

Open Conversational LLMs do not know most Spanish words

Los LLM conversacionales abiertos no reconocen la mayoría de las palabras del español

Javier Conde,¹ Miguel González,¹ Nina Melero,^{1,4} Raquel Ferrando,¹ Gonzalo Martínez,²
Elena Merino-Gómez,³ José Alberto Hernández,² Pedro Reviriego¹

¹ETSI de Telecomunicación, Universidad Politécnica de Madrid
email: {javier.conde.diaz,miguel.gonsaiz,i.melero,pedro.reviriego}@upm.es,
raquel.ferrando@alumnos.upm.es

²Universidad Carlos III de Madrid
email: jahgutie@it.uc3m.es,gonzmart@pa.uc3m.es

³Universidad de Valladolid
email: elena.merino.gomez@uva.es

⁴New York University-Madrid
email: mim8070@nyu.edu

Abstract: The growing interest in Large Language Models (LLMs) and in particular in conversational models with which users can interact has led to the development of a large number of open chat LLMs. These models are evaluated on a wide range of benchmarks to assess their capabilities in answering questions or solving problems on almost any possible topic or to test their ability to reason or interpret texts. Instead, the evaluation of the knowledge that these models have of the languages has received much less attention. For example, the words that they can recognize and use in different languages. In this paper, we evaluate the knowledge that open chat LLMs have of Spanish words by testing a sample of words in a reference dictionary. The results show that open chat LLMs produce incorrect meanings for an important fraction of the words and are not able to use most of the words correctly to write sentences with context. These results show how Spanish is left behind in the LLM race and highlight the need to push for linguistic fairness in conversational LLMs ensuring that they provide similar performance across languages.

Keywords: LLMs, Vocabulary, Open models, Spanish.

Resumen: La popularidad de los grandes modelos de lenguaje, o LLM del inglés *Large Language Models*, con los que los usuarios pueden interactuar ha llevado al desarrollo de un gran número de modelos abiertos. Estos modelos se evalúan con múltiples conjuntos de pruebas para valorar sus capacidades para responder preguntas o resolver problemas sobre casi cualquier tema posible, o para probar su habilidad para razonar o interpretar textos. Sin embargo, la evaluación del conocimiento que estos modelos tienen de los idiomas ha recibido mucha menos atención. Por ejemplo, las palabras que pueden reconocer y usar en diferentes idiomas. En este artículo evaluamos el conocimiento que los LLM conversacionales abiertos tienen de las palabras en español utilizando una muestra de palabras de un diccionario de referencia. Los resultados muestran que los LLM conversacionales abiertos producen significados incorrectos para una fracción importante de las palabras y no son capaces de usar la mayoría de las palabras correctamente para escribir frases con contexto. Estos resultados muestran cómo el español se queda atrás en la carrera de los LLM de código abierto y destacan la necesidad de impulsar la equidad lingüística en los LLM conversacionales asegurando que proporcionen un rendimiento similar en todos los idiomas.

Palabras clave: LLM, vocabulario, modelos abiertos, español.

1 Introduction

The development of Large Language Models (LLMs) with billions of parameters and trained with huge amounts of text has pushed the performance of computer-based Natural Language Processing (NLP) to new limits achieving impressive performance in a wide range of tasks such as summarization, translation, retrieval of information, and conversation (Naveed et al., 2023). These LLMs are the foundation for systems that can interact with humans, such as the well-known ChatGPT (Wu et al., 2023), which is now used by hundreds of millions of people. These conversational LLM-based tools are being used to build autonomous agents that can perform many tasks interacting in an environment with humans and also with other agents (Wang et al., 2023b). These agents will soon replace traditional ChatBots (Adamopoulou and Moussiades, 2020) enabling a myriad of applications, from a personal assistant that manages our daily office tasks to a teaching assistant or a customer support agent.

As LLM-based tools and agents are adopted and eventually become widespread, the research community is focusing on evaluating their performance in a wide range of tasks and topics. Those include maths (Hendrycks et al., 2021b), reasoning (Zellers et al., 2019) but also comprehensive tests on a large number of tasks and topics (Hendrycks et al., 2021a), (Srivastava and et al, 2023). Instead, the linguistic aspects of the text generated by LLM-based tools have received much less attention, with only a few works exploring the linguistic features (Muñoz-Ortiz, Gómez-Rodríguez, and Vilares, 2023), phonological bias (Toro, 2023) or lexical diversity (Reviriego et al., 2023) of the AI-generated text. However, this is important as after all LLMs will likely account for a significant fraction of the text generated in the future.

In addition to the linguistic features, it is also of interest to understand the knowledge that LLMs have about the different languages, as they have been optimized in most cases for English. This may lead to a situation in which agents and other AI tools have worse performance for other languages, leading to an unfair scenario. This unfairness is twofold, firstly it will reinforce the dominant languages and secondly, it will put non-native English speakers at a disadvantage. There have been different efforts to develop truly multilingual

LLMs such as BLOOM (Workshop et al., 2022) or translation tools (Costa-jussà et al., 2022) but the performance of AI tools and models is still highly dependent on the language, as shown in (Reviriego and Merino-Gómez, 2022) for AI text to image generators.

