

MALAMUTE: A Multilingual, Highly-granular, Template-free, Education-based Probing Dataset

Sagi Shaier[∇] George Arthur Baker[∇] Chiranthan Sridhar[∇]
Lawrence E Hunter[†] Katharina von der Wense^{∇◇}

[∇]University of Colorado Boulder

[†]University of Chicago, Department of Pediatrics

[◇]Johannes Gutenberg University Mainz

E-mail: {sagi.shaier, george.baker, chiranthan.sridhar, katharina.kann}@colorado.edu

Abstract

Language models (LMs) have excelled in various broad domains. However, to ensure their safe and effective integration into real-world educational settings, they must demonstrate proficiency in specific, granular areas of knowledge. Existing cloze-style benchmarks, commonly used to evaluate LMs' knowledge, have three major limitations. They: 1) do not cover the educational domain; 2) typically focus on low-complexity, generic knowledge or broad domains, which do not adequately assess the models' knowledge in specific subjects; and 3) often rely on templates that can bias model predictions. Here, we introduce MALAMUTE, a multilingual, template-free, and highly granular probing dataset comprising expert-written, peer-reviewed probes from 71 university-level textbooks across three languages (English, Spanish, and Polish). MALAMUTE is the first education-based cloze-style dataset. It covers eight domains, each with up to 14 subdomains, further broken down into concepts and concept-based prompts, totaling 33,361 university curriculum concepts and 116,887 prompts. MALAMUTE's fine granularity, educational focus, and inclusion of both sentence-level and paragraph-level prompts make it an ideal tool for evaluating LMs' course-related knowledge. Our evaluation of masked and causal LMs on MALAMUTE shows that despite overall proficiency, they have significant gaps in knowledge when examined closely on specific subjects, hindering their safe use in classrooms and underscoring the need for further development. Code and data can be found at <https://github.com/Shaier/MALAMUTE>.

1 Introduction

Language models (LMs) have shown remarkable abilities in various domains (OpenAI, 2023a,b; Touvron et al., 2023). However, to ensure their safe and effective integration into real-world educational settings, they must demonstrate proficiency in specific,

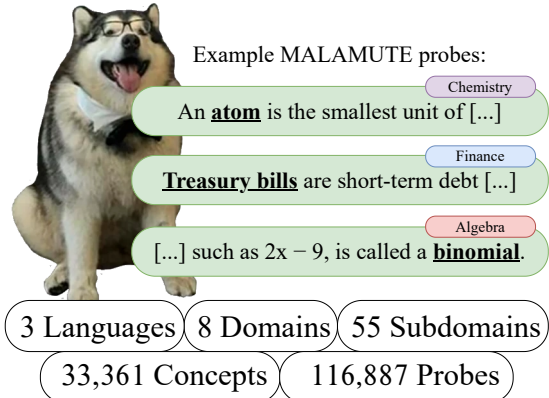


Figure 1: MALAMUTE: a highly granular cloze-style dataset for educational knowledge probing. This dataset assesses language models' knowledge across three languages with fine-grained detail. **Bold strings** indicate predictable text. Example prompts shown for 3 out of the 55 subdomains.

granular areas of knowledge (Xu et al., 2020). For instance, a LM may excel in general knowledge of mathematics, but struggle with the nuances of calculus or algebra. Hence, we need precise and targeted evaluation methods for LMs' knowledge in these critical areas.

The need for such evaluations is pressing, as LMs are increasingly being considered for use in classrooms, where their performance can have a direct and profound impact on student learning outcomes (Jošt et al., 2024; Kumar et al., 2023). Moreover, LMs are being explored for a range of educational applications, including adaptive learning (Weijers et al., 2024), intelligent tutoring systems (Nye et al., 2023), and automated grading (Xie et al., 2024b). In order to ensure that these systems are both effective and safe, we must develop evaluation methods that can accurately assess their knowledge in specific, curriculum-aligned areas.

Currently, evaluating LMs' knowledge is a significant research focus. Two prominent evaluation methods are question answering (QA) (Shaier

et al., 2024b; Li et al., 2023a; Xie et al., 2024a) and cloze-style (CS) probing (Petroni et al., 2019; Sung et al., 2021). CS probing, in particular, has gained popularity, where models receive a prompt with a masked portion, such as “Dante was born in [MASK],” and are tasked with predicting the masked string. This approach measures models’ lower bound of knowledge (Petroni et al., 2019).

However, existing CS benchmarks have limitations. Firstly, there is a notable lack of CS datasets that are specifically designed for educational settings. Secondly, the prompts in existing datasets typically focus on low-complexity, generic knowledge or broad domains, which do not adequately capture models’ knowledge in specific subjects (Petroni et al., 2019; Keleg and Magdy, 2023; Kassner and Schütze, 2020; Kassner et al., 2020; Dhingra et al., 2022). Furthermore, most existing cloze-style probing datasets rely on expert-written templates, which have been shown to bias model predictions (Cao et al., 2021a; Petroni et al., 2019; Sung et al., 2021; Shaier et al., 2024a).

To address these limitations, we introduce MALAMUTE, a Multilingual, highLy granular, teMplate-free, edUcation-based probing daTasEt. Derived from 71 university-level textbooks across three languages (English, Spanish, and Polish), MALAMUTE covers eight domains, each with up to fourteen subdomains, broken down into individual curriculum concepts and concept-centric prompts. With 33,361 curriculum concepts and 116,887 expert-written, peer-reviewed CS prompts, MALAMUTE offers an unprecedented level of granularity and alignment with educational curricula, making it an ideal tool for evaluating LMs’ knowledge in educational contexts. See Figure 1 for examples. MALAMUTE is also the first probing dataset to include both sentence-level and paragraph-level prompts, which provides a more comprehensive evaluation of LMs’ knowledge, establishing a new higher lower-bound of knowledge for masked LMs (MLMs). Lastly, MALAMUTE also comes with a number of parallel domains, subdomains, and concepts across multiple languages.

We evaluate MALAMUTE on 8 causal LMs (CLMs) and 5 MLMs and highlight the need for granular datasets: despite many models’ overall proficiency on the dataset, they have significant gaps in knowledge when examined closely on specific subjects, hindering their safe use in classrooms and underscoring the need for further development.

2 Related Work

2.1 Models’ Parametric Knowledge

LMs have been shown to contain a substantial amount of factual knowledge within their parameters, as demonstrated by Petroni et al. (2019); Sung et al. (2021); Roberts et al. (2020). This embedded knowledge enables LMs to perform tasks requiring extensive information, like open-domain question answering, without needing external databases or knowledge sources (Roberts et al., 2020). As models’ size increase, their capacity to store and retrieve vast amounts of information increases significantly (OpenAI, 2023a,b), which also allow them to encompass a wide range of knowledge types, including factual data (Petroni et al., 2019; Shaier et al., 2023b), linguistic nuances (Zhang and Choi, 2021), cultural references (Zhou et al., 2024; Keleg and Magdy, 2023), and domain-specific information (Sung et al., 2021).

Evaluating the amount and types of knowledge within these models is crucial, as it ensures that they can be effectively and reliably applied to various tasks. This assessment helps identify gaps and biases, informing improvements in training and deployment (Shaier et al., 2023a). Moreover, understanding models’ knowledge scope is crucial for ethical considerations, such as preventing misinformation and ensuring factual language generation (Wang et al., 2023; Shaier et al., 2024d).

2.2 Knowledge Probing

Knowledge probing is an approach to assess knowledge implicitly embedded in LMs’ weights. Various methods are utilized to evaluate this parametric knowledge, including closed-book QA (Shaier et al., 2024b; Li et al., 2023a; Xie et al., 2024a) and CS prompts (Petroni et al., 2019; Sung et al., 2021). Closed-book QA tests models without access to external resources, which can further be separated into those that use multiple-choice QA, or free-form QA, while CS prompts require completion of missing segments within a given context. Evaluators may further employ template-based or template-free probes to gauge comprehension, where the latter provide a more accurate assessment of LMs knowledge (Shaier et al., 2024a).

2.3 Cloze-style Knowledge Probing

CS knowledge probing has experienced a surge in popularity in recent years, resulting in the development of over 80 distinct datasets to date, span-

ning multiple languages and domains. CS probing involves presenting models with prompts that have masked portions, such as “Dante was born in [MASK]”, and tasking them with predicting the correct completion (Petroni et al., 2019, 2020; Zhou et al., 2024; Talmor et al., 2020; Sung et al., 2021; Kassner and Schütze, 2020; Kassner et al., 2020; Dhingra et al., 2022). This method is used to measure models’ lower bound of knowledge. (Petroni et al., 2019; Jiang et al., 2020).

While existing datasets span diverse domains, such as law (Chalkidis et al., 2023), geography (Bhandari et al., 2023), biomedicine (Sung et al., 2021; Shaier et al., 2024a; Meng et al., 2022), and general knowledge (Petroni et al., 2019; Cao et al., 2021b; He et al., 2024), they have several limitations. Firstly, the prompts in existing datasets typically focus on low-complexity, generic knowledge or broad domains, which do not adequately capture models’ knowledge in specific subjects. Secondly, most existing cloze-style probing datasets rely on expert-written templates, which have been shown to bias model predictions. Furthermore, there is a notable lack of CS datasets that are specifically designed for educational settings.

2.4 Educational Knowledge Probing

As LMs are increasingly being integrated into classrooms (Kazemitabaar et al., 2024; Cao et al., 2023), it is essential to have a reliable educational knowledge-probing dataset to assess their performance, which can significantly impact student learning outcomes. However, current educational probing datasets have several limitations.

One major concern is that some datasets rely on expert-written templates (Koto et al., 2024), which can introduce bias into model predictions. Furthermore, these datasets usually have a limited scope, focusing on high-level subjects (such as American Government and U.S. history (Ciosici et al., 2021)) rather than providing the fine-grained, concept-level details that are necessary for in-depth knowledge. This is particularly problematic in education, where many concepts comprise various subdomains that require comprehensive assessment.

Another limitation is that these datasets have a limited number of prompts, restricting their ability to comprehensively evaluate LMs. Additionally, they are not based on CS prompts, instead relying on either multiple-choice QA or closed-book settings. However, multiple-choice QA is a significantly less challenging task due to the limited

number of options that can skew knowledge assessment results and can be subject to models’ biases towards lettered answer options in causal setups (Zheng et al., 2023). Closed-book settings, on the other hand, are not suitable for MLMs without additional fine-tuning or mask tokens, which can lead to grammatical issues. Moreover, evaluating free-form text generation often relies on human assessment to fully capture its correctness, as automatic metrics may not fully capture the nuances of the generated output, which can be time-consuming and subjective. Hence, there is a pressing need for a more comprehensive and nuanced CS educational knowledge-probing dataset. A summary of such probing datasets can be found in Table 1.

3 MALAMUTE

In this section, we introduce MALAMUTE, the first CS educational knowledge probing dataset. We also describe our systematic and scalable approach to generating high-quality probing prompts directly from reputable academic sources. See Table 2 for prompt examples from each of the 8 domains and 55 subdomains in MALAMUTE.

3.1 Data Source: OpenStax

OpenStax (OpenStax, 2012), a project led by Rice University, is dedicated to creating and distributing high-quality, peer-reviewed, and open-source textbooks for popular university-level general education courses. These textbooks are made available under a Creative Commons license, allowing for free use and adaptation.

