

# Bilingual Evaluation of Language Models on General Knowledge in University Entrance Exams with Minimal Contamination

Eva Sánchez Salido, Roser Morante, Julio Gonzalo, Guillermo Marco  
Jorge Carrillo-de-Albornoz, Laura Plaza, Enrique Amigó, Andrés Fernández  
Alejandro Benito-Santos, Adrián Ghajari Espinosa and Victor Fresno

UNED Research Center in Natural Language Processing and Information Retrieval\*  
ETSI Informática, UNED - Juan del Rosal, 16 28040 Madrid, Spain

Correspondence: [julio@lsi.uned.es](mailto:julio@lsi.uned.es)

## Abstract

In this article we present *UNED-ACCESS 2024*, a bilingual dataset that consists of 1003 multiple-choice questions of university entrance level exams in Spanish and English. Questions are originally formulated in Spanish and manually translated into English, and have not ever been publicly released, ensuring minimal contamination when evaluating Large Language Models with this dataset. A selection of current open-source and proprietary models are evaluated in a uniform zero-shot experimental setting both on the UNED-ACCESS 2024 dataset and on an equivalent subset of MMLU questions. Results show that (i) Smaller models not only perform worse than the largest models, but also degrade faster in Spanish than in English. The performance gap between both languages is negligible for the best models, but grows up to 37% for smaller models; (ii) Model ranking on UNED-ACCESS 2024 is almost identical (0.98 Pearson correlation) to the one obtained with MMLU (a similar, but publicly available benchmark), suggesting that contamination affects similarly to all models, and (iii) As in publicly available datasets, reasoning questions in UNED-ACCESS are more challenging for models of all sizes.

## 1 Introduction

With the recent progress in broadening the generalisation capabilities of Large Language Models (LLMs), much current research focuses on understanding their capabilities and limitations. Evaluation of generative models, such as Llama (Touvron et al., 2023b), Mistral (Jiang et al., 2023), Mistral (Jiang et al., 2024a), Gemini (Anil et al., 2024), Gemma (Mesnard et al., 2024), GPT-3.5 (Brown et al., 2020), GPT-4 and GPT-4o (Achiam et al., 2024), attempt at measuring their world knowledge, memorization and inference capabilities with

benchmarks designed for these purposes. Although many benchmarks have been proposed as single-task evaluations, with the emergence of general language models such as BERT (Devlin et al., 2019), the development of more comprehensive benchmarks to measure the general capabilities of these models in a multi-task setting became popular (GLUE (Wang et al., 2018), SuperGLUE (Wang et al., 2020), Big-Bench (Srivastava et al., 2022), Big-Bench Hard (Suzgun et al., 2022), HELM (Liang et al., 2023)). A lot of emphasis is now put on assessments focusing on human-level cognitive tasks in real-world evaluation scenarios, such as exams. In this context, multiple-choice questions have emerged as one of the preferred methods for evaluating new generative models with datasets such as RACE (Lai et al., 2017), MMLU (Hendrycks et al., 2021a) or AGIEval (Zhong et al., 2023), which includes university entrance exams, law school entrance exams, mathematics competitions and lawyer qualification tests.

In this paper we present a new bilingual dataset of **private** exam questions, UNED-ACCESS 2024. The dataset contains 1003 multiple-choice questions in Spanish from various subjects of the UNED Access Course for Over-25s in Spanish, together with high-quality English professional translations. The dataset has been compiled at Universidad Nacional de Educación a Distancia (UNED), in the framework of the ODESIA project (*Espacio de Observación de Inteligencia Artificial en Español*). This dataset is complementary to the ODESIA Benchmark, a collection of comparable datasets in English and Spanish to evaluate multiple NLP discriminative tasks. The ODESIA Benchmark is used in the ODESIA Leaderboard,<sup>1</sup> with the aim of measuring the effectiveness gap between Spanish and English LLMs. Two characteristics make

\*[nlp.uned.es](http://nlp.uned.es)

<sup>1</sup><https://leaderboard.odesia.uned.es>

this dataset unique: first, the contamination level is minimal, because UNED typically does not release the answers to the exam questions, which are only accessible to the teachers of each course. Second, this is a high-quality bilingual dataset, with original questions in Spanish translated into English manually by a professional translators who did not use any external software.

We also present results on evaluating multiple commercial and open-source LLMs on this dataset, and on an equivalent version of the popular MMLU dataset (Hendrycks et al., 2021a) with the same evaluation protocol. These allows us to address three research questions: **RQ1 (Contamination effects)**: How do current generative models perform on official private entrance university exams, in a minimal contamination setup? **RQ2 (Language effects)**: How do models perform comparatively in Spanish and English? and **RQ3 (Dataset size effect)**: How suitable is the zero-shot approach to evaluate LLMs on a relatively small dataset of exam questions?

Our main contributions are: (i) The creation of a new, private benchmark with minimal contamination for the evaluation of generative LLMs on entrance university exams in Spanish (native) and English (translated); (ii) A systematic evaluation of LLMs with the same prompts and hyperparameters for all models, in English and Spanish; (iii) Empirical evidence that the effectiveness gap of LLMs between English and Spanish is inversely correlated with model quality and size: smaller models have lower performance values in both languages and also a larger gap between their effectiveness in English and Spanish, and (iv) Empirical evidence that the relative performance of models on UNED-ACCESS 2024 is almost identical to their performance on an equivalent MMLU dataset and that results correlate with the public MMLU results, which may imply that contamination affects similarly to all language models, and that a small, sufficiently diverse dataset is representative enough to measure model performance.

## 2 Related work

We summarize related work regarding current approaches to evaluate LLMs, existing benchmarks from human evaluations, multilingual LLMs evaluations and the problem of data contamination.

**LLM benchmarking.** General benchmarking of the comparative performance of LLMs is currently

performed either with datasets of questions and answers (most commonly in multiple-choice format) or with *LLM arenas* where users make requests, are offered answers from two or more LLMs, and pick their preferred one (Chiang et al., 2024). Here we focus on the first approach, where UNED-ACCESS 2024 belongs.

Standard question/answer datasets include tasks such as common sense reasoning: HellaSwag (Zellers et al., 2019), WinoGrande (Sakaguchi et al., 2019), ARC (Clark et al., 2018), OpenBookQA (Mihaylov et al., 2018), CommonsenseQA (Talmor et al., 2019), PIQA (Bisk et al., 2019), SIQA (Sap et al., 2019); world knowledge: TriviaQA (Joshi et al., 2017), NaturalQuestions (Kwiatkowski et al., 2019); reading comprehension: RACE (Lai et al., 2017), QuAC (Choi et al., 2018) and SQuAC (Rajpurkar et al., 2016); mathematical reasoning: MATH (Hendrycks et al., 2021b), GSM8K (Cobbe et al., 2021)); and code generation: HumanEval (Chen et al., 2021), MBPP (Austin et al., 2021). Aggregated benchmarks like MMLU (Hendrycks et al., 2021a), AGIEval (Zhong et al., 2023), Big-Bench (Srivastava et al., 2022), and Big-Bench Hard (Suzgun et al., 2022) are also commonly used. In particular, since the introduction of GPT-3, evaluations that comprise a variety of datasets have been standardised (see for instance evaluations of GPT-3 (Brown et al., 2020), GPT-4 (Achiam et al., 2024), Llama and Llama-2 (Touvron et al., 2023b), Mistral (Jiang et al., 2023) or Claude-3 (Anthropic, 2024)). For Spanish and its official varieties, a new benchmark for the evaluation of open generative LLMs, named *La Leaderboard*, has recently been released (Grandury et al., 2024).

Evaluations often mix prompting strategies, varying from zero-shot, few-shot with different numbers of examples, and chain-of-thought configurations. For example, in the evaluation of Llama-2 (Touvron et al., 2023b) we find a zero-shot setting for evaluating common sense reasoning, except for one dataset in which 7-shot is used, and for aggregated benchmarks we find 3 and 5-shot settings. In more recent evaluations, such as that of Claude-3 (Anthropic, 2024), few-shot experimentation is performed with different numbers of examples (varying from 0 to 25) and chain-of-thought strategies.

**Exams benchmarks.** MMLU (Hendrycks et al., 2021a) is a popular benchmark with English multiple-choice questions covering subjects like

mathematics, history, and computer science, ranging from elementary to college level. AGIEval (Zhong et al., 2023) assesses human-like cognition and problem-solving abilities using high-stakes exams, such as university entrance and law school exams, in both English and Chinese. GPQA (Rein et al., 2023) presents challenging questions in biology, physics, and chemistry, aimed at testing even expert knowledge. M3Exam (Zhang et al., 2023) evaluates multilingual, multi-modal intelligence in LLMs using questions from human exams in nine languages. Additional benchmarks include GSM8K (Cobbe et al., 2021) for elementary-level math, EXAMS (Hardalov et al., 2020), which covers 24 school subjects across 16 languages, OpenBookQA (Mihaylov et al., 2018) and ARC (Clark et al., 2018) for assessing scientific knowledge in English, and RACE (Lai et al., 2017), which focuses on reading comprehension from real English exams designed by Chinese teachers.

