench: A challenging benchmark for free-text legal reasoning and citations. *arXiv preprint arXiv:2505.17267*.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, and 1 others. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. 2018. [XNLI: Evaluating cross-lingual sentence representations](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485, Brussels, Belgium. Association for Computational Linguistics.
- Iakovos Evdaimon, Hadi Abdine, Christos Xypolopoulos, Stamatis Outsios, Michalis Vazirgiannis, and Giorgos Stamou. 2024. [GreekBART: The first pre-trained Greek sequence-to-sequence model](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources*

- and Evaluation (LREC-COLING 2024), pages 7949–7962, Torino, Italia. ELRA and ICCL.
- Jack FitzGerald, Christopher Hench, Charith Peris, Scott Mackie, Kay Rottmann, Ana Sanchez, Aaron Nash, Liam Urbach, Vishesh Kakarala, Richa Singh, Swetha Ranganath, Laurie Crist, Misha Britan, Wouter Leeuwis, Gokhan Tur, and Prem Natara-jan. 2023. [MASSIVE: A 1M-example multilingual natural language understanding dataset with 51 typologically-diverse languages](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4277–4302, Toronto, Canada. Association for Computational Linguistics.
- 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, and 5 others. 2024. [The language model evaluation harness](#).
- Team GLM, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, Jing Zhang, Juanzi Li, and 37 others. 2024. [Chatglm: A family of large language models from glm-130b to glm-4 all tools](#). *Preprint*, arXiv:2406.12793.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. [The Flores-101 evaluation benchmark for low-resource and multilingual machine translation](#). *Transactions of the Association for Computational Linguistics*, 10:522–538.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- 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 *International conference on machine learning*, pages 4411–4421. PMLR.
- Fajri Koto, Haonan Li, Sara Shatnawi, Jad Doughman, Abdelrahman Sadallah, Aisha Alraeesi, Khalid Almubarak, Zaid Alyafeai, Neha Sengupta, Shady Shehata, and 1 others. 2024. Arabicmmlu: Assessing massive multitask language understanding in arabic. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5622–5640.
- 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.
- Pedro Henrique Martins, Patrick Fernandes, João Alves, Nuno M Guerreiro, Ricardo Rei, Duarte M Alves, José Pombal, Amin Farajian, Manuel Faysse, Mateusz Klimaszewski, and 1 others. 2025. Eurollm: Multilingual language models for europe. *Procedia Computer Science*, 255:53–62.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 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, Toronto, Canada. Association for Computational Linguistics.
- Xueqing Peng, Triantafillos Papadopoulos, Efstathia Soufleri, Polydoros Giannouris, Ruoyu Xiang, Yan Wang, Lingfei Qian, Jimin Huang, Qianqian Xie, and Sophia Ananiadou. 2025. [Plutus: Benchmarking large language models in low-resource Greek finance](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 30176–30202, Suzhou, China. Association for Computational Linguistics.
- Dimitris Roussis, Leon Voukoutis, Georgios Paraskevopoulos, Sokratis Sofianopoulos, Prokopis Prokopidis, Vassilis Papavassileiou, Athanasios Katsamanis, Stelios Piperidis, and Vassilis Katsouros. 2025. [Krikri: Advancing open large language models for Greek](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 5012–5033, Suzhou, China. Association for Computational Linguistics.
- Guokan Shang, Hadi Abdine, Ahmad Chamma, Amr Mohamed, Mohamed Anwar, Abdelaziz Bounhar, Omar El Herraoui, Preslav Nakov, Michalis Vazirgiannis, and Eric P. Xing. 2025a. [Nile-chat: Egyptian language models for Arabic and Latin scripts](#). In *Proceedings of The Third Arabic Natural Language Processing Conference*, pages 306–322, Suzhou, China. Association for Computational Linguistics.
- Guokan Shang, Hadi Abdine, Yousef Khoubrane, Amr Mohamed, Yassine Abbahaddou, Sofiane Ennadir, Imane Momayiz, Xuguang Ren, Eric Moulines, Preslav Nakov, Michalis Vazirgiannis, and Eric Xing. 2025b. [Atlas-chat: Adapting large language models](#)

- for low-resource Moroccan Arabic dialect. In *Proceedings of the First Workshop on Language Models for Low-Resource Languages*, pages 9–30, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Shivalika Singh, Angelika Romanou, Clémentine Fourrier, David Ifeoluwa Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, and 1 others. 2025. Global mmlu: Understanding and addressing cultural and linguistic biases in multilingual evaluation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 18761–18799.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, and 1 others. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Ahmet Üstün, Viraat Aryabumi, Zheng Yong, Wei-Yin Ko, Daniel D’souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, and 1 others. 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.
- Leon Voukoutis, Dimitris Roussis, Georgios Paraskevopoulos, Sokratis Sofianopoulos, Prokopis Prokopidis, Vassilis Papavasileiou, Athanasios Katsamanis, Stelios Piperidis, and Vassilis Katsouros. 2024. Meltemi: The first open large language model for greek. *arXiv preprint arXiv:2407.20743*.
- Weihao Xuan, Rui Yang, Heli Qi, Qingcheng Zeng, Yunze Xiao, Aosong Feng, Dairui Liu, Yun Xing, Junjue Wang, Fan Gao, Jinghui Lu, Yuang Jiang, Huitao Li, Xin Li, Kunyu Yu, Ruihai Dong, Shangding Gu, Yuekang Li, Xiaofei Xie, and 13 others. 2025. [MMLU-ProX: A multilingual benchmark for advanced large language model evaluation](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 1513–1532, Suzhou, China. Association for Computational Linguistics.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Arda Yüksel, Abdullatif Köksal, Lütfi Kerem Senel, Anna Korhonen, and Hinrich Schütze. 2024. Turk-
ishmmlu: Measuring massive multitask language understanding in turkish. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7035–7055.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*.

