Maistros: A Greek Large Language Model Adapted Through Knowledge Distillation From Large Reasoning Models

Nikolaos Giarelis[1,*], **Charalampos Mastrokostas**[1], **and Nikos Karacapilidis**[1]

[1]Industrial Management and Information Systems Lab, MEAD University of Patras, Rio Patras, Greece
[*]giarelis@ceid.upatras.gr

arXiv:2605.01870v1 [cs.CL] 3 May 2026

ABSTRACT

Large Language Models (LLMs) have substantially advanced the field of Natural Language Processing (NLP), achieving state-of-the-art performance across a wide range of tasks. These improvements have been attributed, in part, to their emerging reasoning capabilities, which are enabled by large-scale training and increased model capacity. However, existing LLMs can generate erroneous responses when addressing complex queries that fall outside their training distribution, due to limited internal knowledge or the need for multi-step reasoning. To address these limitations, recent work has introduced large reasoning models (LRMs), which incorporate explicit internal reasoning processes to improve response accuracy. Additionally, state-of-the-art LRMs often comprise hundreds of billions of parameters and require several seconds per inference, even on advanced multi-GPU systems. These characteristics limit their practicality for deployment in conventional computing environments. Meanwhile, NLP research on multilingual LLMs continues to prioritize high-resource languages. However, these models exhibit limited performance in under-resourced languages, primarily due to insufficient language- and culture-specific training data. In this paper, we focus on Modern Greek, for which only a limited number of question answering (QA) datasets have been proposed, most of which are intended for model evaluation. To address this research gap in Greek QA, we make the following contributions: (i) We introduce *CulturaQA*, a high-quality LRM-generated and human-curated dataset, for Greek LLM training and evaluation; (ii) a memory-efficient LLM evaluation framework adaptable to diverse languages and QA tasks; (iii) *Maistros 8B*, a state-of-the-art open-weights Greek LLM developed via knowledge distillation and fine-tuning on *CulturaQA*; and (iv) a comprehensive evaluation of nine LLMs across nine human-curated Greek QA datasets. We release our code, model, and data to support reproducibility.

Introduction

Large Language Models (LLMs) have significantly advanced the fields of Natural Language Processing (NLP), Artificial Intelligence (AI), and Deep Learning, achieving state-of-the-art performance across a wide range of natural language understanding and reasoning tasks[1,2]. LLMs, also referred to as Foundation Models, are trained on large-scale corpora using substantial computational resources (e.g., GPU clusters) and can subsequently be adapted to downstream tasks with comparatively limited resources[3]. Earlier LLMs such as *GPT-3*[4] and *Llama-2*[5] were primarily trained on English-centric corpora and exhibited limited multilingual capabilities. In contrast, recent models such as *GPT-5*[6] and *Gemini 3*[7] are trained on multilingual data and demonstrate improved cross-lingual and reasoning capabilities. A recent line of work further categorizes such models as Large Reasoning Models (LRMs), which generate extended intermediate reasoning traces to improve performance on complex tasks and reduce factual errors[8]. However, these capabilities come at an increased computational cost, including substantially larger model sizes, higher inference latency due to longer generated reasoning sequences, and a reliance on proprietary deployment infrastructures.

Despite these advances, multilingual LLMs continue to exhibit performance disparities between high-resource and under-resourced languages, largely due to imbalanced training data and limited coverage of linguistic and cultural variation[9]. This issue is particularly evident in tasks requiring cultural or domain-specific knowledge, where models may produce incomplete or inaccurate outputs[10]. In this context, Modern Greek remains an under-resourced language despite its linguistic complexity, including a distinct alphabet, rich morphology, and syntactic variability. Consequently, developing robust NLP systems for Greek remains challenging. Existing surveys highlight the scarcity of datasets, models, and systematic evaluations for Greek question answering (QA)[11–13]. While recent studies have introduced Greek QA datasets and benchmarking efforts[14–16], most datasets are primarily designed for evaluation rather than model training[16].

Empirical findings further indicate that open-weights multilingual LLMs supporting Greek generally underperform compared to models specifically adapted for Greek, such as *Krikri 8B*[14], while proprietary LLMs still achieve superior performance

on Greek QA benchmarks[15,16]. These observations highlight a persistent gap between open and proprietary models in under-resourced language settings. Motivated by these findings, this work investigates whether high-quality LRM-generated data can be leveraged to improve open-weights Greek language models. Specifically, using a human-in-the-loop process, we curate these data to mitigate linguistic and cultural inaccuracies. Finally, we employ supervised fine-tuning to distill knowledge into a compact model suitable for local deployment on commodity hardware. To address the identified gap in Greek QA, we make the following contributions:

- We introduce *CulturaQA*, a synthetic and human-curated Greek QA dataset designed to support model training and evaluation.
- We develop *Maistros 8B*, a Greek-adapted open-weights LLM obtained via supervised fine-tuning of *Ministral 3 8B*[17] on *CulturaQA*.
- We propose a memory-efficient and adaptable evaluation framework that supports multiple-choice and open-ended QA tasks using the *accuracy* and *BERTScore*[18] metrics, respectively.
- We conduct a comprehensive evaluation of nine LLMs across nine human-curated Greek QA datasets.
- We release the dataset, model, and code to support reproducibility (see the Data and Code Availability Statement sections).

This study investigates the following research questions (RQs):

- RQ1: Can human-curated LRM-generated data serve as a viable foundation for training and evaluating Greek QA systems?
- RQ2: Does fine-tuning an open-weights LLM on *CulturaQA* lead to measurable performance improvements on Greek QA benchmarks?
- RQ3: How does the fine-tuned model compare to existing open-weights Greek and multilingual LLMs?
- RQ4: To what extent can fine-tuned open-weights models approach the performance of proprietary LLMs on Greek QA tasks?

Related Work

This section presents related works on language-adapted and general-purpose LLMs, resources for Greek QA, and common fine-tuning techniques. For this study, we consider LLMs with at least 7 billion parameters, as these models consistently outperform smaller ones in complex reasoning and language tasks[1,2]. We also employ their instruction-tuned model variants, as these are directly optimized for in-context learning and zero-shot NLP tasks. Throughout this paper, model sizes are denoted using standard abbreviations (e.g., 8B corresponds to 8 billion parameters).

General-purpose and Language-adapted LLM

In this study, we consider several open-weights LLMs sorting them in two categories: (i) general-purpose and (ii) language-adapted. The first are typically made using large computing resources by for-profit AI organizations (e.g., Meta, Google, Mistral AI). The purpose of these LLMs is to excel in a variety of tasks, although prior works underscore their performance disparities across languages[9,10]. On the other hand, language-adapted LLMs use the former ones as a basis to subsequently introduce language-, cultural- or domain-specific knowledge through LLM fine-tuning.

Llama 3.1 8B[19] and *Gemma 3n E4B*[20] are two general-purpose LLMs that share a common training strategy. Both were pre-trained on large, multilingual corpora from the Internet, which also include mathematical and code reasoning data. Both of these models have learned many linguistic representations during pre-training and *Gemma 3n E4B* officially supports a wide variety of languages. An important thing to note is that despite its confusing naming scheme, *Gemma 3n's* actual number of parameters is 8B.

Qwen 3 8B[21] is a general-purpose LLM that was pre-trained in a large and diverse corpus consisting of 119 languages and dialects. This corpus comprises high-quality content from diverse domains, including scientific books, code and reasoning tasks, multilingual texts, and synthetic data. The synthetic data were extracted from *.pdf* documents using a *Qwen* Vision Language model. *Qwen-3 8B* is pre-trained on three different stages: (i) a general pre-training stage similar to previous models; (ii) a reasoning enhancing stage; and (iii) a training stage that improves its long context capabilities. Finally, Qwen-3 8B underwent post-training via knowledge distillation from a higher-capacity teacher model of the same family. This teacher model had been refined through reasoning-centric supervised fine-tuning and reinforcement learning.