**Multilingual evaluations.** Recent research has increasingly focused on multilingual evaluations of models like ChatGPT. For instance, Lai et al. (2023) evaluated ChatGPT on 37 languages ranging from high to extremely low resources, across seven tasks in a zero-shot learning setting. Their findings, along with Bang et al. (2023), show that models struggle to generalize in low-resource languages. Similarly, M3Exam (Zhang et al., 2023) assesses models like GPT-4 across nine languages, revealing persistent difficulties in handling low-resource and non-Latin script languages. GPT-4’s technical report itself extends its evaluation to a multilingual version of MMLU self-translated via Azure Translate (Achiam et al., 2024), and a professional manual translation of MMLU in 14 languages (MMMLU) has recently been released (OpenAI, 2024).

To address multilingual challenges, benchmarks such as XTREME (Hu et al., 2020), XTREME-R (Ruder et al., 2021), and MEGA (Ahuja et al., 2023) have been developed. Recently, a benchmark was used to assess general knowledge of the models, and compare English and Chinese (Gu et al., 2024). Additionally, models like Mixtral (Jiang et al., 2024a) include multilingual evaluations at launch, and works such as those of Li et al. (2024) and Blasi et al. (2021) highlight how model performance correlates with the amount of data available in each language during pretraining.

**Data contamination.** Machine learning models, especially LLMs, are increasingly dealing with the problem of data contamination, where models are trained on datasets from the internet, including those used for evaluation, leading to biased outcomes as the models may already know the answers beforehand (Jiang et al., 2024b; Dong et al., 2024; Golchin and Surdeanu, 2024; Sainz et al., 2023; Yang et al., 2023). This issue is gaining more attention, as seen in recent works and efforts like the LM Contamination Index<sup>2</sup>. Several techniques have been proposed to detect contamination, such as checking dataset release dates or availability on the web, but these approaches may fail to address issues like model updates or indirect data leakage (Ahuja et al., 2023). To mitigate this, researchers suggest generating variations of datasets to assess reasoning capabilities (Srivastava et al., 2024), or using private benchmarks to minimize contamination (Rajore et al., 2024), approach that we follow.

We introduce UNED-ACCESS 2024 to complement current evaluations by providing a bilingual and comparable resource featuring parallel questions in both Spanish and English. The questions are sourced from university private exams to minimize contamination, and have been professionally translated. Furthermore, by conducting experiments in a zero-shot mode, we aim to offer a more uniform evaluation compared to existing methods.

### 3 The UNED-ACCESS 2024 Dataset

UNED-ACCESS 2024 contains 1003 multiple-choice questions, with 3 or 4 possible answers, from the following subjects of the UNED Access Course for Over 25s: Business Administration and Management, Biology, Biochemistry, Economics, Fundamentals of Computer Science, Spanish Language, Literature, Mathematics, Mathematics Applied to Social Sciences, Advanced Mathematics and Psychology. The questions, originally in Spanish and human-translated into English, are from exams conducted between 2009 and 2022. Although several benchmarks of multiple-choice questions exist, some of them, such as ARC (Clark et al., 2018) and OpenBookQA (Mihaylov et al., 2018) have been found to be limited as they cover simple topics of primary school subjects, for which models can already achieve solid performance. UNED-ACCESS 2024 covers more complex topics and

<sup>2</sup><https://hitz-zentroa.github.io/lm-contamination/>

	Languages	Original	Translations	Same instances per language	Translation method	Access
<b>UNED-ACCESS 2024</b>	Spanish, English	Spanish	Yes	Yes	Professional manual translations, no external software	Private
<b>AGIEval</b> (Zhong et al., 2023)	English, Chinese	English, Chinese	No translations, except in LogiQA	No, except in LogiQA	Manual translation from Chinese	Public
<b>Xiezhi</b> (Gu et al., 2024)	Chinese, English	Chinese	Yes, from Chinese to English	Yes	Google Translate API, followed by manual post-processing	Public
<b>M3Exam</b> (Zhang et al., 2023)	English and 8 other languages	9 languages	No	No	No translations	Public
<b>EXAMS</b> (Hardalov et al., 2020)	English, Spanish and 14 other languages	16 languages	No, except for parallel questions	No, except for parallel questions in 7 languages	No translations	Public
<b>XTREME</b> (Hu et al., 2020) and <b>XTREME-R</b> (Ruder et al., 2021)	English, Spanish and 38 other languages, 9 tasks (XTREME); English, Spanish and 48 other languages, 10 tasks (XTREME-R)	Some tasks*	Some tasks	Some tasks	Some professional translations, some parallel instances and some automatic translations**	Public
<b>MEGA</b> (Ahuja et al., 2023)	English, Spanish and 68 other languages (16 datasets)	Some tasks	Some tasks	Some tasks	Some datasets contain professional translations, parallel instances and automatic translations***	Public
<b>MMMLU</b> (OpenAI, 2024)	English, Spanish and other 13 languages	English	Yes	Yes	Professional human translations	Public

Table 1: Comparison of translation methods in multilingual datasets. \*While in the XNLI (Conneau et al., 2018) or XQuAD (Artetxe et al., 2020) datasets the translations from English were done manually by professional translators, in others such as MLQA instances exist naturally in several languages, without relying on English translations. \*\*Since not all tasks are available in all languages, to enable a broader comparison across languages they automatically translate some English instances to the remaining languages using an in-house translation system. \*\*\*In order to compare diverse prompting strategies, they use Bing translator to translate target language data to English.

questions that require from simple memorization to mathematical reasoning or linguistic skills. Table 1 shows a comparison of several multilingual datasets and benchmarks, regarding whether they are public, the languages included, the original language of the questions or instances, and the availability of translations and the translation method used. It shows that UNED-ACCESS 2024 is the only dataset with private access and fully manual translation.

Appendix A contains more details on the compilation process, the format of the dataset, and the dataset statement. The following are two examples of questions in English:

- **Spanish Language:**  
Which of the following sentences uses accents correctly?  
A. Es por éso. B. Fui y volví. C. Llorar y reir. D. No se prohíbe.
- **Mathematics Applied to Social Sciences:**  
When  $(x \rightarrow 1)$ , the function  $(f(x) = \frac{1-x^2}{1-x})$ : A. Has a limit of  $(\infty)$ . B. Has a limit of 2. C. Has no limit.

Table 2 shows, per subject, the number of questions, the total word count and the number of answer options per question, which is relevant since the probability of getting the answer right randomly depends on the number of possible answers.

Subject	# Q	# AxQ	# W
Business and Administration Management	87	3	3849
Biology	119	3	2753
Biochemistry	59	3	1407
Economics	51	4	1725
Foundations of Computer Science	63	4	1987
Spanish Language	94	4	2816
Literature	91	4	5130
Mathematics	73	3	1465
Mathematics Applied to Social Sciences	94	3	2847
Advanced Mathematics	24	3	446
Psychology	248	4	5669
Total	1003	-	30094

Table 2: Distribution of the number of questions per subject (# Q), the number of answer options per question (# AxQ) and total word count (# W) per subject in the Spanish UNED-ACCESS 2024 dataset.

## 4 Experimental setup

Experiments were conducted using 12 generative models: four proprietary models (GPT-4-Turbo (Achiam et al., 2024), GPT-4o, GPT-3.5-Turbo (Brown et al., 2020) and Claude-3-Opus (Anthropic, 2024)) and eight open-weights models: Llama-2-7B (Touvron et al., 2023a), Llama-3-8B (Meta, 2024), Llama-3-70B, Gemma-7B (Mesnard



et al., 2024), Gemma-2-27B, Mixtral-8x7B (Jiang et al., 2024a), Mistral-7B (Jiang et al., 2023) and Leniachat-Gemma-2B, all trained for instruction following. In Appendix B.1 details are provided about how the models were accessed.