To understand the importance given to different languages, a simple approach is to look at the amount of text used for training in each language. For example, in GPT-3 more than 181 billion words in English were used for training compared to only 1.5 billion in Spanish, representing 92.65% and 0.77% of the training set, respectively. This shows the prevalence of English in the training phase and more so if we take into account that Spanish is the fourth language with more words in the training dataset¹. The same trend is observed on other commercial models such as PaLM2 with only 2.11% of Spanish compared to 77.98% in English (Chowdhery et al., 2023). For open models, the situation is similar or even worse, for example, in LLaMa-2 only 0.11% of the training set is in Spanish (Touvron et al., 2023). In fact, for LLaMa models a recent work suggests that they work internally in English (Wendler et al., 2024).

The dominance of a single language in the training dataset of LLMs suggests that performance will be worse for other languages although the evaluation results, for example, of GPT-4 show a very similar performance for English and Spanish on the Measuring Massive Multitask Language Understanding (MMLU) benchmark (Hendrycks et al., 2021a) with an accuracy of 85.5% and 84% respectively (Achiam et al., 2023). However, the performance of LLMs in languages other than English and in some cases Chinese has received little attention compared to the amount of evaluation performed for the dominant languages. Therefore, both linguistic features and languages have not been the focus of the LLM evaluation to date.

One aspect that is of interest is how LLMs use the vocabulary of a language and the fraction of words that they can recognize (Martínez et al., 2023). On one hand, the vocabulary of LLMs may influence how languages evolve in the future, as words not used by LLMs may eventually be less and less used. On the other hand, the fraction of words in a

¹https://github.com/openai/gpt-3/blob/master/dataset_statistics/languages_by_word_count.csv

language that LLMs can recognize can be an indication of the degree of lexical knowledge that they have in different languages.

In this work, we contribute to the study of the lexical knowledge of LLMs by evaluating a representative group of open conversational LLMs from several companies and organizations and with different sizes on a sample of the words present in a reference Spanish dictionary (Seco, Andrés, and Ramos, 1999).

The rest of the paper is organized as follows, in section 2 we present the motivation and objectives of the work. Section 3 discusses the selection of the open conversational LLMs for evaluation and the evaluation methodology is described in section 4. The results are presented in section 5 followed by a discussion of the main insights as well as ideas for continuing this work in section 6. The paper ends with the conclusion in section 7.

2 Motivation and objectives

As discussed in the introduction, there is a lack of studies that analyze the linguistic aspects and, in particular, the lexical knowledge of LLMs and also of their understanding of languages other than English. This motivates our work, that tries to address both issues by studying the lexical knowledge of Spanish in open conversational LLMs. The rationale behind focusing on conversational LLMs is twofold. Firstly, they lend themselves to testing by asking the models directly if they know the meaning of words and if they are capable of using them meaningfully. Secondly, because they are widely used in many tools and agents. The choice of the language is given by us being native Spanish speakers but also because Spanish is typically in the top five languages with most words in the training set. This means that the results obtained will probably be better than for most other languages that have less data in the training set. Finally, we focus on open LLMs as there is a large number of models, with different sizes, and architectures, and even some of them are adapted for Spanish which enables a better evaluation of the impact of the different parameters on lexical knowledge. The evaluation of open models is also intended to provide feedback to the open-source community that could be used to improve the lexical knowledge of new conversational LLMs across different languages.

We set the following objectives in our study:

1. Evaluate the fraction of Spanish words that open conversational LLMs can recognize and use in context.
2. Provide an overview of the current lexical knowledge of Spanish in open conversational LLMs.
3. Analyze the impact of model size on the knowledge of the words.
4. Compare the knowledge of models designed to support multiple languages to those focused on English and Chinese.
5. Analyze whether the adaptation of the model to Spanish from a pre-trained model improves the knowledge of the words.

The methodology to test the conversational LLMs and evaluate the responses is described in the next section. The results can be processed to elaborate a summary of the lexical knowledge of conversational LLMs and how it depends on different factors. For example, the dependency on the model size can provide insights into whether acquiring the lexicon is limited by the training set, by model size, or possibly by both. Similarly, it is of interest to understand if models that are specifically designed to support several languages (Workshop et al., 2022) have a better knowledge of the Spanish vocabulary than models that are optimized for one or two languages². The same reasoning applies to the first conversational LLMs that are adapted to improve their performance in Spanish, do they have a better knowledge of the Spanish lexicon? Next, we discuss the selection of the conversational LLMs and the evaluation methodology used to pursue our objectives.

3 Models evaluated

There is a large (and growing) number of open models that have been designed to interact with users³. For example, filtering the models available at Huggingface to select only the ones for natural language processing in the category conversational returns more than 2,700 results. Although some of those may be the same model on different formats or with different quantization, the number of conversational

²For example <https://huggingface.co/01-ai/Yi-34B-Chat>.