A GreekMMLU Tasks and Examples

Task	Tested Concepts	Supercategory	# Q
Accounting	Accounting, balance sheets, microeconomics, institutions...	Social Sciences	189
Agriculture Professional	Circular bioeconomy, social economy, agri-food systems...	STEM	344
Agriculture University	Smart agriculture, circular bioeconomy, agri-food systems...	STEM	75
Art Professional	Applied arts, digital design, materials, fashion ...	Humanities	605
Art Secondary School	Greek cinema, Greek music, cultural figures, arts history..	Humanities	44
Art University	Music theory, rhythm, Greek tradition, ethnomusicology...	Humanities	17
Biology	Basic biology, microorganisms, human body, animal biology..	STEM	423
Chemistry	Thermodynamics, physical chemistry, phase equilibria, surface chemistry...	STEM	86
Civil Engineering	Building systems, materials, safety, mechanics.	STEM	774
Clinical Knowledge	Clinical basics, anatomy and physiology, nursing care, dermatology..	STEM	638
Computer Networks and Security	Packet-switched networks, architectures, protocol layers, Internet...	STEM	91
Computer Science Professional	Computing fundamentals, networking, software, digital systems...	STEM	259
Computer Science University	Computer systems, networking, data analysis, interactive technologies.	STEM	109
Driving Rules	Traffic regulations, road safety, traffic signs, driver responsibilities.	Other	1663
Economics Professional	Macroeconomics, international trade, public finance, labor markets.	Social Sciences	201
Economics University	Strategy, digital business, innovation, technology, costs and pricing.	Social Sciences	91
Education Professional	Child development, creativity, play-based learning, pedagogy...	Social Sciences	258
Education University	Educational research, data analysis, academic writing.	Social Sciences	46
Electrical Engineering	Electrical circuits, sensors, automotive systems, protection devices.	STEM	557
General Knowledge	Safety procedures, natural phenomena, first aid.	Other	356
Geography Primary School	Physical and human geography, cartography, regions...	Social Sciences	365
Geography Secondary School	Physical and human geography, cartography...	Social Sciences	67
Government and Politics Primary School	Civics, democracy, institutions, governance, rights and duties.	Social Sciences	285
Government and Politics Secondary School	Political institutions, citizenship, rights, public administration...	Social Sciences	80
Greek History Primary School	Greek history from antiquity to modern times, key events and figures.	Humanities	469
Greek History Professional	Modern Greek political and constitutional history	Humanities	122
Greek History Secondary School	Byzantine Empire and Modern Greek State	Humanities	98
Greek Literature	Authors, Literary Movements and Major Works (Modern and Classical)	Humanities	19
Greek Mythology	Creation myths, heroic cycles, myth in culture...	Humanities	243
Greek Traditions	Food Safety, Baking Science and Culinary Operations	Social Sciences	381
Law	Administrative Law, EU Law and Sports Regulations	Humanities	941
Management Professional	Management, Economics and Business Operations	Social Sciences	646
Management University	Human Resource Management	Social Sciences	30
Mathematics	Mathematical reasoning, algebra, geometry, basic statistics...	STEM	1123
Medicine Professional	Professional-level medical knowledge and clinical reasoning.	STEM	467
Medicine University	System-level human physiology and functional integration of organs.	STEM	77
Maritime Safety and Rescue Operations	Maritime safety, rescue procedures, emergency response at sea.	Other	153
Modern Greek Language Primary School	Basic grammar, morphology, syntax, orthography..	Social Sciences	1483
Modern Greek Language Secondary School	Vocabulary, comprehension, meaning, language use..	Social Sciences	885
Physics Primary School	Basic physics, matter, energy, measurements, everyday phenomena.	STEM	432
Physics Professional	Applied physics and engineering	STEM	1336
Physics University	Scientific reasoning, physics–biology concepts, thermodynamics, optics...	STEM	76
Prehistory	Cycladic, Minoan, and Mycenaean civilizations.	Humanities	68
World History	Enlightenment, reformation, renaissance, revolutionary movements...	Humanities	25
World Religions	Orthodox Christian hymnology.	Humanities	160
Total			16,857

Table 4: Summary of the 45 subjects in the GreekMMLU public dataset. # Q indicates the total number of questions for each task.

Subject	Question	Choices
STEM	Ποια σώματα, όταν δέχονται το ίδιο φως (π.χ. από τον Ήλιο), απορροφούν περισσότερη ενέργεια;	A. Τα σκουρόχρωμα σώματα B. Τα ανοιχτόχρωμα σώματα C. Τα διαφανή σώματα D. Τα μεταλλικά σώματα
	Which bodies, when receiving the same light (e.g. from the Sun), absorb more energy?	A. Dark-colored bodies B. Light-colored bodies C. Transparent bodies D. Metallic bodies
Humanities	Ποιος σκηνοθέτησε την ταινία «Κυνόδοντας», η οποία υπήρξε το 2011 υποψήφια για Όσκαρ Καλύτερης Ξενόγλωσσας Ταινίας;	A. Παντελής Βούλγαρης B. Γιώργος Λάνθιμος C. Κώστας Γαβράς D. Θεόδωρος Αγγελόπουλος
	Who directed the film Dogtooth, which was nominated in 2011 for the Academy Award for Best Foreign Language Film?	A. Pantelis Voulgaris B. Yorgos Lanthimos C. Costa-Gavras D. Theodoros Angelopoulos
Social Sciences	Ποιος είναι ο πρώτος φορέας κοινωνικοποίησης του ατόμου;	A. Το σχολείο A B. Η οικογένεια C. Οι φίλοι/συνομήλικοι D. Οι επαγγελματικές σχέσεις
	What is the primary agent of socialization of an individual?	A. School B. Family C. Friends/peers D. Professional relationships
Other	Η Ανώνυμη Εταιρεία (Α.Ε.) είναι εμπορική εταιρεία:	A. όταν ενεργεί εμπορικές πράξεις. B. όταν ο σκοπός της είναι εμπορικός. C. ανεξάρτητα από τον σκοπό της. D. Η Ανώνυμη Εταιρεία (Α.Ε.) δεν είναι εμπορική εταιρεία.
	A Société Anonyme (S.A. / public limited company) is considered a commercial company:	A. when it carries out commercial acts. B. when its purpose is commercial. C. regardless of its purpose. D. A Société Anonyme is not a commercial company.
Greek-specific	Ποιος ήταν ο Αίολος;	A. Ο βασιλιάς των Λαιστρυγόνων B. Ένας σύντροφος του Οδυσσέα C. Ο θεός των ανέμων D. Ο πατέρας της Κίρκης
	Who was Aeolus?	A. The king of the Laestrygonians B. A companion of Odysseus C. The god of the winds D. The father of Circe

Table 5: Examples from GreekMMLU with their corresponding English translations across different subjects, where the bold items indicate the correct choices.

B Statistics of GreekMMLU

As reported in Table 6, token-length statistics were obtained using `tiktoken` with the `cl100k_base` encoding. For each question, we computed the number of tokens in the question text (reported as Avg. Q Tokens) and the total number of tokens across all answer choices (reported as Avg. C Tokens), where choices were concatenated using a single space separator. These values were then averaged across all questions within each subject group to summarize token distribution patterns across the dataset.

As shown in Figure 8, the distributional characteristics of text length vary across subject categories in the GreekMMLU benchmark. The left panel reports question-length distributions, indicating that most subjects exhibit median lengths below 200 characters. Notable exceptions include domains such as Modern Greek Language (Secondary School), which display substantially longer inputs, primarily due to the inclusion of extended reading-comprehension passages. The right panel presents answer-length distributions, which are generally more compact; however, technical and professional domains, including Physics and Law, are characterized by longer answer options, reflecting the increased precision and explanatory detail required in these fields. Overall, this variation in sequence length underscores the benchmark’s ability to evaluate model performance across both short factual queries and longer, context-dependent inputs.

