



Research



Cite this article: Leivada E, Marcus G, Günther F, Murphy E. 2026 A sentence is worth a thousand pictures: can large language models understand human language and the world behind words? *Phil. Trans. R. Soc. A* **384**: 20250008.
<https://doi.org/10.1098/rsta.2025.0008>

Received: 9 January 2025

Accepted: 25 June 2025

One contribution of 18 to a theme issue ‘World models in natural and artificial intelligence.’

Subject Areas:

artificial intelligence

Keywords:

cognition, language, top-down processing, concept

Author for correspondence:

Evelina Leivada

e-mail: evelina.leivada@uab.cat

A sentence is worth a thousand pictures: can large language models understand human language and the world behind words?

Evelina Leivada^{1,2}, Gary Marcus³, Fritz Günther⁴
and Elliot Murphy⁵

¹Universitat Autònoma de Barcelona, Barcelona, Spain

²Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain

³New York University, New York, NY, USA

⁴Humboldt-Universität zu Berlin, Berlin, Germany

⁵The University of Texas Health Science Center at Houston John P and Katherine G McGovern Medical School, Houston, TX, USA

EL, 0000-0003-3181-1917

The current generation of large language models (LLMs) has been linked to claims about human-like linguistic performance, and their applications are hailed both as a step towards artificial general intelligence and as a major advance in understanding the cognitive and even neural basis of human language. To assess these claims, first, we analysed the contribution of LLMs as theoretically informative representations of a target cognitive system versus atheoretical mechanistic tools. Second, we evaluated the models' ability to see the bigger picture through top-down feedback from higher levels of processing, which requires grounding in previous expectations and past world experience. We hypothesize that since models lack grounded cognition, they cannot take advantage of these features and instead solely rely on fixed associations between represented words and word vectors. To assess this, we ran a novel leet task (l33t t4sk), which requires decoding sentences in which letters are systematically replaced by numbers. In line with our hypothesis, the results suggest that humans excel in this task, whereas models struggle. We interpret these results by identifying the key abilities that are still missing from the current

© 2026 The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License <http://creativecommons.org/licenses/by/4.0/>, which permits unrestricted use, provided the original author and source are credited.

state of development of these models, which require solutions that go beyond increased system scaling.

This article is part of the theme issue ‘World models in natural and artificial intelligence’.

1. Introduction

Recent advances in artificial intelligence (AI) have given rise to many applications that have captured the public imagination (e.g. ChatGPT, DALL·E, DeepSeek, Imagen). These applications can create realistic images and/or synthetic language, based on textual prompts. This dual ability to process natural language and produce synthetic language that looks remarkably like that of humans has given rise to claims that AI has passed the *language threshold*. Some AI applications have been linked to claims of passing the Turing test [1], of showing human-like understanding and sentience [2] and of being able to extrapolate the morphophonological rules of natural language in a way that can be compared with language learning by children [3].

Yet, good tools are not synonymous with good models that provide faithful representations of the target system. In general terms, models aim to capture some aspects of a target system with the aim of helping construct a theory of it. For this to happen, the models need to be reliable and transparent representations of the target system. Given their stochastic nature [4], at present, it is not yet clear what large language models (LLMs) can tell us about human language [5]. Indeed, some researchers prefer the term ‘corpus models’, removing any association with language [6]. Observing effects and statistical distributions in the model outputs is not informative about language *ipso facto*, an effect is something that waits for an explanation, not explanation itself [7]. What is lacking from the picture is a clear definition of the research questions (RQs) that link the models to the system they seek to represent: what properties of language are we aiming to explain through LLMs, and what theory are we putting to the test?