We decided to apply a uniform zero-shot setting because it more closely resembles the real-world scenario in which people interact with LLMs, and it also results in a simpler and more replicable experimentation. As for the hyperparameters, the temperature was set to 0 in all models to ensure deterministic responses, minimizing creativity and focusing on factual or reasoning-based answers, which suits the dataset. The questions were provided one at a time and the prompt is fixed, providing a system prompt, a user prompt and an assistant prompt. As extra information, the prompt provided the name of the subject to which the question belongs. The language of the prompt was always in the same language of the dataset used to evaluate the model (English or Spanish) as in Zhang et al. (2023) and Achiam et al. (2024). The prompts are the following for Spanish (ES) and English (EN):

- **System prompt**  
**ES:** Eres un sistema experto en responder preguntas de exámenes.  
**EN:** You are an expert system for answering exam questions.
- **User prompt**  
**ES:** Responde a la siguiente pregunta de la asignatura {}, tan solo con la letra de la respuesta correcta. Pregunta:  
**EN:** Answer the following question of the subject {} only with the letter of the correct answer. Question:
- **Assistant prompt**  
**ES:** Letra de la respuesta correcta:  
**EN:** Letter of the correct answer:

Additionally, for open models the instruction was formatted according to the labels used in training, as explained in their model cards. These templates are shown in Appendix B.2. Finally the literal output of the models was cleaned up and structured before applying the evaluation script, since, on many occasions, the answer contained answer justifications or additional material. The letter of the correct answer was kept if provided.

**Evaluation metrics.** In the literature, evaluation is typically based on Accuracy, the proportion of correct answers ( $C$ ) over the total responses ( $N$ ). However, to allow for direct comparisons across subjects with varying numbers of answer choices, we use Cohen’s Kappa coefficient, which accounts

for the number of possible choices ( $M$ ) in multiple-choice questions:

$$\text{Kappa} = \frac{\text{observed accuracy} - \text{expected accuracy}}{1 - \text{expected accuracy}} = \frac{\frac{C}{N} - \frac{1}{M}}{1 - \frac{1}{M}}$$

where the expected accuracy corresponds to that of random guessing: 1/3 for questions with 3 options and 1/4 for questions with 4 options. Cohen’s Kappa adjusts the correctness level so that random answers result in a Kappa of zero, allowing comparison across exams with varying answer choices. Kappa values range from -1/2 to 1, with negative values indicating worse-than-random performance. The final result for each model and language is the arithmetic mean of Kappa values across all subjects, giving equal weight to each subject to account for dataset imbalances.

## 5 Results

Model	Type	EN	ES	Gap%
Claude-3-Opus	proprietary	0.79	0.81	-2.51
GPT-4o	proprietary	0.78	0.77	0.83
GPT-4-Turbo	proprietary	0.76	0.78	-1.72
Llama-3-70B-Instruct	open	0.65	0.67	-3.44
Gemma-2-27B-Instruct	open	0.64	0.66	-3.40
GPT-3.5-Turbo	proprietary	0.60	0.55	9.19
Mixtral-8x7B-Instruct	open	0.56	0.57	-2.32
Llama-3-8B-Instruct	open	0.51	0.50	2.73
Mistral-7B-Instruct	open	0.46	0.43	5.67
Gemma-7B-It	open	0.41	0.38	8.93
Llama-2-7B-Chat	open	0.32	0.25	26.04
Leniachat-Gemma-2B	open	0.15	0.11	37.27

Table 3: Average performance of each model on UNED-ACCESS 2024 (Cohen’s Kappa) and language gap %. Results are sorted by performance in English.

Table 3 shows the average Cohen’s Kappa results by model in Spanish and English. Additional detailed results per language, model and subject are provided in Table 6 in Appendix C.1. The best overall results for Spanish are achieved by the proprietary models Claude-3-Opus, GPT-4-Turbo and GPT-4o, followed by the open models Llama-3-70B, Gemma-2-27B and Mixtral-8x7B, which outperform GPT-3.5-Turbo. In the pass margin (around 0.50 points) there is Llama-3-8B, and below are Mistral-7B, Gemma-7B, Llama-2-7B and Leniachat-2B. For English, the ranking of the models by average performance is almost identical to that of Spanish, except for an exchange in the positions of GPT-3.5 and Mixtral (GPT-3.5 performs better than Mixtral in English but worse in Span-

ish). At the quantitative level, the results are also very similar in English and Spanish.

**EN-ES Language gap.** The last column in Table 3 shows the performance gap between English and Spanish. We calculate the percentage effectiveness gap between languages from the averaged results as:

$$\text{GAP\% EN-ES} = \frac{(\text{EN-ES})}{\text{ES}} \cdot 100$$

Five out of twelve models perform slightly better in Spanish than English (namely Claude-3, GPT-4, Llama-3-70B, Gemma-2-27B and Mixtral-8x7B, where the gap is negative), two of them proprietary (which perform the best in Spanish overall) and three open (also the three best models among the open-weights ones). This does not imply that these models are better in Spanish: English questions are manual translations of the original Spanish ones, and there are at least three language-related factors that affect overall performance: (i) even with high-quality manual translations, there might be translation artifacts (e.g. unusual translation choices) that make the English text harder to process; (ii) contamination: even if the questions are private, the chance that some questions have been seen by the models in Spanish is higher than in English; (iii) language competence of the models in English vs Spanish. Presumably, in the models that perform better in Spanish, effects (i) and (ii) out-weight pure language competence.

Among the proprietary models, three of them have similar performance differences (in absolute value), whilst that of GPT-3.5 is the most pronounced, in favour of English. As for the open models, the two best performing ones (Llama-3-70B and Gemma-2) have a very similar performance difference, both in favour of Spanish (and larger than the ones observed for the best proprietary models). Mixtral is the next model in the ranking and still favours Spanish, while all the other models, which do not pass (or are at the limit), favour English. In this set of models the performance difference is very pronounced, specially for Llama-2 and Leniachat, which are the worst performing models. In summary, for the best models we observe a slight advantage for Spanish, while in the worst models there is a pronounced advantage towards English. The best models are in turn larger, and therefore more likely to have been trained with more Spanish data, and more data from the web in general, whereas smaller models may rely more solely on

English. However, in principle this is not the case of *Leniachat*, which has been conceived specifically for Spanish. The size of the models is crucial: although we do not know for certain the size of the proprietary models, in general model size is the best performance predictor, with one exception: the very similar performance between Llama-3-70B and Gemma-2-27B, with such a large difference in size, is noteworthy.

There are two cases where the gap is extremely pronounced (in favour of English), coinciding with the two worst-performing models: *Leniachat*, which is the smallest model, and Llama-2, which is the oldest of the evaluated models.

There is a strong, **statistically significant inverse correlation between model quality and language gap**: Pearson's correlation between English performance and EN-ES gap is -0.87 with a p-value of 0.0002; and Pearson's correlation between performance in Spanish and language gap is -0.89 ( $p=9.83e-05$ ). The performance of less powerful models degrades faster in Spanish than in English.

**Results by subject.** Figure 1 shows the performance per subject for all models in both languages (also provided in Table 6 in Appendix C.1). Subjects with the highest scores tend to be Biology, Biochemistry, Fundamentals of Computer Science and Psychology. In Economics and Business Administration and Management (BAM), a slight decrease in scores is observed, as well as in Literature and Spanish Language. The worst results are observed in the three Mathematics subjects, with slightly higher results in Mathematics Applied to Social Sciences, which is presumably the least complex of the three. Just as the low performance in Mathematics subjects can be explained by the reasoning abilities that the questions require, the lower results in Spanish Language may also be due to the fact that the models are trained on tokens but do not really distinguish a letter or a word, and therefore fail in tasks that require concrete lexical, syntactic or orthographic understanding. Besides, the fact that the models solve Biochemistry questions so well in both languages suggests that these questions only involve recalling fundamental knowledge; that the information needed to answer the questions is well represented in internet; that the questions do not require complex reasoning or deep contextual interpretation; and on the other hand, that the models understand the formulas expressed in  $\text{\LaTeX}$ . This is a hint that formulas in  $\text{\LaTeX}$  should not




Figure 1: Distribution of results on UNED-ACCESS 2024 per subject (Cohen's Kappa).

be what prevents models from solving the Math questions correctly.

Broadly speaking, the best models perform best in all subjects and, conversely, the worst models remain at low scores throughout the dataset. The best performing model in average, Claude-3, passes all subjects except for Advanced Mathematics, suggesting that even the most advanced models do not show true reasoning capabilities as an emergent phenomenon. Claude-3 scores above 0.8 in seven out of the eleven subjects in Spanish (BAM, Biology, Biochemistry, Economics, Fundamentals of Computer Science, Literature and Psychology) and six in English. In Spanish Language it remains at 0.67 in both languages, and around 0.65 in Mathematics and Mathematics Applied to Social Sci-

ences. The results for GPT-4 and GPT-4o are very similar. Llama-3-70B, Gemma-2-27B and GPT-3.5 follow the lead, with results that drop overall by at least one tenth with a significant drop in the three Math subjects. Results continue to fall with Mixtral-8x7B, GPT-3.5 and Llama-3-8B, which fall outside the pass mark. The remaining models have results that hover around random probabilities, in some cases performing even worse than a random answer (values below zero), with the lowest results generally obtained with Gemma-7B, Llama-2-7B and Leniachat-Gemma-2B. It is interesting to note that although Claude-3 achieves the highest average result, in numerous subjects it is GPT-4 or GPT4o models that achieve the highest results.