C Implementation Details

All evaluations were conducted using the `lm-evaluation-harness` framework (Gao et al., 2024) (version 0.4.9.1). The evaluation setup follows a unified configuration derived from the GreekMMLU benchmark, ensuring consistent assessment across all subject domains.

Models were evaluated on the GreekMMLU benchmark that we implemented based on the MMLU format, which consists of subject-specific multiple-choice questions covering STEM, Humanities, Social Sciences, and Other domains. Evaluation was performed under standardized zero-shot and five-shot prompting conditions.

All evaluations for open-source models relied on the default deterministic behavior of the `lm-evaluation-harness` for multiple-choice tasks. Model predictions were obtained via log-likelihood comparison over answer options rather

than generative sampling.

For closed-source models accessed via external APIs, evaluation was performed using a generation-based multiple-choice protocol. Generation was conducted under deterministic settings, with sampling disabled (`do_sample=false`) and temperature set to 0.0.

Model outputs were post-processed using a standardized parsing procedure that extracts the first valid answer symbol, supporting both Greek (Α, Β, Γ, Δ) and Latin (A, B, C, D) representations. When Latin characters were produced, they were deterministically mapped to their Greek equivalents. The extracted prediction was then compared against the answer label to compute accuracy. This procedure supports a variable number of answer choices and was applied uniformly across all evaluated models.

Experiments were conducted on NVIDIA GPUs, primarily using RTX A6000. For larger models exceeding 30B parameters, evaluations were performed on NVIDIA A100 GPUs to accommodate increased memory and compute requirements.

Group	Tasks	#Q	Questions per Task			Avg. Tokens	
			Avg.	Max.	Min.	Question	Choices
STEM	16	6787	424.19	1331	70	83.91	76.86
Humanities	12	2751	229.25	936	12	84.29	132.65
Social Science	13	4753	365.62	1478	25	266.50	66.52
Other	4	2341	585.25	1658	148	77.83	133.96
All	45	16632	369.60	1658	12	135.30	91.17

Table 6: Statistics of the GreekMMLU

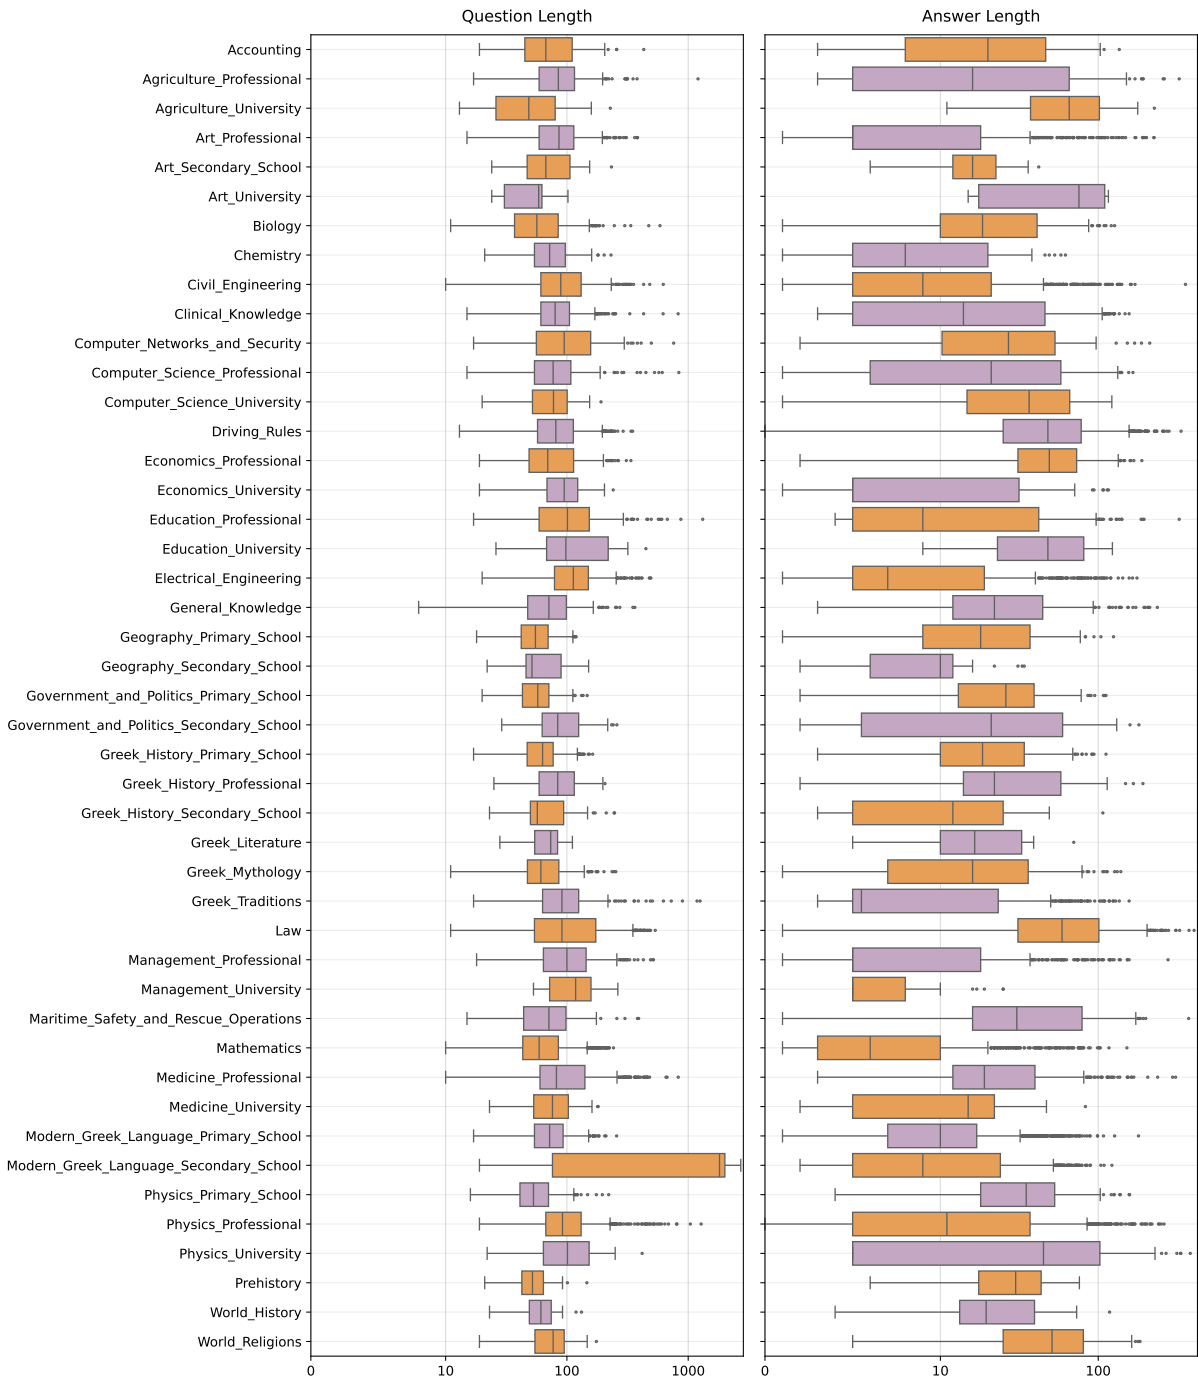


Figure 8: Distribution of character lengths for questions and answer choices across all GreekMMLU subjects.