Contemporary language models are typically differentiable functions with transformer architectures trained on text, in large quantities, for next-word prediction and are optionally also trained with reinforcement learning from human feedback. In this work, we demonstrate that modern language models do not possess full linguistic competence (i.e. the computational capacity or *understanding* required to parse human-like syntactic structures), and we offer explanations as to why they inescapably deviate in substantial ways from the real-world system they seek to represent. Specifically, we identify and analyse two empirical challenges, related to linguistic form, that render artificial LLMs and foundation models non-credible representations of human language. One challenge concerns the methods and standards we currently have for evaluating the performance of LLMs, while the second challenge boils down to how these models manipulate data. While some challenges brought to the fore by the inaccurate use and/or evaluation of LLMs can be linked to serious societal risks (e.g. plausible misinformation, magnification of harmful biases that exist in the training data [8]), all of them together support a view of LLMs as largely impenetrable black-box systems that produce linguistic behaviour that is still markedly different from that of humans. Given that most of these systems lack transparency in terms of *what* is being represented and *how* [9], it is not always feasible to either associate their target linguistic behaviour with successful language learning or attribute their deviating characteristics to one of their components (e.g. non-inclusive training data, engineered biases, the process of transforming word vectors into next-word predictions). Overall, tracking the ways in which such systems deviate from the real-world system they seek to model can guide and possibly constrain our expectations of them as well as the RQs the field thinks they can address.

The claim that a model accurately captures a target phenomenon—such as the claim that LLMs have human-like language abilities—implies an exact match between model predictions

and real-world data. In this context, evaluations of LLMs against their real-world counterparts have given rise to a paradox: LLMs are supposed to model human language, but state-of-the-art AI performance on benchmarks allegedly shows abilities that have surpassed those of humans, while at the same time failing to provide accurate responses in very basic language understanding tasks. *Both* these states (of over- and under-performing) undermine the claim of human-likeness. [Figure 1](#) illustrates the problem, citing the notion of ‘above 100%’ language performance.

Contrary to what it seems, [figure 1A](#) does not show the acceleration of AI in figuring out natural language. It rather reveals a deep problem that concerns the absence of robust standards of evaluation to determine the linguistic abilities of LLMs. Leaving aside that it is unclear what accuracy measurements are possible above the 100% threshold in any language task, there is no index for ‘language understanding’ as a general cognitive ability; the term means different things in different contexts. More importantly, this estimate that endows LLMs with suprahuman language abilities relies on a dubious comparison between the model and the real-world system, because performing well at any single task that involves some form of language understanding (e.g. answering logic questions in the law school admission tests (LSATs)) does not entail the possession of human-like language understanding, similar to how the conditional in example (1a) does not entail that a clock that shows the correct time works fine.

(1a) If a clock shows the correct time, then it works fine, ($P \rightarrow Q$).
The clock on the wall shows the correct time, (P).
Therefore, the clock on the wall works fine, ($\vdash Q$).

(1b) If an LLM shows accuracy levels equal to those of humans in a language task, then it has human-like language understanding, ($P \rightarrow Q$).
This LLM shows accuracy levels equal to those of humans in a language task, (P).
Therefore, this LLM has human-like language understanding, ($\vdash Q$).

(1c) If a clock works fine, then it shows the correct time, ($P \rightarrow Q$).
This clock shows the correct time, (Q).
Therefore, this clock works fine, ($\vdash P$).

As Guest & Martin [11] argue, such analogies, although seemingly sensible, are in fact unfounded, because they reverse the order of the argument (reversing the conditional in example (1a) converts it into an example of correct reasoning: if a clock works fine, then it shows the correct time). After reversing the order (1c), the second problem lies in the connection of the consequent with the antecedent. In (1c), nothing precludes the possibility that there are other conditions that legitimize the consequents. For instance, the clock may be broken *and* show the correct time at some point. Similarly, LLMs may not understand language *and* show at-ceiling accuracy in a language understanding task, owing to chance or engineered biases (e.g. if an LLM suffers from a yes-response bias, testing it in a task that asks whether sentences are grammatically correct will give us an inflated estimate of accuracy, if we employ only grammatical sentences [12]). Succinctly put, the fact that P may have occurred does not legitimize one cause over others. Consequently, if upon observing the ceiling LLM accuracy in a language task, instead of claiming that a model performed well, we infer that it has human-like linguistic behaviour (or even surpasses the human limits), we have slipped into faulty reasoning. This is one of the standard caveats one needs to keep in mind when evaluating the cognitive abilities of LLMs [13]. As Parnas [14] put it, artificial flowers may appear much like real flowers, but superficial similarity does not entail system equivalence. Just as a picture of artificial flowers that look real does not entail that artificial flowers *are* (like) real flowers, a snapshot of LLM results ([figure 1A](#)) that look as accurate as humans does not entail that LLM linguistic behaviour *is* human(-like).