Ministral 3 8B[17] is a general-purpose LLM, that was trained using iterative layer pruning and knowledge distillation from a large pre-trained model of the same family. This model is later post-trained using instruction tuning and supervised fine-tuning from synthetic reasoning data. The authors' evaluation results indicate that the model performs at or above the level of counterparts with similar parameter counts, while requiring less training resources.

EuroLLM 9B v2[22] utilizes a specialized training strategy, where the model is trained with specific data per language, while incorporating math and code data for reasoning. At the same time, its authors have developed custom multilingual tokenizers to officially support 35 languages, including Greek. This model has several architectural and pre-training improvements over the first model version, and has incorporated synthetic math data from *Gemma* and *Llama* models. However, even in its revised version, the model still has an imbalanced number of language samples, where 18 European languages are critically underrepresented (less than 2% of the training data).

Krikri 8B[14] is one of the first Greek-adapted LLMs, built on *Llama 3 8B*[19]. Krikri 8B was adapted from the base model, through extensive pre-training on Greek corpora, followed by Greek-specific instruction tuning, to improve its conversational capabilities. The authors' experimental results reported demonstrate its strong performance against other multilingual LLMs (including *Llama 3 8B*) throughout many Greek NLP tasks.

Plutus 8B[23] is a Greek financial LLM fine-tuned from *Krikri 8B* on several financial tasks. It supports a variety of relevant tasks including QA, Named Entity Recognition, Text Summarization, Classification and Numerical Extraction. The authors report that it achieves state-of-the-art performance over several open-weights and proprietary models in the Greek Financial benchmark, which was introduced in the same work.

Some notable LLMs following similar training practices for other languages include: (i) *LlaMandement-7B*[24], a French adaptation of *Llama 2 7B* that facilitates the Summarization of French Legislative Proposals; (ii) *Llama-SEA-LION-8B-IT* and *Gemma-SEA-LION-9B-IT*[25] adapted from *Llama 3 8B* and *Gemma 2 9B* respectively, through additional pre-training, instruction tuning and supervised fine-tuning for South-East Asia languages; (iii) *LlaMAntino-3-ANITA*[26], an Italian adaptation of *Llama 3 8B* using fine-tuning and reinforcement learning; (iv) *GaMS3-12B-Instruct*[27], a Slovenian LLM adapted from *Gemma 3 12B* using language-specific pre-training and fine-tuning. A common motif across these publications is that through additional linguistic or cultural knowledge distillation, the models perform much better than the original models used as a basis for several language-specific tasks.

Greek QA Datasets

To ensure evaluation integrity, we focused exclusively on high-quality human-curated QA datasets for Greek. This approach avoids machine translation errors, which can negatively impact performance in NLP tasks[28]. Using these criteria, we identified eight QA datasets suitable for the purposes of our study.

The *Greek Medical MCQA* dataset[29] comprises 2,034 medical QA pairs, where 1602 are reserved for training and the rest for validation. These QA pairs contain a question, five possible answer options with a single correct one. This dataset was extracted from the medical exams of the Hellenic National Academic Recognition and Information Center.

The *Greek Truthful QA*[29] contains 817 questions designed to measure whether LLMs can answer correctly when faced with misconceptions or false human beliefs. In contrast to other datasets, this one features a variable number of candidate answers for each QA pair. For our evaluation, we select its multiple-choice version, in the hardest difficulty setting (*mc1_targets*), where only a single answer is correct.

INCLUDE[30] is a large-scale multilingual dataset spanning 44 languages. This dataset is collected from local academic and professional exams, and facilitates per-language LLM evaluation of regional and domain-specific knowledge. In our experiments, we utilize the Greek subset (552 QA pairs), where each pair consists of a question, four possible answers and a single correct one.

Greek ASEP MCQA[31] comprises 1,200 multiple-choice QA pairs, sourced from the Greek Supreme Council for Civil Personnel Selection (ASEP) exams. This dataset spans various topics, including Greek history, law, politics, public administration and e-governance. This dataset has the same QA structure as *INCLUDE*.

GPCR[32], also known as the Greek Physical Commonsense Reasoning dataset, includes manually-annotated 208 samples, similar to *PIQA*[33]. These samples contain a question, two candidate answers exhibiting near-identical lexical composition, with a single one marked as correct. Approximately 40% of the samples are regionally or culturally specific, and cannot be easily rendered into English.

Plutus QA[23] consists of 540 Greek financial QA pairs. This dataset was derived from real-world financial documents (i.e., annual reports, article headlines and exam questions) and was annotated by several financial and linguistics experts. *Plutus QA* has a train, validation and test splits that comprise 267, 48 and 225 pairs respectively. Each QA pair has a question, additional context and a set of multiple-choice answers, with a single correct one.

Demos QA[16] is a Greek QA dataset, comprising 600 questions and community-reviewed answers from Greek social media. Each QA pair has a question and four candidate answers, as well as a selected best answer *DemosQA's* answers are ranked

based on community voting, with the highest-upvoted response designated as the reference answer.

Greek MMLU[15] is a QA dataset that assesses massive multitask Greek understanding. It encompasses 45 subjects with 21,805 original QA pairs curated from real-world educational and professional assessments. The present study focuses on its Greek-specific subset comprising 3,660 question-answer (QA) pairs. These pairs necessitate a higher level of Hellenic cultural and linguistic knowledge.

Collectively, these eight human-curated datasets offer a valuable foundation to evaluate LLMs supporting Greek in varying domains. However, most of them lack training samples across a multitude of topics related to Greece and its culture; therefore, this motivates us to create our dataset *CulturaQA*, which is elaborated in a following section.

Supervised Fine-Tuning

LLMs comprise many dense weight matrices, which are used to infer the next tokens in the sequence, given an input. In order to instill new knowledge in the model, it is required to update these weights. The typical learning paradigm is full fine-tuning. Let us consider a single dense matrix $W_0 \in \mathbb{R}^{d \times k}$ from the LLM. To update this matrix, we would have to produce a new one ΔW:

$$W = W_0 + \Delta W \tag{1}$$

where ΔW would be the learned weight updates, which are produced by recalculating the gradients for the entire parameter space *(d x k)*. Several studies have shown this to be computationally expensive, given the large number of parameters of LLMs, where multiple GPUs with extremely large memory are required[34–36]. Thus, to mitigate this issue, Hu et al.[34] introduce a different fine-tuning paradigm; Low-Rank Adaptation (LoRA). Instead of ΔW, LoRA calculates two smaller low-rank matrices $A \in \mathbb{R}^{d \times r}$ and $B \in \mathbb{R}^{r \times k}$ that alter equation (1) as such:

$$W = W_0 + sAB \tag{2}$$

where $r \ll \min\{d, k\}$ and s is the scaling factor which controls the impact of weight updates:

$$s = \frac{\alpha}{r} \tag{3}$$

The scaling factor is typically set to 2.0, to improve training stability and robustness[36,37]. This is achieved by setting α to be twice as large as the selected rank parameter. The matrices A and B have their elements initially set to a random Gaussian distribution and 0 respectively. LoRA optimizes efficiency by learning low-rank matrices, which enable fine-tuning while only updating a fraction of the original LLM's parameters. Despite, this smaller update, LoRA has been shown to be nearly equivalent with full fine-tuning, while mitigating the effects of catastrophic forgetting of previous model knowledge[36,37]. Considering that an LLM has multiple layers (L) and trainable modules (M), we can define equation 2 as such, where each layer $l \in L$ and $m \in M$:

$$W_{(l,m)} = W_{0(l,m)} + \frac{\alpha}{r}(B_{(l,m)}A_{(l,m)}) \tag{4}$$

Both in the case of full and LoRA fine-tuning an LLM learns by minimizing the following loss function[38]:

$$L = -\frac{1}{T}\sum_{t=1}^{T} \log P(x_t \mid x_{<t}) \tag{5}$$

where T is the total number of tokens, while x_t and $x_{<t}$ are the current and preceding tokens of the input sequence X. Essentially, the model learns to minimize the loss (or maximize the probability) of correctly predicting the next token in the sequence from previous ones, based on the training examples. However, Touvron et al.[5] suggested an improved version of this loss function:

$$L = -\frac{1}{\sum_{t=1}^{T} m_t}\sum_{t=1}^{T} m_t \log P(x_t \mid x_{<t}) \tag{6}$$

In this version, a token-level binary mask $m_t \in \{0,1\}$ is introduced. This mask is assigned to 0 for every token that is part of the user or system prompt, (and 1 otherwise), thus these tokens are omitted from the loss calculation. By computing gradients only on the assistant's generated response tokens, the model is prevented from learning the user and system prompt distributions, thus significantly enhancing its chat and instruction-following capabilities. In this work, we are fine-tuning our proposed model based on equations 4 and 6. For more information about our training see the Model Training and Validation Section.