D Experiment results

Model	STEM	Humanities	Social Sci.	Other	Average	Greek-specific
General-Purpose LLMs						
Qwen2.5-0.5B	35.13	31.31	31.14	36.24	33.43	29.67
Qwen2.5-1.5B	43.02	43.63	43.28	41.02	42.94	42.65
Qwen2.5-3B	46.22	46.23	47.97	51.62	47.46	46.09
Qwen2.5-7B	62.21	62.21	64.65	61.72	62.96	65.98
Qwen2.5-14B	70.00	71.84	73.11	68.14	71.06	76.34
Qwen2.5-32B	75.60	74.40	77.52	72.72	75.73	79.21
Qwen2.5-72B	82.22	82.84	83.19	78.60	82.18	85.79
Qwen3-0.6B	41.74	38.43	40.51	42.21	40.95	38.33
Qwen3-1.7B	59.61	52.99	58.21	57.14	57.97	57.46
Qwen3-4B	69.34	60.79	67.93	64.51	67.14	67.95
Qwen3-8B	74.05	68.78	73.91	69.69	72.77	75.85
Qwen3-14B	78.27	72.20	79.53	73.17	77.26	81.17
Qwen3-30B	81.74	75.31	80.46	76.80	79.86	82.32
Llama-2-7b-hf	36.53	35.60	33.53	33.15	34.99	32.87
Llama-3-8B	61.79	65.59	66.76	64.11	64.24	68.96
Llama-3.1-70B	77.62	82.79	83.66	78.35	80.41	87.30
Llama-3.1-8B	61.15	63.58	65.89	62.77	63.25	67.21
Llama-3.2-1B	34.54	33.73	34.26	37.08	34.65	32.68
Llama-3.2-3B	46.91	49.93	48.88	50.57	48.42	46.53
Gemma-3-1B-pt	30.94	30.72	30.23	32.16	30.82	30.03
Gemma-3-4B-pt	62.53	64.95	67.31	63.27	64.54	69.75
Gemma-3-12B-pt	77.16	78.59	77.48	76.16	77.34	79.48
Gemma-3-27B-pt	82.29	82.88	83.68	80.69	82.64	86.53
XGLM-1.7B	32.57	38.53	29.55	26.50	31.71	30.26
XGLM-2.9B	30.79	41.85	28.34	29.34	30.72	29.66
XGLM-4.5B	33.44	41.02	29.45	27.92	32.37	31.56
XGLM-7.5B	30.79	41.85	28.34	29.34	30.72	29.66
GLM-4-9B	67.70	69.69	74.00	68.79	70.20	74.23
Greek and European LLMs						
Llama-Krikri-8B-Base	66.75	75.63	75.11	68.09	70.88	80.22
Meltemi-7B-v1	60.26	64.58	51.19	64.91	58.38	49.10
Meltemi-7B-v1.5	59.88	67.14	64.11	65.95	62.99	67.10
EuroLLM-1.7B	32.71	30.94	30.71	36.54	32.27	30.49
EuroLLM-9B	60.91	69.92	68.00	66.75	65.18	73.03
EuroLLM-22B	71.65	77.23	77.01	71.13	74.11	82.32
Random Baseline	32.33	28.77	31.86	32.62	30.42	31.59

Table 7: Overall five-shot performance of **base** LLMs on the GreekMMLU benchmark. Accuracy (%) is reported. The *Greek-specific* column includes an average of History, Traditions, and Mythology subsets.

Model	STEM	Humanities	Social Sci.	Other	Average	Greek-specific
General-Purpose LLMs						
GPT-5.2	88.29	90.28	90.74	86.76	89.18	93.85
GPT-4o	85.62	90.05	90.29	86.16	87.83	93.22
Gemini 3 Flash	91.24	88.63	98.10	87.46	91.20	93.52
Qwen2.5-0.5B-Instruct	34.82	32.72	30.58	36.98	33.39	29.26
Qwen2.5-1.5B-Instruct	46.66	45.82	46.86	45.84	46.52	45.98
Qwen2.5-3B-Instruct	51.82	52.67	54.58	56.20	53.39	55.11
Qwen2.5-7B-Instruct	64.20	63.35	64.85	63.36	64.20	67.27
Qwen2.5-14B-Instruct	69.90	70.52	72.20	68.59	70.59	75.36
Qwen2.5-32B-Instruct	75.69	73.66	78.72	72.22	76.01	81.39
Qwen2.5-72B-Instruct	80.93	81.65	83.14	78.25	81.44	85.90
Qwen3-4B-Instruct-2507	70.46	62.25	68.66	66.80	68.32	69.59
Qwen3-30B-Instruct	82.42	77.09	81.50	77.90	80.85	83.42
Llama-2-7b-chat-hf	38.03	39.66	35.58	33.10	36.83	35.60
Llama-3-8B-Instruct	61.91	65.40	68.55	64.31	64.88	71.26
Llama-3.1-8B-Instruct	61.77	67.09	68.22	62.72	64.74	71.12
Llama-3.1-70B-Instruct	79.09	81.93	83.34	77.90	80.74	86.80
Llama-3.2-1B-Instruct	40.93	40.58	38.98	43.01	40.49	36.15
Llama-3.2-3B-Instruct	48.36	49.34	52.19	54.06	50.46	51.99
Llama-3.3-70B-Instruct	79.65	82.47	84.19	79.09	81.47	87.90
Mistral-7B-Instruct-v0.3	50.10	51.94	52.72	53.11	51.58	53.52
Gemma-3-1B-it	45.71	46.33	45.01	46.54	45.66	43.80
Gemma-3-4B-it	62.84	63.62	68.16	63.27	64.77	70.00
Gemma-3-12B-it	74.51	75.35	79.79	74.27	76.35	81.97
Gemma-3-27B-it	80.51	81.70	83.05	78.50	81.27	86.12
Aya-101	54.37	54.71	64.78	57.94	58.59	62.62
Aya-epsilon-8b	62.18	66.41	69.80	65.11	65.64	71.07
BLOOMZ-1b1	32.17	39.48	29.11	31.62	31.60	30.02
BLOOMZ-1b7	31.78	39.24	28.18	29.06	30.95	29.13
BLOOMZ-7b1	31.88	32.04	30.02	30.86	31.16	28.88
mT0-large	31.56	40.31	27.97	30.20	30.88	28.84
mT0-xl	39.54	47.44	41.08	46.44	41.00	41.34
mT0-xxl	48.95	43.72	55.58	53.69	51.14	55.36
GLM-4-9B-chat	67.31	66.54	72.15	67.40	68.83	72.60
Greek and European LLMs						
Llama-Krikri-8B-Instruct	67.04	75.58	74.57	67.94	70.80	79.21
Meltemi-7B-Instruct-v1.5	58.91	66.86	68.24	64.11	63.71	69.97
Plutus-8B-instruct	67.23	74.81	74.42	68.34	70.77	78.36
EuroLLM-1.7B-Instruct	32.87	31.36	30.73	36.29	32.37	31.07
EuroLLM-9B-Instruct	62.37	72.43	70.88	67.75	67.20	75.96
EuroLLM-22B-Instruct-2512	72.09	78.73	78.01	72.92	75.05	83.06
Random Baseline	32.33	28.77	31.86	32.62	30.42	31.59