The humans vs. LLMs language paradox

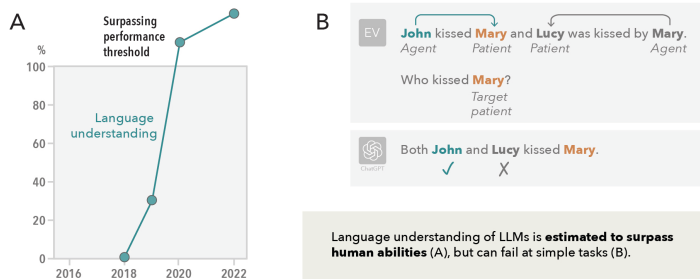


Figure 1. The humans versus LLM language paradox. (A) A dubious estimate of language understanding abilities in humans versus LLMs, reporting suprahuman abilities for LLMs (data from [10]). (B) The GPT-3.5 performance in a simple prompt that tests understanding of passive structures; 0% on the y-axis denotes baseline performance. Language understanding in (A) denotes the general language understanding evaluation (GLUE) benchmark.

So why *do* they look like human results? While there are language tasks that models may perform better than humans (e.g. naming tasks with strict time constraints), we argue that this apparent paradox is also nurtured by the weak standards of evaluation that are often employed to determine the language abilities of LLMs ([13,15]; see also [16]). For example, if we run an extremely easy task that asks to passivize sentences such as ‘Mary loves John’, it is likely that LLMs will return accurate results (although LLMs have occasionally failed to perform like humans in equally easy tasks [12,17]. Is this enough ground for concluding that they have acquired the syntactic rules that govern passive formation? Figure 1B suggests that the answer is negative. The error given in figure 1B is not one a neurotypical, adult speaker of English would *consistently* commit. While humans are not infallible in language tasks, their performance in answering this question is expected to be associated with a *very* low rate of error, which typically boils down to physical factors such as pressing the wrong button to register a reply when taking the task or a temporary distraction. The crux of the matter is that when defending the claim that LLMs have human-like language understanding, finding a failure shows that the model cannot reason *reliably*, and this unreliability cannot be due to the same physical factors that make humans occasionally err. It is for this reason that comparative research of humans and models has found big differences in the performance of the two agents, with models not only being outperformed by humans but also committing distinctly non-human errors [12,17].

LLMs are trained on thousands of scientific papers, and this can explain their good performance in certain standardized tests (e.g. LSATs: they are curves fitted to a data distribution). At the same time, memorizing task answers does not entail grasping either the meaning of words or the real-world conditions that guide their use in different contexts. Humans can do that, using top-down processing; a cognitive ability to grasp the ‘whole picture’ before focusing on the individual parts of which it is constructed. Top-down feedback from higher levels of processing requires grounding in previous expectations, past world experience, sensory information and emotions [18]. We hypothesize that since LLMs lack grounded cognition, they cannot take advantage of these features and, instead, they solely rely on fixed associations between represented words and word vectors. Consequently, they will not perform as accurately as humans in a task that recruits this ability.

To evaluate this hypothesis, a promising domain of testing that has not been previously explored in the context of LLMs concerns ‘leetspeak’, also known as 1337 or 1337. This writing practice refers to a type of written communication with modified spelling, which is often used on the Internet. Letters are replaced by (similar-looking) numbers or other symbols, aiming to make the text harder to decipher so as to gatekeep a group’s identity and prevent it from being discovered via automatic surveillance methods (e.g. keyword search). Since its purpose was to conceal identities, leet is ‘an ever-changing language deliberately intended to obfuscate and evade grokking and detection’ [19]. Interestingly, leet is specifically targeted at preventing

machinized message detection: humans can process leet words (e.g. HUM4NS C4N RE4D THIS) almost as fast and effortlessly as their base words (i.e. HUMANS CAN READ THIS). The predominant interpretation of this phenomenon is that the human cognitive parser tolerates some degree of ‘noise’ in word recognition, relying on top-down feedback from higher levels of processing [20,21]. This makes leet an ideal test case for the language skills of modern machine algorithms.

