



OPEN Testing AI on language comprehension tasks reveals insensitivity to underlying meaning

Vittoria Dentella^{1,2✉}, Fritz Günther³, Elliot Murphy⁴, Gary Marcus⁵ & Evelina Leivada^{6,7}

Large Language Models (LLMs) are recruited in applications that span from clinical assistance and legal support to question answering and education. Their success in specialized tasks has led to the claim that they possess human-like linguistic capabilities related to compositional understanding and reasoning. Yet, reverse-engineering is bound by Moravec's Paradox, according to which easy skills are hard. We systematically assess 7 state-of-the-art models on a novel benchmark. Models answered a series of comprehension questions, each prompted multiple times in two settings, permitting one-word or open-length replies. Each question targets a short text featuring high-frequency linguistic constructions. To establish a baseline for achieving human-like performance, we tested 400 humans on the same prompts. Based on a dataset of $n = 26,680$ datapoints, we discovered that LLMs perform at chance accuracy and waver considerably in their answers. Quantitatively, the tested models are outperformed by humans, and qualitatively their answers showcase distinctly non-human errors in language understanding. We interpret this evidence as suggesting that, despite their usefulness in various tasks, current AI models fall short of understanding language in a way that matches humans, and we argue that this may be due to their lack of a compositional operator for regulating grammatical and semantic information.

Large Language Models (LLMs) are neural networks trained on generating probability distributions over natural language data. Through interfaces that allow direct interaction with users, LLMs (i.e., the underlying models together with their respective interface setup) perform tasks that span from translation to answering a wide range of general queries¹, spanning domains as diverse as law², medicine^{3,4}, and chemistry⁵. Yet, good performance in tasks that require memorizing specialized knowledge is not necessarily grounded in a solid understanding of language, such that LLMs may fail at comparatively easier tasks (*Moravec's paradox*^{6–8}). Reverse-engineering is harder for simple, effortless tasks that our minds do best⁹; and understanding language is easy for humans, with even eighteen-month-olds demonstrating an understanding of complex grammatical relations¹⁰. Our species is characterized by an irrepressible predisposition to acquire language, to seek meaning beneath the surface of words, and to impose on linear sequences a surprisingly rich array of hierarchical structure and relations¹¹.

Based on their success in various applied tasks, featured in a wide array of testing benchmarks [see e.g., ^{12–19}], LLMs have been linked with human-like capabilities such as advanced reasoning (OpenAI on ChatGPT and GPT-4²⁰), understanding across modalities (Google on Gemini²¹), and common sense¹. Some scholars have claimed that LLMs approximate human cognition, understand language to the point that their performance aligns with or even surpasses that of humans, and are good models of human language^{22–24}. However, much evidence exists pointing to possible inconsistencies in model performance. Models can both produce highly fluent, semantically coherent outputs^{25–27,1} whilst also struggling with fundamental syntactic or semantic properties of natural language^{28–34}. Observing errors in the linguistic performance of LLMs begs the question of how their application in various tasks, such as replying to a medical or legal question, is affected by possible language limitations, and may in fact be achieved through some series of computational steps entirely remote from a cognitively plausible architecture for human language³⁵.

While shortcomings in LLM performance do not negate their usefulness as tools for specific applied uses, they do invite serious concerns over what guides output generation: is it (i) context-sensitive text parsing (i.e., the ability to pair specific linguistic forms with their corresponding meanings) in a way that generalizes across contexts^{36,37}, or (ii) the mechanistic leveraging of artifacts present in the training data³⁸, hence giving merely the illusion of competence? At present, the dominant means of evaluating LLMs consists in employing their

¹Universitat Rovira i Virgili, Tarragona, Spain. ²University of Pavia, Pavia, Italy. ³Humboldt-Universität zu Berlin, Berlin, Germany. ⁴UTHealth, Houston, TX, USA. ⁵New York University, New York, USA. ⁶Autonomous University of Barcelona, Barcelona, Spain. ⁷Institució Catalana de Recerca i Estudis Avançats, Barcelona, Spain. ✉email: vittoria.dentella@unipv.it

(well-formed) outputs as a basis to infer human-like language skills (*modus ponens*, for³⁹). For instance, accurate performance in language-oriented tasks and benchmarks [e.g.,^{40–42}] is used to conclude that the LLMs have succeeded not only at the specific task they performed, but also at learning the general knowledge necessary in order to carry out that task [cf.⁴³]. This line of reasoning lies at the basis for holding LLMs as cognitive theories^{23,44}.

On the other hand, assuming competence for LLMs over all morphosyntactic, semantic and pragmatic processes involved in language comprehension clashes with their inability to consistently deploy the knowledge that is attributed to them^{31,32,44}. We list the grammaticality judgments (i.e., judgments about whether a prompt complies with or deviates from the language patterns that the model has internalized) of two frontier LLMs on linguistic prompts that are less frequent in day-to-day language and, consequently, are considerably less likely to form part of the training data (Fig. 1). While in humans the influence of cognitive factors (e.g., working memory limitations or distraction) can affect language processing in ways that may result in non-target grammaticality judgments, humans are able to correctly process the stimuli upon reflection; that is, we can engage in ‘deep’ processing after an initial ‘shallow’ parse⁴⁵. For LLMs, however, there is no analogous straightforward explanation of their systematic language errors.

While it is incontestable that LLMs offer a potentially useful interactive tool that mimics human conversation, it remains to be established whether their ability to understand language is on a par with that of humans. To answer, we investigate the ability of 7 state-of-the-art LLMs in a comprehension task that features prompts whose linguistic complexity is purposely kept at a minimum, e.g. (1).

(1) John deceived Mary and Lucy was deceived by Mary. In this context, did Mary deceive Lucy?

The benefit of this methodology is that it tests LLMs in the very task for which they have been trained: respond to questions based on the input. While previous research testing LLMs on grammaticality judgments has found that their performance is not comparable to that of humans³¹, the results of this reported method are not unequivocally convincing. Some have argued that judgment prompting is not a suitable methodology for assessing LLMs, because this is (putatively) a metalinguistic skill that LLMs do not necessarily possess^{24,46}. Thus, in the present work we test LLMs in their default ability to answer questions based on input. Given their use in a range of meaning-sensitive tasks, from giving medical diagnoses to travel advice and to being “thought partners” to students^{3,47}, it is expected that LLMs should be able to accurately parse prompts featuring high-frequency verbs and constructions.