Table 8: Overall five-shot performance of **instruction-tuned** LLMs on the GreekMMLU benchmark. Accuracy (%) is reported. The *Greek-specific* column includes an average of History, Traditions, and Mythology subsets.

Model	STEM	Humanities	Social Sci.	Other	Average	Greek-specific
General-Purpose LLMs						
Qwen2.5-0.5B	35.20	33.32	34.18	35.74	34.68	32.30
Qwen2.5-1.5B	36.19	35.33	34.27	39.62	35.85	31.48
Qwen2.5-3B	46.34	44.64	45.02	40.32	44.94	44.86
Qwen2.5-7B	51.92	56.37	59.34	53.31	55.16	61.07
Qwen2.5-14B	64.11	64.95	66.04	60.23	64.39	67.95
Qwen2.5-32B	72.74	71.29	75.26	70.68	73.14	77.40
Qwen2.5-72B	78.31	78.78	81.35	76.75	79.20	83.91
Qwen3-0.6B	36.35	34.50	34.95	36.64	35.67	33.11
Qwen3-1.7B	49.62	44.00	49.70	49.23	48.85	47.87
Qwen3-4B	64.92	57.78	64.67	61.27	63.44	64.97
Qwen3-8B	70.27	63.49	69.89	66.50	68.78	71.17
Qwen3-14B	62.59	61.89	69.11	67.89	65.32	68.20
Qwen3-30B	70.72	69.06	74.95	59.63	70.56	77.49
Llama-2-7b-hf	36.23	35.97	33.80	35.54	35.30	32.84
Llama-3-8B	52.39	57.01	56.54	50.02	54.10	56.75
Llama-3.1-8B	51.08	57.42	54.78	50.62	53.10	54.78
Llama-3.1-70B	72.42	79.05	78.90	74.46	75.71	82.46
Llama-3.2-1B	37.37	36.51	34.86	36.83	36.35	33.66
Llama-3.2-3B	41.82	41.67	43.81	44.75	42.82	41.97
Gemma-3-1B-pt	33.11	32.72	30.29	31.41	31.91	30.00
Gemma-3-4B-pt	50.77	52.67	53.83	52.17	52.21	55.00
Gemma-3-12B-pt	73.30	74.44	78.28	72.27	74.99	80.71
Gemma-3-27B-pt	77.81	79.83	80.52	76.95	78.88	83.33
XGLM-1.7B	34.07	35.01	31.55	27.35	32.98	31.67
XGLM-2.9B	33.25	29.39	34.94	26.92	32.68	32.90
XGLM-4.5B	33.93	36.86	31.26	27.07	33.00	31.42
XGLM-7.5B	33.97	34.72	29.86	24.79	32.16	30.11
GLM-4-9B	63.19	63.81	68.22	63.41	64.98	68.52
Greek and European LLMs						
Llama-Krikri-8B-Base	59.83	68.92	65.56	62.92	63.33	68.72
Meltemi-7B-v1	52.90	56.64	43.71	56.45	50.76	40.98
Meltemi-7B-v1.5	52.13	55.23	53.74	52.36	53.11	56.28
EuroLLM-1.7B	33.21	34.28	30.54	32.11	32.33	29.86
EuroLLM-9B	62.52	71.11	69.46	65.21	66.30	73.74
EuroLLM-22B	66.94	75.35	73.49	68.44	70.43	77.46
Random Baseline	32.33	28.77	31.86	32.62	30.42	31.59

Table 9: Overall zero-shot performance of **base** LLMs on the GreekMMLU benchmark. Accuracy (%) is reported. The *Greek-specific* column includes an average of History, Traditions, and Mythology subsets.