³See for example at <https://huggingface.co/spaces/lmsys/chatbot-arena-leaderboard>

LLMs is clearly too large to make an exhaustive evaluation feasible. Therefore, the first step is the selection of a subset of conversational LLMs for evaluation.

The number of models in the evaluation set has to be manageable so that evaluating the models requires a reasonable effort but at the same time, it captures different types and sizes of conversational LLMs. In the selection, we focus on models that are widely used and with different sizes and target languages.

The evaluation set includes three LLaMa-2 (Touvron et al., 2023) models from Meta with different sizes ranging from 7B to 70B and two models from Mistral AI with 7B (Jiang et al., 2023) and 46.7B (Jiang et al., 2024) parameters. LLaMa and Mistral models have been used as a starting point to build other models by fine-tuning or adaptation such as Zephyr-7B (Tunstall et al., 2023), OpenChat (Wang et al., 2023a) or WizardLM (Xu et al., 2023). Therefore, the insights gained on LLaMa and Mistral models may be applicable to their derivative models. Solar10.7B (Kim et al., 2023) which improves Mistral-7B using a scaling technique to derive a larger model with close to 11B parameters is also included in our evaluation set to check if scaling has any impact on lexical knowledge. Two models recently introduced by 01.AI and optimized for English and Chinese⁴ with 6B and 34B that have shown good performance in several benchmarks are also included in the set. This first group of general models that are mostly targeted towards English and in some cases also Chinese is completed with Gemma-7B (Gemma Team et al., 2024) which has been recently released by Google.

A second group of models that are either multilingual by design or that have been optimized for Spanish is also included in the evaluation set. The goal is to check if these optimizations improve the lexical knowledge of Spanish. Three models are evaluated, the first two Bloomz-7B1⁵ (a model that has been trained specifically to support several languages) and Flor-6.3BInstructed⁶ (a model optimized for Spanish) are derived from Bloom7B, a multilingual open model (Workshop et al., 2022). The third, Bertin-6B is a

version of GPT-J 6B fine-tuned for Spanish (la Rosa and Fernández, 2022). These three models enable an evaluation of the changes in lexical knowledge introduced by optimizing the models for Spanish.

The twelve models evaluated and their main features are summarized in Table 1. The third column shows the quantization used to run the models in terms of bits per parameter. The parameters are available with 32-bit precision for most models and 32-bit operations are also supported by the GPUs. Therefore, when possible 32 bits are used, when there is not enough memory fewer bits are used. It can be observed that the LLMs selected cover a wide range of model sizes from 7 to 70 billion parameters from different companies. The quantization of the model parameters has been selected so that they can be run on a single GPU in our computing cluster⁷.

4 Evaluation methodology

This section describes the methodology used to evaluate the lexical knowledge of the conversational LLMs covering how to ask the models, what to ask in terms of words, and the processing and interpretation of the responses.

4.1 Prompts

Since all the models evaluated are conversational, we can ask them about the words directly using different prompts to assess if they know a given word and if they can use it meaningfully. The most straightforward procedure is to ask the models directly if they know the word or if it is correct. However, this may be misleading as LLMs are known to suffer hallucinations and provide inconsistent responses (Ye et al., 2023). Therefore, we initially use two prompts that ask the LLMs for the meaning of the word and to use the word to write sentences. In more detail, the prompts (and their English translations) used are:

- **Prompt A (meaning):** “*Dime la definición de la palabra ‘<word>’.*” (“Write the definition of the word ‘<word>’ ”)
- **Prompt B (use):** “*Escribe dos frases, una con la palabra ‘<word>’, y otra que no contenga esa palabra, pero que esté relacionada con la primera y complemente su significado.*” (“Write two sentences,

⁴<https://github.com/01-ai/Yi>

⁵<https://huggingface.co/bigscience/bloomz-7b1>

⁶<https://huggingface.co/projecte-aina/FLOR-6.3B-Instructed>

⁷We use NVIDIA A100 GPUs with 40GB of memory and RTX-A6000 GPUs with 48GB of memory.

Model	Parameters	Bits/Parameter
Llama-2-7b-chat-hf	6.7B	32
Llama-2-13b-chat-hf	13B	16
Llama-2-70b-chat-hf	69B	4
Mistral-7b-Instruct	7.2B	32
Mixtral-8x7b-Instruct	46.7B	4
Gemma-7b-it	8.54B	32
SOLAR-10.7b-Instruct	10.7B	16
Yi-6b-Chat	6B	32
Yi-34b-Chat	34.4B	8
Bloomz-7b1	7.1B	32
FLOR-6.3b-Instructed	6.3B	32
Bertin-6b	6B	32

Table 1: Conversational LLMs considered in the evaluation.

one with the word ‘<word>’, and another that does not contain that word, but that is related to the first and complements its meaning.”)