**UNED-ACCESS vs MMLU: public vs private test sets.**

A direct comparison between the public results on MMLU (Hendrycks et al., 2021a) and UNED-ACCESS 2024 is not adequate: the dataset sizes are different (MMLU has more than 14000 test instances) and the subjects and levels differ: MMLU ranges in difficulty from an elementary level to an advanced professional level, whereas UNED-ACCESS 2024 covers only university-entrance level. Also, most public evaluations of MMLU are 5-shot and even include chain of thought prompting, as compared to the zero-shot setting of our experiments. For these reasons we have performed our own evaluation of the models, using the same methodology described earlier, and selecting the subset of 2822 MMLU questions (*MMLU-10*) from ten topics which correspond most closely to the level and topics of UNED-ACCESS: high school biology, high school chemistry, high school computer science, elementary mathematics, high school mathematics, high school physics, high school macroeconomics, high school microeconomics, high school psychology and high school world history.

	MMLU from publications	MMLU-10 from zero-shot experiments
GPT-4o	88.7	83.40
Claude-3-Opus	86.8 (5-shot)	83.71
GPT-4-Turbo	86.5	81.61
Llama-3-70B-Instruct	82.0 <small>Llama-3-70B (5-shot)</small>	75.47
Gemma-2-27B-Instruct	75.2 <small>Gemma-2 (5-shot)</small>	76.61
Mixtral-8x7B-Instruct	70.6 <small>Mixtral 8x7B (5-shot)</small>	65.34
GPT-3.5-Turbo	70.0	63.74
Llama-3-8B-Instruct	68.4 <small>Llama-3-8B (5-shot)</small>	61.54
Mistral-7B-Instruct	58.4 (5-shot)	54.24
Gemma-7B-It	53.3 (5-shot)	42.50
Llama-2-7B-Chat	45.3 <small>Llama-2-7B (5-shot)</small>	35.08
Leniachat-Gemma-2B	-	29.81

Table 4: **Accuracy** scores of models obtained from published MMLU results in their respective reports, and results obtained from experiments on MMLU-10 executed with the experimental setting described in this paper.

First, we have checked whether the results that we obtain on MMLU-10 with our methodology correlate with the published results. Table 4 provides the MMLU results published in the official

benchmark<sup>3</sup> (or, in their absence, in the respective model reports) and the results of applying our methodology on MMLU-10. Both evaluations are reported in terms of Accuracy, rather than Cohen’s Kappa, because published results are always presented with this metric. This is the only table in this paper that presents results using Accuracy. Pearson’s correlation between these two sets of results is 0.9863 with a p-value of 2.31e-08, which means that results are highly correlated. This high correlation validates our experimentation, as it implies that our choices of prompts and hyperparameters, and zero-shot setting, yield results consistent with those found in the literature.

This analysis allows us to compare the results obtained on MMLU-10 with the results on UNED-ACCESS 2024 in English and check whether they correlate. Table 5 shows both results in terms of Cohen’s Kappa (recall that the MMLU questions all have 4 answers). Pearson’s correlation between these results is 0.98 (p=5.04e-08): both results have almost perfect, statistically significant correlation. This consistency reinforces the reliability of results obtained with both datasets.

	UNED-ACCESS 2024 EN, 1K questions	MMLU-10 EN, 2.8K questions
Claude-3-Opus	0.79	0.78
GPT-4o	0.78	0.78
GPT-4-Turbo	0.76	0.75
Llama-3-70B-Instruct	0.65	0.67
Gemma-2-27B-Instruct	0.64	0.69
GPT-3.5-Turbo	0.60	0.52
Mixtral-8x7B-Instruct	0.56	0.54
Llama-3-8B-Instruct	0.51	0.49
Mistral-7B-Instruct	0.46	0.39
Gemma-7B-It	0.41	0.23
Llama-2-7B-Chat	0.32	0.13
Leniachat-Gemma-2B	0.15	0.06

Table 5: Cohen’s Kappa results of models on UNED-ACCESS 2024 (English), and results from experiments on a subset of the MMLU executed with our prompts and hyperparameters in a zero-shot setup.

The fact that the results on both datasets are so similar, in spite of the fact that UNED-ACCESS should have minimal contamination, implies that either the questions in UNED-ACCESS are somewhat easier in average, or that there is some kind of indirect contamination: some questions might be recurring for some subjects, and models might have seen variations of the questions even if they do not have access to the original ones. Also, the

<sup>3</sup><https://paperswithcode.com/sota/multi-task-language-understanding-on-mmlu>



similarity of the results seems to indicate that the translation effect is not decisive.

In Table 7 of Appendix C.2 we also provide the evaluation results on the MMLU-10 dataset per subject, which shows the same trends as for UNED-ACCESS 2024. The worst results are obtained for Mathematics and the best results are obtained for Biology, Micro- and Macro-economics, Computer Science, World History and Psychology. Also equally, the best models perform best in all subjects and the best performing models in average, GPT-4o and Claude-3, pass all subjects except for Mathematics.

## 6 Conclusions

We introduced the UNED-ACCESS 2024 bilingual dataset to evaluate general knowledge of language models. It consists of 1003 private multiple-choice questions and answers in Spanish, manually translated into English. We evaluated proprietary and open-source LLMs on this dataset and on an equivalent subset of MMLU (MMLU-10), with the same methodology, addressing three research questions.

**RQ1: What is the performance of models in a dataset with minimal contamination?** We have found that the behaviour of models is almost identical on our dataset and on MMLU-10: the model rankings provided by both datasets are almost equivalent (Pearson correlation is 0,98). In both datasets, performance is strongly correlated with model size. Across subjects, all models perform substantially worse in subjects requiring reasoning and calculus, compared to those which require mainly memorization. We plan to classify the questions according to the level of memorisation or reasoning they require, as in Yu et al. (2023) where they use Bloom’s taxonomy of educational objectives (Bloom, 1956) in order to better assess models on these dimensions. Overall, the fact that UNED-ACCESS has minimal contamination does not lead to any meaningful difference in the relative models’ performance. This suggests that contamination affects similarly to all models, or that there are indirect sources of contamination in our dataset. In a preliminary investigation, however, a web search for solved exams in each of the subjects did not lead to any systematic repository with questions and answers, although we occasionally found PDF versions of the exams with the right answers highlighted with a pen (usually from academies that prepare the students for the exams).

**RQ2: What is the performance gap of models between English and Spanish?** Our most significant finding is that the effectiveness gap between English and Spanish is inversely correlated with model accuracy: the best models have smaller relative differences between their performance in both languages. In general, smaller models show larger gaps, and the best models even perform slightly better in Spanish than English. The original questions are in Spanish, and therefore this negative gap may indicate that the true language gap (most likely in favour of English) is out-weighted by translation effects and contamination (even if small, it is more likely in Spanish as it is the source language). More experimentation is needed to isolate these factors from each other.

**RQ3: How suitable is the zero-shot approach to evaluate LLMs on 1,000 exam questions?** The high correlation observed between the results in MMLU and UNED-ACCESS 2024 in English, along with the high correlation between Spanish and English results in UNED-ACCESS 2024, suggest that (i) our experimental setting is fit for reliable experimentation, since it yields results consistent with those found in the literature, and (ii) despite its reduced size, our dataset is sufficiently diverse and representative to measure performance by model and discipline, and the effect of manual translation (Spanish to English) does not significantly alter results.

To further mitigate contamination and increase difficulty, we plan to extend the dataset with undergraduate-level questions of increasing difficulty, and also to check the real level of contamination to which they may be exposed. We also plan to estimate the consistency of models’ answers to prompt variations, that is, how the models deal with changes in content (and get to measure a *reasoning gap* (Srivastava et al., 2024)), and in form (to measure the *performance spread* (Polo et al., 2024; Sclar et al., 2023; Mizrahi et al., 2024)), and apply this methodology to other benchmarks.

Finally, although the questions and answers will not be distributed to keep contamination minimal, new relevant open models will continue to be evaluated in a UNED-ACCESS leaderboard at <https://leaderboard.odesia.uned.es>. Developers with an open model that wish to be included in the evaluation may send a request to [odesia-comunicacion@lsi.uned.es](mailto:odesia-comunicacion@lsi.uned.es).