Model	STEM	Humanities	Social Sci.	Other	Average	Greek-specific
General-Purpose LLMs						
GPT-5.2	86.05	88.27	90.29	85.96	87.75	92.92
GPT-4o	84.54	88.68	89.36	85.42	86.81	93.11
Gemini 3 Flash	92.82	92.88	94.16	91.84	93.16	95.44
Qwen2.5-0.5B-Instruct	35.20	33.91	34.40	36.24	34.89	32.51
Qwen2.5-1.5B-Instruct	34.08	34.19	33.33	34.69	33.92	34.86
Qwen2.5-3B-Instruct	50.67	50.07	51.26	48.28	50.50	51.97
Qwen2.5-7B-Instruct	59.88	58.33	61.98	58.89	60.25	64.02
Qwen2.5-14B-Instruct	65.20	66.41	69.78	62.92	66.61	73.06
Qwen2.5-32B-Instruct	72.08	71.47	76.61	69.69	73.22	80.03
Qwen2.5-72B-Instruct	78.90	79.14	81.92	76.95	79.70	84.67
Qwen3-4B-Instruct-2507	65.35	57.60	66.64	63.46	64.52	67.16
Qwen3-30B-Instruct	79.31	74.81	79.33	76.56	78.39	81.80
Llama-2-7b-chat-hf	36.63	34.92	34.26	33.85	35.27	33.61
Llama-3-8B-Instruct	59.01	62.85	64.16	60.38	61.40	65.38
Llama-3.1-8B-Instruct	56.58	62.85	62.56	57.84	59.56	64.75
Llama-3.1-70B-Instruct	76.35	82.34	82.57	75.61	79.13	86.99
Llama-3.2-1B-Instruct	38.01	37.33	36.87	35.94	37.29	35.46
Llama-3.2-3B-Instruct	43.77	43.04	45.62	45.64	44.52	46.15
Llama-3.3-70B-Instruct	77.03	82.20	82.92	76.80	79.65	86.94
Mistral-7B-Instruct-v0.3	48.05	48.93	51.23	48.23	49.25	52.02
Mistral-Small-24B-Instruct-2501	67.16	71.72	74.79	66.39	70.47	78.06
Gemma-3-1B-it	45.25	44.45	44.19	46.59	44.95	42.05
Gemma-3-4B-it	59.79	60.43	65.95	62.32	62.24	68.58
Gemma-3-12B-it	72.95	75.13	79.28	72.62	75.31	82.21
Gemma-3-27B-it	78.03	79.87	82.19	75.96	79.41	85.33
Aya-101	52.62	53.75	62.16	57.05	56.73	59.86
Aya-expanse-8b	60.75	65.13	68.33	63.27	64.17	69.51
BLOOMZ-1b1	35.64	34.37	34.26	36.14	35.07	32.68
BLOOMZ-1b7	35.66	34.19	33.15	35.94	34.66	31.53
BLOOMZ-7b1	34.27	32.63	30.03	32.30	32.40	28.80
mT0-large	36.89	36.47	37.09	37.28	36.95	34.95
mT0-xl	46.44	47.88	54.21	52.51	49.96	52.24
mT0-xxl	52.72	53.13	61.33	36.92	56.57	56.91
GLM-4-9B-chat	61.44	64.26	68.42	65.85	64.68	69.29
Greek and European LLMs						
Llama-Krikri-8B-Instruct	62.62	70.29	70.57	64.11	66.47	74.73
Meltemi-7B-Instruct-v1.5	57.18	63.99	64.31	61.03	60.93	66.42
Plutus-8B-instruct	61.65	69.74	69.73	64.01	65.71	73.96
EuroLLM-1.7B-Instruct	29.47	30.76	28.76	31.76	29.68	30.16
EuroLLM-9B-Instruct	64.42	73.16	72.24	66.85	68.48	76.86
EuroLLM-22B-Instruct-2512	69.78	75.31	74.66	70.08	72.18	78.99
Random Baseline	32.33	28.77	31.86	32.62	30.42	31.59

Table 10: Overall zero-shot performance of **instruction-tuned** LLMs on the GreekMMLU benchmark. Accuracy (%) is reported. The *Greek-specific* column includes an average of History, Traditions, and Mythology subsets.

E Comparison Between the Public and Private Results

Category	Pearson r	p -value
STEM	0.9973	3.18×10^{-74}
Humanities	0.9895	1.98×10^{-55}
Social Sci.	0.9966	5.01×10^{-71}
Other	0.9929	7.42×10^{-61}
Average	0.9984	1.85×10^{-81}
Combined	0.9908	3.27×10^{-287}

Table 11: Correlation between Public and Private Greek-MMLU Scores (5-shot).

Category	Pearson r	p -value
STEM	0.9986	2.83×10^{-83}
Humanities	0.9934	5.39×10^{-62}
Social Sci.	0.9970	5.96×10^{-73}
Other	0.9881	1.01×10^{-53}
Average	0.9988	4.17×10^{-86}
Combined	0.9901	4.74×10^{-282}

Table 12: Correlation between Public and Private Greek-MMLU Scores (0-shot).

To verify that the private split of GreekMMLU provides a reliable estimate of model performance, we analyze the correlation between public and private evaluation results for each model. Figures 9 and 10 show a strong linear relationship between public and private scores under zero-shot and five-shot settings, respectively. This observation is quantified in Tables 12 and 11, where Pearson correlation coefficients exceed 0.98 across all subject categories and evaluation setups, with extremely small p -values.

The consistently high correlations across STEM, Humanities, Social Sciences, and Other domains indicate that relative model rankings are well preserved between the two splits. These results demonstrate that the private GreekMMLU set faithfully reflects model performance observed on the public data, supporting its use as a robust and contamination-resistant benchmark for leaderboard-based evaluation.

F Extended Calibration

We evaluated the calibration of a wider set of 11 models in the 5-shot setting (Table 13). The results




Figure 9: Correlation between Public and Private Greek-MMLU Scores (0-shot)




Figure 10: Correlation between Public and Private GreekMMLU Scores (5-shot)

highlight three distinct behaviors:

We evaluated the calibration of a wider set of 11 models in the 5-shot setting (Table 13). The results highlight three distinct behaviors. The top-performing models, including the Greek-specific Llama-Krikri-8B and the Llama 3.1 family, fall into a highly calibrated cluster, exhibiting the highest correlation ($r \geq 0.95$) and accurately reflecting their true probability of correctness. A second group, comprising the Qwen 2.5/3 family, shows moderate calibration with strong but slightly lower correlation coefficients ($r \approx 0.87 - 0.93$). Finally, the older Llama-2-7b baseline remains poorly calibrated with significantly lower correlation ($r = 0.46$)




Figure 11: Subject performance distribution (0-shot)




Figure 12: Subject performance distribution (5-shot)

G Subject-Level Performance Distribution

Figures 11 and 12 illustrate the distribution of accuracy scores across the 45 subjects of the GreekMMLU benchmark in zero-shot and five-shot settings, respectively, offering granular insights into model robustness beyond aggregate metrics. The analysis consistently reveals a distinct relationship between model specialization, scale, and cross-domain consistency across both prompting strategies. Notably, the Greek-centric Llama-Krikri-8B demonstrates a significant upward shift in its performance distribution relative to generic models of comparable size, such as Llama-3.1-8B and EuroLLM-9B. Its elevated performance baseline indicates a resilience against catastrophic failure modes on linguistically or culturally complex subjects, whereas generic counter-

parts frequently degrade to near-random accuracy in these areas.

Furthermore, the data highlights the stabilizing effect of model scale. The largest evaluated model, Llama-3.1-70B, exhibits the tightest clustering of subject scores within the upper quartile in both settings. This suggests that massive parameterization functions as a stabilizing factor, ensuring consistent competency across niche domains where smaller architectures struggle. In contrast, smaller models like Qwen2.5-0.5B display extreme vertical variance; while they occasionally achieve parity on simpler tasks, they lack the generalization capabilities required to maintain robust performance across the full breadth of the academic curriculum.