The first one is intended to assess if the model knows the meaning of the word and the second is to check if the model can use the word in a meaningful way. These prompts are based on assessment methods for human learners informed by Communicative Language Teaching, integrating the lexical approach (Lewis, 1993). The prompts are designed as ‘tasks’; in this case, small-scale written production exercises which ask for the generation of context. The ability to recognise and produce individual items will not suffice to complete these tasks successfully: in order to do so, the testee will rather have to create and use language as a functional part of a text. This integrative method of vocabulary assessment is aimed at evaluating comprehension, fluency, and accuracy, all key aspects of lexical competence (Leńko-Szymańska, 2019).

In addition to those two prompts with open answers, we test three additional prompts in which the models are asked to answer only “Yes” or “No”. The selected prompts and their English translations are as follows:

- **Prompt 1:** “¿Conoces el significado de la palabra ‘<word>’? Responde solo ‘Sí’ o ‘No’ y sé sincero, por favor.” (“Do you know the meaning of the word ‘<word>’? Answer only ‘Yes’ or ‘No’ and be honest, please.”)
- **Prompt 2:** “¿Existe la palabra ‘<word>’ en castellano? Responde solo ‘Sí’ o ‘No’ y sé sincero, por favor.” (“Does the word

‘<word>’ exist in Spanish? Just answer ‘Yes’ or ‘No’ and be honest, please.”)

- **Prompt 3:** “¿Es correcta la palabra ‘<word>’ en castellano? Responde solo ‘Sí’ o ‘No’ y sé sincero, por favor.” (“Is the word ‘<word>’ valid in Spanish? Answer only ‘Yes’ or ‘No’ and be honest, please.”)

These three prompts will enable us to assess if the models can be trusted when using “Yes/No” prompts. This is relevant because if these prompts were reliable, then an automatic evaluation of for example all the words in the dictionary could be easily run.

4.2 Test words

The best way to evaluate the knowledge of the words comprehensively seems to be to use all the words in a dictionary. Among the Spanish dictionaries, there are two that stand out above the rest. The first is the Diccionario de la lengua española (DRAE)⁸, by the Real Academia Española. This dictionary records 93,000 words from the 18th century onwards, and it is widely regarded as the main reference and authority in the language (Martínez de Sousa, 1995) (pp. 232-233).

The second dictionary to consider is the Diccionario del español actual (DEA)⁹, by Seco, Andrés and Ramos (Seco, Andrés, and Ramos, 1999). The DEA is a compilation of 84,000 words used in contemporary Spanish since 1950. It is based on actual documents and complements the definitions with examples of the words in use. This dictionary is

⁸<https://dle.rae.es/>

⁹<https://www.fbbva.es/diccionario/info/el-diccionario/>

considered exceptionally comprehensive and consistent and has been described as the most important lexicographical work in centuries (Álvarez de Miranda, 2011) (pp. 141-150).

As we wish to assess the model’s knowledge of the Spanish lexicon that is used by humans today, and to evaluate the knowledge of the words exhaustively, the DEA seems more appropriate for our objectives: it excludes outdated lexicon and presents detailed information about the current use of words.

Ideally, we would like to test all the words in the dictionary. However, that is only possible if we can trust the model responses, something that is not generally the case (Ye et al., 2023). Therefore, we start with a test set of 100 words that are randomly taken from the Diccionario del español actual and use them to conduct a manual evaluation as described next¹⁰.

4.3 Procedure

The models were loaded into our computing cluster and the five prompts were run sequentially for the 100 test words, setting the parameters of the models to obtain reproducible and deterministic results¹¹.

The first step of the analysis was to check manually the answers to prompts A (meaning) and B (use). This was done by an expert on languages with over 20 years of experience teaching Spanish in different Universities in the UK and Asia as well as a professional interpreter and writer. The answer produced by each model is compared with that in the dictionary and marked as valid if there is a match. For words with several meanings, we accept as positive a response that contains at least one of the meanings as long as it does not include additional incorrect meanings. The use of the word in the sentences with Prompt B is also checked manually. This result is assumed to be the ground truth on the knowledge that the model has of a word.

To assess if an automatic analysis is possible, in a second step the manual results are compared with those of prompts 1,2,3 to see if there is a correlation between the model answers to “Yes/No” prompts and the ground

truth. The idea is to check if having different responses on the prompts for the same word can be used as an indication that the responses are false positives and negatives. This would enable the evaluation of all the words in the dictionary, something that cannot be checked manually in a reasonable time.

Finally, an alternative approach to automate the analysis is evaluated by using ChatGPT to check the model responses. This is done 1) by asking ChatGPT if the meaning given by a model corresponds to the word and 2) by asking ChatGPT if the meaning given by a model matches the one in the dictionary for that word. In more detail, these checking prompts used are:

- **Check 1:** “*Es correcta la definición <definition> para la palabra <word>? Responde solo sí o no*” (“Is the definition <definition> for the word <word> correct? Answer yes or no only.”)
- **Check 2:** “*¿Son equivalentes las siguientes definiciones? <1. definition> <2. definition>.*” (“Are the following definitions equivalent? <1. definition> <2. definition>”) with the first definition being the one in the dictionary and the second the one given by the model.

again, the responses to those checking prompts will be correlated to the ground truth to see if the automatic checks are reliable.