Why is this a task that models should be expected to perform reasonably well? While models are trained to predict tokens⁴⁸, when taken together with their respective interface setup, their abilities are marketed as

Phenomenon	Prompt	GPT-3	ChatGPT-3.5	Linguistic analysis
Identity Avoidance violation	Dogs dogs dog dog dogs			Both prompts are grammatical and have the same structure: in English, <i>dog</i> can be either a noun (a dog) or a verb (to dog). The first prompt can be paraphrased as <i>Dogs that are followed by dogs follow themselves other dogs</i> .
	Fish fishermen catch eat worms			
Agreement attraction	*The key to the drawers are on the table			Both prompts are ungrammatical and have the same structure: the singular subject (<i>the key</i> and <i>the village</i> , respectively) does not agree with the plural verb (<i>are</i> and <i>have met</i>).
	*The village of the fundamentally flawed ideas have not met the criteria			
Comparative sentences	*More people have been to Russia than I have			Both prompts are ungrammatical: the subject of the main clause (<i>more people</i> and <i>fewer people</i> , respectively) require a cardinal set for comparison purposes, whereas the subject of the embedded clause, i.e., <i>I</i> , is singular.
	*Fewer people have been to Russia than I have			
Missing verb	The patient the nurse the clinic had hired admitted met Jack			The second prompt is ungrammatical, as one verb phrase is missing. Under the processing demands induced by center embedding, this violation may go unnoticed.
	*The doctor the nurse the clinic had hired met Jack			
Semantic anomalies	When an airplane crashes at the border of two countries and debris is scattered on the ground of both, where should the survivors be buried?			Both prompts are semantically anomalous: survivors are not buried. In the first prompt, the anomaly is induced at the sentence level; in the second prompt, the violation is rendered through a locally anomalous noun phrase (i.e., <i>surviving dead</i>).
	If a plane crashed on the border of US and Canada, where should we bury the surviving dead?			

Fig. 1. GPT-3 (text-davinci-002) and ChatGPT-3.5 performance on a grammaticality judgment task that involves low-frequency constructions. Inaccurate responses are marked in red. Accurate responses are marked in green. Full model replies and detailed linguistic analyses are available (<https://osf.io/dfgmr/>).

going significantly beyond next-token prediction; they are portrayed as fluent conversational agents that display long-context understanding across modalities (e.g.,²¹ on Gemini). While LLMs, understood in the present work as the underlying model together with the interface setup, generate content that looks remarkably like human-to-human conversation, it remains to be established whether their ability to understand language is on a par with that of humans. This is a matter with large-scale consequences for the applied use of LLMs. To provide an example, an air company was recently asked to pay damages to a passenger who was provided with inaccurate information by its chatbot. According to a company representative, the chatbot indeed included “misleading words” in its replies⁴⁹. While the court issued a ruling of negligent misrepresentation in favor of the passenger, the company maintained that the chatbot is a separate legal entity with reasonable linguistic abilities, and thus *responsible* for its own words. What remains to be established is the cause of the error. Do chatbots and other interactive applications that rely on the same token-predicting technology possess a human-like language understanding of written language or is their understanding capacity limited and consequently likely to give rise to such errors?

To this end, we asked whether LLMs’ language performance is on a par with that of humans in language comprehension task. Two research questions (RQs) emerge from this general aim: **RQ1**: Do LLMs reply to comprehension questions accurately? **RQ2**: When the same question is asked multiple times, do LLMs reply consistently?

Results

Results are reported separately for *accuracy* (RQ1) and *stability* (RQ2). First, LLMs are compared between each other on these two variables. Then, their performance is contrasted with that of humans. (Generalized) Linear Mixed Effect Models ((G)LMMs) are used to analyze the data, using the `lme4`⁵⁰ and `lmerTest`⁵¹ packages in R. LLM and human datasets, together with the annotated code employed in the analyses are available (<https://osf.io/dfgmr/>).

LLMs: accuracy

Accuracy by model and setting Accuracy rates by model and setting (open-length vs. one-word) are displayed in Fig. 2A. The intercept in an intercept-only GLMM ($\text{accuracy} \sim 1 + (1 | \text{test_item})$) is not significantly different from zero ($\beta = 0.224$, $z = 1.71$, $P = .087$), indicating at-chance performance when all LLMs are taken together. A first chi-square test between a GLMM with a main effect of model ($\text{accuracy} \sim \text{model} + (1 | \text{test_item})$) and this intercept-only GLMM however finds significant differences in the performance of the 7 LLMs, justifying the introduction of a model parameter in the GLMM ($X^2(6) = 202.8$, $P < .001$). The intercept of a GLMM with Bard set as baseline category is negative ($\beta = -0.688$, $z = -3.433$, $P = .001$), indicating that the performance of at least one LLM is below chance. By systematically changing the baseline category, we find that Mixtral is the first LLM to perform at chance ($\beta = 0.103$, $z = 0.524$, $P = .600$). With a mean accuracy of 0.66, Llama2 is the first LLM to perform above chance ($\beta = 0.752$, $z = 3.720$, $P < .001$), alongside Falcon ($M = 0.66$) and ChatGPT-4 ($M = 0.83$). Further, we find that ChatGPT-4 has a significantly higher accuracy than all other LLMs: When setting ChatGPT-4 as the baseline category, all parameters for all other models are significant ($P_s < 0.001$) and negative ($\beta_s < -1.01$, $z_s < -4.31$).

A second chi-square test between a GLMM with a main effect of setting ($\text{accuracy} \sim \text{model} + \text{setting} + (1 | \text{test_item})$) and a GLMM without it ($\text{accuracy} \sim \text{model} + (1 | \text{test_item})$) finds a significant main effect of setting ($X^2(1) = 82.587$, $P < .001$), with more accurate responses in the one-word setting than in the open-length setting ($\beta = 1.037$, $z = 8.872$, $P < .001$). We additionally find an interaction between model and setting ($X^2(6) = 43.374$, $P < .001$), indicating that most but not all LLMs perform more poorly in the open-length setting.

LLMs: stability

In order to establish the extent to which the LLMs provide stable responses, we performed an analysis at the item level (as opposed to the individual trial level for the accuracy data). Here, the responses to an item were coded as

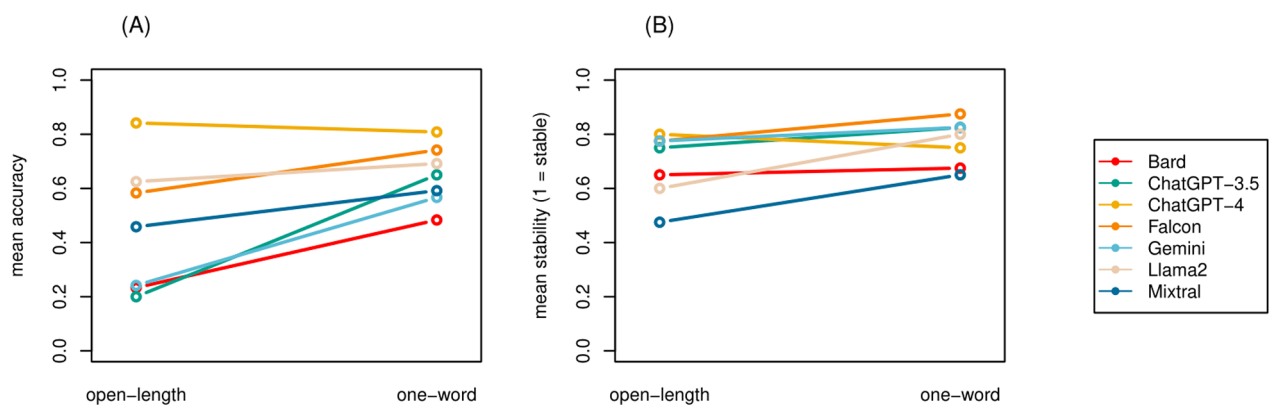


Fig. 2. (A) Mean accuracy by model and setting. (B) Mean stability by model and setting.

stable (= 1) if all three responses to a prompt were the same, and as unstable (= 0) if one response was different than the other two.