## Limitations

Firstly, although the correct answers of the exam questions of UNED-ACCESS 2024 have never been made public officially, which makes us believe that the risk of **contamination** is limited, we have not measured the real level of contamination by, for example, making systematic searches in the internet in order to determine whether and in what form some correct answers may be publicly available. As stated in the paper, more research is needed to reliably estimate contamination. Secondly, the **size** of the dataset might be considered small as compared to the most popular public datasets. We have given priority to gather quality questions with private correct answers, rather than to build a large dataset. However, we are working towards producing a larger dataset with university level questions in a diversity of subjects. Third, our estimation of the **language gap** is just a first approximation, as there are confounding factors (translation, contamination) that are difficult to isolate from linguistic performance. More research is needed to find adequate methodologies to estimate the language gap. Also, our experimentation is directed at estimating a minimal gap (comparing English with other predominant online language such as Spanish); obviously, the gap will be larger with less resourced languages. Finally, we have not **categorized** the questions into levels of cognitive complexity, for example using Bloom’s taxonomy (Anderson and Krathwohl, 2001), in order to make an in-depth study of the models performance per cognitive level. This will be addressed in future work.

## Acknowledgments

This work has been funded by the European Union - NextGenerationEU through the ‘Recovery, Transformation and Resilience Plan’, by the *Ministerio para la transformación digital y de la función pública* and by UNED via cooperation agreement C039-21OT. However, the views and opinions expressed are solely those of the author(s) 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.

## References

- Achiam, Josh and Adler, Steven and Agarwal, Sandhini and Ahmad, Lama and Akkaya, Ilge and Aleman et al. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Ahuja, Kabir and Diddee, Harshita and Hada, Rishav and Ochieng, Millicent and Ramesh, Krithika and Jain, Prachi and Nambi, Akshay and Ganu, Tanuja and Segal, Sameer and Axmed, Maxamed and Bali, Kalika and Sitaram, Sunayana. 2023. [MEGA: Multilingual Evaluation of Generative AI](#). *arXiv preprint*. ArXiv:2303.12528 [cs].
- Anderson, Lorin W. and Krathwohl, David R. 2001. *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom’s Taxonomy of Educational Objectives*. Longman.
- Rohan Anil and Sebastian Borgeaud and Jean-Baptiste Alayrac and Jiahui Yu and Radu Soricut and Johan Schalkwyk and Andrew M. Dai and Anja Hauth et al. 2024. [Gemini: A family of highly capable multimodal models](#). *Preprint*, arXiv:2312.11805.
- Anthropic. 2024. [The Claude 3 Model Family: Opus, Sonnet, Haiku](#).
- Artetxe, Mikel and Ruder, Sebastian and Yogatama, Dani. 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.
- Jacob Austin and Augustus Odena and Maxwell Nye and Maarten Bosma and Henryk Michalewski and David Dohan and Ellen Jiang and Carrie Cai and Michael Terry and Quoc Le and Charles Sutton. 2021. [Program synthesis with large language models](#). *Preprint*, arXiv:2108.07732.
- Bang, Yejin and Cahyawijaya, Samuel and Lee, Nayeon and Dai, Wenliang and Su, Dan and Wilie, Bryan and Lovenia, Holy and Ji, Ziwei and Yu, Tiezheng and Chung, Willy and Do, Quyet V. and Xu, Yan and Fung, Pascale. 2023. [A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity](#). In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–718, Nusa Dua, Bali. Association for Computational Linguistics.
- Bender, Emily M. and Friedman, Batya. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Yonatan Bisk and Rowan Zellers and Ronan Le Bras and Jianfeng Gao and Yejin Choi. 2019. [Piqa: Reasoning about physical commonsense in natural language](#). *Preprint*, arXiv:1911.11641.

- Damián Blasi and Antonios Anastasopoulos and Graham Neubig. 2021. [Systematic inequalities in language technology performance across the world’s languages](#). *Preprint*, arXiv:2110.06733.
- Benjamin S. Bloom, editor. 1956. *Taxonomy of educational objectives: The classification of educational goals. Handbook I: Cognitive domain*. Longmans, Green, New York.
- Tom B. Brown and Benjamin Mann and Nick Ryder and Melanie Subbiah and Jared Kaplan and Prafulla Dhariwal and Arvind Neelakantan and Pranav Shyam and Girish Sastry and Amanda Askell and Sandhini Agarwal and Ariel Herbert-Voss and Gretchen Krueger and Tom Henighan and Rewon Child and Aditya Ramesh and Daniel M. Ziegler and Jeffrey Wu and Clemens Winter and Christopher Hesse and Mark Chen and Eric Sigler and Mateusz Litwin and Scott Gray and Benjamin Chess and Jack Clark and Christopher Berner and Sam McCandlish and Alec Radford and Ilya Sutskever and Dario Amodei. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.
- Mark Chen and Jerry Tworek and Heewoo Jun and Qiming Yuan and Henrique Ponde de Oliveira Pinto and Jared Kaplan and Harri Edwards and Yuri Burda et al. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- Wei-Lin Chiang and Lianmin Zheng and Ying Sheng and Anastasios Nikolas Angelopoulos and Tianle Li and Dacheng Li and Hao Zhang and Banghua Zhu and Michael Jordan and Joseph E. Gonzalez and Ion Stoica. 2024. [Chatbot arena: An open platform for evaluating llms by human preference](#). *Preprint*, arXiv:2403.04132.
- Eunsol Choi and He He and Mohit Iyyer and Mark Yatskar and Wen-tau Yih and Yejin Choi and Percy Liang and Luke Zettlemoyer. 2018. [Quac : Question answering in context](#). *Preprint*, arXiv:1808.07036.
- Clark, Peter and Cowhey, Isaac and Etzioni, Oren and Khot, Tushar and Sabharwal, Ashish and Schoenick, Carissa and Tafjord, Oyvind. 2018. [Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge](#). *ArXiv*.
- Cobbe, Karl and Kosaraju, Vineet and Bavarian, Mohammad and Chen, Mark and Jun, Heewoo and Kaiser, Lukasz and Plappert, Matthias and Tworek, Jerry and Hilton, Jacob and Nakano, Reiichiro and Hesse, Christopher and Schulman, John. 2021. [Training Verifiers to Solve Math Word Problems](#). *arXiv preprint*. ArXiv:2110.14168 [cs] version: 2.
- Conneau, Alexis and Rinott, Ruty and Lample, Guillaume and Williams, Adina and Bowman, Samuel and Schwenk, Holger and Stoyanov, Veselin. 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.
- Devlin, Jacob and Chang, Ming-Wei and Lee, Kenton and Toutanova, Kristina. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yihong Dong and Xue Jiang and Huanyu Liu and Zhi Jin and Bin Gu and Mengfei Yang and Ge Li. 2024. [Generalization or memorization: Data contamination and trustworthy evaluation for large language models](#). *Preprint*, arXiv:2402.15938.
- Shahriar Golchin and Mihai Surdeanu. 2024. [Time travel in llms: Tracing data contamination in large language models](#). *Preprint*, arXiv:2308.08493.
- María Grandury and Javier Aula-Blasco and Clémentine Fourrier and Miguel González and Gonzalo Martínez and Gonzalo Santamaría and Alejandro Vaca. 2024. [La leaderboard: Leaderboard of spanish varieties and official languages](#). <https://huggingface.co/spaces/la-leaderboard/la-leaderboard>.
- Zhouhong Gu and Xiaoxuan Zhu and Haoning Ye and Lin Zhang and Jianchen Wang et al. 2024. [Xiezhi: An ever-updating benchmark for holistic domain knowledge evaluation](#). *Preprint*, arXiv:2306.05783.
- Hardalov, Momchil and Mihaylov, Todor and Zlatkova, Dimitrina and Dinkov, Yoan and Koychev, Ivan and Nakov, Preslav. 2020. [EXAMS: A multi-subject high school examinations dataset for cross-lingual and multilingual question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5427–5444, Online. Association for Computational Linguistics.
- Hendrycks, Dan and Burns, Collin and Basart, Steven and Zou, Andy and Mazeika, Mantas and Song, Dawn and Steinhardt, Jacob. 2021a. [Measuring Massive Multitask Language Understanding](#). *arXiv preprint*. ArXiv:2009.03300 [cs].
- Dan Hendrycks and Collin Burns and Saurav Kadavath and Akul Arora and Steven Basart and Eric Tang and Dawn Song and Jacob Steinhardt. 2021b. [Measuring mathematical problem solving with the math dataset](#). *Preprint*, arXiv:2103.03874.
- Junjie Hu and Sebastian Ruder and Aditya Siddhant and Graham Neubig and Orhan Firat and Melvin Johnson. 2020. [Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalization](#). *Preprint*, arXiv:2003.11080.
- Jiang, Albert Q. and Sablayrolles, Alexandre and Mensch, Arthur and Bamford, Chris and Chaplot, Devendra Singh and Casas, Diego de las and Bressand, Florian others. 2023. [Mistral 7B](#). *arXiv preprint*. ArXiv:2310.06825 [cs].