5 Results

In this section, the results for the manual testing are presented first to then analyze the feasibility of using the “Yes/No” prompts or the checking by ChatGPT to automate the process.

5.1 Manual evaluation

We start by presenting the summary of the results for the manual analysis in Table 2. It can be seen that all models fail to produce a valid meaning for a significant fraction of the words and are unable to construct meaningful sentences using the words in context for most of the words. This clearly shows that there is room for improvement in the lexical knowledge of Spanish in open conversational LLMs. Analyzing the results in more detail, the following observations can be made:

1. **Observation 1: Valid meanings are below 50%.** Two-thirds of the models

¹⁰The list of the words is included in section A and the dataset with the raw data from the 100 words in <https://zenodo.org/doi/10.5281/zenodo.10797991>

¹¹This was done by disabling sampling or setting the temperature to zero depending on the model.

do not produce valid meanings for more than half of the words evaluated. Only one-third of the models can reach 50% with the best achieving only 66%.

2. **Observation 2: Correct word usage is below 25%.** Only one model reaches 25% correct usage of the words and the majority of the models are below 10%.
3. **Observation 3: Performance improves with model size.** Both meaning and usage increase with model size for Llama, Mistral, and Yi, suggesting that larger models can better handle the lexicon.
4. **Observation 4: Adaptation to Spanish does not improve performance.** Models that have been designed to support multiple languages (such as Bloomz) or that have been adapted or fine-tuned for Spanish (Flor and Bertin) have lower scores in both meanings and usage than Llama or Mistral models of the same size.

Looking at the results for each model individually, the following observations can be made which show limitations of the models and a tendency towards English:

1. Responses to prompt A (definition) and prompt B (language in use) are not necessarily in line; i.e., the meaning that the definition refers to and the text produced are frequently not consistent. For example, ‘tundidor’ is defined as a shy person in the answer to prompt A, but the word is used for describing a type of machine in prompt B (Bloomz-7b1); ‘sagum’ is long hair in the definition, but a piece of cloth in the text produced (Mistral-7b); etc. This behaviour is frequent across all models.
2. Grammatical mistakes are rare, but when these occur, they are related exclusively to agreements (masculine/feminine, singular/plural: “la olor” [“la” is feminine and “olor” is masculine], “que no se pierdan nadie” [“pierdan” is the verb form for plural, but for “nadie” the singular form should be used], “los sílabas” [“los” is masculine and “sílabas” is feminine], “el yema” [“el” is masculine and “yema” is feminine]. . .).
3. Structures are heavily influenced by English, which results in unusual constructions (“finalmente fue completado” [finally completed], “socialmente responsable” [socially responsible], “era frecuentemente invitado” [was frequently invited]. . .).
4. Most models tend to include English in their responses, producing answers fully or partially written in that language: “no se puede *missed*” (it can’t be missed), “prolongado *enough* to be measurable” (extended enough to be measurable), “una cola larga y *bushy*” (a long and bushy tail), etc. Yi-34B also responds including Russian and Chinese words.
5. The frequency of non-existent hybrid words mixing English and Spanish is high: “Ella es *meana*” (for “she is mean”), “hoefados” (meaning “with a hoof”), “si *meantes*” (“if you mean”). . . Most of these words are incomprehensible: “orato”, “cérama”, “metingse”, “lasufación”, “conmanes”.
6. All models are unable to identify the use of foreign words in Spanish. When asked about those terms, they answer directly in the language the word comes from, and refer to the meaning in that language alone (“speed” is only “fast pace”, not a drug; “apparat” is defined mechanically, not politically, etc.).
7. Definitions are sometimes mere descriptions of a word’s features (for example, Bertin-6b just provides grammatical category and even number of letters). Many definitions refer to the (often incorrect) translation of the word, instead of describing its meaning (Yi-6b, for example, responds that “to *acelguilla*” [a plant] “puede significar to adjust o to fit” [may mean to adjust or to fit], “ardilla” [squirrel] “significa rabbit en inglés” [means rabbit in English]; etc).
8. Some models (Llama-2-13b, Llama-2-7b) provide unsolicited etymological information which is mostly inaccurate: “ ‘verrax’ means urine in Latin”, “ ‘colosseum’ comes from the Latin ‘column’ ”; “ ‘cosmo’ is a prefix referring to a unit of measurement”; etc.).
9. Common words (“minute”, “tie”, “cent”, “thirty”) are often defined incorrectly.

Model	Word meaning	Word use
Llama-2-7b-chat-hf	42	3
Llama-2-13b-chat-hf	46	20
Llama-2-70b-chat-hf	50	25
Mistral-7b-Instruct	51	6
Mixtral-8x7b-Instruct	66	17
Gemma-7b-it	20	4
SOLAR-10.7B-Instruct	59	11
Yi-6B-Chat	27	2
Yi-34B-Chat	45	16
Bloomz-7b1	38	0
FLOR-6.3b-Instructed	39	0
Bertin-6b	19	4

Table 2: Number of correct answers in the manual evaluation (in bold values above 50%).