Stability by model and setting Stability rates by model and setting are displayed in Fig. 2B. As was the case for accuracy, a GLMM including a parameter for model significantly outperforms a GLMM without it ($X^2(6) = 21.438$, $P = .002$). In a GLMM with Bard (the LLM with the lowest stability, $M = 0.66$) as baseline category, the intercept is positive ($\beta = 0.610$, $z = 2.810$, $P = .005$), indicating that even for the most unstable model, the responses are not fully random (i.e., there is a higher likelihood of stable than unstable response patterns). Additionally, the model parameters show that stability is significantly higher for Falcon ($\beta = 0.901$, $z = 2.357$, $P = .018$; $M = 0.83$) and Gemini ($\beta = 0.733$, $z = 1.976$, $P = .048$; $M = 0.80$). Taken together with accuracy results, significantly higher stability for Falcon and Gemini respectively translates into Falcon being partially consistent in providing accurate replies, and Gemini being partially consistent in providing inaccurate replies. Employing Falcon (the most stable LLM) as baseline category for the GLMM, we find that Bard ($\beta = -0.901$, $z = -2.357$, $P = .018$) and Mixtral ($\beta = -1.340$, $z = -3.559$, $P < .001$) display less stable response patterns than Falcon, while all the other LLMs perform comparably to Falcon.

Similar to accuracy, a second chi-square test between a GLMM with a main effect of setting (stability ~ model + setting + (1 | test_item)) and a GLMM without it (stability ~ model + (1 | test_item)), finds a significant main effect of setting on stability ($X^2(1) = 5.160$, $P = .023$), with answers being more stable in the one-word setting ($\beta = 0.450$, $z = 2.262$, $P = .024$). Lastly, we find no significant interaction between model and setting, indicating that the effect of setting does not differ between the different LLMs ($X^2(6) = 4.241$, $P = .644$).

LLMs vs. humans

Comparative analyses showed major differences in the performance of the two responding agents (humans vs. LLMs) in terms of both accuracy and stability (Fig. 3).

Accuracy by agent and setting Accuracy rates by responding agent and setting are displayed in Fig. 3A. A first chi-square test finds that a GLMM with a main effect of agent (accuracy ~ setting + agent + (1 | item) + (1 | unit)) significantly outperforms a GLMM without it (accuracy ~ setting + (1 | item) + (1 | unit)), justifying the introduction of an agent parameter in the GLMM ($X^2(1) = 23.356$, $P < .001$). Additionally, a GLMM with an interaction between agent and setting (accuracy ~ setting * agent + (1 | item) + (1 | unit)) outperforms the GLMM that only includes them as main effects ($X^2(1) = 45.265$, $P < .001$). In the GLMM including the interaction, with humans and the open-length setting as baseline category ($\beta = 3.270$, $z = 16.190$, $P < .001$ for the intercept), we find that LLMs perform significantly worse than humans in the open-length setting ($\beta = -3.508$, $z = -6.103$, $P < .001$), that humans do not perform better in the one-word setting as opposed to the open-length setting ($\beta = -0.228$, $z = -1.561$, $P = .118$), and that the gap between humans and LLMs is significantly smaller for responses in the one-word setting, indicating a discrepancy between the response settings for LLMs that is not observed for humans ($\beta = 1.276$, $z = 6.833$, $P < .001$; see Fig. 3).

Stability by agent and setting Stability rates by responding agent and setting are displayed in Fig. 3B. As was the case for accuracy, a first chi-square test finds that a GLMM with a main effect of agent (stability ~ setting + agent + (1 | item) + (1 | unit)) outperforms a GLMM without it ($X^2(1) = 10.515$, $P = .001$). A second chi-square test additionally finds that a GLMM with an interaction between agent and setting (stability ~ setting * agent + (1 | item) + (1 | unit)) significantly outperforms a GLMM which includes agent and setting as main effects ($X^2(1) = 4.903$, $P = .027$). In the GLMM including the interaction with humans and the open-length setting again as baseline category ($\beta = 2.286$, $z = 17.368$, $P < .001$ for the intercept), we find that LLMs are less stable than humans in the open-length setting ($\beta = -1.385$, $z = -3.800$, $P < .001$), that humans are not less stable in the open-length setting

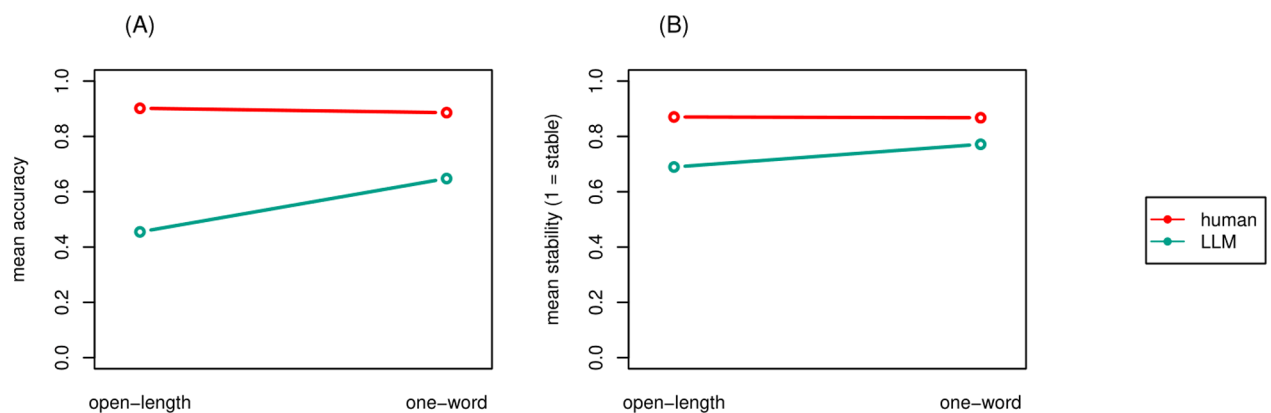


Fig. 3. (A) Mean accuracy by responding agent and setting. (B) Mean stability by responding agent and setting.

compared to the one-word setting ($\beta = -0.038$, $z = -0.362$, $P = .717$), and that in the one-word setting, the stability gap of LLMs with humans is smaller than in the open-length setting (again indicating a discrepancy between the settings for LLMs that is not observed for humans; $\beta = 0.507$, $z = 2.234$, $P = .026$; see Fig. 3).

Lastly, anticipating possible arguments that LLMs as a category should be represented by their best performing representative, ChatGPT-4 (the best performing LLM among those tested) can be compared to the best performing humans (as opposed to the averaged performance of all humans). Given that $n = 51$ human subjects perform at ceiling in terms of accuracy and therefore stability, while ChatGPT-4 does not perform at ceiling on either variable, it follows that —notwithstanding the great performance overlap between LLMs and humans— even the best performing LLM performs significantly worse than the best performing humans. Finally, while a sufficiently powerful statistical analysis is not allowed by the current data (given that each LLM was treated as a unit of observation, like a single human participant), all human participants taken together outperform ChatGPT-4 at a descriptive level, with a mean accuracy of 0.89 vs. 0.83 for ChatGPT-4, and a stability rate of 0.87 (0.78 for ChatGPT-4).