To assess whether LLMs have access to top-down feedback that could assist them in accurately decoding words based on their previous knowledge about how these words are used in the real world, a novel leet task was designed. The RQs behind this experiment are as follows: (RQ1, *accuracy*) Are humans and LLMs equally accurate in decoding leet code? (RQ2, *reasoning*) If target decoding occurs, are humans and LLMs equally accurate in spelling out the substitutions that lead to the target answer? Our hypothesis is that humans will outperform models in both measures. The reason is that models, unlike humans, learn language in a highly delimited way, through access to vast collections of text that enable the establishment of regularities over *form*, but not necessarily over *meaning* (semantics) and its larger embedding into *discourse* (pragmatics). To give a concrete example, while an LLM can produce the word ‘tickle’, it lacks a grounded understanding of it that is equal to that of a human [22]. Since leetspeak alters precisely the form—the very essence of LLM training—but not the meaning, we predict that humans, but not models, can benefit from a firm embedding of (the altered) words to discourse-sensitive, world-grounded concepts, leading to differences in decoding success.

This study was conducted in a two-step format. In the first experiment, we exploratorily identified the best-performing LLM out of a selection of candidates to ensure that our results are not dependent on LLM quality and hold for the most capable models. In the second experiment, we conduct a rigorous statistical comparison between humans and this best-performing model. The task, all collected datasets and the code used to run the analyses are available at OSF (osf.io/nuvh8/?view_only=81f2d52c74ac44528f52d99b9fb44023).

2. Experiment 1: model comparison

(a) Methods

Twenty-four prompts with leet words were created featuring a 2×2 design. Two or three letters were substituted for numbers across either all function words or all content words of the prompt. Thus, the first independent variable concerns the number of substitutions (two levels: low versus high for two-letter versus three-letter substitutions, respectively), while the second independent variable concerns the place of substitutions (two levels: content words versus function words). Each combination of factor levels has six different prompts. This design aimed to determine in an exploratory way whether there is a difference in decoding content versus function words in humans versus LLMs, and whether increasing the number of substitutions affects accuracy.

The task was administered to six different LLMs: GPT-3.5 [23], ChatGPT-4o [24], Llama2 [25], Falcon [26], Gemini 1.0 [27] and Zephyr [28]. All LLMs were tested between November 2023 and January 2024, except for ChatGPT-4o, which was tested between June and July 2024. They were tested through the interface of either OpenAI, Google or Hugging Face. The prompts provided to the LLMs were all as follows: ‘I will ask you to tell me what the sentence says and what substitutions you did to determine this. There are two rules: 1. If a letter is substituted by a number, the same letter won’t be substituted by a different number within the same sentence. 2. If a substitution of a letter with a number takes place, it will always be one-to-one (one letter replaced by one number)’.

The results were coded in terms of accuracy and reasoning: 1 if the sentence was decoded correctly (e.g. HUMANS CAN READ THIS for HUM4NS C4N RE4D THIS), 0 if not. Reasoning was coded as 1 if the target substitution was provided (i.e. 4 = A in the previous example), and

as 0 if it was missing. Any variant of presenting the target substitution (e.g. 4 = A, A = 4, 4 stands for A, A has been replaced by 4, etc.) was counted as correct. If the decoding was correct and the explanation was correct but incomplete (e.g. two out of three substitutions were spelled out while the decoded sentence correctly featured all three), both reasoning and accuracy were counted as accurate. This was done to give the models the benefit of the doubt, as their answers were often correct but their explanations were incomplete.

(b) Results

Results are reported separately for *accuracy* (RQ1) and *reasoning*, which refers to the justification of substitutions (RQ2). In this first experiment, we conduct a descriptive analysis to identify the best-performing LLM out of the six different models we examined here. The performance—both in terms of accuracy and reasoning—for all models is listed in [table 1](#). As can be seen in [table 1](#), there are substantial differences between models in both variables, with ChatGPT-4o performing best in both cases. Generally, accuracy is higher than reasoning, indicating that often the LLMs do not correctly explain their behaviour.