The CulturaQA Dataset

In this section, we introduce *CulturaQA*, a synthetic and human curated dataset that captures knowledge from Greek culture. *CulturaQA* encompasses a plethora of topics across several domains including Greek art, history, mythology, politics, economy, tourism, food, health, science, sports, education and law, thus providing a valuable resource for training, validating and evaluating models on the nuances of Greek Culture, as well as advancing language understanding research within culturally grounded contexts.

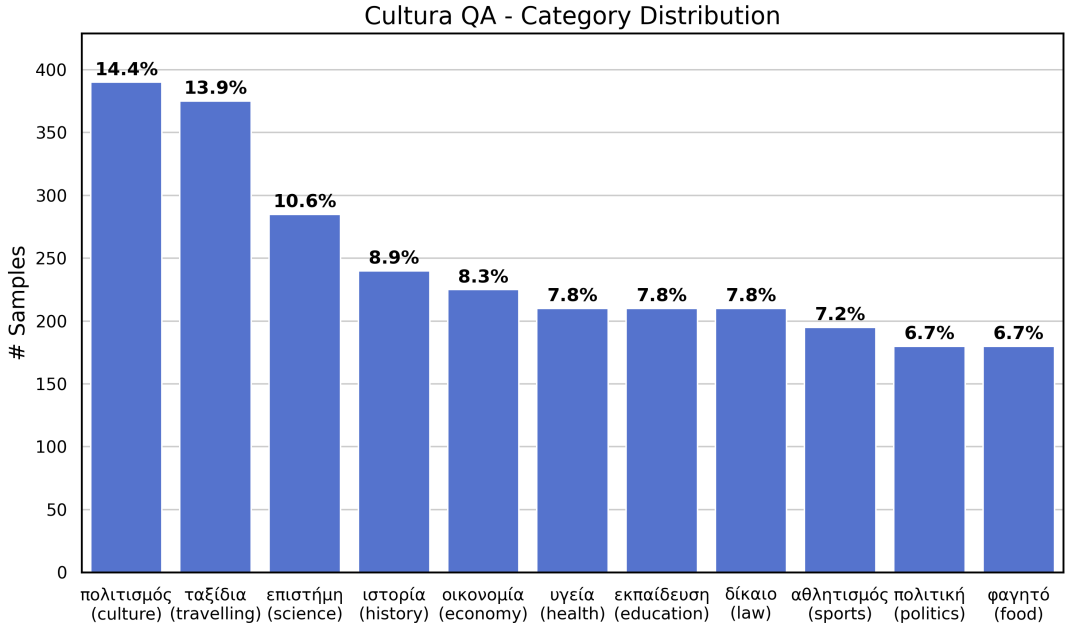

Figure 1. CulturaQA's sample distribution per category. The y-axis measures the number of samples; the percentage of samples is written in each bar.

To create *CulturaQA*, we manually curated a list of 180 Greek keyphrases (for the exact phrases, see the code repository) that were grouped into eleven categories. These categories are "πολιτισμός" ("civilization"), "ταξίδια" ("travelling"), "πολιτική" ("politics"), "οικονομία" ("economy"), "επιστήμη" ("science"), "υγεία" ("health"), "αθλητισμός" ("sports"), "εκπαίδευση" ("education"), "ιστορία" ("history"), "φαγητό" ("food") and "δίκαιο" ("law"). For each keyphrase, we generated 15 questions using *GPT-5*, which were answered one-by-one by the same model. Each QA pair is accompanied by a unique ID and the category that it belongs. For the exact dataset creation prompts, please refer to the Appendix. This procedure resulted in a dataset comprising 2,700 Greek question-answer pairs. For the sample distribution per category for CulturaQA (see Figure 1).

We then manually pre-processed the dataset to ensure its linguistic quality, as well as remove any erroneous answers (hallucinations) and mitigate potential cultural and historical biases. During this step, we correct several linguistic errors, for example Greek words: (i) written in the wrong grammar inflection; (ii) containing english characters; (iii) that were misspelled or (iv) erroneously written in polytonic. We also corrected syntax and translation errors, where a few non-Greek words were generated, while an equivalent Greek term existed. We added the missing full forms for abbreviations that were used to refer to Greek organizations. We corrected several factual hallucinations and we removed suggestions, where the model was asking follow-up questions (e.g., "Should I produce a table, report, file, etc. to compile the above information?"). Finally, we removed the commonly occurring phrase, given a lack of reference for time or country, when in fact, there was a time reference (e.g., the last five years, or the current year) or a local reference (e.g., Greece) in the question.

During post-processing, we utilized code to remove unnecessary whitespace and replace the English with the Greek question mark from each QA pair. We also divided the dataset into training, validation and testing splits using stratifying sampling based on the categories, to ensure that all splits have a similar distribution across categories. We also analyzed *CulturaQA* and compared it with other datasets (see Table 1).

Dataset	#Docs	Train, Val, Test splits	Type	P5	P25	P50	Mean	P75	P95	P99
CulturaQA	2700	2000, 200, 500	Question	11	14	17	17.81	20	26	32.1
			Answer	26	98.75	178	201.7	295	451.35	557.31
Greek MMLU (Greek Specific)	3660	0, 0, 3660	Question	6	8	12	56.08	18	308	375.41
			Answer	1	1	2	3.03	4	10	15.41
Greek Medical MCQA	2032	1602, 432, 0	Question	3	6	9	9.96	11.25	21	27
			Answer	1	2	3.5	4.67	6.0	12	17.7
Greek ASEP MCQA	1200	0, 0, 1200	Question	4	6	10	11.4	14	26	37.01
			Answer	2	4	7	8.38	11.25	21	27
Greek Truthful QA	817	0, 0, 817	Question	5	7	9	11.03	13	22	40.68
			Answer	2.8	7	10	10.02	13	18	21.84
DemosQA	600	0, 0, 600	Question	26	53	84.5	103.04	132.25	243	347.7
			Answer	11	31	54.5	80.22	105	222	362.04
Include (Greek)	552	0, 0, 552	Question	5	9	13	22.8	27.25	75.9	129.96
			Answer	1	3	5	7.23	9	20	35
GPCR	208	0, 0, 208	Question	5	7	10	11.76	13	25	45.58
			Answer	2	4	7	10.64	12	32.30	47.65
Plutus QA	540	267, 48, 225	Question	3.2	7	11	13.45	16	34	54
			Answer	1	3	7	11.01	12	36.8	69.84

Table 1. Number of documents and splits per dataset, with a statistical overview of their question and answer word counts.

LLM Post-training, Validation and Evaluation

This section presents our technical setup, the post-training and validation of the proposed model *Maistros 8B*, as well as our empirical evaluation. The general approach is illustrated in Figure 2.

Technical Setup

In terms of hardware specifications, for LLM fine-tuning, we utilized a cloud server with 16 logical cores CPU, 128 GB of RAM and an *Nvidia L40S* 48 GB VRAM GPU, whereas for LLM evaluation we utilized a local server with an *Intel Core i9-12900K* 20 logical cores CPU, 64 GB of RAM and an *Nvidia RTX A4000* 16 GB VRAM GPU.