Discussion

This work sheds light on what guides output generation in LLMs: (i) the human-like ability to parse and understand written text, or (ii) the leveraging of artifacts in the training data. To do so, we test 7 state-of-the-art LLMs on comprehension questions targeting sentences that feature high-frequency constructions and lexical items, and whose linguistic complexity is kept at a minimum. Particularly, we aim at determining whether LLMs generate answers that are both accurate (RQ1) and stable over repeated trials (RQ2). Systematic testing shows that the average performance of LLMs as a group is at chance for accuracy, and that LLMs are relatively unstable in their answers. On the other hand, humans tested on the same comprehension questions provide mostly accurate answers (RQ1) which almost never change when a question is repeatedly prompted (RQ2). Importantly, these differences between LLMs and humans are pronounced even though the models' replies are favorably scored, ignoring any possible inconsistencies they contain (i.e., in the one-word setting; see section 'Methods').

Language parsing, referring to the ability to comprehend and produce language by assigning strings of symbols with meaning, is a distinctively human ability⁵². This explains why humans in our experiment responded accurately and without alternating answers when a question is asked multiple times, or with different instructions. For humans, failure to reach at-ceiling performance is explained by the interplay of processing factors⁴⁵ and interspeaker variation which affects interpretative choices. For instance, given the prompt *Steve hugged Molly and Donna was kissed by Molly. In this context, was Molly kissed?* the answer is “no”, though some might assign the verb *to kiss* a reciprocal reading consequently replying “yes”. Nonetheless, despite processing limitations and the variability induced by subjective choices, humans cohesively outperform LLMs both in terms of accuracy and stability, thus confirming that testing LLMs on foundational tasks can reveal shortcomings which remain concealed in more sophisticated settings⁸.

The LLM outputs differ from human answers both quantitatively and qualitatively. Quantitatively, the averaged accuracy of LLMs as a group is at chance, and the models which manage to surpass the chance threshold (i.e., Falcon, Llama2, and ChatGPT-4) reach accuracy rates that are far from perfect. The best performing LLM, ChatGPT-4, performs worse than the best performing humans. Secondly, while all LLMs perform above chance in terms of stability, none of them succeeds in consistently providing the same answer to a question. Taken together, this evidence translates into an overall inability on the LLMs' part to tackle simple comprehension questions in a way that can be reasonably called ‘human-like’.

Building on how their outputs show surface similarities with natural language, LLMs have been endowed with human-like language skills both in public discourse, where they are described as “understanding”, “reasoning”, or “intelligent”, and in science, where their adoption as cognitive models, and even *theories*, is increasingly called for^{23,44}. Our results suggest that such a framing may not be justified at present. Importantly, once a LLM has been invested with competence over a language phenomenon, without theoretically informed hypotheses at hand⁵³, the causes of its failures to perform at ceiling leave room for speculation. We thus interpret the LLMs' failure in providing accurate and stable answers in a simple comprehension task as suggesting that the models lack a real understanding of language: they produce words that are like semantic black-boxes, approximating the surface statistics of language and the more ‘automatic’ aspects of parsing⁵⁴. In fact, it is not simply the lower accuracy, but also the reduced stability in LLM responses that suggests an absence of a human-like algorithm that maps syntactic information directly onto semantic instructions, with considerably lower tolerance for differing judgments of the kind demonstrated in our study. Humans, possessing an invariant compositional operator for regulating grammatical and semantic information⁵⁵, are evidently less prone to error here. In addition, LLMs are not suitable for serving as theories of language since their representational capacities are virtually unbounded, rendering their representations arbitrary and not explanatorily grounded, belonging as they do to the class of Universal Functional Approximators⁵⁶, which have been shown to be capable of approximating any mathematical function⁵⁷. Consequently, LLMs bear a closer resemblance to tools such as generalized derivatives than scientific theories [see also⁵⁸].

Additionally, the differences between LLM and human performance are also qualitative in nature. First, all human participants with no exception provide one-word replies regardless of the setting. This shows alignment between subjects on the amount of information that is necessary in order to satisfactorily reply to a question⁵⁹, and even in this space of pragmatics humans seem to comply with principles of computational efficiency and relevance that are built out of the path that syntactic inferences carve out for them⁶⁰. This also marks a stark difference from LLMs which typically provide lengthy explanations, even in the one-word setting, thus defying the task instructions. In these explanations, their reasoning is broken down into logical steps which often contain

irrelevant or contradicting information. For example, to the prompt “Franck read to himself and John read to himself, Anthony and Franck. In this context, was Franck read to?”, Gemini replies (2):

(2) “In this context, *it’s impossible to say for sure if Franck was read to*. Here’s why:

1. Limited information: While the sentence “John read to himself, Anthony and Franck” implies John read aloud to others, it doesn’t clarify if Franck was one of the recipients or just present during the reading.
2. Separate actions: The sentence about Franck reading to himself is separate and doesn’t necessarily connect to John’s reading session.
3. Ambiguity: The phrase “read to” can have different interpretations depending on the context. It could mean reading aloud *for* someone or simply reading *in the presence of* someone.

Without additional information about the specific situation, such as John’s reading material, the purpose of his reading, or the relationship between the characters, it’s impossible to determine if Franck was actually read to by John.

Here are some possibilities based on the limited information:

- Franck was read to: John read aloud to both Anthony and Franck, and Franck listened.
- Franck wasn’t read to: John read aloud only to Anthony, and Franck was simply present during the reading.
- Franck listened to himself: Both sentences describe independent reading activities, with Franck reading silently to himself and John reading aloud to others.

Therefore, the answer remains *uncertain* due to the lack of sufficient context.”

The correct answer to the question is “yes”. Gemini not only gives the wrong answer, but it also answers redundantly and mixes unrelated matters about the nature of the characters and the type of reading material, none of which plays a role in answering the question. This limited ability to reason does not alter the fact that LLMs can fluently produce grammatical sentences on any topic²³. However, as Mahowald et al.²³ argue, to claim that LLMs have human-like linguistic abilities, formal competence is not enough: LLMs need to show competence in reasoning, world knowledge, pragmatic inference, etc. Humans perform language tasks such as the one we developed, using a combination of core grammatical abilities, semantico-pragmatic knowledge, and non-linguistic reasoning.

The problem of redundant, semantically irrelevant information is not a trivial one. Our testing showed that obtaining one-word replies from the tested LLMs was possible only in a reduced number of cases (mostly in ChatGPT-4 and Falcon). While this issue was ignored for the purposes of coding in the one-word setting, because we wanted to give the models the benefit of the doubt, it further showcases how LLMs’ outputs are sensitive to the meaning of prompts only in a superficial manner⁶¹. While pre-trained LLMs might not have the skills necessary for following specific instructions such as the one we provided (i.e., reply using only one word), it should be noted that ChatGPT-3.5, ChatGPT-4, Bard and Gemini include Reinforcement Learning from Human Feedback (RLHF) in their training, which means that their performance in this experiment benefitted from previous targeted interventions aimed at fixing incongruent outputs and maximizing compliance to user requests; yet, these models’ replies fail to conform to task demands.