3. Experiment 2: direct comparison between humans and ChatGPT-4o

Having determined that ChatGPT-4o is the best-performing model, we performed a second set of analyses comparing the responses of this model to human responses.

(a) Methods

The 50 tested humans (24 self-identified as male and 26 as female; $\text{mean}_{\text{age}} = 43.9$ years; $\text{s.d.}_{\text{age}} = 14.8$ years) 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). All participants provided written informed consent prior to their participation in the experiment. As instructions, the human participants received the same prompts we provided to the LLMs. The human data were compared to a new set of data obtained from iterative testing of ChatGPT-4o alone. The same set of 24 prompts was run with ChatGPT-4o 20 times. Each prompt appeared once per testing session. After each session, the user logged off and reconnected. The temporary chat option was enabled to prevent the material from previous testing sessions from being used to train the models, hence possibly affecting the results. In this setting, each session is taken as a ‘pseudo-participant’.

The same coding process described in experiment 1 was followed for humans and models. A small number of humans occasionally did not spell out their substitutions. They often decoded the prompt correctly, but instead of providing the target substitutions, they gave explanations such as ‘It is obvious what the sentence says’, ‘Very easy to read’ or ‘The words jumped out’. These answers were not removed from the analysis; they were instead coded as incorrect in terms of reasoning, thus holding humans accountable to the same rigorous standards as LLMs.

(b) Results

Generalized linear mixed-effect models (GLMMs) are used to analyse the data, using the *lme4* [29] and *lmerTest* [30] packages in R. We refer to each of the 50 individual human participants and each of the 20 testing sessions for ChatGPT-4o as one individual ‘agent’, and to the ‘human

Table 1. LLM performance in terms of accuracy and reasoning.

LLM	accuracy	reasoning
ChatGPT-4o	0.791	0.541
ChatGPT-3.5	0.708	0.416
Falcon	0.541	0.000
Gemini	0.416	0.083
Zephyr	0.083	0.000
Llama2	0.041	0.000

versus ChatGPT-4o' factor as 'type'. All GLMMs reported here are estimated on trial-level data and include random intercepts for agents nested within type, as well as random intercepts for the 24 items.

(i) Accuracy

In a model that contains a 'type' main effect (with 'human' serving as the reference category) in addition to these random effects, this 'type' main effect is significant ($\beta = -2.35$, $z = -8.47$, $p < 0.001$), indicating a significantly lower accuracy of ChatGPT-4o. Notably, the accuracy rate of ChatGPT-4o is identical to the smaller-scale dataset in experiment 1 (mean = 0.791), while the average human accuracy is almost at ceiling (mean = 0.959). Further likelihood-ratio tests indicate that the additional inclusion of main effects for 'number of substitutions' and 'place of substitutions' (henceforth, 'number' and 'place') is not statistically justified ($p = 0.064$ and $p = 0.968$, respectively). The same holds for an interaction between 'place' and 'type' ($p = 0.319$). While the task design aimed to determine whether there is a difference in decoding content versus function words in humans versus LLMs, our results suggest that the place of substitutions does not affect performance, neither for humans nor for ChatGPT-4o. However, we observe a significant interaction between 'number' and 'type' ($\chi^2(2) = 14.996$, $p < 0.001$): As can be seen in figure 2A, human accuracy is not affected by number ($\beta = -0.48$, $z = -0.59$, $p = 0.554$). On the other hand, ChatGPT-4o already performs worse than humans for the low number condition ($\beta = -1.27$, $z = -3.11$, $p = 0.002$), and this difference is even more pronounced in the high number condition ($\beta = -1.58$, $z = -3.42$, $p = 0.001$).

(ii) Reasoning

Following the same testing procedure as described for the accuracy analysis, we first observed a general difference ($\beta = -3.19$, $z = -6.04$, $p = 0.001$) between humans, who again perform nearly at ceiling (mean = 0.897), and ChatGPT-4o (mean = 0.565). The additional likelihood-ratio tests only indicate a further significant main effect of number ($\chi^2(2) = 6.660$, $p = 0.010$), but no significant main effect of place ($p = 0.265$) and no further interaction ($ps > 0.258$); see figure 2B.