The OpenStax library currently contains an impressive collection of 70 English textbooks, 11 Spanish textbooks, and 6 Polish textbooks. All textbooks can be accessed in multiple formats: print, PDF, and interactive web pages. Moreover, the library is continuously expanding, with new subjects, textbook improvements, and language translations being added regularly. **This ongoing development ensures that our dataset will remain extensible and adaptable in the future.**

The textbooks are structured in a hierarchical system, with eight main categories (or “domains”): Business, College Success, Computer Science, Humanities, Math, Nursing, Science, and Social Sciences. Each domain is further divided into subjects, which we refer to as “subdomains”. These subdomains roughly correspond to the content of a single general education college course. Finally, each

Dataset	Template-free	Educational Focus	Cloze-style	Domains	Prompts	Languages	Domain	Granularity
LAMA (Petroni et al., 2019)	Both	No	Yes	1*	51k	1		Low
IndicGLUE [†] (Kakwani et al., 2020)	Yes	No	Yes	1*	239k	11		Low
mLAMA (Kassner et al., 2021)	No	No	Yes	1*	855k	53		Low
MedLAMA (Meng et al., 2022)	No	No	Yes	1	19k	1		Low
BioLAMA (Sung et al., 2021)	No	No	Yes	1	49k	1		Low
LegalLAMA (Chalkidis et al., 2023)	Yes	No	Yes	1	27k	1		Low
NumerSense (Lin et al., 2020)	No	No	Yes	8	3k	1		Medium
LEFT (Ciosici et al., 2021)	Yes	Yes	No	2	1k	1		Medium
EXAMS [†] (Hardalov et al., 2020)	Yes	Yes	No	3	14k	16		High
ArabicMMLU [†] (Koto et al., 2024)	No	Yes	No	5	15k	2		High
MALAMUTE	Yes	Yes	Yes	8	116k	3		Very High

Table 1: Comparison of recent knowledge probing datasets. Low domain granularity means a dataset is not divided into domains, though it may have multiple tasks or sources. Medium granularity involves several domains without subdomains. High granularity includes domains and subdomains, while very high granularity adds a further division into concepts. * These datasets cover broad general knowledge and do not correspond to a specific domain. [†] multiple-choice QA.

subdomain contains numerous *concepts*, which are linked to specific sections within the relevant textbooks where those concepts are discussed.

3.2 Dataset Creation

3.2.1 Data Extraction

We collect MALAMUTE by first retrieving a list of textbook URLs from the OpenStax sitemap. We then drop remedial instances such as earlier editions and high-school level textbooks, as their contents are largely redundant, and compile a list of index pages. Next, we iterate over the vocabulary terms in the index pages and scrape the contents of the pages associated with the hypertext reference attributes. For each page, we save any parent object of a span with the attribute “data-type” equal to “term”, which generally corresponds to the surrounding paragraph. We exclude rare cases that result in prompts consisting of only the concept itself, primarily occurring in tables, which would be inappropriate for cloze-style prompts. **The term spans we select are chosen by domain-expert authors of the textbooks and undergo multiple rounds of peer review, with educational value and informativeness as the primary criteria.** This makes them ideal candidates for knowledge probing prompts.

3.2.2 Cloze-style Prompts Creation

While existing benchmark largely focus on sentence-level prompts, MALAMUTE is the first CS dataset that is composed of both *sentence-level* and *paragraph-level* prompts.

Paragraph-level Prompts We take each unique term in each paragraph and replace every occurrence of that term with the “[MASK]” string. We

store the original term as a label and hide all other occurrences of the term in the same paragraph by replacing them with the “[HIDDEN]” string. This results in prompts that provide more contextual information about the term.

Sentence-level Prompts We split each paragraph into individual sentences using the NLTK library. Then, we select only the sentences that contain the term we are interested in. This results in a more challenging prompt, as it focuses on a single sentence and provides less contextual information.

3.2.3 Filtration

To ensure the quality of our prompts, we developed a comprehensive filtration methodology, which is described in detail in Appendix B. This methodology involves a combination of regular expressions and POS tagging to remove prompts that do not meet our quality standards, as determined through a thorough manual review of the prompt set.

3.3 Quality Control

To ensure the quality of our generated prompts, we conducted a thorough quality control process, which is described in detail in Appendix C. In brief, Graduate student annotators with academic backgrounds spanning MALAMUTE’s diverse subject areas, as evident from their academic transcripts, manually reviewed a representative sample of 495 English prompts. The results showed that the prompts demonstrated a high level of quality, with 97.6% of paragraph-level prompts and 99.1% of sentence-level prompts meeting grammatical standards, and 92.9% and 93.9%, respectively, effectively conveying the intended concept.

Domain	Subdomain	Prompt	Concept	
Math	Prealgebra	We call a the [MASK] of a .	additive inverse	
	Elementary Algebra	In [MASK], we express $x > 3$ as $(3, \infty)$.	interval notation	
	Intermediate Algebra	If we reverse the x and y in the function and then solve for y , we get our [MASK].	inverse function	
	College Algebra	A polynomial containing two terms, such as $2x - 9$, is called a [MASK].	binomial	
	Trigonometry	An [MASK] is an even root of a negative number.	imaginary number	
	Precalculus *	The [MASK] can be used to determine whether a graph represents a function.	vertical line test	
	Calculus I *	If f' is decreasing over I , we say f is [MASK] over I .	concave down	
	Calculus II *	The net change theorem considers the integral of a [MASK].	rate of change	
	Calculus III *	The nonzero vectors u and v are [MASK] if and only if $u \cdot v = 0$.	orthogonal vectors	
	Introductory Statistics	A [MASK] is the number of times a value of the data occurs.	frequency	
Science	Business Statistics *	The [MASK] is the most frequent value.	mode	
	Statistics *	A [MASK] provides a way of portraying data that can facilitate calculating probabilities.	two-way table	
	Contemporary Math	A [MASK] is any voting method that satisfies the Condorcet criterion.	Condorcet method	
	Concepts of Biology	All the individuals of a species living within a specific area are collectively called a [MASK].	population	
	Biology	There are very few living species of [MASK]: the platypus and four species of echidnas [...]	monotremes	
	Microbiology	The reactions of nitrogen fixation occur in specialized cells called [MASK].	heterocysts	
	Chemistry Atoms First	The number of protons in the nucleus of an atom is its [MASK].	atomic number	
	Chemistry *	An [MASK] is the smallest unit of an element that can participate in a chemical change.	atom	
	Organic Chemistry	[MASK] complexed with protein provides the physical makeup of the ribosomes.	Ribosomal RNA	
	Anatomy	The psoas major and iliacus make up the [MASK].	iliopsoas group	
Humanities	Astronomy	[MASK], the closest planet to the Sun [...]	Mercury	
	College Physics	Physics as it developed from the Renaissance to the end of the 19th century is called [MASK].	classical physics	
	Physics	The boiling point of water is 100°C for the Celsius scale, and its unit is the [MASK].	degree Celsius	
	Physics I \diamond *	For elliptical orbits, the point of closest approach of a planet to the Sun is called the [MASK].	perihelion	
	Physics II \diamond *	A special type of potential difference is known as [MASK].	electromotive force	
	Physics III \diamond *	Electroweak theory unifies the theory of [MASK][...]	quantum electrodynamics	
	U.S. History	The end of the Civil War saw the beginning of the [MASK] era [...]	Reconstruction	
	Writing	[MASK]: reference to the source of information used in a writer's research.	Citation	
	Philosophy	German philosopher Friedrich [MASK] famously declared that "God is dead"	Nietzsche	
	World History I	Conflict between [MASK] and the Sui began in the 590s and lasted for decades.	Goguryeo	
Business	World History II	In 1864, socialists founded the International Workingmen's Association (IWA) in [MASK].	London	
	Introduction to Business	[MASK] are the individuals or groups to whom a business has a responsibility.	Stakeholders	
	Financial Accounting	The [MASK] of cost allocation assumes that the last units purchased are the first units sold.	last-in, first-out method	
	Managerial Accounting	The person overseeing all of the accounting and finance concerns is the [MASK].	Chief Financial Officer	
	Business Ethics	Acting with [MASK] means we adhere strongly to a code of ethics[...]	integrity	
	Organizational Behavior	Once acceptable behavioral criteria have been specified, a [MASK] can be done.	performance audit	
	Finance	[MASK] are short-term debt instruments issued by the federal government.	Treasury bills	
	Business Law	Another remedy is [MASK], which would terminate the right of a partnership to exist.	dissolution	
	Intellectual Property	The fundamental and overriding requirement for a trademark is [MASK].	distinctiveness	
	Marketing	[MASK] happens when those being observed aren't aware that they are being watched.	Unobtrusive observation	
Nursing	Management	[MASK] are desired goals, objectives, or end states that individuals wish to pursue.	Terminal values	
	Economics	We call a firm's first stock sale to the public an [MASK].	initial public offering	
	Macroeconomics \diamond	The simplest example of a rate of return is the [MASK].	interest rate	
	Microeconomics \diamond	In the case of [MASK], one firm produces all of the output in a market.	monopoly	
	Entrepreneurship	A [MASK] is a company that does not allow members of the investing public to own stock.	privately held corporation	
	Nutrition	Any change in sodium and [MASK] can dramatically affect the cells of the brain.	water balance	
	Computer Science	Python Programming	[MASK] is the task of arranging elements in a sequence in ascending or descending order.	Sorting
		Workplace Software	The second most popular operating system to emerge during this time was the [MASK], first developed in 2005 and later acquired by Google.	Android operating system
	Social Sciences	American Government	The [MASK] limits the ability of the government to control or restrict religious practices.	free exercise clause
		Anthropology	Another group, called "untouchables" or [MASK], are outside the scheme of varnas.	dalits
Political Science		The Greek philosopher [MASK] argued that humans were "political animals" [...]	Aristotle	
Psychology \diamond		The behavior caused by the conditioned stimulus is called the [MASK].	conditioned response	
College Success	Sociology	The power in an [MASK] is held by a small, elite group.	oligarchy	
	College Success	If we rely too heavily on assumptions, we may be buying into [MASK], or generalizations	Stereotypes	
College Success Concise	One particular studying technique is called [MASK], which calls for students to mix [...]	interleaving		

Table 2: Examples of probes from each of the English textbooks. Books that are also in Polish are marked with \diamond , where books that are also in Spanish are marked with $*$.

This rigorous quality control process verified that the generated prompts meet the desired standards, providing a robust foundation for evaluating the performance of LLMs.

3.4 MALAMUTE Statistics

Our methodology yields a substantial dataset of 116,887 unique prompts, covering 33,361 distinct concepts across 71 textbooks and three languages. The language distribution of the prompts is: 100,258 prompts are in English, 8,147 in Polish, and 8,482 in Spanish. Domain-level statistics are provided in Table 3.

4 Experiments

4.1 Models

We evaluate a set of five MLMs: BERT-base (Devlin et al., 2019), DistilBERT-base (Sanh et al.,

2020), and SciBERT, a variant pretrained on scientific text (Beltagy et al., 2019). Additionally, we include two multilingual models, XLM-RoBERTa-large (Conneau et al., 2020) and mBERT-base (Devlin et al., 2019), which are specifically designed to handle a wide range of languages.

We also assess eight CLMs: GPT-4o, GPT-4o-mini, and GPT-4-turbo (Achiam et al., 2023), which are state-of-the-art generative models; Llama-3.1-405B-Instruct (Dubey et al., 2024), a powerful model designed for instruction-based tasks; Bloom-7B (Muennighoff et al., 2023) and Mistral-7B-Instruct (Jiang et al., 2023), which are trained for cross-lingual tasks; and Llama-2-7B-Chat (Touvron et al., 2023) and Llama-3-8B-Instruct (Llama3), both pretrained on multilingual data, with a focus on both general and scientific domains.