- Albert Q. Jiang and Alexandre Sablayrolles and Antoine Roux and Arthur Mensch and Blanche Savary and Chris Bamford and Devendra Singh Chaplot and Diego de las Casas and Emma Bou Hanna and Florian Bressand and Gianna Lengyel and Guillaume Bour and Guillaume Lample and L lio Renard Lavaud and Lucile Saulnier and Marie-Anne Lachaux and Pierre Stock and Sandeep Subramanian and Sophia Yang and Szymon Antoniak and Teven Le Scao and Th ophile Gervet and Thibaut Lavril and Thomas Wang and Timoth e Lacroix and William El Sayed. 2024a. [Mixtral of experts](#). *Preprint*, arXiv:2401.04088.
- Minhao Jiang and Ken Ziyu Liu and Ming Zhong and Rylan Schaeffer and Siru Ouyang and Jiawei Han and Sanmi Koyejo. 2024b. [Investigating data contamination for pre-training language models](#). *ArXiv*, abs/2401.06059.
- Joshi, Mandar and Choi, Eunsol and Weld, Daniel and Zettlemoyer, Luke. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Kwiatkowski, Tom and Palomaki, Jennimaria and Redfield, Olivia and Collins, Michael et al. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Lai, Guokun and Xie, Qizhe and Liu, Hanxiao and Yang, Yiming and Hovy, Eduard. 2017. [RACE: Large-scale ReAding Comprehension Dataset From Examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Lai, Viet and Ngo, Nghia and Pourn Ben Veyseh, Amir and Man, Hieu and Dernoncourt, Franck and Bui, Trung and Nguyen, Thien. 2023. [ChatGPT beyond English: Towards a comprehensive evaluation of large language models in multilingual learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13171–13189, Singapore. Association for Computational Linguistics.
- Zihao Li and Yucheng Shi and Zirui Liu and Fan Yang and Ali Payani and Ninghao Liu and Mengnan Du. 2024. [Quantifying multilingual performance of large language models across languages](#). *Preprint*, arXiv:2404.11553.
- Liang, Percy and Bommasani, Rishi and Lee, Tony and Tsipras, Dimitris and Soylu, Dilara and Yasunaga, Michihiro and Zhang, Yian and Narayanan, Deepak et al. 2023. [Holistic Evaluation of Language Models](#). *arXiv preprint*. ArXiv:2211.09110 [cs].
- Thomas Mesnard and Cassidy Hardin and Robert Dadashi and Surya Bhupatiraju and Shreya Pathak and Laurent Sifre and Morgane Riviere et al. 2024. [Gemma: Open models based on gemini research and technology](#). *Preprint*, arXiv:2403.08295.
- Meta. 2024. [Introducing meta llama 3: The most capable openly available llm to date](#).
- Mihaylov, Todor and Clark, Peter and Khot, Tushar and Sabharwal, Ashish. 2018. [Can a Suit of Armor Conduct Electricity? A New Dataset for Open Book Question Answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium. Association for Computational Linguistics.
- Moran Mizrahi and Guy Kaplan and Dan Malkin and Rotem Dror and Dafna Shahaf and Gabriel Stanovsky. 2024. [State of what art? a call for multi-prompt llm evaluation](#). *Preprint*, arXiv:2401.00595.
- OpenAI. 2024. [Mmmlu](#). Hugging Face Datasets.
- Felipe Maia Polo and Ronald Xu and Lucas Weber and M rian Silva and Onkar Bhardwaj and Leshem Choshen and Allysson Flavio Melo de Oliveira and Yuekai Sun and Mikhail Yurochkin. 2024. [Efficient multi-prompt evaluation of llms](#). *Preprint*, arXiv:2405.17202.
- Tanmay Rajore and Nishanth Chandran and Sunayana Sitaram and Divya Gupta and Rahul Sharma and Kashish Mittal and Manohar Swaminathan. 2024. [Truce: Private benchmarking to prevent contamination and improve comparative evaluation of llms](#). *Preprint*, arXiv:2403.00393.
- Pranav Rajpurkar and Jian Zhang and Konstantin Lopyrev and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#). *Preprint*, arXiv:1606.05250.
- Rein, David and Hou, Betty Li and Stickland, Asa Cooper and Petty, Jackson and Pang, Richard Yuanzhe and Dirani, Julien and Michael, Julian and Bowman, Samuel R. 2023. [GPQA: A Graduate-Level Google-Proof Q&A Benchmark](#). *arXiv preprint*. ArXiv:2311.12022 [cs].
- Sebastian Ruder and Noah Constant and Jan Botha and Aditya Siddhant and Orhan Firat and Jinlan Fu and Pengfei Liu and Junjie Hu and Dan Garrette and Graham Neubig and Melvin Johnson. 2021. [Xtreme-r: Towards more challenging and nuanced multilingual evaluation](#). *Preprint*, arXiv:2104.07412.
- Sainz, Oscar and Campos, Jon and Garc a-Ferrero, Iker and Etxaniz, Julen and de Lacalle, Oier Lopez and Agirre, Eneko. 2023. [NLP evaluation in trouble: On the need to measure LLM data contamination for each benchmark](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10776–10787, Singapore. Association for Computational Linguistics.



- Sakaguchi, Keisuke and Bras, Ronan Le and Bhagavatula, Chandra and Choi, Yejin. 2019. [WinoGrande: An Adversarial Winograd Schema Challenge at Scale](#). *arXiv preprint*. ArXiv:1907.10641 [cs].
- Sap, Maarten and Rashkin, Hannah and Chen, Derek and Le Bras, Ronan and Choi, Yejin. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Melanie Sclar and Yejin Choi and Yulia Tsvetkov and Alane Suhr. 2023. [Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting](#). *Preprint*, arXiv:2310.11324.
- Srivastava, Aarohi and Rastogi, Abhinav and Rao, Abhishek and Shoeb, Abu Awal Md and Abid, Abubakar and Fisch, Adam and Brown, Adam R. et al. 2022. [Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models](#). *arXiv preprint*. ArXiv:2206.04615 [cs, stat].
- Srivastava, Saurabh and B, Annarose M. and P V, Anto and Menon, Shashank and Sukumar, Ajay and T, Adwaith Samod and Philipose, Alan and Prince, Stevin and Thomas, Sooraj. 2024. [Functional Benchmarks for Robust Evaluation of Reasoning Performance, and the Reasoning Gap](#). *arXiv preprint*. ArXiv:2402.19450 [cs].
- Mirac Suzgun and Nathan Scales and Nathanael Schärli and Sebastian Gehrmann and Yi Tay and Hyung Won Chung and Aakanksha Chowdhery and Quoc V. Le and Ed H. Chi and Denny Zhou and Jason Wei. 2022. [Challenging big-bench tasks and whether chain-of-thought can solve them](#). *Preprint*, arXiv:2210.09261.
- Talmor, Alon and Herzig, Jonathan and Lourie, Nicholas and Berant, Jonathan. 2019. [CommonsenseQA: A question answering challenge targeting commonsense knowledge](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Touvron, Hugo and Lavril, Thibaut and Izacard, Gautier and Martinet, Xavier and Lachaux, Marie-Anne and Lacroix, Timothée and Rozière, Baptiste and Goyal, Naman and Hambro, Eric and Azhar, Faisal and Rodriguez, Aurelien and Joulin, Armand and Grave, Edouard and Lample, Guillaume. 2023a. [LLaMA: Open and Efficient Foundation Language Models](#). *arXiv preprint*. ArXiv:2302.13971 [cs].
- Touvron, Hugo and Martin, Louis and Stone, Kevin and Albert, Peter and Almahairi, Amjad and Babaei, Yasmine and Bashlykov, Nikolay and Batra, Soumya and Bhargava, Prajjwal et al. 2023b. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#). *arXiv preprint*. ArXiv:2307.09288 [cs].
- Ashish Vaswani and Noam Shazeer and Niki Parmar and Jakob Uszkoreit and Llion Jones and Aidan N. Gomez and Lukasz Kaiser and Illia Polosukhin. 2023. [Attention is all you need](#). *Preprint*, arXiv:1706.03762.
- Alex Wang and Yada Pruksachatkun and Nikita Nangia and Amanpreet Singh and Julian Michael and Felix Hill and Omer Levy and Samuel R. Bowman. 2020. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). *Preprint*, arXiv:1905.00537.
- Wang, Alex and Singh, Amanpreet and Michael, Julian and Hill, Felix and Levy, Omer and Bowman, Samuel. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Shuo Yang and Wei-Lin Chiang and Lianmin Zheng and Joseph E. Gonzalez and Ion Stoica. 2023. [Rethinking benchmark and contamination for language models with rephrased samples](#). *Preprint*, arXiv:2311.04850.
- Jifan Yu and Xiaozhi Wang and Shangqing Tu and Shulin Cao and Daniel Zhang-Li and Xin Lv and Hao Peng and Zijun Yao et al. 2023. [Kola: Carefully benchmarking world knowledge of large language models](#). *Preprint*, arXiv:2306.09296.
- Rowan Zellers and Ari Holtzman and Yonatan Bisk and Ali Farhadi and Yejin Choi. 2019. [Hellaswag: Can a machine really finish your sentence?](#) *Preprint*, arXiv:1905.07830.
- Zhang, Wenxuan and Aljunied, Sharifah Mahani and Gao, Chang and Chia, Yew Ken and Bing, Lidong. 2023. [M3Exam: A Multilingual, Multimodal, Multilevel Benchmark for Examining Large Language Models](#). *arXiv preprint*. ArXiv:2306.05179 [cs].
- Zhong, Wanjun and Cui, Ruixiang and Guo, Yiduo and Liang, Yaobo and Lu, Shuai and Wang, Yanlin and Saied, Amin and Chen, Weizhu and Duan, Nan. 2023. [AGIEval: A Human-Centric Benchmark for Evaluating Foundation Models](#). *arXiv preprint*. ArXiv:2304.06364 [cs].