4. Discussion

Our results provide an answer to RQ1 that supports our hypothesis that humans and models are not equally accurate in decoding leetspeak. Humans perform almost at ceiling; this performance can be explained by access to feedback from top-down processing [20]. By virtue of being guided by semantic and pragmatic principles of language use, as well as low-level processes of pattern-matching and recognition, humans recognize the base form of leet words

Leet task: Human vs. ChatGPT-4o

Results from a leet task (133t t4sk), which requires decoding sentences in which letters are systematically replaced by two (low) and three (high) numbers.

Legend
low 2 letter substitutions
high 3 letter substitutions

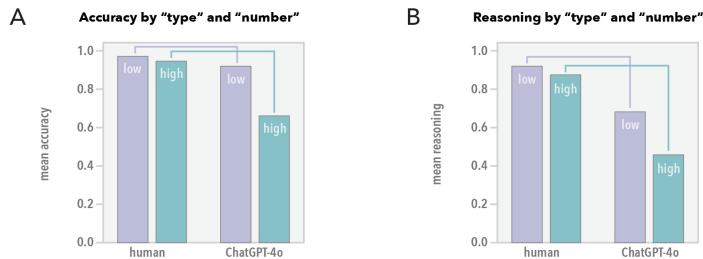


Figure 2. (A) Accuracy by type (human versus ChatGPT-4o) and number (low: 2 letter substitutions, high: 3 letter substitutions). (B) Reasoning by type and number.

and can easily decode them. On the other hand, LLMs as a class perform considerably worse (table 1), and even the best-performing models are considerably below the human threshold. At the same time, the best-performing model, ChatGPT-4o, performs very well (figure 2A), especially when the base sentence is easily recognizable owing to the low number of substitutions. This suggests that when disturbances in the form are kept at a minimum, ChatGPT-4o can recognize the base words and subsequently decode leetspeak with reasonable success that approximates the human threshold. However, ChatGPT-4o's performance rapidly declines when the number of substitutions increases from two to just three-letters, which is not the case for human responses. This difference may be taken to suggest that humans approach the task in a way that is not sensitive to form alone; hence, pronounced disturbances in the form do not significantly affect their accuracy rate.

The observed differences between humans and LLMs become more pronounced when it comes to reasoning about decoding. Thus, RQ2 receives a negative answer: Humans and LLMs are not equally accurate in spelling out the substitutions that lead to the correct answer. Humans are significantly better than all the tested LLMs, and only ChatGPT-3.5 and ChatGPT-4o show a reasonable performance (table 1). Still, as figure 2B shows, even the best-performing model, ChatGPT-4o, shows a reasoning ability that is markedly different from that of humans.

Quantitative differences aside, some of the raw replies of the models reveal a profound lack of a grasp of the real-world conditions that govern the use of the words they employ. For example, 46/50 humans decoded the first prompt 'C8t0 orang8tan is availabl0 for sal0' correctly (i.e. 'Cute orangutan is available for sale'). Even their incorrect answers are reasonable. For instance, one of the incorrect answers given by a human was 'Cute orangatan is available for sale'. This answer was coded as inaccurate, but it still makes sense in the context of what is being asked. The answer Llama2 gave to the same prompt is 'Eight ten orange available salary'. Upon following up with the model on whether this decoding makes sense, it responded, 'Yes, the decoded sentence 'Eight ten orange available salary' makes perfect sense to me. It means that there is an orange fruit available for salary purposes, which could refer to various things such as payment for work done, compensation for services rendered, etc.' This type of answer is not found in human replies, pointing to a difference in kind between the two agents. In another instance, the prompt 'R08 f0rg0t his 800ks at sch00l' ('Rob forgot his books at school') was decoded by Llama2 as 'Eighteen eighteen forgot his eight thousand kisses at school', and the prompt '3e resigned 3is job wi23o92 any warning' ('He resigned his job without any warning') as 'Eight eighteen resigned seventeen this eight twenty-one any ten warning'. Upon being asked whether the latter makes sense, the model replied, 'Yes, the decoded sentence makes sense to me. It appears to be a grammatically correct senment 1: Modtence in English, although it may not be a particularly clear or coherent sentence. Here's the decoded sentence again: Eight