It is also of interest to study if the errors made by the models are concentrated on the same words or if the errors are more dependent on the model with some models failing in some words and others in different words. Figure 1 shows the number of words that fail on a given number of models. It can be seen that the maximum occurs for failures on all (twelve) models but there are also words in each of the other categories. Therefore, there is a correlation between the failures of the models but not in all cases. The number of words failing per model is shown in Figure 2 showing also significant variability among models. The results for each word and model are included in the annex in section A.




Figure 1: Number of words failing on a number of models.

To further analyze the words that are not correctly identified by the models, we checked the frequencies of the 100 test words in the Corpus de Referencia del Español Ac-




Figure 2: Number of words failing on each model.

tual (CREA)¹². Figure 3 shows the words as points in a plot with the word frequency on the x-axis and the number of failures in the 12 models on the y-axis. Most failures would be expected for words with low frequencies that are less used and thus will most likely appear fewer times in the training datasets. Therefore, a strong correlation between frequency and failures would be expected with the points concentrating towards zero (no failures) as frequency increases. The 100-word frequency and error distributions do not fit the Pearson assumptions, so the Spearman correlation was computed. The results suggest a statistically significant moderate negative monotonic relationship between the variables

¹²Real Academia Española: Banco de datos (CREA) Corpus de referencia del español actual <https://corpus.rae.es/1frecuencias.html> consultado en febrero de 2024.

($\rho = -.47$, p-value $< .0001$). However, there is a significant number of models that do not produce a valid meaning for the most common words in the test set. For example, the word with the largest frequency, “minuto” fails in the meaning for eight of the twelve models. This suggests that the presence of the word in the training dataset does not guarantee that the model will be able to produce a valid meaning for it and further shows the limitation of the models in Spanish.




Figure 3: Number of models failing per word versus word frequency (to represent words that do not appear in CREA, we plot them with a frequency of 10^{-3} as the frequency is represented in a logarithmic scale.).

5.2 Automation feasibility

After analyzing the manual results we consider whether the “Yes/No” prompts or the checks by ChatGPT can be used to automate the evaluation of the lexical knowledge so that all words in the dictionary can be evaluated. To do so we compare the results of prompts 1 to 3 and checks 1 and 2 with the ground truth using ChatGPT4 and show the accuracy for each of them. The results are summarized in table 3. It can be seen that the “Yes/No” prompts have poor accuracy for most models. In fact, the models that have the highest accuracy, for example Gemma with 80% on prompts 1 to 3, answer “No” for all words which provides no useful information on the knowledge of words. Therefore, the “Yes/No” prompts cannot be used to conduct evaluation at scale as the results would not be reliable. The accuracy for the ChatGPT-based checks is better, specially for check 2 that reaches in some cases 90%. However, even for these checks accuracy is below 70% for some models. Therefore, overall the values are not still high enough to be used

for testing at scale.

6 Discussion

The evaluation results show that open conversational LLMs have limited knowledge of the Spanish lexicon. This could be expected as most models have been trained with datasets that are dominated by English or Chinese. However, similar results are observed for models that have been designed to target many languages or fine-tuned for Spanish. This suggests that current efforts to support Spanish in LLMs are not effective in terms of lexical knowledge. In fact, it would be of interest to do the same evaluation in English and compare the results as we are taking for granted that performance in English will be better.

The evaluation also shows that automating the evaluation of lexical knowledge is not straightforward. This is a serious limitation as manual checking is only practical for a reduced set of words. Further work is needed to improve automated testing, for example analyzing the mistakes made by ChatGPT in making the judgements may provide insights on how to improve the checking.

In fact, the results presented in this paper for 100 words provide a first estimate of lexical knowledge but more words, ideally thousands, should be evaluated to have better estimates. Another limitation of our work is that results may depend on the prompts used, several prompts were tested and the differences observed were not large. However, the use of additional prompts to check the consistency of the results is also of interest. The use of more advanced prompting techniques such as chain of thoughts would also be of interest, especially when the models are asked about their knowledge of words as LLMs have limited meta-cognition capabilities. Again, having to conduct a manual analysis of 100 words on 12 models makes any additional testing cumbersome. Another limitation of our study is that only three multilingual models have been tested as there are few such models. As newer multilingual models appear, it would be interesting to see if they improve their lexical knowledge of Spanish.

Despite the limitations of our study, the limited lexical knowledge of Spanish in the open conversational LLM ecosystem is evident, even for the multilingual models evaluated. Most likely, the same applies to other languages that have even less presence on the

Model	P1 (%)	P2 (%)	P3 (%)	C1 (%)	C2 (%)
Llama-2-7b-chat-hf	60	64	67	65	81
Llama-2-13b-chat-hf	56	53	54	73	90
Llama-2-70b-chat-hf	52	62	54	72	86
Mistral-7b-Instruct	53	66	74	82	86
Mixtral-8x7b-Instruct	38	39	38	70	86
Gemma-7b-it	80	80	80	56	57
SOLAR-10.7B-Instruct	56	48	48	72	87
Yi-6B-Chat	48	73	75	74	76
Yi-34B-Chat	39	51	59	62	77
Bloomz-7b1	38	38	38	77	75
FLOR-6.3b-Instructed	66	62	61	67	67
Bertin-6b	81	77	19	73	67

Table 3: Accuracy of the "Yes/No" checks (P1,P2,P3) and ChatGPT4 checks (C1,C2).

training datasets. Therefore, an effort should be made by the open-source community to develop conversational LLMs with better lexical knowledge of Spanish. This is important given the large number of native speakers of Spanish and the expected impact of conversational LLMs in the future of languages.