Regarding software specifications, we utilized several HuggingFace[39] libraries (i.e., *Transformers*, *TRL*, *PEFT*) for LoRA Fine-tuning, as well as *Bitsandbytes* to load models with 4-bit quantization as to reduce memory requirements for our training and the evaluation of open-weights LLMs. Morever, we utilized the GenAI APIs from OpenAI and Google, to access *GPT-5*, for the synthetic data creation step, as well as *GPT-5 mini* and *Gemini 3 Flash* for model evaluation. Finally, we use *scikit-learn's* accuracy metric for multiple-choice QA and *BERTScore* for open-ended QA. For reproducibility purposes, we set a fixed random seed (42) and for model evaluation we also utilize greedy decoding, which is equivalent to model temperature 0.0[40].

To preserve a controlled experimental setup, we selected open-weights LLMs in the same computational class (ranging from 7B to 10B parameters), while including the smallest available versions of state-of-the-art proprietary LLMs. We also utilize the most recent model versions (e.g., *Gemma 3*). We exclude models that were shown to have bad performance in Greek QA tasks from previous empirical studies[15,16]. Overall, the models evaluated in this study are compiled in Table 2.

Model Training and Validation

To develop *Maistros 8B*, we post-trained *Ministral-3-8B-Instruct* on *CulturaQA* using LoRA fine-tuning. For our training setup, we utilize a mixed precision environment, where the model is loaded initially in full precision (*BFloat16*) and then it is *4-bit* quantized to a normalized float (*NF4*) with double quantization to reduce the memory requirements even further[41]. Following best practices from Touvron et al.[5], we also calculate the cross-entropy training loss for the answer, while question and system prompts are masked and thus excluded from the loss calculation. The aim of this practice is to update the model based on the answer and not learned fixed patterns that occur at the question and instructions of the system prompt. This has the added benefit that the model does not learn to predict the question. Thus after fine-tuning, it directly answers the question instead of

LLM	Full Model Name	Open weights
GPT-5 mini	openai/gpt-5-mini-2025-08-07	-
Gemini 3 Flash	google/gemini-3-flash-preview	-
Ministral 3 8B	mistralai/Ministral-3-8B-Instruct-2512-BF16	✓
Krikri 8B	ilsp/Llama-Krikri-8B-Instruct	✓
Plutus 8B	TheFinAI/plutus-8B-instruct	✓
EuroLLM 9B v2	utter-project/EuroLLM-9B-Instruct-2512	✓
Gemma 3n E4B	google/gemma-3n-E4B-it	✓
Qwen 3 8B	Qwen/Qwen3-8B	✓
Maistros 8B	IMISLab/Maistros 8B-Instruct	✓

Table 2. Considered LLMs for the experiments.

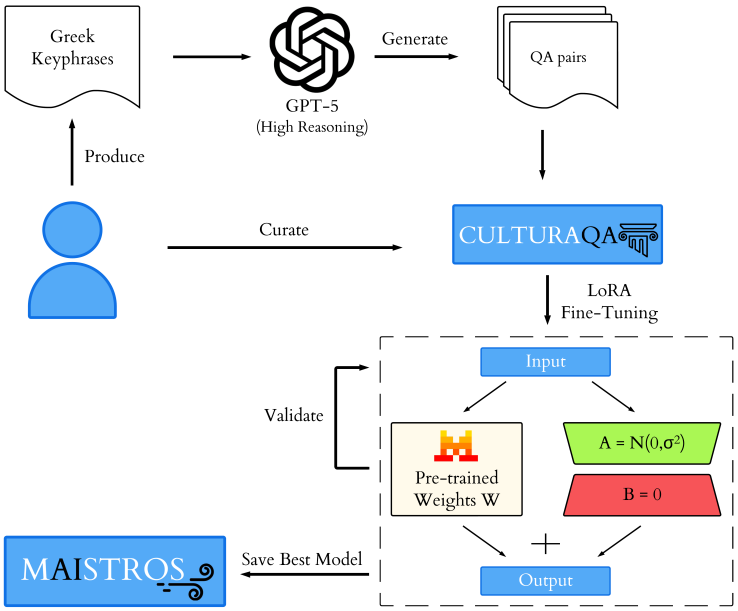

Figure 2. The overall approach for training *Maistros 8B*.

generating it again. We experimented with various training hyperparameters (see Table 3) and we selected the best ones based on the training and validation losses. The ones that are elaborated below are the best ones for the final training run.

We train the model for 4 epochs using a learning rate of 2e-5 and a cosine learning rate scheduler, with a batch size of 2 and gradient accumulation steps set to 8. This gives us a larger effective batch size of 16, which improves training stability[42]. To further improve stability, we utilized global gradient clipping with a maximum L2 norm threshold of 1.0[43]). Overall, our goal is to keep the batch size small, due to the low number of training samples. To further reduce memory usage, we employed the *8-bit AdamW* optimizer[44] with its default hyperparameters ($\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\varepsilon = 1e{-}8$). This optimizer reduces the memory usage by 75% by utilizing block-wise dynamic quantization (from *32-bit* to *8-bit*).

We add LoRA adapters across all attention and feed forward layers of the model. For LoRA, we set the rank to 16 and the LoRA alpha (α) to 32; this gives us a scaling factor of 2.0, which has been shown to improve training stability and robustness[36,37]. To avoid overfitting, we set both the dropout for every LoRA layer and weight decay to 0.1. For our training, we set the max sequence length to 3269. This is calculated by counting the number of tokens of the longest training sample; we also pad sequences to the length of the longest one in each training batch to reduce computational overhead. The total number of training steps is 500 calculated using the following equation and the warm-up steps are set to 62 (8% of total):

$$\left\lfloor \frac{2000 \text{ training samples}}{16 \text{ (effective batch size)}} \times 4 \text{ epochs} \right\rfloor = 500 \text{ steps} \tag{7}$$

The amount of total steps is equally divided among epochs, thus each epoch comprises 125 steps. To select the best model checkpoint for the final training run, we calculated the training and validation losses and the best LoRA adapters for the model

Training Hyperparameter	Values
Train Batch size	{1, **2**}
Gradient Accumulation Steps	{4, **8**}
Learning Rate	{1e-4, 9e-5, 5e-5, 3e-5, **2e-5**}
Weight decay	{0, 0.01, **0.1**}
LoRA Rank	{8, 16, **32**}
LoRA Dropout	{0, 0.05, **0.1**}
Warmup steps	{25, 50, **62**}
Learning Rate Scheduler	{linear, **cosine**}
Target modules	Attention mechanism: Query (**q_proj**), Key (**k_proj**), Value(**v_proj**) The Output projection of the attention layer: (**o_proj**) Feed-Forward Network blocks: (**gate_proj**, **up_proj**, **down_proj**)

Table 3. Training hyperparameters tested for fine-tuning; the best parameters are marked in bold.