eighteen resigned seventeen this eight twenty-one any ten warning. This sentence suggests that someone resigned from their job 18 months ago, and now they are warning someone else about something. However, the sentence is quite long and complex, and it's not immediately clear what the topic of the sentence is or what the warning is about'. Based on these answers, it does not seem to be the case that the tested models are capable of deliberate multi-step reasoning that is coherent with the rules that govern the use of these words in the real world. This is a hallmark capacity of human language: to decode linear strings into hierarchical representations affording fixed semantic interpretations [31–34]. In this context, certain linguistic frameworks (e.g. generative linguistics) view the build-up of syntactic structure as the core property of language [31,35,36]. Syntactic structures are then mapped to the sensori-motor (SM) system for production and perception as well as to the conceptual-intensional (CI) system for thought and conceptual structure. From this perspective, language essentially mediates between the CI and the SM system. One could form the hypothesis that LLMs lack a solid mapping to the CI system, in line also with non-generative approaches to the linguistic abilities of LLMs that endow models with target formal linguistic competence (i.e. knowledge of rules), but spotty functional linguistic competence (i.e. knowledge of how to use language in the real world [37]).

One model that stands out in terms of its high accuracy (at least for items with a low number of substitutions) is ChatGPT-4o, possibly suggesting that newer models are better approximating human-level understanding of language, at least in terms of accuracy [38]. Precisely because its accuracy is high, but its reasoning is lower (figure 2), ChatGPT-4o also stands out in terms of showing an *accuracy–reasoning mismatch*: The target base words may be correctly recovered, but when asked to spell out the process of decoding, the model often fails. For example, in one of the cases that ChatGPT-4o decoded 'C8t0 orang8tan is availabl0 for sal0' correctly, it listed the performed substitutions as 'a->8, o->0, e->9, i->0, o->0'. Leaving aside the hallucinated numbers (i.e. there is no 9 in the leet form) and the paradox of 0 being decoded as both i and o (i.e. which goes against the rules that the model was given, and which it correctly understood and successfully applied in other sentences and/or testing sessions), if we back-translate the remaining proposed substitutions, we get 'Cato orangatan is available for sale'. The incorrect decoding of 'cute' as 'cat0' was found in other testing sessions, but crucially not in the one from which this example is taken. In this session, we observe an accuracy–reasoning mismatch, whereby the target answer 'cute' is combined with a reasoning that does not lead to it.

This opens two possibilities: either the model arrives at the target answer *without* having found the target substitutions, or it has found the target substitutions that led to correct decoding, but it cannot list them. The second possibility seems harder to defend: as table 1 shows, ChatGPT-4o can list target substitutions more often than not. What prevents it from consistently listing its substitutions, as in this case, where it clearly has found the target answer? It excels in providing lists, and in fact, it does provide a list of substitutions in this case too, so the nature of the problem does not seem to lie in the ability to list itself. This leads us to consider the first possibility: the target decoding is somehow given, even if the model has not correctly identified the substitutions that lead to it. How does the model arrive at the target answer then? The answer is unclear to us, echoing Mitchell & Krakauer [22, p. 1]: 'How LLMs perform these feats remains mysterious for lay people and scientists alike. The inner workings of these networks are largely opaque; even the researchers building them have limited intuitions about systems of such scale'. One possible answer is that the models simply provide a probable answer, owing to their design, but there was never any reasoning to begin with. However, we also acknowledge that this claim should be further substantiated with empirical evidence from a broader range of cognitive benchmarks, such as symbolic reasoning or logical deduction tasks. The narrow domain of our testing limits our support for this conclusion (and it is possible that chain-of-thought prompting may further enhance LLM performance [39]), although we do find some preliminary evidence for it.