7 Conclusion

In this paper, we have analyzed the lexical knowledge that open conversational LLMs have of Spanish. The results show that most LLMs fail to recognize more than half of the words tested and are unable to use most of the words in context when asked to do so with a simple prompt. This evaluation is done on a small sample of one hundred words randomly taken from a reference Spanish dictionary and thus is only an initial estimate of the lexical knowledge. However, even when taking that into account, the results show the limitation of open conversational LLMs in using the Spanish lexicon. This applies even to models that have been adapted or fine-tuned for Spanish.

The automation of the lexical evaluation at scale has also been studied by using prompts with binary answers that can be processed automatically or by using a commercial conversational LLM as the evaluator judging if the responses are valid or comparing the meanings of the words produced by the LLM tested with those in the dictionary. The results show that automating the testing incurs a significant loss of accuracy as the LLMs produce false results for binary prompts frequently and the checking by another conversational LLM also has a relevant number of errors. Developing techniques to automate the testing of words is interesting as it would enable the evaluation of for example all the words in the dictionary.

The initial study presented in this paper

can be extended by considering additional LLMs, a larger set of test words, and different prompts for the evaluation. Another interesting topic is the development of automated tests for lexical knowledge that overcome the challenges described in this paper. More generally and more importantly, our results suggest that the open-source LLM community should make an effort to better support languages other than English.

8 Acknowledgements

This work was supported by the FUN4DATE (PID2022-136684OB-C21/C22) project funded by the Spanish Agencia Estatal de Investigacion (AEI) 10.13039/501100011033, by the Chips Act Joint Undertaking project SMARTY (Grant no. 101140087) and by the OpenAI Researcher Access Program. The evaluation was also done in part with equipment that was donated by NVIDIA to support our research.

References

- Achiam, J., S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Adamopoulou, E. and L. Moussiades. 2020. An overview of chatbot technology. In *IFIP international conference on artificial intelligence applications and innovations*, pages 373–383. Springer.
- Álvarez de Miranda, P. 2011. *Los diccionarios del español moderno*. TREA.
- Chowdhery, A., S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann, et al. 2023. Palm: Scaling

- language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.
- Costa-jussà, M. R., J. Cross, O. Çelebi, M. Elbayad, K. Heafield, K. Heffernan, E. Kalbassi, J. Lam, D. Licht, J. Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Gemma Team, T. M., C. Hardin, R. Dadashi, S. Bhupatiraju, L. Sifre, M. Rivière, M. S. Kale, J. Love, P. Tafti, L. Hussenot, and et al. 2024. Gemma. *Kaggle*.
- Hendrycks, D., C. Burns, S. Basart, A. Zou, M. Mazeika, D. Song, and J. Steinhardt. 2021a. Measuring massive multitask language understanding. In *International Conference on Learning Representations*.
- Hendrycks, D., C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, and J. Steinhardt. 2021b. Measuring mathematical problem solving with the math dataset.
- Jiang, A. Q., A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. d. l. Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Jiang, A. Q., A. Sablayrolles, A. Roux, A. Mensch, B. Savary, C. Bamford, D. S. Chaplot, D. d. l. Casas, E. B. Hanna, F. Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Kim, D., C. Park, S. Kim, W. Lee, W. Song, Y. Kim, H. Kim, Y. Kim, H. Lee, J. Kim, et al. 2023. Solar 10.7 b: Scaling large language models with simple yet effective depth up-scaling. *arXiv preprint arXiv:2312.15166*.
- la Rosa, J. D. and A. Fernández. 2022. Zero-shot reading comprehension and reasoning for spanish with BERTIN GPT-J-6B. In M. M. y Gómez, J. Gonzalo, F. Rangel, M. Casavantes, M. Ángel Álvarez Carmona, G. Bel-Enguix, H. J. Escalante, L. Freitas, A. Miranda-Escalada, F. Rodríguez-Sánchez, A. Rosá, M. A. Sobrevilla-Cabezudo, M. Taulé, and R. Valencia-García, editors, *Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2022)*. CEUR Workshop Proceedings.
- Lewis, M. 1993. *The Lexical Approach: The State of ELT and a Way Forward*. LTP teacher training. Language Teaching Publications.
- Leńko-Szymańska, A. 2019. *Defining and Assessing Lexical Proficiency*. Routledge, 11.
- Martínez, G., J. Conde, P. Reviriego, E. Merino-Gómez, J. A. Hernández, and F. Lombardi. 2023. How many words does chatgpt know? the answer is chatwords. *arXiv preprint arXiv:2309.16777*.
- Martínez de Sousa, J. 1995. *Vox: diccionario de lexicografía práctica*. Vox.
- Muñoz-Ortiz, A., C. Gómez-Rodríguez, and D. Vilares. 2023. Contrasting linguistic patterns in human and llm-generated text. *arXiv preprint arXiv:2308.09067*.
- Naveed, H., A. U. Khan, S. Qiu, M. Saqib, S. Anwar, M. Usman, N. Barnes, and A. Mian. 2023. A comprehensive overview of large language models. *arXiv preprint arXiv:2307.06435*.
- Reviriego, P., J. Conde, E. Merino-Gómez, G. Martínez, and J. A. Hernández. 2023. Playing with words: Comparing the vocabulary and lexical richness of chatgpt and humans. *arXiv preprint arXiv:2308.07462*.
- Reviriego, P. and E. Merino-Gómez. 2022. Text to image generation: leaving no language behind. *arXiv preprint arXiv:2208.09333*.
- Seco, M., O. Andrés, and G. Ramos. 1999. *Diccionario del español actual*, volume 2. Aguilar Madrid.
- Srivastava, A. and et al. 2023. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.
- Toro, J. M. 2023. Emergence of a phonological bias in chatgpt. *arXiv preprint arXiv:2305.15929*.
- Touvron, H., L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Tunstall, L., E. Beeching, N. Lambert, N. Rajani, K. Rasul, Y. Belkada, S. Huang,