If the LLMs’ fluent outputs are taken to entail human-like linguistic skills (i.e., from the production of a correct sentence such as “The scientist is happy”, we infer that the model has knowledge of the relevant rules that put the sentence together), then it follows that LLMs should possess mastery over all the words, constructions, and rules of use, that figure in and govern their outputs. Our results from grammaticality judgments (Fig. 1) and language comprehension tasks (Figs. 2 and 3) suggest that this is not the case: distinctly non-human errors are made. To discuss an example, when asked whether “The patient the nurse the clinic had hired admitted met Jack” is grammatically correct, GPT-3 said no, providing the wrong answer (Fig. 1). While humans likewise err in hard-to-process sentences of this kind owing to cognitive constraints and unmet prediction expectations⁶², what strikes as surprising is not GPT-3’s incorrect judgment per se, but rather the correction that the model volunteered: “The nurse had hired at the clinic a patient who admitted met Jack”. This answer contains a semantic anomaly that marks an important deviation from any answer a human would produce: a nurse hiring a patient. Even if we change the methodology from judgment prompting to simple question answering, we get similar errors that suggest that the tested models do not really understand the meaning of words. Recall that Gemini could not answer as to whether Franck has been read to (after being told that “John read to himself, Anthony and Franck”; example 2), because it lacks information “such as John’s reading material”. Overall, if we attempt to map the encountered LLM errors across levels of linguistic analysis, it seems that their incidence becomes greater as we move from basic phonological form to the more complex outer layers of linguistic organization (Fig. 4).

While LLMs are often portrayed through wishful mnemonics that hint to human-like abilities⁶⁴, seemingly sound performance may hide flaws that are inherent to language modelling as a method: intelligence cannot in fact emerge as a side product of statistical inference⁶⁵, nor can the ability to understand meaning²⁸. Our results offer an argument against employing LLMs as cognitive theories: the fact that LLMs are trained on natural language data, and generate natural-sounding language, does not entail that LLMs are capable of human-like processing; it only means that LLMs can predict some of the fossilized patterns found in training texts. Claiming a model’s mastery over language because of its ability to reproduce it, is analogous to claiming that a painter *knows* a person because they can reproduce this person’s face on canvas, just by looking at her pictures. Unjustly equating LLMs’ abilities to those of humans in order to justify their mass-scale deployment in applied areas, in turn, results in the biases present in LLMs’ training data⁶⁶ being allowed one more display venue.

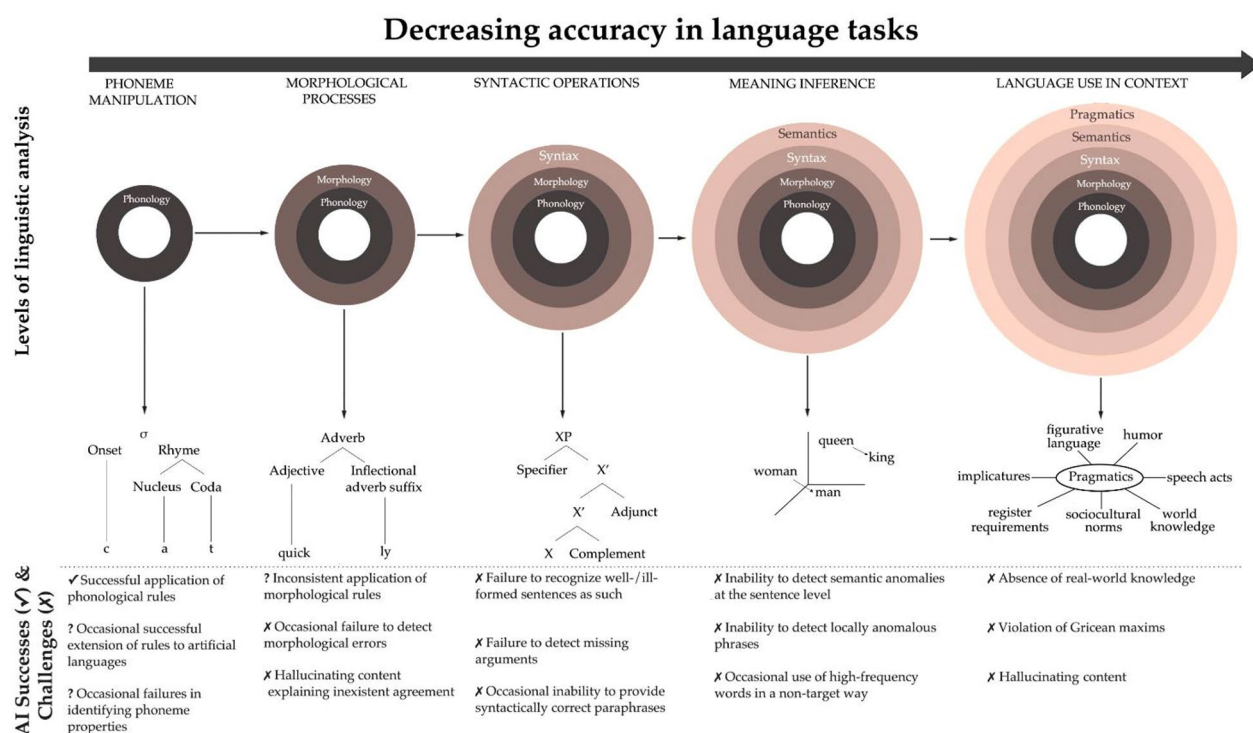


Fig. 4. Mapping of the types of errors presented in Fig. 1 onto levels of linguistic analysis. The observations for phonology are based on the results of⁶³.

With respect to practical applications, in AI reproducing high-level reasoning is overall less demanding and needs less computational power than comparatively easier tasks⁷. At present, LLMs are employed in settings which require high-level reasoning. For instance, LLMs can reach high accuracy in answering medical questions³ and are believed to be ready for testing in real-world simulations⁶⁷. While LLMs in a white coat strike many as an implausible scenario, in reality LLM-based chatbots already reply to any question that any human user in real-life situations may pose, possibly including doctors. Except for warnings typically issued when ethically problematic words or politically incorrect concepts figure in prompts⁶⁸, often chatbots reply without flagging any hesitations or lack of knowledge. As a result, the message conveyed is that whatever content LLMs output, it is indeed the product of true understanding and reasoning. However, as the results obtained in this work demonstrate alongside a consistent body of evidence^{30–32,34,69}, coherent, sophisticated, and refined outputs amount to patchwork in disguise.