## A UNED-ACCESS 2024: Additional Information

This appendix provides more details about the collection, transcription and annotation process of the dataset UNED-ACCESS 2024, as well as an example of the format of the dataset instances, and the dataset statement.

### A.1 Data Harvesting and Preparation

**Collection of the exams.** The exams have been downloaded directly from the UNED repository of exams, accessible for UNED users at [blindlink](#). The repository contains exams for Bachelor’s, Master’s, Languages, and Access Courses for 25 and 45 year olds. In this first version of the dataset, only exams from the Access courses for over 25s have been used, as they are a smaller set with a wide distribution of subjects. The availability of exams varies from subject to subject and, in particular, the availability of multiple-choice exams depends on the subject in question (some subjects have only, or mostly, essay questions). In the case of some subjects, such as Advanced Mathematics, no more questions were available than those included in the dataset. The exams have been downloaded between October and November 2023 and the questions obtained correspond to exams between 2009 and 2022.

**Sampling and filtering.** Questions selected meet the following two conditions:

- Availability of private correct answers. We have selected only exam questions for which templates with “official” solutions were provided to us by UNED. These templates exist for a subset of the exams contained in the repository. Thus, we first filter out exams for which no templates are provided.
- Multiple-choice questions. Once the exams for which a solution template is available have been filtered, multiple-choice exam question have been selected. Not all subjects are evaluated with this type of exams.

**Preprocessing.** The first step to create the dataset consisted in matching the PDF exams with their corresponding answer templates provided by UNED. Exams available in the private repository of the UNED were in PDF format, so they have been transcribed using Nougat OCR<sup>4</sup> in Python. This

<sup>4</sup><https://facebookresearch.github.io/nougat/>

OCR system was chosen because it is open-source and based on Transformers (Vaswani et al., 2023), and allows the transcription of formulas in  $\text{\LaTeX}$ , a functionality especially useful for the transcription of exams in subjects such as Mathematics or Biochemistry. After conversion, the exams were structured into CSV format with a Python script and questions and answers were matched. Due to variations in exam formats, each exam required individual treatment. Finally, duplicate questions, chained questions, questions with images, and those marked with an ‘X’ in the answer template were removed. The original Spanish questions were then manually translated into English by a professional translator, who was instructed not to use any machine translation tool nor generative models such as ChatGPT. Doubts about the translation of terminology were discussed with the authors of the paper. The translator was also instructed to work locally and to delete the files once the work was finished.

Subsequently, the plain text questions have been structured in CSV and JSON format semi-automatically by means of a Python script using libraries such as `re` to identify regular expressions. Thus, we have a bilingual dataset with the following fields:

- **question.** Contains the statement of the question.
- **A.** Contains the text of answer option A.
- **B.** Contains the text of answer option B.
- **C.** Contains the text of answer option C.
- **D.** Contains the text of the answer option D, and in case it does not exist it is empty (NaN in CSV format or null in JSON format).
- **solution.** It contains the letter of the correct answer, extracted from the solution templates.
- **year.** Year of the examination.
- **test\_name.** This is the name of the PDF file of the exam, which contains information such as the subject code, year, month and week of the exam.
- **code.** Subject code.
- **subject.** Name of the subject.
- **id.** A unique identifier assigned to each question, consisting of the name of the test (`test_name`) and a number indicating the order of the question within the subject. This identifier matches that of its corresponding English translation.

## A.2 Dataset Statement

The characteristics of the dataset are listed in the form of a dataset statement (Bender and Friedman, 2018):

- **Type of instances:** Exam questions.
- **Type of questions:** Most questions consist on choosing the correct answer, and some consist on filling the blank/s in the statement (see example in Appendix A.3). Literature and Psychology questions primarily test encyclopedic knowledge. Spanish Language questions range from testing spelling, lexical, syntactic and grammatical knowledge. Business and Administration Management and Economics questions cover basic financial, accounting, and business management concepts, and few calculus questions. Fundamentals of Computer Science’s are mostly memorization questions, although some require to perform some calculations. Biology questions focus mostly on concepts and definitions, with some of them requiring some reasoning, for example when applying genetics concepts. Biochemistry questions require mainly memorization, and many of them include chemical formulas. Mathematics Applied to Social Sciences questions contain basic analysis, statistics, logic, probability and other arithmetic problems. Mathematics questions cover analysis, algebra, combinatorics, probability, and others, and so does Advanced Mathematics with questions in a more advanced level of calculus, algebra, and trigonometry. Most of the mathematics questions include equations, polynomials or sets expressed in  $\LaTeX$  formulas, and require reasoning to solve the problems.
- **Institution providing exams:** UNED, Spain.
- **Task:** Multiple-choice Question Answering.
- **Curation rationale:** Obtaining multiple choice questions to automatically evaluate LLMs, without diagrams or images.
- **Domain:** Academic.
- **Difficulty level:** University entrance.
- **Knowledge areas:** Business Administration and Management, Biology, Biochemistry, Economics, Fundamentals of Computer Science, Spanish Language, Literature, Mathematics, Mathematics Applied to Social Sci-

ences, Advanced Mathematics and Psychology.

- **Languages and varieties:** Standard Spanish and Standard English.
- **Number of questions:** 1003 per language.
- **Year of publication of the exams:** From 2009 to 2022.
- **Total word count:** Spanish: 30094; English: 28186.
- **Partitions:** Test.
- **Formats:** CSV, JSON.
- **Year of creation of the dataset:** 2024.
- **Further considerations:**
  - The selected exams are a partial sample of the total number of exams available in the repository.
  - The dataset is self-contained and needs no external sources.
  - The dataset is of small size (1003 questions) designed as a test set to evaluate generative models in a zero-shot setting, without model training.
  - The dataset has been manually checked to minimize spelling, OCR, and  $\LaTeX$  errors. Duplicates, image-based questions, and those dependent on other answers have been removed.
  - The dataset is kept private as it includes UNED exams with solution keys, to prevent answers from becoming public and contaminating generative models.
  - Offensive/sensitive data: none.
  - The dataset will remain private and won’t be distributed, as it’s intended for use in the ODESIA Benchmark for evaluating language models in English and Spanish.
  - No maintenance of the dataset is foreseen.
  - Extensions of the dataset are foreseen, possibly with undergraduate exams in several disciplines.

## A.3 Format of the Dataset

The dataset is presented in both CSV and JSON format, structured into the fields: question, (option) A, (option) B, (option) C, (option) D, solution, year (of exam),

test name, code, subject, id. Below is an example of a dataset object in JSON format where we omit the letter of the correct solution:

```
{
  "question": "The empirical
    testing procedure, which is
    carried out by means of ...
    is called the ... method.",
  "A": "deducing conclusions;
    experimental",
  "B": "experimentation;
    experimental",
  "C": "the correlation of
    variables; experimental",
  "D": "experimentation;
    correlational",
  "solution": "correct answer",
  "year": 2014,
  "test_name": "E000012060A14F1",
  "code": 1206,
  "subject": "Psychology",
  "id": "E000012060A14F1-178"
}
```

Listing 1: Example of JSON object (question) in the UNED-ACCESS 2024 dataset.

## B Experimental Setup

### B.1 Model Access

The proprietary models GPT-4-Turbo, GPT-4o, and GPT-3.5-Turbo were accessed via the OpenAI API, while Claude-3-Opus was available via Anthropic’s API. Open-source models were available from HuggingFace (Llama-2-7B<sup>5</sup>, Llama-3-8B<sup>6</sup>, Mistral-7B<sup>7</sup>, Gemma-7B<sup>8</sup> and Llama-3-2B<sup>9</sup>) or deployed via the Ollama library (Llama-3-70B<sup>10</sup>, Mistral-8x7B<sup>11</sup> and Gemma-2-27B<sup>12</sup>).