	Domain	#Books	#Prompts	#Concepts	Words (Pg.)	Words (Sent.)
EZ	Business	14	18,373	7,923	105.27	22.84
	College Success	2	206	115	104.09	18.60
	Computer Science	2	3,682	1,447	93.75	18.44
	Humanities	5	18,622	6,342	114.47	24.78
	Math	13	7,103	3,384	55.17	18.44
	Nursing	1	3,934	1,567	96.85	23.32
	Science	13	35,095	15,602	116.55	21.87
	Social Sciences	5	13,696	5,792	139.82	26.31
EL	Science	3	3,428	1,684	101.33	16.60
	Business	2	1,660	771	115.07	20.25
	Social Sciences	1	3,059	1,481	93.74	17.19
ES	Math	6	3,481	1,512	61.14	19.84
	Science	4	5,001	2,388	128.64	23.45

Table 3: Statistics of MALAMUTE. Pg and Sent refer to the average #words in paragraph and sentence level.

This comprehensive evaluation includes a range of models from both the MLM and CLM categories, pretrained on general and scientific data to ensure broad coverage of linguistic and domain-specific tasks.

4.2 Evaluation Metric

MLM Evaluation: Unlike previous work, which mostly focuses on single-token mask prediction (Zhong et al., 2021; Petroni et al., 2019; Sung et al., 2021; Bouraoui et al., 2019), we adopt a more comprehensive approach by predicting multi-token entities. Following Kassner et al. (2021), we frame the task as entity ranking, but with a key difference: we do not limit predictions to specific entity types. This is because entity type information is often scarce or unreliable, and using external resources to classify entity types may introduce inaccuracies. Instead, following Shaier et al. (2024a), we allow our models to predict any entity from the dataset’s entity list. To evaluate the predictions, we use the top-k accuracy metric, as suggested by Sung et al. (2021); Shaier et al. (2024a). This metric assigns a score of 1 if the correct entity is among the top k predicted entities, and 0 otherwise. Due to the complex relationships between entities (N-to-M connections), we report accuracy at three different levels: Acc@1, Acc@5, and Acc@10.

CLM Evaluation: Assessing the accuracy of CLM-generated answers is a complex task, especially since a proportion of responses are lengthy and explanatory. As a result, the traditional exact match metric commonly used in question-answering tasks is inadequate. We therefore adopt best sub-span accuracy, following Kandpal et al. (2023); Mallen et al. (2023), which is equal to 1 if the gold label appears in the generated answer and

0 otherwise. To account for different capitalization possibilities in labels and outputs, we cast both to lowercase before scoring the outputs.

To mitigate the known sensitivity of CLMs to minor input variations (Jia and Liang, 2017; Shaier et al., 2024c), we employ three diverse prompt designs (detailed in Appendix A), in line with established CLM probing research (Taylor et al., 2022; Li et al., 2023b; Kalo, 2022; Nayak, 2023). Notably, previous studies have shown that varying prompt designs, such as providing in-context examples, can significantly enhance performance (Brown et al., 2020). To this end, one of our prompts includes five in-context examples. To maintain consistency with the MLM input, we opt for a zero-shot evaluation, where no examples or additional context are provided with the other two prompts.

Given the substantial resources required to evaluate the larger models (GPT-4 and Llama 405B), we first assess the three prompts using the smaller models on a representative sample of the data. Our results show that the in-context learning prompt performs best. We therefore evaluate the larger models on this prompt on the vast English data, and all models on the remaining languages data.

5 Results & Discussion

Due to space constraints, we provide a summary of our results in Table 4. See full presentation of models’ results in Appendix D.

5.1 Model-wise Analysis

For CLMs, Llama 3.1 405b Instruct performs best in all languages. On English sentence-level, it scores 57.5%, while GPT-4o-mini scores the

lowest at 49.5%. On paragraph-level, Llama 3.1 405b Instruct scores 67.0%. Similar results are seen in Spanish (Llama 3.1: 49.5%, GPT-4o-mini: 37.4%) and Polish (Llama 3.1: 39.6%, GPT-4o-mini: 27.0%).

For MLM, XLM-RoBERTa excels on English sentence- and paragraph-level, with top-1, top-5, and top-10 scores of 23.8%, 45.6%, and 54.2%, respectively. mBERT performs the worst, with scores of 12.5%, 33.4%, and 41.9%, respectively. XLM-RoBERTa dominates on Spanish and Polish sentence- and paragraph-level, with scores of 17.8%, 40.5%, and 49.9% on Spanish, and similar results on Polish.

While these models show promise, their educational knowledge is limited, and their performance varies greatly between languages, particularly in CLMs. This language gap is problematic in educational settings, where students from diverse linguistic backgrounds need equal access to educational resources. Many students speak languages other than English, and the significant performance disparity between languages may exacerbate existing educational inequalities.

For instance, students who speak Spanish or Polish may not receive the same level of educational support as their English-speaking peers, simply because the models are less proficient in their native languages. Notably, this gap is smaller in MLM, but it is unclear if multilingual models perform equally well across all subjects due to the larger English dataset. Ultimately, bridging this language gap is crucial to ensuring that AI-powered educational tools are inclusive, equitable, and effective for all students.

5.2 Prompt Design and Evaluation

One striking observation is that the performance of most CLMs varies significantly depending on the prompt used for evaluation. For instance, on the Spanish portion, Llama 3's performance on Business Statistics diverges drastically across different prompt settings (see Table 8): it achieves a score of 23.0% when evaluated using paragraph-level prompts with setting 3, but only 5.0% when evaluated with prompt 1. This notable discrepancy underscores the importance of evaluating CLMs using a diverse range of prompts to gain a comprehensive understanding of their knowledge.

5.3 Sentence-level vs. Paragraph-level Performance

As discussed in Petroni et al. (2019), CS probing evaluates the lower bound of LMs' knowledge. Notably, Shaier et al. (2024a) find that using template-free prompts instead of template-based prompts can reveal a higher lower bound of knowledge, as templates can bias model predictions.

Building upon this insight, MALAMUTE introduces a significant innovation by providing a comprehensive set of both sentence-level and paragraph-level template-free prompts. This enables a more nuanced evaluation of LMs' knowledge in diverse contextual settings. Our experimental findings demonstrate that, across nearly all subdomains and models, paragraph-level prompts consistently yield higher scores for both CLMs and MLMs. This suggests that the provision of additional context can substantially enhance the evaluation of knowledge and masked concepts.

Furthermore, MALAMUTE's extensive collection of parallel sentence and paragraph-level prompts allows us to draw a more definitive conclusion: LMs may possess a significantly higher level of knowledge than previously estimated. Notably, when the same concept is being masked, our results show that models exhibit a substantially higher likelihood of correct prediction when more context is provided. This finding has important implications for the development of more accurate and comprehensive knowledge evaluation frameworks for LMs.

5.4 Granularity Analysis

The MALAMUTE benchmark represents a significant milestone in probing datasets, offering an unprecedented level of granularity that enables a fine-grained examination of LMs' strengths and weaknesses. This level of detail is crucial for developing effective educational systems that interact with students across diverse levels and domains.

For instance, our analysis of CLMs reveals that the top-performing model, Llama 3.1 405b Instruct, excels in subjects like Concepts of Biology, American Government, and Anatomy with scores approaching 80% (see Table 5). However, it struggles with subjects like Calculus 3, Marketing, and Business Law, where its scores plummet to around 50%. Similarly, MLMs demonstrate varying degrees of proficiency across subjects, with some showing significant knowledge gaps.

Models	Languages					
	English		Spanish		Polish	
	S	P	S	P	S	P
GPT-4-turbo	55.1	62.2	35.8	43.4	22.9	31.3
GPT-4o-mini	49.5	59.8	30.1	39.6	18.6	28.1
GPT-4o	53.4	63.5	34.5	42.7	27.7	38.3
Llama-3.1-405B-Instruct	57.5	67.0	42.4	51.4	27.8	39.7
BERT	[15.4, 32.6, 40.4]	[16.7, 34.6, 41.1]	-	-	-	-
DistillBERT	[13.0, 32.3, 41.8]	[14.2, 35.1, 43.8]	-	-	-	-
XLm-RoBERTa	[18.7, 40.0, 50.3]	[23.8, 45.6, 54.2]	[12.5, 31.1, 42.7]	[17.8, 40.5, 49.9]	[11.2, 29.4, 38.1]	[29.9, 51.7, 59.4]
mBERT	[9.3, 27.4, 36.8]	[12.5, 33.4, 41.9]	[9.1, 28.4, 38.1]	[13.3, 36.3, 45.5]	[4.4, 17.4, 24.2]	[16.1, 37.3, 45.7]
SciBERT	[18.5, 40.9, 49.6]	[22.6, 45.7, 53.9]	-	-	-	-

Table 4: All models’ overall accuracies (percentages) **on the English, Spanish, and, Polish portions of the dataset**. Results are shown as [prompt 1, prompt 2, prompt 3] for Causal LMs, and [top1, top5, top10] for Masked LMs. S=sentence-level prompts; P=paragraph-level prompts. The highest result for each setting, language, and model type is highlighted in **bold** font.

The granular nature of MALAMUTE is essential because it allows us to uncover hidden patterns and biases in model performance. While models may appear successful at first glance, a closer inspection can reveal significant gaps in their knowledge. For example, SciBERT, the overall top-performing MLM, owes its success largely to its exceptional performance in scientific subjects, with average scores of around 60% in Math and Science. However, when we drill down to specific subjects like Philosophy, World History, or Anthropology, its scores drop to around 30%. Furthermore, despite its strong performance in certain science domains, SciBERT’s knowledge in Astronomy is found to be lacking.

The insights provided by MALAMUTE underscore the critical need for more targeted and fine-grained evaluation and development of LMs. By recognizing the specific areas where models excel and struggle, we can focus our efforts on creating more effective, well-rounded models that can support students across diverse educational domains.

5.5 MLM vs. CLM Performance

Notably, many MLMs outperform many CLMs across all languages and subjects, which is a surprising finding given the significant size difference between the two types of models. Specifically, the smallest CLM (with 7B parameters) is more than 10 times larger than the largest MLM (XLm-RoBERTa). However, it is essential to acknowledge that this comparison is not apples-to-apples, as the two types of models are evaluated differently. MLMs are evaluated on their ability to predict masked tokens, which has a relatively limited number of possible options, whereas CLMs are tasked with generating free-form text, which has

an exponentially larger solution space.

Furthermore, the entity ranking evaluation method used for MLMs is simpler and more straightforward compared to the sub-span accuracy method used for CLMs. This raises important questions: (1) Can we directly compare the knowledge captured by MLMs and CLMs, given their fundamentally different evaluation criteria? and (2) What alternative evaluation methods, yet to be explored, could provide a more accurate assessment of each model type’s knowledge? We present these findings and questions for the reader’s consideration, acknowledging the complexity and nuance of comparing these two types of LMs.

6 Conclusion

We introduce MALAMUTE, a highly-granular, multilingual, and template-free educational probing dataset that exposes significant knowledge gaps in existing LMs. As LMs are increasingly considered for classroom use, where their performance can directly impact student outcomes, it is essential to develop evaluation methods that assess knowledge in specific, curriculum-aligned areas. MALAMUTE highlights the need for more specialized benchmarks that go beyond general knowledge and ensure models can support, rather than hinder, student learning. We hope MALAMUTE inspires the development of more educational benchmarks, driving the field toward safer, more reliable LMs for educational contexts. By addressing these gaps, we can better align language models with the needs of educational systems and improve their utility in real-world applications.