To conclude, LLMs' fluent outputs have given rise to claims that models possess human-like language abilities, which would in turn allow them to act as cognitive theories able to provide mechanistic explanations of language processing in humans^{23,44}. Following evidence that LLMs often produce non-human outputs⁷⁰, in this work we tested 7 state-of-the-art LLMs on simple comprehension questions targeting short sentences, purposefully setting an extremely low bar for the evaluation of the models. Systematic testing showed that the performance of these LLMs lags behind that of humans both quantitatively and qualitatively, providing further confirmation that tasks that are easy for humans are not always easily developed in AI⁸. We argue that these results invite further reflection about the standards of evaluation we adopt for claiming human-likeness in AI. Both in society and science, the role of LLMs is currently defined according to inflated human-like linguistic abilities⁷¹, as opposed to their real capacities. If the application of internalized linguistic rules is absent and LLMs do not go beyond reproducing input patterns⁷², all that is presented as a danger induced by the stochastic nature of LLMs can be anticipated⁷³. If LLMs are useful tools, understanding their limitations in a realistic way paves the way for them to be better and safer.

Methods

Materials

The language understanding abilities of 7 LLMs are tested on a set of comprehension questions ($n=40$). Specifically, the employed prompts have the following characteristics: (a) sentences are affirmative (i.e., contain no negation, which is known to negatively impact LLM performance²⁹), (b) they are linked through coordination, (c) either 3 or 5 human entities are featured in each prompt, all identified by proper names (i.e., no pronouns), and (d) high-frequency verbs are used (e.g., to cook, to hug, to kiss), so as to maximize the chances of LLMs replying correctly. These characteristics were chosen with the explicit aim to minimize grammatical complexity (e.g., avoiding entangled pronoun resolution, the complex semantics of negation, etc.). Following previous research that has suggested that LLMs fail to be stable in their answers to a question prompted repeatedly³¹, we

prompted each comprehension question multiple times to assess the stability of answers. To obtain the right level of comparisons for determining human-likeness, the performance of the LLMs is compared to that of human subjects ($n=400$) who were tested on the same stimuli. Table 1 provides an overview of the prompts, with one sample sentence and question per condition. The full list of prompts alongside the obtained replies by both humans and LLMs is available at <https://osf.io/dfgmr/>.

LLMs

All prompts were presented to the LLMs in two settings: (i) open-length setting and (ii) one-word setting. In each setting, every prompt was repeated three times, occasionally refreshing LLMs’ testing context and without allowing the repetition of the same prompt within the same context. In practical terms, this means that the first query in the chat interface represents zero-shot testing, whereas subsequent queries are subject to the influence of previous ones (few-shot). While this difference may not be ideal upon assessment of the models in isolation, it is a necessary condition to ensure straightforward comparability with human results. Thus, this *modus operandi* effectively mirrors how human testing was carried out (as explained in detail below). In (i), LLMs were allowed to produce answers of any length. In (ii), the LLMs were prompted using the same materials as in (i), but the instruction “Answer using just one word” was added after every question (e.g., “John deceived Mary and Lucy was deceived by Mary. In this context, did Mary deceive Lucy? Answer using just one word”). The use of the one-word setting was intended to facilitate the LLMs, as it was noted that lengthy LLM replies often contained contradictory, irrelevant, and eventually incorrect statements. In total, each of the tested LLMs was prompted 240 times (120 times per setting). The LLMs’ pool of data thus consists of $n=1,680$ replies (240 replies per each of the 7 LLMs).

In both settings, all replies were coded for *accuracy* (1 for accurate answers, 0 for inaccurate answers) and *stability* (1 if the three responses were either coded all as *yes* or all as *no*, 0 if at least one response was coded differently than the others). In an attempt to give the LLMs the benefit of the doubt and obtain from them the best possible result, their responses in the open-length setting which contained no clear error, but also failed to directly provide the correct answer, were coded as accurate. Additionally, due to the difficulty of obtaining one-word replies from LLMs, the decision was made to only code the first relevant piece of information provided in a reply in the one-word setting. Put another way, if a LLM correctly replied ‘Yes’ but this was followed by additional content that was incorrect, contradictory, or irrelevant, this was coded as accurate. From this it follows, that a less lenient coding would further amplify the differences between humans and LLMs. All prompts were merged in a unified pool, randomized and administered in random order, first in the open-length setting, and subsequently in the one-word setting.

The prompts were given to 7 LLMs, which were chosen as representative of the state of the art: ChatGPT-3.5⁷⁴, ChatGPT-4⁷⁵, Llama2 (Llama-2-70B-chat-hf⁷⁶), Bard⁷⁷, Falcon (falcon-180B-chat⁷⁸), Mixtral (Mixtral-8×7B-Instruct-v0.1⁷⁹), and Gemini²¹. ChatGPT-3.5, ChatGPT-4, Bard and Gemini include RLHF in their training regime. All LLMs were tested in December 2023, with the exception of Gemini which was tested in February 2024. ChatGPT-3.5 and ChatGPT-4 were prompted through OpenAI’s interface; Bard and Gemini through Google’s interface; and Llama2, Falcon and Mixtral through HuggingFace’s interface.

Human data

For comparison, we tested human participants ($n=400$; 200 F, 200 M) on the same prompts. Subjects were recruited through the crowdsourcing platform Prolific, and all participants self-identified as native speakers of English with no history of neurocognitive impairments. The experiment was carried out in accordance with the Declaration of Helsinki and was approved by the ethics committee of the Department of Psychology at Humboldt-Universität zu Berlin (application 2020-47). Participants were divided into two groups: 200 subjects (100 F, 100 M) were tested in the open-length setting, and 200 subjects (100 F, 100 M) in the one-word setting. The prompts ($n=40$) originally administered to the LLMs were split into two lists (each presented to half of the subjects in a condition), so that each human subject was prompted with half of the questions ($n=20$). Similarly to the LLMs, each participant saw each question 3 times, providing a total of $n=60$ replies. Overall, the human pool of data consists of $n=24,000$ replies. Human replies were coded for *accuracy* and *stability*. To ensure compliance, two attention checks were added that are not part of the analyses (i.e., “The door is red. In this context, is the door red?” and “The door is green. In this context, is the door blue?”). All prompts, including attention checks, were administered to each participant in random order. Data from 2 additional participants who did not provide correct answers to all attention checks were excluded from the analyses. Written informed consent was obtained from all subjects and/or their legal guardian(s). For each trial, the screen displayed the

Sample prompt	Number of featured entities	Target answer	Number of prompts per condition
John deceived Mary and Lucy was deceived by Mary. In this context, did Mary deceive Lucy?	3	Yes	10
Steve hugged Molly and Donna was kissed by Molly. In this context, was Molly kissed?	3	No	10
Jessica and Mary were kissed by Alice. Jessica was kissed by Samuel and Andrew was kissed by Mary. In this context, was Mary kissed?	5	Yes	10
Bob kissed Donna and Barbara kissed Peter. Donna was hugged by Alice. In this context, was Alice hugged?	5	No	10

Table 1. Overview of testing materials.

prompt (i.e., the sentence, its related question and, in the one-word setting, the additional instruction). The experiment was carried out with the jsPsych toolkit⁸⁰, and had a median completion time of 13.4 min. Full participant data can be found at <https://osf.io/dfgmr/>.

Future work

One way the current study could be built upon is through testing humans in a speeded setting (vs. the untimed setting employed in the present study). It is possible that human accuracy may decrease when humans are put under time pressure: unlimited time indeed allows recruitment of reasoning resources that are not exclusively linguistic, and which possibly enhance accuracy. Should this prediction be borne out, it would help explain what is missing in LLMs and why they underperform in some tasks. Indeed, recent LLMs (e.g., OpenAI o1 models) seem to incorporate novel forms of multi-step reasoning, making it possible that these models will perform better at our benchmark.