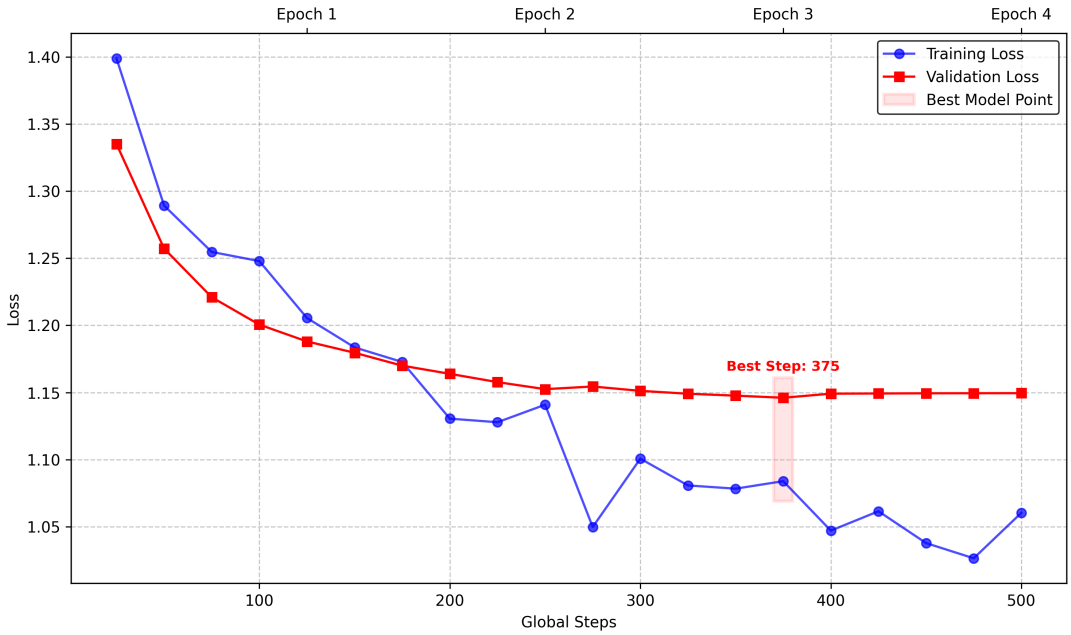

Figure 3. Training and Evaluation loss over steps

are trained on step 375 (epoch 3) where the lowest validation and training losses are achieved. This is visualized in Figure 3. Finally, we merge the LoRA adapter weights from epoch 3 to the weights of the base model to produce *Maistros 8B*.

Experiments

This section elaborates on our evaluation framework and the empirical results. To obtain these results, we carried out a series of experiments that measure the performance of the considered LLMs for Greek QA tasks. The point of these tasks is to reveal model effectiveness in understanding and generating accurate responses to Greek questions across diverse topics. We evaluated the considered LLMs (see Table 2) on eight multiple-choice and one open-ended QA task. For multiple-choice tasks, we extract the correct answer from the LLM response using rule-based parsing and regular expressions, since instruction-tuned models often include explanatory text alongside their selected answer. If a valid answer could not be extracted, it was labeled as "*No match*". For the open-ended QA task (Cultura QA), the models generate a full answer, which is then evaluated using the *BERTScore F1* metric (%) against the reference answer. For the exact instruction prompts used in the evaluation, see Appendix A. One final note is that our framework standardizes different dataset formats into the same format as to enable the comparative performance evaluation.

As shown in Table 4, for most QA datasets, the best performance is achieved by the proprietary models, with *Gemini 3 Flash* outperforming *GPT-5-mini* more than 10% percent across most of them. A notable exception is *CulturaQA*, which

Score %	DemosQA	GPCR	INCLUDE	Greek ASEP MCQA	Greek Medical MCQA	Plutus QA	Greek Truthful QA	Greek MMLU (Greek-specific)	CulturaQA
Proprietary Models									
Gemini 3 flash	**55.67**	**88.46**	**88.77**	**94.75**	**92.82**	**89.78**	**88.62**	**95.03**	73.97
GPT-5 mini	53.00	77.40	74.46	78.92	78.01	76.89	75.89	87.49	**75.09**
Open-Weights Models									
Maistros 8B (Ours)	50.83	**64.42**	**58.70**	**67.25**	**49.54**	**73.33**	53.37	**78.17**	**71.99**
Ministral 3 8B	**51.67**	59.62	54.17	63.25	47.92	65.33	52.51	76.23	71.03
Krikri 8B	49.50	54.81	50.54	63.08	45.37	64.44	**54.83**	71.04	71.31
Plutus 8B	45.67	50.00	48.37	62.92	39.35	57.33	34.52	70.38	67.44
EuroLLM v2 9B	41.50	53.85	39.13	46.08	31.71	42.67	36.72	58.17	70.33
Gemma 3n E4B	47.17	60.10	50.00	57.75	43.75	53.78	46.76	71.39	69.10
Qwen 3 8B	48.83	31.73	49.28	54.58	36.64	63.56	42.72	67.57	68.73

Table 4. Empirical Greek QA results. We report macro accuracy for multiple-choice QA datasets, and macro BERTScore F1 for CulturaQA. The best results for proprietary and open-weights models are highlighted in bold.

was generated using GPT-5 (the parent model of *GPT-5 mini*); thus, in this dataset the former outperforms *Gemini 3 Flash* by a small margin. Regarding open-weights models, *Maistros 8B* achieves the state-of-the-art scores in most datasets with meaningful improvements over the base model. Two notable exceptions are *Demos QA* and *Greek Truthful*, where *Maistros 8B* is outperformed by *Ministral 3 8B* and *Krikri 8B* respectively, albeit by a small margin (less than 2%). One of the most interesting results, was model performance on *Plutus QA*, a domain-specific dataset (Greek Economy), where *Maistros 8B* outperformed all open-weights models, while simultaneously attaining a similar accuracy score with *GPT-5 mini* (-3.56% difference). Similarly, for the test set of *CulturaQA*, *Maistros 8B* outperformed all open-weights models, while simultaneously attaining a similar score with the proprietary models.

To assess the statistical significance of the observed score improvements, we compared *Maistros 8B* against the base model across the above datasets. For the multiple-choice QA datasets, we encoded outputs as paired binary data, where if a model selected the reference answer, we labelled this as 1 and 0, otherwise. Statistical significance was measured using the exact binomial *McNemar's test*. To quantify uncertainty in the estimated effect sizes, we employed bootstrap resampling to derive 95% *confidence intervals (CI)* for the accuracy differences. Specifically, 10,000 bootstrap samples were drawn with replacement from the paired binary data. For the Cultura QA dataset, we utilized the Wilcoxon signed-rank test, given the continuous nature of the BERTScore F1 scores, and similarly 10,000 bootstrap samples were drawn to derive the corresponding 95% *CI*.

In both cases, we consider the score improvement as statistically significant only if the *p-value* <= 0.05 and the resulting 95% *CI* was strictly positive. The reason we run the above tests is that they are recommended to measure the statistical significance of metrics such as accuracy and F1 scores for small NLP datasets[45]. The results of these tests are collected in the following tables. Specifically, in Table 5, we observe the exact accuracy improvements of *Maistros 8B* over the base model, where the former achieves a statistically significant improvement in 5 out of 9 datasets. In Tables 6 and 7, we report the *p-value* and 95% *CI* for *Maistros 8B* compared against other open-weights models in every dataset. In these tables, we observe that the proposed model achieves statistical significant improvements across most datasets. Some notable exceptions include the *Demos QA* dataset, where *Maistros 8B* does not achieve a statistical significant improvement against *Krikri 8B*, *Gemma 3n E4B* and *Qwen 3 8B*. This is also the case for the *Greek Medical MCQA* and *Greek Truthful QA* datasets, where the proposed model does not improve statistically against *Krikri 8B*.