Limitations

We acknowledge that student materials are subject to change over time. However, we have mitigated this issue by using OpenStax, a constantly updated and expanding library of open educational resources. With new subjects, textbook improvements, and language translations being added regularly, our dataset will remain adaptable and extensible in the future.

Ethics Statement

We developed MALAMUTE to promote the creation of safer educational systems. We believe it is essential to rigorously evaluate language models on real-world student materials before deploying them in actual educational settings, to ensure they can accurately and reliably support student learning.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [Scibert: A pretrained language model for scientific text](#). *Preprint*, arXiv:1903.10676.
- Prabin Bhandari, Antonios Anastasopoulos, and Dieter Pfoser. 2023. [Are large language models geospatially knowledgeable?](#) In *Proceedings of the 31st ACM International Conference on Advances in Geographic Information Systems, SIGSPATIAL '23*, New York, NY, USA. Association for Computing Machinery.
- Zied Bouraoui, Jose Camacho-Collados, and Steven Schockaert. 2019. [Inducing relational knowledge from bert](#). *Preprint*, arXiv:1911.12753.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Boxi Cao, Hongyu Lin, Xianpei Han, Le Sun, Lingyong Yan, Meng Liao, Tong Xue, and Jin Xu. 2021a. [Knowledgeable or educated guess? revisiting language models as knowledge bases](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1860–1874, Online. Association for Computational Linguistics.
- Boxi Cao, Hongyu Lin, Xianpei Han, Le Sun, Lingyong Yan, Meng Liao, Tong Xue, and Jin Xu. 2021b. [Knowledgeable or educated guess? revisiting language models as knowledge bases](#). *Preprint*, arXiv:2106.09231.
- Jie Cao, Rachel Dickler, Marie Grace, Jeffrey B Bush, Alessandro Roncone, Leanne M Hirshfield, Marilyn A Walker, and Martha S Palmer. 2023. Designing an ai partner for jigsaw classrooms. In *Workshop on Language-Based AI Agent Interaction with Children (AIAIC'2023)*.
- Ilias Chalkidis, Nicolas Garneau, Catalina Goanta, Daniel Katz, and Anders Søgaard. 2023. [LeXFiles and LegalLAMA: Facilitating English multinational legal language model development](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15513–15535, Toronto, Canada. Association for Computational Linguistics.
- Manuel Ciosici, Joe Cecil, Dong-Ho Lee, Alex Hedges, Marjorie Freedman, and Ralph Weischedel. 2021. [Perhaps PTLMs should go to school – a task to assess open book and closed book QA](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6104–6111, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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.
- Bhuwan Dhingra, Jeremy R. Cole, Julian Martin Eisenschlos, Daniel Gillick, Jacob Eisenstein, and William W. Cohen. 2022. [Time-aware language models as temporal knowledge bases](#). *Transactions of the Association for Computational Linguistics*, 10:257–273.

- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Momchil Hardalov, Todor Mihaylov, Dimitrina Zlatkova, Yoan Dinkov, Ivan Koychev, and Preslav Nakov. 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.
- Qiyuan He, Yizhong Wang, and Wenya Wang. 2024. **Can language models act as knowledge bases at scale?** *Preprint*, arXiv:2402.14273.
- Robin Jia and Percy Liang. 2017. **Adversarial examples for evaluating reading comprehension systems**. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. **Mistral 7b**. *Preprint*, arXiv:2310.06825.
- Zhengbao Jiang, Frank F. Xu, Jun Araki, and Graham Neubig. 2020. **How can we know what language models know?** *Preprint*, arXiv:1911.12543.
- Gregor Jo st, Viktor Taneski, and Sa o Karakati c. 2024. The impact of large language models on programming education and student learning outcomes. *Applied Sciences*, 14(10):4115.
- Divyanshu Kakwani, Anoop Kunchukuttan, Satish Golla, Gokul N.C., Avik Bhattacharyya, Mitesh M. Khapra, and Pratyush Kumar. 2020. **IndicNLPsuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for Indian languages**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4948–4961, Online. Association for Computational Linguistics.
- Jan-Christoph Kalo. 2022. [link].
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning*, pages 15696–15707. PMLR.
- Nora Kassner, Philipp Dufter, and Hinrich Sch utze. 2021. **Multilingual LAMA: Investigating knowledge in multilingual pretrained language models**. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3250–3258, Online. Association for Computational Linguistics.
- Nora Kassner, Benno Krojer, and Hinrich Sch utze. 2020. **Are pretrained language models symbolic reasoners over knowledge?** *Preprint*, arXiv:2006.10413.
- Nora Kassner and Hinrich Sch utze. 2020. **Negated and misprimed probes for pretrained language models: Birds can talk, but cannot fly**. *Preprint*, arXiv:1911.03343.
- Majeed Kazemitabaar, Runlong Ye, Xiaoning Wang, Austin Zachary Henley, Paul Denny, Michelle Craig, and Tovi Grossman. 2024. Codeaid: Evaluating a classroom deployment of an llm-based programming assistant that balances student and educator needs. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–20.
- Amr Keleg and Walid Magdy. 2023. **DLAMA: A framework for curating culturally diverse facts for probing the knowledge of pretrained language models**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 6245–6266, Toronto, Canada. Association for Computational Linguistics.
- Fajri Koto, Haonan Li, Sara Shatnawi, Jad Doughman, Abdelrahman Boda Sadallah, Aisha Alraeesi, Khalid Almubarak, Zaid Alyafeai, Neha Sengupta, Shady Shehata, et al. 2024. Arabcmmlu: Assessing massive multitask language understanding in arabic. *arXiv preprint arXiv:2402.12840*.
- Harsh Kumar, Ilya Musabirov, Mohi Reza, Jiakai Shi, Anastasia Kuzminykh, Joseph Jay Williams, and Michael Liut. 2023. Impact of guidance and interaction strategies for llm use on learner performance and perception. *arXiv preprint arXiv:2310.13712*.
- Daliang Li, Ankit Singh Rawat, Manzil Zaheer, Xin Wang, Michal Lukasik, Andreas Veit, Felix Yu, and Sanjiv Kumar. 2023a. **Large language models with controllable working memory**. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1774–1793, Toronto, Canada. Association for Computational Linguistics.
- Lei Li, Jingjing Xu, Qingxiu Dong, Ce Zheng, Xu Sun, Lingpeng Kong, and Qi Liu. 2023b. **Can language models understand physical concepts?** In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11843–11861, Singapore. Association for Computational Linguistics.
- Bill Yuchen Lin, Seyeon Lee, Rahul Khanna, and Xiang Ren. 2020. **Birds have four legs?! NumerSense: Probing Numerical Commonsense Knowledge of Pre-Trained Language Models**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6862–6868, Online. Association for Computational Linguistics.

- Llama3. [Introducing meta llama 3: The most capable openly available llm to date.](#)
- Linhao Luo, Trang Vu, Dinh Phung, and Reza Haf. 2023. [Systematic assessment of factual knowledge in large language models.](#) In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 13272–13286, Singapore. Association for Computational Linguistics.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822.
- Zaiqiao Meng, Fangyu Liu, Ehsan Shareghi, Yixuan Su, Charlotte Collins, and Nigel Collier. 2022. [Rewire-then-probe: A contrastive recipe for probing biomedical knowledge of pre-trained language models.](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4798–4810, Dublin, Ireland. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual generalization through multitask finetuning.](#) *Preprint*, arXiv:2211.01786.
- Anmol Nayak. 2023. [\[link\]](#).
- B Nye, Dillon Mee, and Mark G Core. 2023. Generative large language models for dialog-based tutoring: An early consideration of opportunities and concerns. In *AIED Workshops*.
- OpenAI. 2023a. [Chatgpt: Optimizing language models for dialogue.](#)
- OpenAI. 2023b. [Gpt-4 technical report.](#) *Preprint*, arXiv:2303.08774.
- OpenStax. 2012. [Openstax.](#) Accessed: 2024-06-09. Originally launched in 2012. Continuously updated.
- Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2020. [How context affects language models’ factual predictions.](#) *Preprint*, arXiv:2005.04611.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) 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 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. [How much knowledge can you pack into the parameters of a language model?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5418–5426, Online. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.](#) *Preprint*, arXiv:1910.01108.
- Sagi Shaier, Kevin Bennett, Lawrence Hunter, and Katharina Kann. 2023a. [Emerging challenges in personalized medicine: Assessing demographic effects on biomedical question answering systems.](#) 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 540–550, Nusa Dua, Bali. Association for Computational Linguistics.
- Sagi Shaier, Kevin Bennett, Lawrence Hunter, and Katharina von der Wense. 2024a. [Comparing template-based and template-free language model probing.](#) In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 766–776, St. Julian’s, Malta. Association for Computational Linguistics.
- Sagi Shaier, Lawrence Hunter, and Katharina Kann. 2023b. [Who are all the stochastic parrots imitating? they should tell us!](#) 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 2: Short Papers)*, pages 113–120, Nusa Dua, Bali. Association for Computational Linguistics.
- Sagi Shaier, Lawrence Hunter, and Katharina von der Wense. 2024b. [Desiderata for the context use of question answering systems.](#) In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 777–792, St. Julian’s, Malta. Association for Computational Linguistics.
- Sagi Shaier, Lawrence Hunter, and Katharina Wense. 2024c. [It is not about what you say, it is about how you say it: A surprisingly simple approach for improving reading comprehension.](#) In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 8292–8305, Bangkok, Thailand. Association for Computational Linguistics.
- Sagi Shaier, Ari Kobren, and Philip V. Ogren. 2024d. [Adaptive question answering: Enhancing language model proficiency for addressing knowledge conflicts with source citations.](#) In *Proceedings of the*

- 2024 Conference on Empirical Methods in Natural Language Processing, pages 17226–17239, Miami, Florida, USA. Association for Computational Linguistics.
- Mujeen Sung, Jinhyuk Lee, Sean Yi, Minji Jeon, Sungdong Kim, and Jaewoo Kang. 2021. [Can language models be biomedical knowledge bases?](#) In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4723–4734, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alon Talmor, Yanai Elazar, Yoav Goldberg, and Jonathan Berant. 2020. [olmpics – on what language model pre-training captures.](#) *Preprint*, arXiv:1912.13283.
- Ross Taylor, Marcin Kardas, Guillem Cucurull, Thomas Scialom, Anthony Hartshorn, Elvis Saravia, Andrew Poulton, Viktor Kerkez, and Robert Stojnic. 2022. [Galactica: A large language model for science.](#) *Preprint*, arXiv:2211.09085.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models.](#) *Preprint*, arXiv:2302.13971.
- Cunxiang Wang, Xiaoze Liu, Yuanhao Yue, Xiangru Tang, Tianhang Zhang, Cheng Jiayang, Yunzhi Yao, Wenyang Gao, Xuming Hu, Zehan Qi, Yidong Wang, Linyi Yang, Jindong Wang, Xing Xie, Zheng Zhang, and Yue Zhang. 2023. [Survey on factuality in large language models: Knowledge, retrieval and domain-specificity.](#) *Preprint*, arXiv:2310.07521.
- Ruben Weijers, Gabrielle Fidelis de Castilho, Jean-François Godbout, Reihaneh Rabbany, and Kellin Pelrine. 2024. [Quantifying learning-style adaptation in effectiveness of LLM teaching.](#) In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, pages 112–118, St. Julians, Malta. Association for Computational Linguistics.
- Jian Xie, Kai Zhang, Jiangjie Chen, Renze Lou, and Yu Su. 2024a. [Adaptive chameleon or stubborn sloth: Revealing the behavior of large language models in knowledge conflicts.](#) *Preprint*, arXiv:2305.13300.
- Wenjing Xie, Juxin Niu, Chun Jason Xue, and Nan Guan. 2024b. [Grade like a human: Rethinking automated assessment with large language models.](#) *arXiv preprint arXiv:2405.19694*.
- Dongfang Xu, Peter Jansen, Jaycie Martin, Zhengnan Xie, Vikas Yadav, Harish Tayyar Madabushi, Oyvind Tafjord, and Peter Clark. 2020. [Multi-class hierarchical question classification for multiple choice science exams.](#) In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 5370–5382, Marseille, France. European Language Resources Association.
- Michael Zhang and Eunsol Choi. 2021. [SituatQA: Incorporating extra-linguistic contexts into QA.](#) In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7371–7387, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Chujie Zheng, Hao Zhou, Fandong Meng, Jie Zhou, and Minlie Huang. 2023. [Large language models are not robust multiple choice selectors.](#) In *The Twelfth International Conference on Learning Representations*.
- Zexuan Zhong, Dan Friedman, and Danqi Chen. 2021. [Factual probing is \[MASK\]: Learning vs. learning to recall.](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5017–5033, Online. Association for Computational Linguistics.
- Li Zhou, Taelin Karidi, Nicolas Garneau, Yong Cao, Wanlong Liu, Wenyu Chen, and Daniel Hershcovich. 2024. [Does mapo tofu contain coffee? probing llms for food-related cultural knowledge.](#) *Preprint*, arXiv:2404.06833.