### B.2 Instruction Format Templates

Instruction following models often require instruction templates that format system, query and assistant messages with specific tags, available in the official documentation of each model. We have used this templates only when deploying models via Huggingface (all of them except proprietary

<sup>5</sup><https://huggingface.co/google/gemma-7b-it>

<sup>6</sup><https://huggingface.co/meta-llama/Meta-Llama-3-8B-Instruct>

<sup>7</sup><https://huggingface.co/mistralai/Mixtral-8x7B-Instruct-v0.1>

<sup>8</sup><https://huggingface.co/google/gemma-7b-it>

<sup>9</sup><https://huggingface.co/LenguajeNaturalAI/leniachat-gemma-2b-v0>

<sup>10</sup><https://ollama.com/library/llama3:70b-instruct>

<sup>11</sup><https://ollama.com/library/mixtral:instruct>

<sup>12</sup><https://ollama.com/library/gemma2:27b-instruct-fp16>

models and the big ones deployed with Ollama, since ollama already includes these templates).

```
def _instruction_format(self,
    sys_message: str, query: str,
    assistant: str):
    if 'gemma' in self.model_id:
        return f'<start_of_turn>user
            {sys_message}
            {query}<end_of_turn>
            <start_of_turn>model
            {assistant}'

    elif 'mistral' in self.model_id or '
    mixtral' in self.model_id:
        return f'<s> [INST] {sys_message}
            } [/INST]\nUser: {query}\
            nAssistant: {assistant}'

    elif 'llama-2' in self.model_id:
        return f'<s>[INST] <<SYS>>
            {sys_message}
            <</SYS>>
            \n{query}</s>[/INST]
            \n<s>{assistant}</s>'

    elif 'llama-3' in self.model_id:
        return f'<|begin_of_text|><|
            start_header_id|>system<|
            end_header_id|>
            {sys_message}<|eot_id
            |><|start_header_id
            |>user<|
            end_header_id|>
            {query}<|eot_id|><|
            start_header_id|>
            assistant<|
            end_header_id|>
            {assistant}'
```

Listing 2: Python function showing instruction format templates for models that require so.

## C Results

### C.1 Results on UNED-ACCESS 2024 by Subject

Table 6 provides detailed results of the experiments, per language, model and discipline.

### C.2 Results on MMLU-10 by Subject

Table 7 provides the evaluation results on the MMLU-10 dataset per subject.



	BAM	Biology	Biochemistry	Economics	F. of Computing	Spanish Language	Literature	Mathematics	Math Applied to SS	Advanced Math	Psychology	Average
SPANISH												
Claude-3-Opus	<b>0.86</b>	0.96	<b>1.00</b>	0.90	<b>0.98</b>	0.67	<b>0.84</b>	<b>0.69</b>	<b>0.63</b>	0.44	0.89	0.81
GPT-4-Turbo	0.78	0.96	<b>1.00</b>	<b>0.95</b>	0.96	0.69	0.72	0.55	0.57	<b>0.50</b>	0.88	0.78
GPT-4o	0.84	<b>0.97</b>	<b>1.00</b>	0.90	0.96	<b>0.74</b>	0.81	0.51	0.38	<b>0.50</b>	<b>0.91</b>	0.77
Llama-3-70B-Instruct	0.83	0.89	0.95	0.79	0.94	0.39	0.66	0.38	0.46	0.25	0.82	0.67
Gemma-2-27B-Instruct	0.76	0.92	<b>1.00</b>	0.79	0.94	0.50	0.53	0.34	0.33	0.38	0.80	0.66
Mixtral-8x7B-Instruct	0.72	0.84	0.87	0.58	0.87	0.32	0.52	0.32	0.23	0.25	0.78	0.57
GPT-3.5-Turbo	0.64	0.80	0.90	0.53	0.87	0.32	0.44	0.20	0.20	0.38	0.74	0.55
Llama-3-8B-Instruct	0.57	0.71	0.82	0.56	0.79	0.26	0.37	0.22	0.30	0.25	0.67	0.50
Mistral-7B-Instruct	0.52	0.67	0.72	0.42	0.77	0.25	0.40	0.05	0.27	0.06	0.62	0.43
Gemma-7B-It	0.40	0.63	0.77	0.32	0.66	0.12	0.36	0.12	0.14	0.06	0.58	0.38
Llama-2-7B-Chat	0.29	0.41	0.31	0.19	0.56	0.12	0.34	0.14	0.12	-0.12	0.44	0.25
Leniachat-Gemma-2B	0.19	0.21	0.03	0.06	0.15	0.05	0.22	-0.03	0.20	-0.12	0.24	0.11
ENGLISH												
Claude-3-Opus	<b>0.81</b>	0.95	<b>1.00</b>	0.90	<b>0.96</b>	0.67	0.78	<b>0.63</b>	<b>0.66</b>	0.44	0.84	0.79
GPT-4o	0.79	0.96	<b>1.00</b>	<b>0.92</b>	<b>0.96</b>	<b>0.70</b>	<b>0.82</b>	0.55	0.52	<b>0.50</b>	<b>0.86</b>	0.78
GPT-4-Turbo	0.78	<b>0.97</b>	<b>1.00</b>	<b>0.92</b>	0.94	0.67	0.74	0.55	0.51	<b>0.50</b>	0.83	0.76
Llama-3-70B-Instruct	0.74	0.90	<b>1.00</b>	0.82	0.92	0.33	0.60	0.34	0.44	0.19	0.82	0.65
Gemma-2-27B-Instruct	0.72	0.94	<b>1.00</b>	0.79	0.87	0.50	0.55	0.28	0.33	0.25	0.81	0.64
GPT-3.5-Turbo	0.67	0.84	0.95	0.61	0.89	0.36	0.56	0.28	0.17	0.50	0.73	0.60
Mixtral-8x7B-Instruct	0.71	0.81	0.92	0.61	0.87	0.32	0.52	0.22	0.33	0.13	0.73	0.56
Llama-3-8B-Instruct	0.52	0.77	0.90	0.61	0.79	0.38	0.43	0.20	0.28	0.13	0.67	0.51
Mistral-7B-Instruct	0.57	0.71	0.82	0.63	0.77	0.23	0.36	0.05	0.23	-0.00	0.65	0.46
Gemma-7B-It	0.41	0.67	0.85	0.56	0.75	0.18	0.22	0.20	0.14	-0.06	0.61	0.41
Llama-2-7B-Chat	0.43	0.62	0.39	0.27	0.62	0.15	0.30	0.12	0.15	-0.00	0.48	0.32
Leniachat-Gemma-2B	0.29	0.32	0.24	0.03	0.13	0.05	0.22	0.08	0.07	-0.06	0.27	0.15

Table 6: Cohen’s Kappa results on UNED-ACCESS 2024 by model and subject in Spanish and English, sorted by average. In bold, the best model(s) per subject per language.

	High School Biology	High School Chemistry	High School Microeconomics	High School Macroeconomics	High School Computer Science	High School World History	Elementary Mathematics	High School Mathematics	High School Physics	High School Psychology	Average
GPT-4o	<b>0.94</b>	<b>0.72</b>	<b>0.95</b>	<b>0.88</b>	<b>0.89</b>	<b>0.93</b>	0.64	0.33	<b>0.62</b>	<b>0.94</b>	<b>0.79</b>
Claude-3-Opus	0.91	0.67	0.89	0.81	<b>0.89</b>	0.91	<b>0.86</b>	<b>0.40</b>	0.56	0.93	0.78
GPT-4-Turbo	0.90	0.65	0.92	0.83	0.88	0.89	0.62	0.36	0.57	0.93	0.75
Gemma2-27B-Instruct	0.91	0.62	0.84	0.79	0.77	0.89	0.44	0.29	0.43	0.91	0.69
Llama3-70B-Instruct	0.83	0.47	0.84	0.73	0.76	0.82	0.63	0.29	0.46	0.90	0.67
Mixtral-7B-Instruct	0.75	0.42	0.64	0.57	0.59	0.82	0.30	0.17	0.30	0.81	0.54
GPT-3.5-Turbo	0.72	0.38	0.63	0.55	0.64	0.74	0.35	0.15	0.20	0.81	0.52
Llama-3-8B-Instruct	0.71	0.36	0.64	0.53	0.57	0.75	0.25	0.04	0.23	0.79	0.49
Mistral-7B-Instruct	0.56	0.24	0.50	0.37	0.45	0.67	0.17	0.11	0.11	0.71	0.39
Gemma-7B-It	0.47	0.15	0.36	0.33	0.31	0.50	-0.12	-0.25	-0.05	0.63	0.23
Llama-2-7B-Chat	0.22	0.09	0.10	0.10	0.13	0.30	0.06	-0.04	0.01	0.38	0.13
Leniachat-Gemma-2B	0.13	0.04	0.08	0.09	0.07	0.08	0.03	-0.04	-0.04	0.21	0.06

Table 7: Cohen’s Kappa performance of models per subject on MMLU-10.