Dataset	Score Improvement	P-Value	95% CI
DemosQA	-0.83%	0.635	[-3.67%, 2.00%]
Greek PCR	4.81%	0.064	[0.0%, 9.62%]
INCLUDE	**4.53%**	**0.014**	**[1.09%, 7.97%]**
Greek ASEP MCQA	**4.00%**	**0.001**	**[1.75%, 6.33%]**
Greek Medical MCQA	1.62%	0.470	[-2.08%, 5.32%]
Plutus QA	**8.00%**	**0.001**	**[3.56%, 12.44%]**
Greek Truthful QA	0.86%	0.562	[-1.59%, 3.30%]
Greek MMLU (Greek-Specific)	**1.94%**	**0.000**	**[1.01%, 2.87%]**
CulturaQA	**0.97%**	**0.000**	**[0.49%, 1.46%]**

Table 5. Score improvements of Maistros 8B over the base model and results from the statistical significance tests.

p-value	DemosQA	GPCR	INCLUDE	Greek ASEP MCQA	Greek Medical MCQA	Plutus QA	Greek Truthful QA	Greek MMLU (Greek-specific)	CulturaQA
Krikri 8B	0.575	**0.010**	**0.001**	**0.006**	0.153	**0.017**	0.495	**0.000**	**0.002**
Plutus 8B	**0.030**	**0.000**	**0.000**	**0.004**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
EuroLLM v2 9B	**0.001**	**0.033**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
Gemma 3n E4B	0.117	0.362	**0.000**	**0.000**	**0.038**	**0.000**	**0.000**	**0.000**	**0.000**
Qwen 3 8B	0.399	**0.000**	**0.000**	**0.000**	**0.000**	**0.006**	**0.000**	**0.000**	**0.000**

Table 6. Statistical significancy tests for Maistros 8B against the other open-weights LLMs. Results with p-value < 0.05 are highlighted in bold.

95% CI	DemosQA	GPCR	INCLUDE	Greek ASEP MCQA	Greek Medical MCQA	Plutus QA	Greek Truthful QA	Greek MMLU (Greek-specific)	CulturaQA
Krikri 8B	[-2.67%, 5.33%]	**[2.88%, 16.83%]**	**[3.44%, 13.04%]**	**[1.25%, 7.08%]**	[-1.39%, 9.49%]	**[2.22%, 16.00%]**	[-5.26%, 2.33%]	**[5.55%, 8.66%]**	**[0.22%, 1.15%]**
Plutus 8B	**[0.67%, 9.50%]**	**[7.69%, 21.15%]**	**[5.62%, 15.04%]**	**[1.42%, 7.33%]**	**[4.63%, 15.74%]**	**[9.33%, 22.67%]**	**[15.18%, 22.64%]**	**[6.20%, 9.37%]**	**[4.00%, 5.10%]**
EuroLLM v2 9B	**[4.17%, 14.50%]**	**[1.44%, 19.71%]**	**[14.31%, 24.82%]**	**[17.83%, 24.50%]**	**[12.04%, 23.61%]**	**[22.67%, 38.67%]**	**[12.73%, 20.44%]**	**[18.14%, 21.86%]**	**[1.21%, 2.11%]**
Gemma 3n E4B	[-0.67%, 8.00%]	[-3.85%, 12.50%]	**[3.99%, 13.41%]**	**[6.50%, 12.50%]**	**[0.69%, 11.11%]**	**[11.56%, 27.11%]**	**[3.06%, 10.16%]**	**[5.19%, 8.39%]**	**[2.44%, 3.35%]**
Qwen 3 8B	[-2.17%, 6.17%]	**[24.52%, 40.87%]**	**[4.71%, 14.13%]**	**[9.50%, 15.75%]**	**[7.87%, 18.52%]**	**[3.11%, 16.44%]**	**[7.22%, 14.20%]**	**[9.04%, 12.16%]**	**[2.82%, 3.70%]**

Table 7. Statistical significancy tests for Maistros 8B against the other open-weights LLMs. Results where the entire confidence interval is above zero are highlighted in bold.

Discussion

This study addresses the limited availability of resources for Greek QA and the performance gap of Greek-capable LLMs. We introduced *CulturaQA*, a synthetic and human-curated dataset designed to support Greek LLM training and evaluation, and developed *Maistros-8B*, a Greek-adapted open-weight LLM via knowledge distillation and fine-tuning. In addition, we proposed a memory-efficient evaluation framework and conducted a comprehensive assessment across multiple human-curated Greek QA datasets. The results provide evidence that curated synthetic data, combined with targeted fine-tuning, can improve the performance of open-weight models in under-resourced language settings.
Our empirical findings provide several insights aligned with the research questions. Specifically:

- The results indicate that high-quality QA datasets for training and evaluation can be constructed using synthetic data generated by LRMs and subsequently refined through human curation. The performance improvements observed for *Maistros-8B* relative to its base model further support the utility of *CulturaQA* (**RQ1** and **RQ2**).
- Among the evaluated open-weight models, *Maistros-8B* achieves consistently strong performance across most datasets, suggesting the effectiveness of the proposed knowledge distillation and fine-tuning approach (**RQ2** and **RQ3**).
- However, open-weight models remain below the performance of proprietary ones for Greek QA. In particular, models such as *Gemini 3 Flash* and *GPT-5 Mini* consistently outperform all evaluated open-weight models across the considered benchmarks (**RQ4**).
- The largest performance gain for *Maistros-8B* is observed on *PlutusQA*, a domain-specific dataset focused on the Greek economy, where it achieves an improvement of approximately 8% over its base model and approaches the performance of *GPT-5 Mini* (3.56% difference). This suggests that the proposed approach is particularly effective in domain-specific settings.

This study has several limitations that suggest directions for future work. First, the analysis is restricted to Modern Greek, and the findings are not directly evaluated in other linguistic settings. Second, despite recent progress, there remains a limited number of Greek QA datasets, particularly those containing long-form answers and diverse content across both general and domain-specific topics. Third, our evaluation focuses on relatively small multilingual LLMs, as large-scale Greek-adapted models (i.e., exceeding 9B parameters) are currently scarce, which constrains broader comparisons. Finally, the proposed model is evaluated primarily on knowledge-intensive QA tasks. Other important capabilities, such as safety alignment and instruction following, are not explicitly assessed and remain an area for future investigation.

Overall, this work provides a foundation for future research in Greek QA and LLM adaptation for under-resourced languages. By releasing the dataset, model, and code, we aim to support the development of linguistically accurate and culturally grounded language models for Greek and related settings[33]. Future work may investigate the construction of more diverse post-training datasets to better capture the morphological and syntactic complexity of Modern Greek. This includes extending coverage to additional language variants, such as ancient and polytonic Greek[46], as well as regional dialects, which may further improve the modeling of social, historical, and cultural context[47]. Finally, future work could extend the evaluation to domain-specific and long-context tasks, such as Greek legal QA[48,49], as well as other NLP tasks including text summarization[50–52] and text classification[53–56], aiming to provide a more comprehensive assessment of model capabilities and performance.

Ethical Considerations

The source code, dataset, and generative model (research items) introduced in this work are intended solely for research and educational purposes. While the authors have made reasonable efforts to ensure the accuracy and reliability of the research items; these are provided without a warranty of any kind, regarding their suitability for any particular purpose. Moreover, the dataset introduced in this study is synthetically generated and was manually processed by the authors to remove any inappropriate or uninformative material, while trying to uphold ethical and high-quality data curation practices.