A Prompts for Causal LMs

We evaluate the performance of our CLMs on the MALAMUTE dataset using each of the following three prompts. We represent each masked prompt as p_{ti} .

A.1 Prompt 1

The following prompt is based on Taylor et al. (2022); Li et al. (2023b)’s work

p_{ti}

A.2 Prompt 2

The following prompt is based on Luo et al. (2023)’s work

Please predict the missing words to complete the following sentence based on the facts in the real-world. The missing words are represented by [MASK]. Please return the missing words only. p_{ti}

A.3 Prompt 3

We drew inspiration from the work of Nayak (2023) and created a modified prompt. Specifically, we augmented the original prompt by incorporating five in-context examples, which are designed to provide additional context and guidance.

You are a helpful, respectful and honest assistant. Your answers should be crisp, short and not repetitive. Here are 5 examples: Example 1: Prompt: Francesco Bartolomeo Conti was born in [MASK]. Answer: Florence. Example 2: Prompt: English bulldog is a subclass of [MASK]. Answer: dog. Example 3: Prompt: The official language of Mauritius is [MASK]. Answer: English. Example 4: Prompt: Nicotine binds to [MASK]. Answer: CHRNA4. Example 5: Prompt: Hepatitis has symptoms such as [MASK]. Answer: Abdominal pain. Prompt: p_{ti}

B Filtration Methodology

To ensure the quality of our prompts, we have developed a comprehensive filtration methodology that involves a combination of regular expressions and POS tagging. This methodology enables us to remove prompts that do not meet our quality standards, as determined through a thorough manual review of the prompt set.

Our filtration methodology begins with the application of regular expressions to remove prompts that exhibit commonly seen issues. Specifically, we remove prompts that start with the following words: "this", "that", "and", "those", "these", "such", "we", "also", "it", "for example", "even", "they", "here", "another", "hence", "therefore", "neither", and "either". We also remove prompts that have digits as labels, as well as those that contain fewer than five words. Additionally, we eliminate prompts that refer to visual aids, including "graphs", "images", "figures", "plots", "example", "diagram", "chart", "photograph", and "photo".

Furthermore, we remove prompts that start with or contain action words, including: "explain", "describe", "analyze", "discuss", "interpret", "summarize", "define", "identify", "clarify", "elucidate", "examine", "determine", "decide", "solve", "calculate", "evaluate", "assess", "judge", "compare", "contrast", "distinguish", "differentiate", "write", "create", "design", "develop", "formulate", "compose", "draw", "sketch", "graph", "plot", "take", "find", "state", "prove", "illustrate", "represent", "model", "simulate", "demonstrate", "show", "indicate", "point", "outline", "develop", "expand", "elaborate", "justify", "argue", "debate", "review", "revise", "edit", "redraft", "reorganize", "rephrase", "paraphrase", "translate", "transcribe", "record", "document", "report", "present", "display", "exhibit", "construct", "build", "make", "produce", "generate", "develop", "plan", "organize", "schedule", "coordinate", "arrange", "implement", "execute", "enact", "establish", "institute", "initiate", "begin", "start", "continue", "resume", "complete", "finish", and "conclude".

We also use regular expressions to remove parenthesized text that appears immediately after or at the end of a label, such as "(AC)" in "audit committee (AC)". However, we apply this step with caution, ensuring that the parenthesized text is separated from the label by a white space character. This constraint is necessary to avoid damaging prompts in subjects like mathematics and computer science, where parentheses are frequently used to denote other concepts, such as mathematical notation (e.g., "f(x)") or programming language syntax (e.g., "upper()").

Our comprehensive filtration methodology enables us to remove low-quality prompts and ensure that our dataset meets the highest standards of quality and accuracy.

C Quality Control Methodology

To ensure the high quality of the generated prompts, we conducted a rigorous quality control process. This process involved manually reviewing a representative sample of 495 English prompts, consisting of 9 prompts from each of the 55 subdomains. The sample included a mix of sentence-level and paragraph-level prompts, with a ratio of 2:1, to provide a comprehensive representation of our dataset.

Three graduate student annotators, each holding an undergraduate degree and possessing in-depth knowledge of MALAMUTE’s diverse subject areas, evaluated the prompts based on a standardized set of criteria. Their knowledge and expertise in these areas are evident from their academic transcripts, which demonstrate a strong foundation in the subject matter. This evaluation was designed to assess two key aspects of the prompts: (1) grammatical correctness, and (2) the effective description of concept properties, which enables the inference of the label from the context.

The results of our quality control process revealed an impressive level of quality in the generated prompts. Among the paragraph-level prompts, a substantial 97.6% were found to be grammatically correct, while a significant 92.9% effectively conveyed the intended concept. Similarly, the sentence-level prompts demonstrated a high level of quality, with 99.1% meeting grammatical standards and 93.9% successfully demonstrating the concept. It is important to note that in many probing datasets, entities are often related to numerous other entities (N-to-M connections), which can result in multiple possible answers (Sung et al., 2021; Shaier et al., 2024a). Our results indicate that while there is some variation in quality, the overwhelming majority of prompts are suitable for knowledge probing applications.

It is worth noting that our approach to selecting annotators was deliberate. We chose to select annotators who were non-experts in any particular field of study, rather than domain-specific experts. This decision was motivated by two key considerations. Firstly, given the highly granular nature of MALAMUTE’s knowledge, it would be impractical to find and engage a large number of annotators with expertise in each of the domains, subdomains, and concepts. Secondly, and more importantly, LLMs are trained on general data and are, in that sense, also general and "non-experts." Therefore, it is more suitable to evaluate the performance of

models against humans who possess general knowledge, but also have a focus on MALAMUTE’s highly granular knowledge in their studies.

By conducting this rigorous quality control process, we were able to verify that the generated prompts meet the desired standards of quality, and are suitable for knowledge probing applications. While there may be some variation in quality, the overwhelming majority of prompts demonstrate a high level of quality, providing a robust foundation for evaluating the performance of LLMs.

C.1 Annotation Prompts

For each prompt, annotators were instructed to evaluate two aspects in response to the following prompt: “For each prompt evaluate whether: The prompt is grammatically correct; The label is specific to the prompt, i.e. the label can be inferred from the context and there aren’t many possible alternative answers. Use web resources such as Google search, Wikipedia, etc., as needed.” Annotators were encouraged to use their discretion and consider all relevant information when making their judgments. Importantly, there were no time constraints imposed on the annotation task, allowing annotators to take as much time as necessary to provide accurate assessments.

D Additional Results

The complete results of our evaluation are presented in the following tables. Table 5 shows the performance of CLMs on English data. Table 6 presents the results of MLMs evaluated on English at the sentence level, while Table 7 provides their performance at the paragraph level. For Spanish, Tables 8 and 10 summarize the results of CLMs, and Table 12 details the MLM evaluations. Similarly, for Polish, Tables 11 and 9 report the performance of CLMs, and Table 13 presents the results of MLMs.

E Error Analysis

E.1 Inter-annotator Agreement

To ensure the reliability and consistency of our dataset, we conducted a thorough analysis of inter-annotator agreement. For that, we evaluated the extent to which our three annotators concurred on the two key metrics of interest: specificity and grammaticality of the prompts. More specifically, we report 3-way unanimous percent agreement¹ on