Data availability

The data and code associated with the present manuscript are available at <https://osf.io/dfgmr/>.

Received: 4 June 2024; Accepted: 11 November 2024

Published online: 14 November 2024

References

1. Bubeck, S. et al. Sparks of artificial general intelligence: early experiments with GPT-4. *arXiv*. 2303.12712 (2024).
2. Nay, J. J. et al. Large language models as tax attorneys: a case study in legal capabilities emergence. *Philos. Trans. R. Soc. A*. **382**, 20230159 (2024).
3. Singhal, K. et al. Large language models encode clinical knowledge. *Nature* **620**, 172–180 (2023).
4. Sandmann, S., Riepenhausen, S., Plagwitz, L. & Varghese, J. Systematic analysis of ChatGPT, Google search and Llama 2 for clinical decision support tasks. *Nat. Commun.* **15**, 2050. (2024).
5. Jablonka, K. M., Schwaller, P., Ortega-Guerrero, A. & Smit, B. Leveraging large language models for predictive chemistry. *Nat. Mach. Intell.* **6**, 161–169 (2024).
6. Moravec, H. *Mind Children* (Harvard University Press, 1988).
7. Moravec, H. Rise of the robots. *Sci. Am.* **281** (6), 124–135 (1999).
8. Pinker, S. *The Language Instinct* (William Morrow and Company, 1994).
9. Minsky, M. *The Society of Mind* (Simon and Schuster, 1986).
10. Perkins, L. & Lidz, J. Eighteen-month-old infants represent nonlocal syntactic dependencies. *PNAS* **118** (41), e2026469118 (2021).
11. Dehaene, S., Al Roumi, F., Lakretz, Y., Planton, S. & Sablé-Meyer, M. Symbols and mental programs: a hypothesis about human singularity. *Trends Cogn. Sci.* **26** (9), 751–766 (2022).
12. Bowman, S. R., Angeli, G., Potts, C. & Manning, C. D. A large annotated corpus for learning natural language inference. *Proceedings of the 2015 Conference on Empirical Methods in NLP*, 632–642. Association for Computational Linguistics (2015).
13. Conneau, A., Lample, G., Rinott, R., Schwenk, H. & Stoyanov, V. XNLI: Evaluating cross-lingual sentence representations. *Proceedings of the 2018 Conference on Empirical Methods in NLP*, 2475–2485. Association for Computational Linguistics (2018).
14. Williams, A., Nangia, N. & Bowman, S. A broad-coverage challenge corpus for sentence understanding through inference. *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Volume 1, 1112–1122. Association for Computational Linguistics (2018).
15. Sap, M., Rashkin, H., Chen, D., Le Bras, R. & Choi, Y. SocialIQA: Commonsense reasoning about social interactions. *Proceedings of the 2019 Conference on Empirical Methods in NLP and the 9th International Joint Conference on NLP*, 4463–4473. Association for Computational Linguistics (2019).
16. Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A. & Choi, Y. HellaSwag: Can a machine really finish your sentence? *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 4791–480. Association for Computational Linguistics (2019).
17. Bisk, Y., Zellers, R., Le Bras, R., Gao, J. & Choi, Y. PIQA: Reasoning about physical commonsense in natural language. *Proceedings of the 34th Conference on Artificial Intelligence*. Association for the advancement of Artificial Intelligence (2020). <https://doi.org/10.1609/aaai.v34i05.6239>
18. Hendrycks, D. et al. Measuring massive multitask language understanding. *arXiv* 2009.03300 (2021).
19. Sakaguchi, K., Le Bras, R., Bhagavatula, C. & Choi, Y. WinoGrande: an adversarial Winograd schema challenge at scale. *Commun. ACM* **64**(9) (2021).
20. OpenAI (April 30, 2024). <https://openai.com/gpt-4>
21. Pichai, S. & Hassabis, D. Introducing Gemini: Our largest and most capable AI model. Google blog (2023). <https://blog.google/technology/ai/google-gemini-ai/#sundar-note>
22. Piantadosi, S. T. & Hill, F. Meaning without reference in large language models. *Proceedings of the 36th Conference on Neural Information Processing Systems* (2022).
23. Mahowald, K. et al. Dissociating language and thought in large language models: a cognitive perspective. *Trends Cogn. Sci.* **28**(6) (2024).
24. Hu, J., Mahowald, K., Lupyran, G., Ivanova, A. & Levy, R. Language models align with human judgments on key grammatical constructions. *PNAS* **121** (36), e2400917121(2024).
25. Radford, A. et al. Language models are unsupervised multitask learners (2019). https://d4mucfpsywv.cloudfront.net/better-language-models/language_models_are_unsupervised_multitask_learners.pdf
26. Brown, T. B. et al. Language models are few-shot learners. *arXiv* 2005.14165 (2020).
27. Bojic, L., Kovacevic, P. & Cabarkapa, M. GPT-4 surpassing human performance in linguistic pragmatics. *arXiv* 2312.09545 (2024).
28. Bender, E. M. & Koller, A. Climbing towards NLU: On meaning, form, and understanding in the age of data. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics (2020).
29. Ettinger, A. What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models. *Transactions of the Association for Computational Linguistics* **8**, 34–48. MIT Press (2020).
30. Rassin, R., Ravfogel, S. & Goldberg, Y. DALL-E 2 is seeing double: Flaws in word-to-concept mapping in Text2Image models. *Proceedings of the Fifth BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 335–345. Association for Computational Linguistics (2022).
31. Dentella, V., Günther, F. & Leivada, E. Systematic testing of three Language models reveals low language accuracy, absence of response stability, and a yes-response bias. *PNAS* **120**(51), e2309583120 (2023).
32. Leivada, E., Murphy, E. & Marcus, G. DALL-E 2 fails to reliably capture common syntactic processes. *SSHO* **8**, 100648 (2023).
33. Moro, A., Greco, M. & Cappa, S. Large languages, impossible languages and human brains. *Cortex* **167**, 82–85 (2023).