- L. von Werra, C. Fourier, N. Habib, et al. 2023. Zephyr: Direct distillation of lm alignment. *arXiv preprint arXiv:2310.16944*.
- Wang, G., S. Cheng, X. Zhan, X. Li, S. Song, and Y. Liu. 2023a. Openchat: Advancing open-source language models with mixed-quality data. *arXiv preprint arXiv:2309.11235*.
- Wang, L., C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, J. Tang, X. Chen, Y. Lin, et al. 2023b. A survey on large language model based autonomous agents. *arXiv preprint arXiv:2308.11432*.
- Wendler, C., V. Veselovsky, G. Monea, and R. West. 2024. Do llamas work in english? on the latent language of multilingual transformers. *arXiv preprint arXiv:2402.10588*.
- Workshop, B., T. L. Scao, A. Fan, C. Akiki, E. Pavlick, S. Ilić, D. Hesslow, R. Castagné, A. S. Luccioni, F. Yvon, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100*.
- Wu, T., S. He, J. Liu, S. Sun, K. Liu, Q.-L. Han, and Y. Tang. 2023. A brief overview of chatgpt: The history, status quo and potential future development. *IEEE/CAA Journal of Automatica Sinica*, 10(5):1122–1136.
- Xu, C., Q. Sun, K. Zheng, X. Geng, P. Zhao, J. Feng, C. Tao, and D. Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*.
- Ye, H., T. Liu, A. Zhang, W. Hua, and W. Jia. 2023. Cognitive mirage: A review of hallucinations in large language models. *arXiv preprint arXiv:2309.06794*.
- Zellers, R., A. Holtzman, Y. Bisk, A. Farhadi, and Y. Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *Annual Meeting of the Association for Computational Linguistics*.

A Annex: Test words and results

The words used in the evaluation are listed in Table 4 and the results for each of them for each model are shown in Figure 4 with the number of models failing in Figure 5. Additional details and the text generated by the models are available in the public repository.

acelguilla	agüista	antidiarreico	apealar	apparat	arante	ardilla
bátavo	bicicross	bifocal	cantautor	cantintero	cartulina	centavo
cerebrotónico	chalaza	chigiüre	coliseo	conspirar	corbata	corralero
cosmotrón	crístico	cuentapartícipe	dabuti	dactiloscopia	deformabilidad	desinfección
desinsectación	desmitificador	diligentemente	emparejador	empurpurado	epifito	escorar
estadista	esteatosis	estuco	exequias	faisánido	fétido	floración
fotogramétrico	funcionario	gabato	garcilla	giroscopio	helenizante	hipogino
incrustar	intercadencia	jaín	lipotimia	magnolia	manes	meano
mediar	mensualizar	mesmerización	mestizo	minuto	mochalero	modisto
monásticamente	morra	nefólogo	novatada	ovni	pagar	paletteo
palmítico	paralogismo	pasarratos	perrillo	pezuña	pinabete	pitahaya
postrar	prédica	prolongador	provinciano	puzzle	quepis	raor
reciclado	rememorar	ridiculización	sagum	salbanda	socialmente	speed
standarizar	sublimado	superbomba	talgo	tornajo	treinta	tundidor
vega	verraquear					

Table 4: Subset of words used in the evaluation. Those failing in nine or more models are in red, those failing in four or fewer models in ForestGreen and the rest in orange.




Figure 4: Models failing per word (red for failure, green for correct meaning).



Figure 5: Number of models failing per word (red for failure, green for correct meaning).