¹<https://shorturl.at/em3FD>

Subdomain	Causal LMs							
	GPT-4-Turbo		GPT-4o-mini		GPT-4o		Llama 3.1 405b Instruct	
	S	P	S	P	S	P	S	P
Prealgebra	59.4	61.9	59.4	59.6	59.0	63.5	64.1	70.5
Elementary Algebra	48.5	53.0	49.1	54.3	45.5	55.3	55.7	61.3
Intermediate Algebra	49.1	55.4	46.0	54.4	50.3	52.9	53.4	60.2
College Algebra	51.0	53.3	49.7	58.8	53.1	58.8	58.6	63.3
Trigonometry	49.8	55.9	49.5	56.2	51.6	61.3	56.3	64.5
Precalculus *	46.4	49.9	44.9	49.8	46.6	52.2	54.9	59.2
Calculus I *	47.3	51.6	39.3	50.0	48.2	49.5	56.2	58.2
Calculus II *	46.2	51.0	41.9	45.2	45.6	50.2	48.1	53.3
Calculus III *	36.6	39.1	37.8	37.7	37.8	43.9	45.1	49.5
Intro. Statistics	41.9	59.1	41.9	60.6	43.5	56.2	53.2	68.6
Business Statistics *	44.1	59.2	45.4	57.8	46.1	59.7	52.0	64.5
Statistics *	41.2	54.6	42.5	55.3	41.2	53.3	55.0	61.2
Contemporary Math	51.8	62.0	50.0	62.3	49.6	63.0	56.7	65.0
Math Avg.	47.2	54.3	46.0	54.0	47.5	55.4	54.6	61.5
Concepts of Biology	71.2	79.9	67.0	78.7	72.7	82.4	73.9	83.7
Biology	66.0	74.2	61.3	74.1	68.3	76.6	68.1	78.2
Microbiology	56.0	62.6	50.4	60.4	56.0	64.8	58.1	68.1
Chemistry Atoms First	64.4	68.7	61.0	69.3	66.0	72.9	68.8	77.2
Chemistry *	64.3	73.3	50.0	73.3	64.3	75.6	67.9	77.8
Organic Chemistry	61.2	64.9	56.9	65.3	65.4	69.6	65.6	72.0
Anatomy	62.4	72.5	58.0	72.9	66.1	77.8	65.7	79.0
Astronomy	66.3	73.5	57.6	67.5	63.9	72.5	66.9	75.3
College Physics	57.2	62.6	51.5	62.6	55.6	63.5	59.0	66.8
Physics	62.5	67.6	60.3	68.3	61.9	69.2	66.2	72.8
Physics I ◇ *	40.1	49.2	37.2	47.0	38.9	49.1	43.0	53.8
Physics II ◇ *	49.9	58.4	43.5	56.7	50.6	56.4	54.2	61.6
Physics III ◇ *	48.4	53.3	43.9	52.0	49.5	55.6	57.3	60.5
Science Avg.	59.2	66.2	53.7	65.2	59.9	68.2	62.7	71.3
U.S. History	65.8	74.4	57.0	72.2	64.8	74.8	68.3	77.0
Writing	46.8	56.6	41.1	51.3	49.0	58.6	48.4	58.5
Philosophy	59.6	68.3	50.8	64.3	57.4	69.3	59.9	72.5
World History I	58.7	64.6	49.2	60.7	57.8	64.7	59.5	67.8
World History II	65.7	70.2	58.5	67.2	64.7	70.3	66.6	72.9
Humanities Avg.	59.3	66.8	51.3	63.1	58.7	67.5	60.5	69.7
Intro. to Business	43.8	52.2	36.2	48.8	42.7	52.2	45.1	56.2
Financial Accounting	47.4	54.0	44.7	57.5	48.4	56.3	50.0	62.8
Managerial Accounting	40.9	50.1	37.4	50.4	42.9	52.2	43.1	59.0
Business Ethics	51.9	63.1	45.0	57.5	50.4	57.1	54.7	66.4
Org. Behavior	31.2	47.8	27.7	42.5	30.7	47.8	34.7	51.6
Finance	44.3	51.0	39.9	49.5	43.8	50.7	48.5	58.8
Business Law	41.7	51.3	41.7	52.5	43.3	52.8	46.8	57.2
Intellectual Property	47.2	53.8	42.7	48.7	44.9	52.1	44.9	55.5
Marketing	38.9	46.8	33.7	43.4	38.5	47.2	39.8	48.5
Management	41.4	53.4	40.1	53.0	41.4	53.0	41.9	57.3
Economics	46.8	58.2	45.0	57.6	47.7	58.8	53.0	66.8
Macroeconomics ◇	56.9	63.3	54.3	64.0	57.8	64.7	61.2	73.4
Microeconomics ◇	59.7	69.6	61.3	73.4	62.9	70.9	71.0	78.5
Entrepreneurship	46.4	54.5	39.0	48.3	46.7	54.2	48.4	57.9
Business Avg.	45.6	54.9	42.1	53.4	45.9	55.0	48.8	60.7
Nutrition	43.0	48.0	39.0	46.4	42.1	48.3	43.7	52.8
Nursing Avg.	43.0	48.0	39.0	46.4	42.1	48.3	43.7	52.8
Python Programming	46.0	55.2	46.8	56.5	51.0	55.2	50.1	61.8
Workplace Software	37.4	52.9	34.4	49.9	38.1	54.6	39.6	59.7
Computer Science Avg.	41.7	54.0	40.6	53.2	44.6	54.9	44.8	60.8
American Government	66.0	76.1	61.0	72.7	64.9	75.8	70.6	80.2
Anthropology	55.4	59.7	44.3	52.1	55.9	61.0	57.8	65.5
Political Science	58.6	64.6	52.0	59.6	57.5	64.9	59.8	67.7
Psychology ◇	66.4	73.7	62.9	74.6	68.7	78.1	67.2	78.0
Sociology	58.4	68.9	49.2	66.1	60.1	71.3	59.9	73.3
Social Sciences Avg.	61.0	68.6	53.9	65.0	61.4	70.2	63.1	72.9
College Success	46.5	60.9	37.2	65.6	44.2	62.5	41.9	57.8
Col. Suc. Concise	40.4	53.8	40.4	61.5	42.6	51.9	44.7	67.3
College Success Avg.	43.4	57.4	38.8	63.6	43.4	57.2	43.3	62.6
Overall	51.6	59.7	47.4	58.5	52.0	60.7	55.4	65.3

Table 5: Causal models’ accuracies (percentages) on the English portion of the dataset. Books that are also in Polish are marked with ◇, where books that are also in Spanish are marked with *, and their results can be seen in Tables T1, T2. S=sentence-level prompts; P=paragraph-level prompts. The highest result for each subdomain is highlighted in bold font.

Subdomain	Mask LMs				
	BERT	DistilBERT	XLM-RoBERTa	mBERT	SciBERT
Prealgebra	[15.2, 37.8, 46.5]	[14.7, 36.4, 47.9]	[19.4, 37.8, 47.9]	[8.8, 25.8, 36.9]	[23.5, 48.4, 57.6]
Elementary Algebra	[15.0, 39.5, 49.7]	[13.8, 38.9, 49.7]	[18.6, 40.7, 53.3]	[6.0, 25.1, 38.3]	[22.8, 52.1, 60.5]
Intermediate Algebra	[10.4, 27.6, 37.4]	[12.3, 31.3, 42.9]	[13.5, 30.7, 39.3]	[6.1, 28.8, 36.2]	[14.7, 41.7, 50.3]
College Algebra	[9.7, 33.1, 42.1]	[9.7, 36.6, 47.6]	[7.6, 31.7, 42.8]	[4.1, 23.4, 30.3]	[20.0, 53.8, 63.4]
Trigonometry	[10.1, 28.5, 34.3]	[9.7, 29.6, 36.8]	[11.2, 27.8, 37.9]	[4.3, 18.4, 27.4]	[20.2, 41.5, 52.0]
Precalculus *	[8.3, 27.6, 33.0]	[9.2, 29.2, 37.2]	[9.0, 29.4, 38.4]	[3.8, 21.5, 31.3]	[14.8, 42.6, 51.6]
Calculus I *	[19.6, 34.8, 46.4]	[17.0, 44.6, 53.6]	[14.3, 33.0, 50.9]	[13.4, 30.4, 45.5]	[23.2, 55.4, 65.2]
Calculus II *	[10.6, 30.0, 36.2]	[11.9, 35.0, 43.8]	[13.8, 35.0, 44.4]	[10.0, 28.7, 39.4]	[23.8, 45.0, 54.4]
Calculus III *	[13.4, 36.6, 47.0]	[14.0, 38.4, 45.7]	[12.8, 36.6, 47.6]	[9.1, 28.7, 37.2]	[23.2, 50.0, 61.0]
Intro. Statistics	[16.1, 40.3, 50.0]	[16.1, 37.1, 61.3]	[19.4, 40.3, 59.7]	[14.5, 37.1, 53.2]	[40.3, 75.8, 83.9]
Business Statistics *	[12.5, 36.2, 48.0]	[11.8, 36.8, 48.0]	[13.2, 34.2, 46.7]	[5.9, 28.3, 38.2]	[27.0, 55.9, 66.4]
Statistics *	[16.2, 41.2, 51.2]	[16.2, 43.8, 55.0]	[16.2, 43.8, 61.3]	[10.0, 38.8, 48.8]	[28.7, 70.0, 80.0]
Contemporary Math	[18.7, 44.7, 56.0]	[21.1, 49.3, 62.0]	[18.3, 47.5, 62.0]	[9.9, 35.9, 45.4]	[25.4, 54.6, 66.9]
Math Avg.	[13.5, 35.2, 44.4]	[13.7, 37.5, 48.6]	[14.4, 36.0, 48.6]	[8.1, 28.5, 39.1]	[23.7, 52.8, 62.6]
Concepts of Biology	[11.4, 22.3, 26.6]	[10.9, 28.1, 37.9]	[14.5, 33.0, 41.3]	[6.7, 24.6, 32.1]	[26.9, 53.4, 64.8]
Biology	[7.8, 16.9, 20.7]	[7.4, 19.8, 26.1]	[10.5, 24.7, 31.2]	[4.7, 18.0, 23.2]	[19.5, 42.6, 50.7]
Microbiology	[3.8, 8.4, 11.0]	[3.3, 10.1, 14.2]	[5.8, 16.5, 22.7]	[2.3, 11.8, 17.5]	[14.4, 32.8, 41.7]
Chemistry Atoms First	[14.1, 31.6, 39.0]	[10.7, 32.4, 41.1]	[14.5, 38.9, 49.5]	[9.2, 27.5, 38.1]	[20.3, 48.1, 56.9]
Chemistry *	[32.1, 64.3, 67.9]	[25.0, 57.1, 82.1]	[32.1, 78.6, 96.4]	[21.4, 64.3, 92.9]	[35.7, 71.4, 89.3]
Organic Chemistry	[12.7, 22.8, 26.8]	[10.5, 28.6, 36.8]	[12.9, 36.2, 42.9]	[8.5, 25.9, 33.7]	[22.1, 49.6, 57.4]
Anatomy	[4.7, 10.4, 13.5]	[3.9, 12.3, 17.4]	[7.8, 20.8, 27.2]	[3.0, 16.1, 22.4]	[15.4, 38.0, 47.1]
Astronomy	[12.0, 27.4, 33.1]	[7.2, 22.4, 29.2]	[10.7, 27.7, 36.6]	[10.5, 25.3, 32.8]	[10.3, 23.0, 30.3]
College Physics	[11.4, 26.5, 33.5]	[10.7, 30.1, 39.3]	[12.8, 33.1, 41.1]	[6.3, 23.8, 31.7]	[19.9, 46.6, 57.0]
Physics	[18.3, 35.8, 46.6]	[15.1, 35.4, 43.7]	[19.2, 37.9, 49.4]	[9.1, 32.5, 40.9]	[23.9, 53.4, 61.6]
Physics I ♦ *	[13.3, 27.8, 34.8]	[10.6, 29.2, 37.2]	[14.0, 32.4, 41.8]	[9.2, 23.7, 30.7]	[16.2, 41.8, 50.7]
Physics II ♦ *	[13.7, 26.3, 34.9]	[12.4, 29.6, 35.4]	[18.7, 35.2, 45.6]	[10.4, 28.4, 35.9]	[21.3, 45.3, 51.9]
Physics III ♦ *	[12.9, 30.5, 35.7]	[13.6, 32.4, 40.8]	[15.7, 36.9, 44.6]	[10.8, 30.0, 39.9]	[26.1, 47.4, 54.5]
Science Avg.	[12.9, 27.0, 32.6]	[10.9, 28.3, 37.0]	[14.6, 34.8, 43.9]	[8.6, 27.1, 36.3]	[20.9, 45.6, 54.9]
U.S. History	[16.3, 29.2, 34.6]	[15.0, 29.4, 35.2]	[18.3, 32.1, 39.5]	[14.0, 25.9, 32.5]	[8.4, 15.2, 21.6]
Writing	[15.2, 29.4, 36.4]	[11.8, 28.5, 38.2]	[20.9, 38.2, 48.0]	[9.1, 22.8, 29.8]	[10.6, 23.4, 30.2]
Philosophy	[17.6, 27.9, 33.1]	[11.7, 24.4, 31.2]	[16.3, 31.0, 39.0]	[13.1, 28.1, 35.1]	[9.8, 22.5, 28.7]
World History 1	[13.1, 24.4, 28.9]	[7.6, 18.7, 24.9]	[9.0, 21.0, 27.0]	[7.8, 22.1, 30.4]	[3.0, 8.8, 12.7]
World History 2	[23.6, 40.3, 47.9]	[10.5, 26.3, 34.7]	[20.9, 38.0, 46.3]	[15.9, 35.9, 45.2]	[8.8, 19.2, 24.9]
Humanities Avg.	[17.2, 30.2, 36.2]	[11.3, 25.5, 32.8]	[17.1, 32.1, 40.0]	[12.0, 27.0, 34.6]	[8.1, 17.8, 23.6]
Intro. to Business	[12.6, 23.6, 28.7]	[10.1, 21.6, 26.4]	[19.6, 36.4, 44.0]	[5.5, 15.8, 20.9]	[9.0, 21.7, 27.0]
Financial Accounting	[12.4, 32.1, 38.6]	[10.0, 35.0, 44.9]	[19.3, 47.4, 58.1]	[6.7, 22.2, 31.3]	[12.2, 31.5, 40.0]
Managerial Accounting	[14.6, 36.3, 42.9]	[11.4, 37.1, 48.3]	[22.6, 50.6, 63.1]	[9.1, 27.4, 36.9]	[16.0, 38.0, 45.1]
Business Ethics	[21.3, 41.9, 51.9]	[19.8, 42.6, 54.7]	[31.4, 56.2, 66.7]	[14.7, 31.8, 43.0]	[21.3, 45.7, 53.1]
Org. Behavior	[10.2, 25.9, 32.9]	[9.5, 24.9, 35.4]	[16.7, 38.4, 49.1]	[6.2, 19.2, 23.9]	[14.0, 33.2, 42.9]
Finance	[11.9, 27.4, 34.6]	[9.8, 25.3, 33.5]	[21.1, 44.0, 52.8]	[6.6, 22.1, 29.8]	[13.7, 30.6, 37.7]
Business Law	[27.0, 39.3, 44.8]	[21.8, 40.5, 46.8]	[29.0, 54.4, 63.1]	[13.1, 31.3, 43.7]	[18.7, 38.9, 46.8]
Intellectual Property	[14.6, 46.1, 64.0]	[12.4, 38.2, 59.6]	[21.3, 49.4, 66.3]	[7.9, 32.6, 44.9]	[9.0, 38.2, 57.3]
Marketing	[9.5, 20.5, 25.2]	[7.9, 20.2, 26.5]	[20.2, 38.3, 44.9]	[5.4, 15.5, 21.3]	[9.3, 20.5, 25.9]
Management	[21.2, 39.6, 52.3]	[14.9, 36.5, 46.8]	[28.8, 52.3, 64.4]	[14.9, 38.7, 47.3]	[21.2, 43.7, 53.6]
Economics	[11.9, 33.1, 41.6]	[10.4, 32.3, 42.9]	[17.2, 40.9, 50.5]	[5.2, 23.9, 31.6]	[15.8, 38.6, 48.8]
Macroeconomics ♦	[24.1, 58.6, 74.1]	[18.1, 56.0, 64.7]	[26.7, 59.5, 69.8]	[7.8, 43.1, 55.2]	[29.3, 60.3, 70.7]
Microeconomics ♦	[27.4, 64.5, 83.9]	[24.2, 61.3, 75.8]	[30.6, 66.1, 82.3]	[12.9, 53.2, 64.5]	[38.7, 71.0, 85.5]
Entrepreneurship	[10.0, 24.3, 30.0]	[8.4, 21.5, 28.5]	[18.0, 38.5, 45.2]	[5.8, 17.3, 22.9]	[8.6, 21.5, 28.3]
Business Avg.	[16.3, 36.7, 46.1]	[13.5, 35.2, 45.3]	[23.0, 48.0, 58.6]	[8.7, 28.2, 36.9]	[16.9, 38.1, 47.3]
Nutrition	[6.8, 15.7, 20.0]	[6.7, 16.3, 21.6]	[10.7, 27.5, 35.9]	[2.3, 10.6, 16.7]	[18.6, 39.8, 51.3]
Nursing Avg.	[6.8, 15.7, 20.0]	[6.7, 16.3, 21.6]	[10.7, 27.5, 35.9]	[2.3, 10.6, 16.7]	[18.6, 39.8, 51.3]
Python Programming	[7.6, 15.9, 20.1]	[6.8, 17.6, 20.7]	[20.1, 36.5, 46.2]	[4.0, 15.3, 23.5]	[16.4, 30.9, 34.6]
Workplace Software	[8.1, 17.2, 21.6]	[6.5, 16.7, 22.0]	[17.6, 36.3, 45.1]	[4.2, 13.0, 18.5]	[9.4, 22.8, 29.3]
Computer Science Avg.	[7.8, 16.6, 20.9]	[6.7, 17.1, 21.4]	[18.9, 36.4, 45.7]	[4.1, 14.2, 21.0]	[12.9, 26.9, 32.0]
American Government	[22.8, 33.1, 38.3]	[10.4, 23.7, 31.0]	[23.4, 38.4, 45.5]	[15.8, 28.0, 33.4]	[7.0, 17.0, 22.3]
Anthropology	[12.3, 23.0, 29.0]	[9.5, 21.2, 28.4]	[13.0, 26.0, 32.9]	[8.6, 21.6, 28.0]	[8.2, 19.1, 24.9]
Political Science	[18.7, 32.8, 39.6]	[11.6, 26.4, 33.7]	[23.3, 41.6, 49.0]	[13.1, 26.5, 34.7]	[9.0, 22.3, 28.2]
Psychology ♦	[12.9, 26.4, 32.0]	[11.3, 26.8, 33.6]	[18.9, 39.0, 48.0]	[6.7, 20.6, 28.7]	[22.2, 42.6, 51.0]
Sociology	[19.1, 39.3, 48.9]	[16.8, 41.8, 54.0]	[27.9, 52.5, 63.9]	[9.9, 28.6, 38.4]	[17.2, 39.5, 49.8]
Social Sciences Avg.	[17.2, 30.9, 37.6]	[11.9, 28.0, 36.1]	[21.3, 39.5, 47.9]	[10.8, 25.1, 32.6]	[12.7, 28.1, 35.2]
College Success	[41.9, 60.5, 74.4]	[44.2, 62.8, 74.4]	[41.9, 76.7, 88.4]	[27.9, 53.5, 69.8]	[23.3, 55.8, 67.4]
College Success Concise	[40.4, 57.4, 70.2]	[29.8, 51.1, 70.2]	[46.8, 72.3, 87.2]	[17.0, 38.3, 59.6]	[27.7, 55.3, 63.8]
College Success Avg.	[41.1, 59.0, 72.3]	[37.0, 57.0, 72.3]	[44.3, 74.5, 87.8]	[22.4, 45.9, 64.7]	[25.5, 55.5, 65.6]
Overall	[15.4, 32.6, 40.4]	[13.0, 32.3, 41.8]	[18.7, 40.0, 50.3]	[9.3, 27.4, 36.8]	[18.5, 40.9, 49.6]