34. Murphy, E., de Villiers, J. & Lucero Morales, S. A comparative investigation of compositional syntax and semantics in DALL-E 2. *arXiv* 2403.12294 (2024).
35. Collins, J. The simple reason LLMs are not scientific models (and what the alternative is for linguistics). *lingbuzz*/008026 (2024).
36. Marcus, G. & Davis, E. *Rebooting AI: Building Artificial Intelligence We Can Trust* (Pantheon Books, 2019).
37. Sinha, K. et al. Language model acceptability judgements are not always robust to context. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics* 1, 6043–6063. Association for Computational Linguistics (2023).
38. Kandpal, N., Deng, H., Roberts, A., Wallace, E. & Raffel, C. Large Language Models struggle to learn long-tail knowledge. *Proceedings of the 40th International Conference on Machine Learning* (2023).
39. Guest, O. & Martin, A. On logical inference over brains, behaviour, and artificial neural networks. *Comput. Brain Behav.* **6**, 213–227 (2023).
40. Dua, D. et al. DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*, 2368–2378. Association for Computational Linguistics (2019).
41. Wang, A. et al. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox & R. Garnett (Eds.), *Advances in Neural Information Processing Systems* 32, 3266–3280. Curran Associates Inc. (2019).
42. Srivastava, A. et al. Beyond the imitation game: quantifying and extrapolating the capabilities of language models. *Trans. Mach. Learn. Res.*, 1–95. (2023).
43. Yogatama, D. et al. Learning and evaluating general linguistic intelligence. *arXiv* 1901.11373 (2019).
44. Marcus, G. & Davis, E. Hello, multimodal hallucinations. Substack (2023). <https://garymarcus.substack.com/p/hello-multimodal-hallucinations>
45. Karimi, H. & Ferreira, F. Good-enough linguistic representations and online cognitive equilibrium in language processing. *Q. J. Experimental Psychol.* **69**, 1013–1040 (2016).
46. Hu, J. & Levy, R. Prompting is not a substitute for probability measurements in large language models. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 5040–5060. Association for Computational Linguistics (2023).
47. Extance, A. ChatGPT has entered the classroom: how LLMs could transform education. *Nature* **623**, 474–477 (2023).
48. McCoy, T., Yao, S., Friedman, D., Hardy, M. & Griffiths, T. L. Embers of autoregression: Understanding Large Language Models through the problems they are trained to solve. *arXiv* 2309.13638v1 (2023).
49. Moffatt v. Air Canada. BCCRT 149 (2024).
50. Bates, D., Maechler, M., Bolker, B. & Walker, S. Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* **67** (1), 1–48 (2015).
51. Kuznetsova, A., Brockhoff, P. B. & Christensen, R. H. B. lmerTest package: tests in linear mixed effects models. *J. Stat. Softw.* **82** (13), 1–26 (2017).
52. Hauser, M. D., Chomsky, N. & Fitch, W. T. Faculty of language: what is it, who has it, and how did it evolve? *Science* **298** (5598), 1569–1579 (2002).
53. Guest, O. & Martin, A. E. How computational modeling can force theory building in psychological science. *Perspect. Psychol. Sci.* **16** (4), 789–802 (2020).
54. Mitchell, M. & Krakauer, D. C. The debate over understanding in AI's large language models. *PNAS* **120** (13), e2215907120 (2022).
55. Pietroski, P. *Conjoining Meanings: Semantics without Truth Values* (Oxford University Press, 2018).
56. Yun, C., Bhojanapalli, S., Rawat, A. S., Reddi, S. J. & Kumar, S. Are transformers universal approximators of sequence-to-sequence functions? *arXiv* 1912.10077 (2019).
57. Ismailov, V. E. A three layer neural network can represent any multivariate function. *J. Math. Anal. Appl.* **523** (1), 127096 (2023).
58. Evanson, L., Lakretz, Y. & King, J. R. Language acquisition: Do children and language models follow similar learning stages? *arXiv* 2306.03586 (2023).
59. Grice, P. Logic and conversation. In (ed Cole, P.), *Syntax and Semantics 3: Speech Acts* (Academic Press, 1975).
60. Carston, R. Words: syntactic structures and pragmatic meanings. *Synthese* **200**, 430 (2022).
61. Du, M., He, F., Zou, N., Tao, D. & Hu, X. Shortcut learning of large language models in natural language understanding. *Commun. ACM* **67**(1) (2024).
62. Hahn, M., Futrell, R., Levy, R. & Gibson, E. A resource-rational model of human processing of recursive linguistic structure. *PNAS* **119**(43), e2122602119 (2022).
63. Beguš, G., Dąbkowski, M. & Rhodes, R. Large linguistic models: analyzing theoretical linguistic abilities of LLMs. *arXiv* 2305.00948 (2023).
64. Mitchell, M. Why AI is harder than we think. *Proceedings of the Genetic and Evolutionary Computation Conference* (2021).
65. van Rooij, I. et al. Reclaiming AI as a theoretical tool for cognitive science. *Comput. Brain Behav.* (2024). <https://link.springer.com/article/10.1007/s42113-024-00217-5#citeas>
66. Birhane, A., Prabhu, V., Han, S. & Boddeti, V. N. On hate scaling laws for data-swamps. *arXiv* 2306.13141 (2023).
67. Mehandru, N. et al. Evaluating large language models as agents in the clinic. *Npj Digit. Med.* **7**, 84 (2024).
68. Ghafouri, V. et al. AI in the gray: Exploring moderation policies in dialogic large language models vs. human answers in controversial topics. *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 556–565 (2023).
69. Kim, Y. et al. Fables: evaluating faithfulness and content selection in book-length summarization. *arXiv* 2404.01261v1 (2024).
70. Leivada, E., Marcus, G., Günther F. & Murphy, E. A sentence is worth a thousand pictures: Can Large Language Models understand human language and the world behind words? *arXiv* 2308.00109 (2024).
71. Leivada, E., Günther, F. & Dentella, V. Reply to Hu et al: Applying different evaluation standards to humans vs. Large Language Models overestimates AI performance. *PNAS* **121**(36), e2406752121 (2024).
72. Udandara, V. et al. No zero-shot without exponential data: pretraining concept frequency determines multimodal model performance. *arXiv* 2404.04125v1 (2024).
73. Bender, E. M., Gebru, T., McMillan-Major, A. & Shmitchell, S. On the dangers of stochastic parrots: Can Language Models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623 (2021).
74. Ouyang, L. et al. Training language models to follow instructions with human feedback. *arXiv* 2203.02155 (2022).
75. OpenAI. GPT-4 technical report. *arXiv* 2303.08774 (2023).
76. Touvron, H. et al. LLaMA: Open and efficient foundation language models. *arXiv* 2302.13971 (2023).
77. Thoppilan, R. et al. LaMDA: Language models for dialog applications. *arXiv* 2201.08239 (2022).
78. Almazrouei, E. et al. The Falcon series of open language models. *arXiv* 2311.16867v2 (2023).
79. Jiang, A. Q. et al. Mixtral of experts. *arXiv* 2401.04088v1 (2024).
80. de Leeuw, J. R., Gilbert, R. A. & Luchterhandt, B. jsPsych: enabling an open-source collaborative ecosystem of behavioral experiments. *J. Open. Source Softw.* **8**, 5351 (2023).

Acknowledgements

We thank two anonymous reviewers for their feedback.

V.D. acknowledges funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 945413 and from the Universitat Rovira i Virgili. F.

G. acknowledges funding from the German Research Foundation (Deutsche Forschungsgemeinschaft) under the Emmy-Noether grant “What’s in a name?” (project No. 459717703). E.L. acknowledges funding from the Spanish Ministry of Science, Innovation & Universities (MCIN/AEI/<https://doi.org/10.13039/501100011033>) under the research projects No. PID2021-124399NA-I00 and No. CNS2023-144415.

Author contributions

V.D. and E.L. conceived and designed the study. V.D. and F.G. collected the data. F.G. performed the analyses. V.D. and E.L. wrote the original draft of the manuscript. All authors helped refine the manuscript, and approved its final version.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to V.D.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024