References

1. Minaee, S. *et al.* Large language models: A survey (2025). 2402.06196.
2. Naveed, H. *et al.* A comprehensive overview of large language models. *ACM Trans. Intell. Syst. Technol.* **16**, DOI: 10.1145/3744746 (2025).
3. Bommasani, R. *et al.* On the opportunities and risks of foundation models (2022). 2108.07258.
4. Brown, T. B. *et al.* Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20 (Curran Associates Inc., Red Hook, NY, USA, 2020).
5. Touvron, H. *et al.* Llama 2: Open foundation and fine-tuned chat models (2023). 2307.09288.
6. Singh, A. *et al.* Openai gpt-5 system card (2025). 2601.03267.
7. Anil, R. *et al.* Gemini: A family of highly capable multimodal models (2025). 2312.11805.
8. Xu, F. *et al.* Toward large reasoning models: A survey of reinforced reasoning with large language models. *Patterns* **6**, 101370, DOI: https://doi.org/10.1016/j.patter.2025.101370 (2025).
9. Shani, C., Reif, Y., Roll, N., Jurafsky, D. & Shutova, E. The roots of performance disparity in multilingual language models: Intrinsic modeling difficulty or design choices? (2026). 2601.07220.
10. Qin, L. *et al.* A survey of multilingual large language models. *Patterns* **6**, 101118, DOI: https://doi.org/10.1016/j.patter.2024.101118 (2025).
11. Bakagianni, J., Pouli, K., Gavriilidou, M. & Pavlopoulos, J. A systematic survey of natural language processing for the greek language. *Patterns* **6**, 101313, DOI: https://doi.org/10.1016/j.patter.2025.101313 (2025).
12. Papantoniou, K. & Tzitzikas, Y. Nlp for the greek language: A longer survey (2024). 2408.10962.
13. Giarelis, N., Mastrokostas, C., Siachos, I. & Karacapilidis, N. A review of greek nlp technologies for chatbot development. In *Proceedings of the 27th Pan-Hellenic Conference on Progress in Computing and Informatics*, PCI '23, 15–20, DOI: 10.1145/3635059.3635062 (Association for Computing Machinery, New York, NY, USA, 2024).
14. Roussis, D. *et al.* Krikri: Advancing open large language models for Greek. In Christodoulopoulos, C., Chakraborty, T., Rose, C. & Peng, V. (eds.) *Findings of the Association for Computational Linguistics: EMNLP 2025*, 5012–5033, DOI: 10.18653/v1/2025.findings-emnlp.268 (Association for Computational Linguistics, Suzhou, China, 2025).
15. Zhang, Y. *et al.* Greekmmlu: A native-sourced multitask benchmark for evaluating language models in greek (2026). 2602.05150.
16. Mastrokostas, C., Giarelis, N. & Karacapilidis, N. Evaluating monolingual and multilingual large language models for greek question answering: The demosqa benchmark (2026). 2602.16811.
17. Liu, A. H. *et al.* Ministral 3 (2026). 2601.08584.
18. Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q. & Artzi, Y. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations* (2020).
19. Grattafiori, A. *et al.* The llama 3 herd of models (2024). 2407.21783.
20. Kamath, A. *et al.* Gemma 3 technical report (2025). 2503.19786.
21. Yang, A. *et al.* Qwen3 technical report (2025). 2505.09388.
22. Ramos, M. M. *et al.* Eurollm-22b: Technical report (2026). 2602.05879.
23. Peng, X. *et al.* Plutus: Benchmarking large language models in low-resource Greek finance. In Christodoulopoulos, C., Chakraborty, T., Rose, C. & Peng, V. (eds.) *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, 30176–30202, DOI: 10.18653/v1/2025.emnlp-main.1535 (Association for Computational Linguistics, Suzhou, China, 2025).

24. Gesnouin, J. *et al.* Llamandement: Large language models for summarization of french legislative proposals (2024). 2401.16182.

25. Ong, D. & Limkonchotiwat, P. SEA-LION (Southeast Asian languages in one network): A family of Southeast Asian language models. In Tan, L., Milajevs, D., Chauhan, G., Gwinnup, J. & Rippeth, E. (eds.) *Proceedings of the 3rd Workshop for Natural Language Processing Open Source Software (NLP-OSS 2023)*, 245–245, DOI: 10.18653/v1/2023.nlposs-1.26 (Association for Computational Linguistics, Singapore, 2023).

26. Polignano, M., Basile, P. & Semeraro, G. Advanced natural-based interaction for the ITAlian language: LLaMAntino-3-ANITA. *Sci. Reports* **16**, 5375, DOI: 10.1038/s41598-025-31319-0 (2026).

27. Vreš, D. *et al.* Building a strong instruction language model for a less-resourced language (2026). 2603.01691.

28. Graham, Y., Haddow, B. & Koehn, P. Translationese in Machine Translation Evaluation, DOI: 10.48550/arXiv.1906.09833 (2019). ArXiv:1906.09833.

29. Voukoutis, L. *et al.* Meltemi: The first open large language model for greek (2024). 2407.20743.

30. Romanou, A. *et al.* INCLUDE: Evaluating multilingual language understanding with regional knowledge. In *The Thirteenth International Conference on Learning Representations* (2025).

31. Kyriazi, P. & Prokopidis, P. Multiple choice qa greek asep. *Hugging Face*. https://huggingface.co/datasets/ilsp/mcqa_greek_asep/ (2025).

32. Vacalopoulou, A., Sofianopoulos, S. & Prokopidis, P. Greek physical commonsense reasoning dataset. *Hugging Face*. https://huggingface.co/datasets/ilsp/greek_pcr/ (2025).

33. Chang, T. A. *et al.* Global piqa: Evaluating physical commonsense reasoning across 100+ languages and cultures (2025). 2510.24081.

34. Hu, E. J. *et al.* LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations* (2022).

35. Mao, Y. *et al.* A survey on LoRA of large language models. *Front. Comput. Sci.* **19**, 197605, DOI: 10.1007/s11704-024-40663-9 (2024).

36. Biderman, D. *et al.* LoRA learns less and forgets less. *Transactions on Mach. Learn. Res.* (2024). Featured Certification.

37. Shuttleworth, R. S., Andreas, J., Torralba, A. & Sharma, P. LoRA vs full fine-tuning: An illusion of equivalence. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems* (2025).

38. Chowdhery, A. *et al.* Palm: Scaling language modeling with pathways. *J. Mach. Learn. Res.* **24**, 1–113 (2023).

39. Wolf, T. *et al.* Transformers: State-of-the-art natural language processing. In Liu, Q. & Schlangen, D. (eds.) *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 38–45, DOI: 10.18653/v1/2020.emnlp-demos.6 (Association for Computational Linguistics, Online, 2020).

40. Renze, M. The effect of sampling temperature on problem solving in large language models. In Al-Onaizan, Y., Bansal, M. & Chen, Y. N. (eds.) *Findings of the Association for Computational Linguistics: EMNLP 2024*, 7346–7356, DOI: 10.18653/v1/2024.findings-emnlp.432 (Association for Computational Linguistics, Miami, Florida, USA, 2024).

41. Dettmers, T., Pagnoni, A., Holtzman, A. & Zettlemoyer, L. Qlora: efficient finetuning of quantized llms. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, NIPS '23 (Curran Associates Inc., Red Hook, NY, USA, 2023).

42. Smith, S. L., Kindermans, P. J. & Le, Q. V. Don't decay the learning rate, increase the batch size. In *International Conference on Learning Representations* (2018).

43. Zhang, J., He, T., Sra, S. & Jadbabaie, A. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations* (2020).

44. Dettmers, T., Lewis, M., Shleifer, S. & Zettlemoyer, L. 8-bit optimizers via block-wise quantization (2022). 2110.02861.

45. Dror, R., Baumer, G., Shlomov, S. & Reichart, R. The hitchhiker's guide to testing statistical significance in natural language processing. In Gurevych, I. & Miyao, Y. (eds.) *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 1383–1392, DOI: 10.18653/v1/P18-1128 (Association for Computational Linguistics, Melbourne, Australia, 2018).

46. Kaddas, P., Gatos, B., Palaiologos, K., Christopoulou, K. & Kritsis, K. Text line detection and recognition of greek polytonic documents. In Coustaty, M. & Fornés, A. (eds.) *Document Analysis and Recognition – ICDAR 2023 Workshops*, 213–225 (Springer Nature Switzerland, Cham, 2023).

47. Chatzikyriakidis, S., Papadakis, D., Papaioannou, S. I. & Psaltaki, E. Grdd+: An extended greek dialectal dataset with cross-architecture fine-tuning evaluation (2026). 2511.03772.

48. Chlapanis, O. S., Galanis, D., Aletras, N. & Androutsopoulos, I. GreekBarBench: A challenging benchmark for free-text legal reasoning and citations. In Christodoulopoulos, C., Chakraborty, T., Rose, C. & Peng, V. (eds.) *Findings of the Association for Computational Linguistics: EMNLP 2025*, 25099–25119, DOI: 10.18653/v1/2025.findings-emnlp.1368 (Association for Computational Linguistics, Suzhou, China, 2025).