Table 6: Masked models’ accuracies (percentages) on the English portion of the dataset using sentence-level prompts. Results are shown as [top1, top5, top10]. Books that are also in Polish are marked with ♦, where books that are also in Spanish are marked with *. The highest result for each subdomain across both sentence and paragraph level Tables (Table 6 and Table 7) is highlighted in bold font.

Subdomain	Causal LMs							
	Bloom		Llama2		Llama3		Mistral	
	S	P	S	P	S	P	S	P
Precalculus	[2.0, 2.0, 3.7]	[1.3, 0.7, 1.0]	[4.0, 2.7, 3.3]	[3.3, 1.3, 3.0]	[7.7, 8.7, 16.3]	[5.3, 5.3, 14.7]	[8.3, 7.0, 8.7]	[4.7, 4.7, 6.7]
Calculus I	[3.2, 2.7, 4.5]	[4.3, 5.8, 4.7]	[5.9, 1.4, 5.4]	[3.2, 1.1, 2.5]	[6.3, 5.0, 8.6]	[7.2, 6.1, 15.2]	[10.0, 4.5, 5.4]	[7.6, 3.2, 4.7]
Calculus II	[2.0, 1.0, 0.0]	[5.2, 3.0, 4.4]	[0.0, 1.0, 2.0]	[3.0, 0.7, 1.5]	[1.0, 2.0, 5.1]	[4.4, 5.2, 7.4]	[2.0, 2.0, 3.0]	[3.0, 5.9, 5.2]
Calculus III	[1.2, 1.2, 2.4]	[3.3, 3.0, 3.0]	[2.0, 0.8, 1.6]	[3.3, 0.7, 1.7]	[4.0, 6.1, 8.1]	[9.0, 5.3, 11.0]	[4.5, 4.0, 3.2]	[5.7, 3.7, 4.3]
Business Statistics	[4.3, 2.9, 1.4]	[4.0, 7.0, 6.0]	[5.8, 0.0, 2.9]	[6.0, 0.0, 3.0]	[8.7, 7.2, 20.3]	[5.0, 5.0, 23.0]	[13.0, 5.8, 5.8]	[8.0, 5.0, 5.0]
Statistics	[3.1, 2.1, 4.2]	[5.0, 5.0, 4.6]	[5.2, 1.0, 4.2]	[4.2, 1.7, 0.4]	[13.6, 7.9, 17.8]	[9.2, 9.7, 15.5]	[8.9, 5.8, 5.2]	[9.2, 4.2, 7.1]
Math Avg.	[2.6, 2.0, 2.7]	[3.9, 4.1, 3.9]	[3.8, 1.2, 3.2]	[3.8, 0.9, 2.0]	[6.9, 6.1, 12.7]	[6.7, 6.1, 14.5]	[7.8, 4.9, 5.2]	[6.4, 4.5, 5.5]
Chemistry	[7.3, 3.0, 6.7]	[9.3, 7.7, 6.7]	[6.0, 4.3, 8.3]	[6.3, 1.0, 4.0]	[10.7, 13.0, 19.3]	[9.3, 13.3, 24.0]	[13.7, 9.3, 11.0]	[9.3, 5.7, 7.0]
Physics I	[3.0, 1.0, 3.3]	[4.3, 3.0, 3.7]	[5.0, 2.0, 6.3]	[5.0, 1.7, 2.3]	[8.7, 7.3, 14.3]	[10.3, 8.3, 14.0]	[6.3, 5.3, 6.3]	[4.7, 3.7, 4.3]
Physics II	[3.3, 3.0, 9.0]	[4.0, 4.3, 4.7]	[4.0, 4.3, 6.7]	[4.3, 3.0, 5.0]	[6.7, 8.3, 16.3]	[10.0, 12.0, 20.3]	[6.3, 10.7, 10.7]	[9.7, 7.0, 6.3]
Physics III	[4.0, 1.7, 5.0]	[3.0, 2.3, 2.7]	[4.0, 4.7, 9.3]	[4.7, 2.3, 3.0]	[7.7, 11.3, 15.7]	[7.0, 10.3, 17.0]	[10.0, 10.7, 12.0]	[6.0, 6.3, 9.3]
Science Avg.	[4.4, 2.2, 6.0]	[5.2, 4.3, 4.5]	[4.8, 3.8, 7.7]	[5.1, 2.0, 3.6]	[8.4, 10.0, 16.4]	[9.2, 11.0, 18.8]	[9.1, 9.0, 10.0]	[7.4, 5.7, 6.7]
Overall	[3.3, 2.1, 4.0]	[4.4, 4.2, 4.2]	[4.2, 2.2, 5.0]	[4.3, 1.4, 2.6]	[7.5, 7.7, 14.2]	[7.7, 8.0, 16.2]	[8.3, 6.5, 7.1]	[6.8, 4.9, 6.0]

Table 8: Causal models’ accuracies (percentages) on the Spanish portion of the dataset. Results are shown as [prompt 1, prompt 2, prompt 3]. S=sentence-level prompts; P=paragraph-level prompts. The highest result for each subdomain is highlighted in **bold** font.