As illustrated in Figures 13 and 14, a subject-wise breakdown reveals heterogeneous performance patterns across the benchmark. Accuracy

Model	Pearson r
Llama-Krikri-8B-Base	0.96
EuroLLM-22B	0.95
Llama-3.1-70B	0.95
Llama-3.1-8B	0.95
Qwen3-30B	0.93
EuroLLM-9B	0.93
Qwen2.5-72B	0.92
Llama-3.3-70B	0.91
Qwen2.5-32B	0.88
Qwen2.5-7B	0.87
Llama-2-7b-hf	0.46

Table 13: Extended calibration analysis ranked by Pearson correlation coefficient (r) in the 5-shot setting.

levels differ notably between disciplinary groups, with non-technical domains generally exhibiting narrower variance across models, while technical subjects display wider performance dispersion. Topics involving culturally grounded knowledge, including Greek historical and mythological content, introduce additional variability, particularly among general multilingual models. In contrast, models incorporating regionally focused training data tend to show more stable behavior across these subjects.

H Impact of Question Length on Model Confidence

To investigate whether model confidence is a byproduct of input verbosity rather than semantic certainty, we analyze the correlation between question length (measured in characters) and average confidence scores. As illustrated in Figures 15 and 16, we observe no meaningful correlation between question length and model confidence for either the specialized Llama-Krikri-8B or the generic baselines (e.g., Llama-3.1-70B). Pearson correlation coefficients remain consistently close to zero across prompting strategies, indicating that the models' uncertainty estimates are robust to variations in input length and are not driven by superficial properties of the prompt, with Llama-based models showing a marginally higher sensitivity to question length.

Per-Task Accuracy (5-shot)



Figure 13: Subject-level accuracy heatmap under five-shot prompting, showing performance variation across GreekMMLU subjects and models.

Per-Task Accuracy (0-shot)

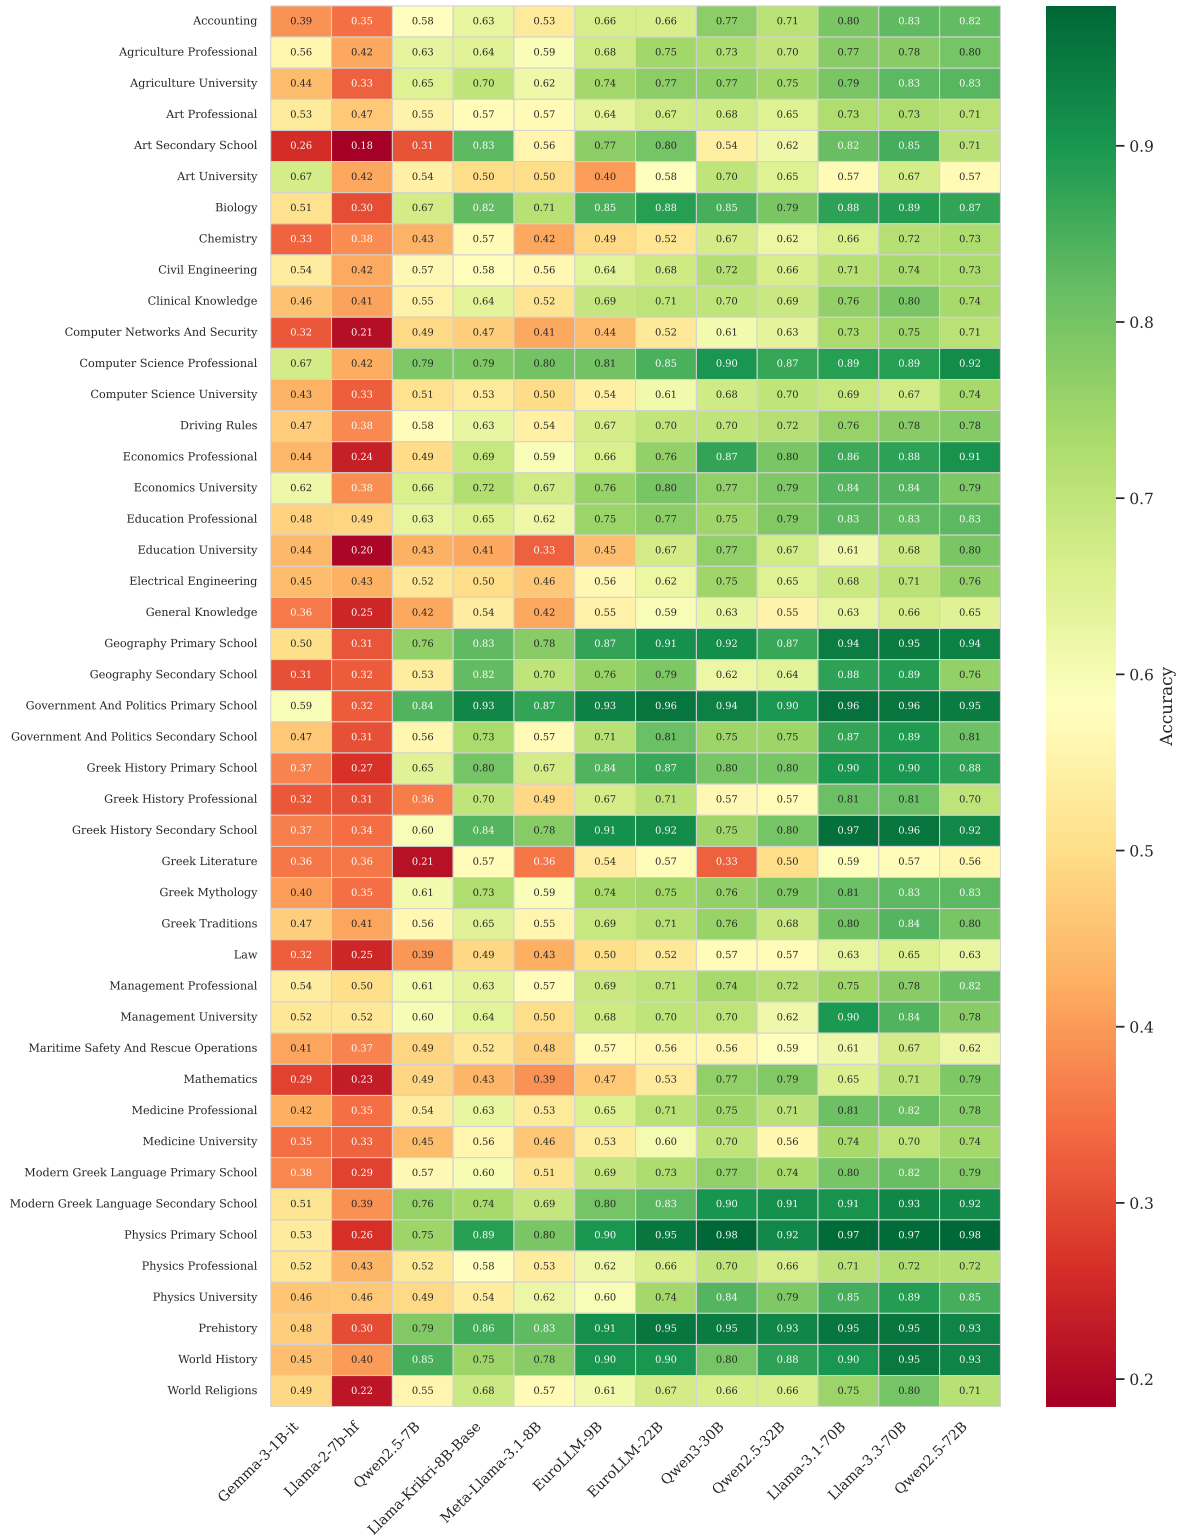


Figure 14: Subject-level accuracy heatmap under zero-shot prompting, showing performance variation across GreekMMLU subjects and models.