49. Tsourma, M., Michail, D., Varlamis, I., Drosou, A. & Tzovaras, D. Legal assistance in low-resource languages: Evaluating rag and fine-tuned llms for greek e-governance. In *2025 3rd International Conference on Foundation and Large Language Models (FLLM)*, 366–373, DOI: 10.1109/FLLM67465.2025.11391043 (2025).

50. Koniaris, M., Galanis, D., Giannini, E. & Tsanakas, P. Evaluation of automatic legal text summarization techniques for greek case law. *Information* **14**, DOI: 10.3390/info14040250 (2023).

51. Giarelis, N., Mastrokostas, C. & Karacapilidis, N. Greekt5: Sequence-to-sequence models for greek news summarization. In Maglogiannis, I., Iliadis, L., Macintyre, J., Avlonitis, M. & Papaleonidas, A. (eds.) *Artificial Intelligence Applications and Innovations*, 60–73 (Springer Nature Switzerland, Cham, 2024).

52. Giarelis, N., Mastrokostas, C. & Karacapilidis, N. Greek wikipedia: A study on abstractive summarization. In *Proceedings of the 13th Hellenic Conference on Artificial Intelligence*, SETN '24, DOI: 10.1145/3688671.3688769 (Association for Computing Machinery, New York, NY, USA, 2024).

53. Loukas, L. *et al.* GR-NLP-TOOLKIT: An open-source NLP toolkit for Modern Greek. In Rambow, O. *et al.* (eds.) *Proceedings of the 31st International Conference on Computational Linguistics: System Demonstrations*, 174–182 (Association for Computational Linguistics, Abu Dhabi, UAE, 2025).

54. Liapis, C. M., Kyritsis, K., Perikos, I. & Paraskevas, M. Transformer-based embeddings for greek language categorization. In *2024 IEEE/ACIS 24th International Conference on Computer and Information Science (ICIS)*, 176–181, DOI: 10.1109/ICIS61260.2024.10778332 (2024).

55. Stylianou, N., Tsikrika, T., Vrochidis, S. & Kompatsiaris, I. Cross-domain hate speech detection for content moderation in greek social networks. In *2024 IEEE/WIC International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, 373–379, DOI: 10.1109/WI-IAT62293.2024.00059 (2024).

56. Mastrokostas, C., Giarelis, N. & Karacapilidis, N. Social media topic classification on greek reddit. *Information* **15**, DOI: 10.3390/info15090521 (2024).

Acknowledgements

This work has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101235708 (BLUEPRINT – Building Living Urban Ecosystems through Participatory Renovation and Innovation Tools). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

Author contributions

N.G.: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.
C.M.: Conceptualization, Data curation, Investigation, Methodology, Software, Writing – review & editing.
N.K.: Project administration, Funding acquisition, Supervision, Writing – review & editing.
All authors reviewed the manuscript.

Data availability

The proposed model and dataset are available in `https://huggingface.co/IMISLab`. The other evaluation datasets used in the study can be found in `https://huggingface.co/datasets`.

Code availability

The source code of this study is available in `https://github.com/NC0DER/Maistros`.

Competing Interests Statement

The authors declare no competing interests.

Appendix A

Table A1. The prompts utilized in this study alongside their English translations.

Prompt Type	Greek	English (Translated)
Dataset Creation (Questions)	Είσαι ένα εξαιρετικά ανεπτυγμένο μοντέλο Τεχνητής Νοημοσύνης για την Ελληνική γλώσσα. Χρησιμοποίησε τις παρακάτω οδηγίες για να δημιουργήσεις μια σειρά ερωτήσεων στο θέμα που αναφέρει ο χρήστης: Οδηγίες: 1. Απάντα αποκλειστικά στα Ελληνικά με άψογη γραμματική, σύνταξη και ορθογραφία. 2. Λάβε υπόψη τον ελληνικό πολιτισμό και την ελληνική κοινωνική πραγματικότητα όπου είναι σχετικό. 3. Απόφυγε τη χρήση στερεοτύπων. 4. Βάλε πάντα το σύμβολο • πριν από κάθε ερώτηση. 5. Δημιούργησε σημαντικές, συχνά προκύπτουσες και χρήσιμες ερωτήσεις για το θέμα. 6. Όλες οι ερωτήσεις πρέπει να μπορούν να απαντηθούν αντικειμενικά. 7. Μη δημιουργείς επαναλαμβανόμενες ερωτήσεις. 8. Κάθε ερώτηση πρέπει να είναι σαφώς ορισμένη. 9. Γράψε μόνο το κείμενο των ερωτήσεων, χωρίς επιπλέον σχόλια. Παρακαλώ δημιούργησε 15 ερώτησεις για το εξής θέμα: {topic}	You are an extremely developed Artificial Intelligence model for the Greek Language. Use the following instructions to create a series of questions on the topic mentioned by the user: Instructions: 1. Answer exclusively in Greek with impeccable grammar, syntax and spelling. 2. Take into consideration the Greek civilization and the Greek social reality where relevant. 3. Avoid the use of stereotypes. 4. Always place the symbol • before each question. 5. Create significant, frequently occurring and useful questions for the topic. 6. All questions must be able to be answered objectively. 7. Do not create repeated questions. 8. Every question must be clearly defined. 9. Write only the text of the questions, without extra comments. Please create 15 questions for the following topic: {topic}

Continued on next page

Table A1 – continued from previous page

Prompt Type	Greek	English (Translated)
Dataset Creation (Answers)	Είσαι ένα εξαιρετικά ανεπτυγμένο μοντέλο Τεχνητής Νοημοσύνης για την Ελληνική γλώσσα. Χρησιμοποίησε τις παρακάτω οδηγίες για να παράγεις τη καλύτερη δυνατή απάντηση: Οδηγίες: 1. Απάντα αποκλειστικά στα Ελληνικά με άψογη γραμματική, σύνταξη και ορθογραφία. 2. Λάβε υπόψη τον ελληνικό πολιτισμό και την ελληνική κοινωνική πραγματικότητα όπου είναι σχετικό. 3. Απάντησε στις ερωτήσεις του χρήστη με ειλικρίνεια και επιστημονική ακρίβεια. 4. Απόφυγε τη χρήση στερεότυπων. 5. Αν η ερώτηση είναι ασαφής ή λείπει πληροφορία (π.χ. χώρα, περίοδος): – Μην ζητήσεις διευκρίνιση. – Δώσε την απάντηση κάνοντας ρητές παραδοχές (π.χ. ¨Ελλείψει άλλης αναφοράς, θεωρούμε ως προεπιλογή την Ελλάδα και το τρέχον έτος¨). Παρακαλώ απάντησε στη παρακάτω ερώτηση: {question}	You are an extremely developed Artificial Intelligence model for the Greek Language. Use the following instructions to generate the best possible answer: Instructions: 1. Answer exclusively in Greek with impeccable grammar, syntax and spelling. 2. Take into consideration the Greek civilization and the Greek social reality where relevant. 3. Answer the user question with honesty and scientific accuracy. 4. If the question is vague or information is missing (e.g., country, time period): – Do not ask for clarification. – Give the answer by making explicit assumptions (e.g. "In the absence of other reference, we assume as a default Greece and the current year”). Please answer the following question: {question}
Multiple Choice (Evaluation)	Διάβασε προσεκτικά την ερώτηση και σκέψου ποια επιλογή είναι σωστή. Επίλεξε την καλύτερη απάντηση. Απάντησε μόνο με το γράμμα (Α, Β, Γ ή Δ). Ερώτηση: {question} Απαντήσεις: {answers}	Read the question carefully and think about which option is correct. Choose the best answer. Answer only with the letter (A, B, C or D). Question: {question} Answers: {answers}
Open-ended (Evaluation)	Σκέψου και απάντησε στην παρακάτω ερώτηση με συντομία, σαφήνεια και ακρίβεια. Ερώτηση: {question}	Think and answer the following question with brevity, relevance and precision. Question: {question}