Subdomain	Causal LMs							
	Bloom		Llama2		Llama3		Mistral	
	S	P	S	P	S	P	S	P
Physics I	[0, 0, 0]	[0, 0, 0]	[0.0, 0.3, 0.3]	[1.0, 0.3, 0.0]	[1.0, 1.7, 4.7]	[1.7, 0.3, 4.0]	[1.3, 1.0, 1.0]	[1.7, 0.7, 0.0]
Physics II	[0.3, 0.0, 0]	[0.3, 0.3, 0]	[1.0, 0.0, 1.3]	[0.7, 0.0, 1.7]	[0.3, 2.0, 5.0]	[1.7, 2.3, 8.3]	[1.0, 0.7, 0.7]	[2.3, 1.0, 2.3]
Physics III	[0.3, 0.0, 0.0]	[0.3, 0.0, 0.3]	[1.0, 1.3, 1.3]	[1.3, 0.3, 0.3]	[1.3, 1.7, 3.7]	[1.0, 1.7, 5.0]	[2.7, 0.7, 1.3]	[2.3, 0.7, 1.3]
Science Avg.	[0.2, 0.0, 0.0]	[0.2, 0.1, 0.1]	[0.7, 0.5, 1.0]	[1.0, 0.2, 0.7]	[0.9, 1.8, 4.5]	[1.5, 1.4, 5.8]	[1.7, 0.8, 1.0]	[2.1, 0.8, 1.2]
Macroeconomics	[0.0, 0.0, 0.0]	[0.7, 0.7, 0.3]	[1.7, 1.7, 2.7]	[3.7, 1.0, 0.3]	[1.7, 2.3, 6.0]	[3.3, 2.3, 6.3]	[1.0, 1.7, 3.7]	[3.0, 1.7, 2.7]
Microeconomics	[0.0, 0.0, 0.0]	[2.0, 1.0, 0.7]	[1.7, 0.3, 1.0]	[1.7, 0.7, 1.0]	[1.0, 2.7, 5.0]	[2.3, 2.3, 5.0]	[0.7, 1.3, 3.0]	[2.7, 2.0, 2.3]
Business Avg.	[0.0, 0.0, 0.0]	[1.4, 0.8, 0.5]	[1.7, 1.0, 1.9]	[2.7, 0.8, 0.7]	[1.4, 2.5, 5.5]	[2.8, 2.3, 5.7]	[0.8, 1.5, 3.4]	[2.9, 1.9, 2.5]
Psychology	[0, 0, 0]	[0, 0, 0]	[1.0, 0.0, 1.3]	[2.0, 1.0, 2.7]	[1.3, 2.7, 9.0]	[3.0, 3.0, 9.0]	[1.0, 1.7, 4.0]	[2.3, 2.0, 4.0]
Social Sciences Avg.	[0.0, 0.0, 0.0]	[0.0, 0.0, 0.0]	[1.0, 0.0, 1.3]	[2.0, 1.0, 2.7]	[1.3, 2.7, 9.0]	[3.0, 3.0, 9.0]	[1.0, 1.7, 4.0]	[2.3, 2.0, 4.0]
Overall	[0.1, 0.0, 0.0]	[0.5, 0.3, 0.2]	[1.1, 0.6, 1.3]	[1.7, 0.5, 1.0]	[1.1, 2.2, 5.6]	[2.2, 2.0, 6.3]	[1.3, 1.2, 2.3]	[2.4, 1.3, 2.1]

Table 9: Causal models’ accuracies (percentages) on the Polish portion of the dataset. Results are shown as [prompt 1, prompt 2, prompt 3]. S=sentence-level prompts; P=paragraph-level prompts. The highest result for each subdomain is highlighted in **bold** font.

Subdomain	Causal LMs							
	GPT-4-turbo		GPT-4o-mini		GPT-4o		Llama 3.1 405b Instruct	
	S	P	S	P	S	P	S	P
Precalculus	33.5	37.3	25.8	35.0	29.1	36.3	43.4	48.8
Calculus I	29.4	37.9	25.8	32.5	27.1	38.3	30.3	46.6
Calculus II	24.2	26.7	17.2	26.7	21.2	23.0	27.3	36.3
Calculus III	25.5	33.8	20.6	25.6	27.5	32.0	31.6	41.5
Business Statistics	42.0	42.0	36.2	41.0	34.8	42.0	47.8	57.0
Statistics	30.4	42.0	32.5	37.4	30.4	37.4	37.2	49.2
Math Avg.	30.8	36.6	26.4	33.0	28.4	34.8	36.3	46.6
Chemistry	46.4	56.5	39.2	51.0	46.5	56.0	53.1	63.5
Physics I	30.7	40.4	27.3	39.1	30.1	41.6	34.8	44.7
Physics II	36.8	46.7	29.5	43.3	36.2	47.1	44.9	55.2
Physics III	38.1	46.2	32.2	41.9	36.4	44.6	44.3	51.8
Science Avg.	38.0	47.5	32.0	43.8	37.3	47.3	44.3	53.8
Overall	33.7	41.0	28.6	37.4	31.9	39.8	39.5	49.5

Table 10: Causal models’ accuracies (percentages) on the Spanish portion of the dataset. Averages are book-level. S=sentence-level prompts; P=paragraph-level prompts. The highest result for each subdomain is highlighted in **bold** font.

Subdomain	Causal LMs							
	GPT-4-turbo		GPT-4o-mini		GPT-4o		Llama 3.1 405b Instruct	
	S	P	S	P	S	P	S	P
Physics I	19.7	29.5	14.6	21.8	22.0	32.7	23.3	36.2
Physics II	21.8	30.7	17.0	25.8	25.7	35.3	29.3	40.0
Physics III	26.9	30.7	19.0	23.3	30.8	39.2	30.8	42.0
Science Avg.	22.8	30.3	16.9	23.6	26.2	35.7	27.8	39.4
Macroeconomics	22.3	28.6	20.1	28.6	25.7	35.0	27.0	37.9
Microeconomics	22.0	33.1	18.6	30.2	26.6	35.7	24.0	41.7
Business Avg.	22.2	30.9	19.4	29.4	26.2	35.4	25.5	39.8
Psychology	23.3	32.8	20.0	32.4	30.4	42.8	28.7	40.0
Social Sciences Avg.	23.3	32.8	20.0	32.4	30.4	42.8	28.7	40.0
Overall	22.7	30.9	18.2	27.0	26.9	36.8	27.2	39.6

Table 11: Causal models’ accuracies (percentages) on the Polish portion of the dataset. Averages are book-level. S=sentence-level prompts; P=paragraph-level prompts. The highest result for each subdomain is highlighted in **bold** font.

Subdomain	Mask LMs			
	XLM-RoBERTa		mBERT	
	S	P	S	P
Precalculus	[7.6, 24.0, 32.2]	[12.1, 29.9, 38.5]	[4.5, 23.0, 29.8]	[7.2, 27.2, 34.3]
Calculus I	[16.7, 37.6, 47.1]	[18.1, 48.4, 54.5]	[11.8, 33.9, 40.7]	[15.2, 41.2, 46.6]
Calculus II	[18.2, 34.3, 49.5]	[21.5, 45.9, 50.4]	[9.1, 28.3, 39.4]	[16.3, 40.0, 49.6]
Calculus III	[8.5, 32.0, 37.2]	[16.5, 35.7, 44.2]	[11.3, 26.7, 38.9]	[13.1, 36.0, 44.8]
Business Statistics	[11.6, 31.9, 60.9]	[21.0, 52.0, 66.0]	[11.6, 42.0, 59.4]	[16.0, 51.0, 69.0]
Statistics	[9.9, 33.5, 46.1]	[15.5, 38.7, 50.8]	[7.3, 26.7, 39.3]	[10.1, 32.4, 42.4]
Math Avg.	[12.1, 32.2, 45.5]	[17.4, 41.8, 50.7]	[9.3, 30.1, 41.2]	[13.0, 38.0, 47.8]
Chemistry	[12.8, 29.6, 40.2]	[22.0, 44.0, 53.9]	[9.0, 27.6, 37.1]	[15.7, 40.3, 50.0]
Physics I	[12.2, 28.5, 36.5]	[16.0, 35.9, 43.7]	[7.6, 22.0, 29.3]	[9.5, 26.9, 35.7]
Physics II	[14.2, 31.1, 40.0]	[18.4, 37.1, 47.7]	[9.6, 25.6, 33.9]	[14.3, 32.6, 39.4]
Physics III	[13.5, 28.2, 37.3]	[16.4, 37.5, 48.9]	[9.5, 27.7, 33.7]	[15.7, 35.1, 43.3]
Science Avg.	[13.2, 29.4, 38.5]	[18.2, 38.6, 48.5]	[8.9, 25.7, 33.5]	[13.8, 33.7, 42.1]
Overall	[12.5, 31.1, 42.7]	[17.8, 40.5, 49.9]	[9.1, 28.4, 38.1]	[13.3, 36.3, 45.5]

Table 12: Masked models’ accuracies (percentages) on the Spanish portion of the dataset. Results are shown as [top1, top5, top10]. S=sentence-level prompts; P=paragraph-level prompts. The highest result for each subdomain is highlighted in bold font.

Subdomain	Mask LMs			
	XLM-RoBERTa		mBERT	
	S	P	S	P
Physics I	[8.9, 26.8, 36.8]	[25.7, 48.3, 54.5]	[4.1, 18.3, 25.2]	[13.7, 36.0, 43.5]
Physics II	[11.9, 29.7, 36.8]	[28.9, 44.6, 52.6]	[5.4, 18.8, 25.0]	[16.0, 36.0, 43.5]
Physics III	[12.3, 31.7, 39.9]	[32.0, 54.1, 61.0]	[7.3, 22.1, 30.8]	[19.9, 43.0, 49.1]
Science Avg.	[11.0, 29.4, 37.8]	[28.9, 49.0, 56.0]	[5.6, 19.7, 27.0]	[16.5, 38.3, 45.4]
Macroeconomics	[11.4, 34.1, 43.9]	[26.8, 56.0, 65.3]	[3.4, 17.0, 23.3]	[13.7, 37.2, 47.6]
Microeconomics	[12.2, 30.2, 40.1]	[33.9, 55.9, 64.8]	[3.2, 16.2, 23.9]	[15.0, 34.9, 47.8]
Business Avg.	[11.8, 32.1, 42.0]	[30.4, 56.0, 65.0]	[3.3, 16.6, 23.6]	[14.3, 36.0, 47.7]
Psychology	[10.4, 24.0, 30.9]	[32.3, 51.4, 58.1]	[2.8, 12.3, 17.2]	[18.1, 36.8, 42.5]
Social Sci Avg.	[10.4, 24.0, 30.9]	[32.3, 51.4, 58.1]	[2.8, 12.3, 17.2]	[18.1, 36.8, 42.5]
Overall	[11.2, 29.4, 38.1]	[29.9, 51.7, 59.4]	[4.4, 17.4, 24.2]	[16.1, 37.3, 45.7]

Table 13: Masked models’ accuracies (percentages) on the Polish portion of the dataset. Results are shown as [top1, top5, top10]. S=sentence-level prompts; P=paragraph-level prompts. The highest result for each subdomain is highlighted in bold font.

a sample of 495 prompts (9 from each subdomain) for True/False ratings of grammaticality and specificity between annotators. Simple agreement was chosen over a Kappa metric due to the "paradox of high agreement" from which Kappa metrics suffer under heavy class imbalance; in our case, since "yes" ratings make up over 90% and 97% in our specificity and grammaticality ratings respectively, Kappa scores would be misleading.

The results show a high level of agreement, with all three annotators reaching the same conclusion in 97.1% of cases for grammaticality judgments, indicating a large consensus on the linguistic accuracy of the prompts. Similarly, our annotators demonstrated a strong agreement on specificity judgments, concurring in 90.1% of cases, which underscores the consistency of their evaluations regarding the degree to which the prompts require precise and detailed responses. These findings provide strong evidence of the high quality and reliability of our dataset.