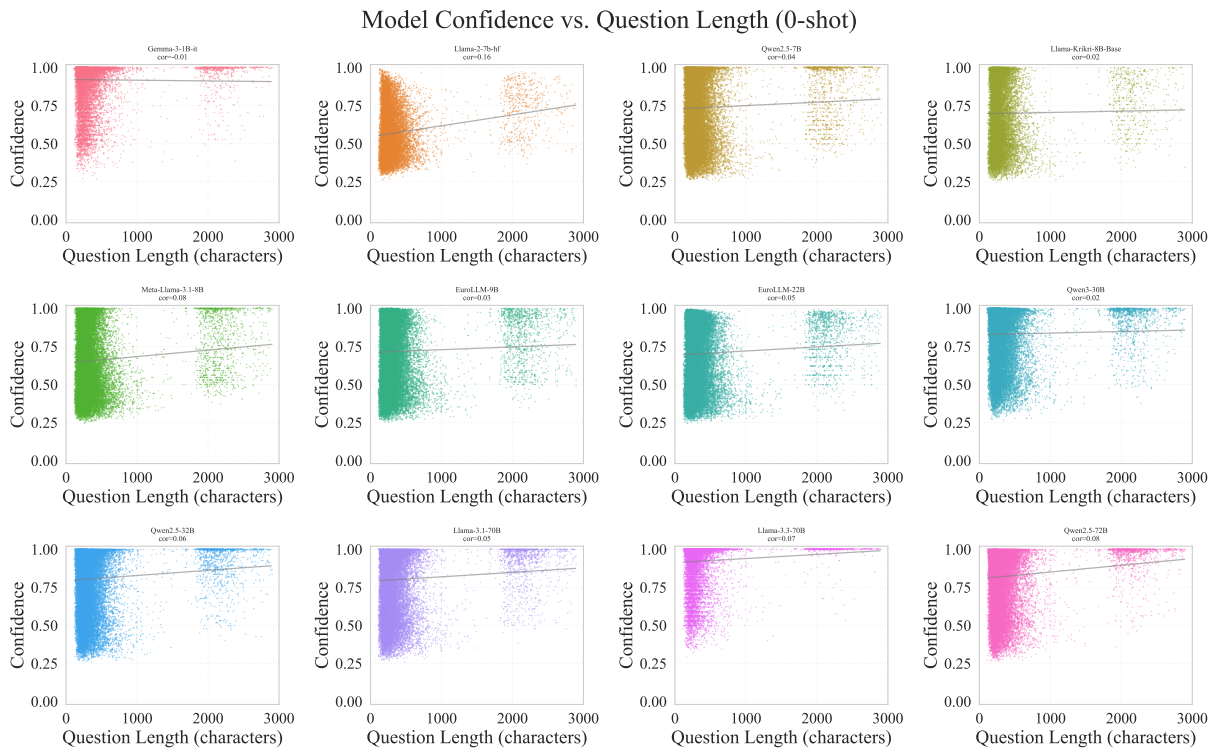


Figure 15: Correlation between Question Length and Confidence (Zero-Shot).

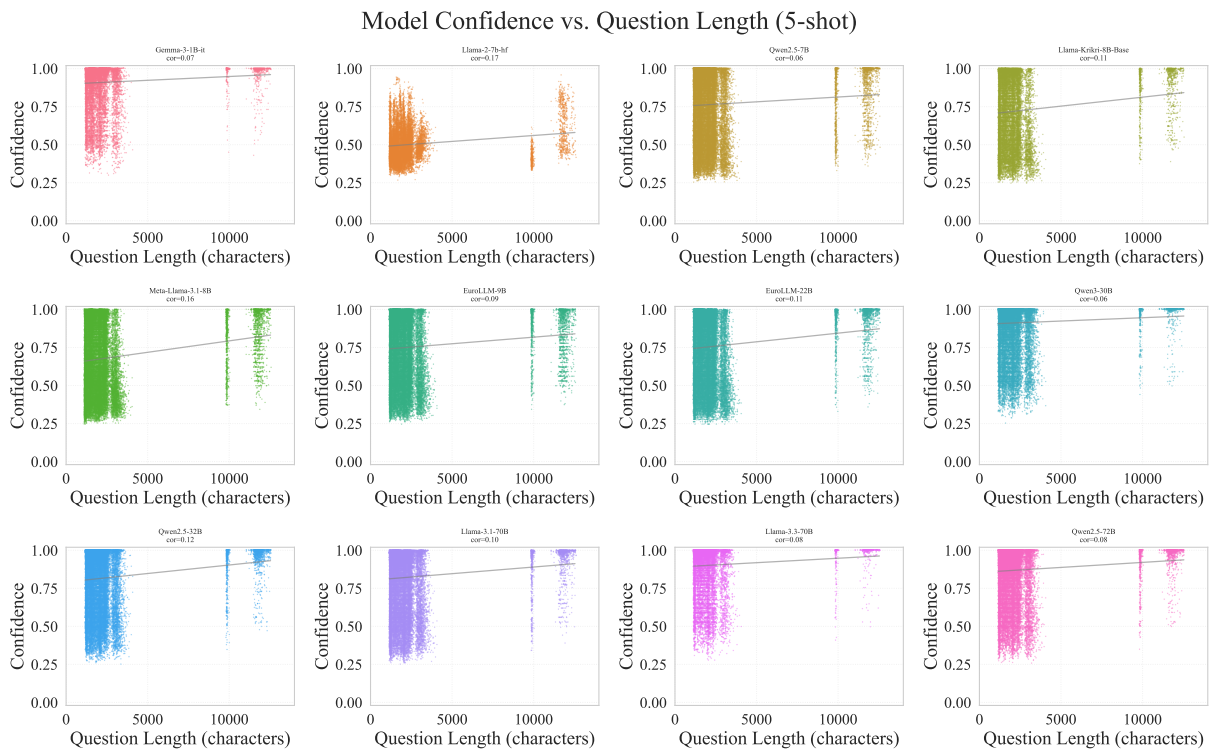


Figure 16: Correlation between Question Length and Confidence (Five-Shot).