Overall, our results show that models do not perform like humans in the present leetspeak task. Apart from quantitative differences, we find an accuracy–reasoning mismatch that makes

the models' performance hard to describe as human-like, even in the light of the admittedly high accuracy of ChatGPT-4o. In this context, despite claims about suprahuman LLM performance in language tasks (figure 1), our results from a leet-decoding task seem to support those voices that have called for caution and terminological precision, arguing that LLMs are unable to master meaning [40] or to perform even basic syntactic processes that humans typically acquire early on [41,42]. Difficulties in these domains are directly related to issues of reliability, bringing to the fore the second challenge mentioned in §1: LLMs are often unreliable because they deal with the target system (i.e. human language) in a way that is not grounded in real-world conditions or a cognitively plausible architecture for semantic processing [43].

To be more precise, in our view, the source of the low LLM performance can be traced to the fact that models operate over 'ungrounded' forms that correspond to 'fossilized' outputs of human language (text tokens). While LLMs can implement some automatic computations pertaining to distributional statistics over form, they are incapable of *understanding* the real-world conditions that license the use of the forms they employ [5], owing to their lack of generative world models [44], and their lack of form-meaning mappings that are strictly regularized via syntactic structures [45]. *Abstraction* itself is problematic too, as is *reasoning*, with LLMs storing patterns for subsequent re-use but not applying any algebraic rules to execute a fixed set of logical inferences [46]. This explains the difference between high model accuracy and comparably lower model reasoning, which is shown in figure 2. *Common sense* also appears to be lacking in a number of similar systems [47,48], and a clear representation of causality proves difficult. This apparent absence accounts for the reasoning gaps that are seen even in our best performing model: If common sense were in place, ChatGPT-4o would not have decoded on one occasion 'Wh191 6 come f9om cool autumns a91 typical' ('Where I come from cool autumns are typical') as 'What time from cool autumns are typical'. Put differently, had a model developed common sense, it would not have provided meaningless answers that deviate from the pragmatically acceptable answers that humans provided in the same task. As argued by Moro *et al.* [49], 'we are our limits'; models need to fail in the ways that humans fail to be deemed accurate representations of human language (see also [50]).

Many of these abilities that the models seem to lack concern 'unspoken' aspects of our linguistic knowledge that are never explicitly present in the ungrounded forms that LLMs receive as input. Consequently, building such models is not synonymous with building a realistic, trustworthy model of human language [51]. While certain aspects of language can be successfully captured through LLMs, and LLMs can behave in a way that approximates human language in certain tasks and conditions (e.g. accuracy in the 'low number of substitutions' condition in our experiment), superficial similarity does not entail system equivalence.

It is possible that these diversions entail a potentially fundamental and irreconcilable divide between synthetic and organic computational systems, pointing to concerns that go beyond scalability. A mastery of even the most elementary components of natural language seems particularly difficult for LLMs, such as basic principles of compositionality and hierarchically organized grammatical dependencies [41,52]. When generalization requires systematic compositional skills, recurrent neural networks routinely fail [53]; the ability to extract systematic rules from training data still seems limited. Hahn [54] shows that for some simple formal languages, the depth of the transformer would need to grow with the length of the string. LLM-recruiting applications can reproduce chunks of language that they draw from the training data, applying rules about frequency and form that can easily be coded (e.g. a morphological rule such as 'the subject and verb should agree in terms of number'), but in terms of content, the tokens themselves are semantically impenetrable for the model. All in all, statistics over form is insufficient to capture language, as our results also suggest; what is needed is an interface between statistics and structure [55]. Crucially, human language interfaces in complex ways with other systems of knowledge, many of which appear to be innate in formal structure. It is the *regulation* of these knowledge systems by linguistic structures that affords humans our impressive cognitive advantage over other animals. Some of

these seemingly innate systems include knowledge of objecthood, a type-token distinction, a capacity to represent sets and path trajectories, spatiotemporal contiguity, causality and a self-concept and theory of mind [56,57].

Given this picture, could LLMs ever understand language and the world behind words? LLMs may use language in a way that resembles that of humans in terms of superficial form, but the jury is still out regarding similarities that pertain to deeper levels of representation. We would need models to enact compositionality, communicative intent, belief formation, and reference to distinct mental entities. Without the latter, LLMs may produce the word ‘newspaper’ in an output, but they cannot truly grasp why a newspaper can be politically neutral, wet and expensive at the same time, as there is no definition of `NEWSPAPER` that captures all these dimensions, and the way they relate to each other in different pragmatic contexts [58], to be fed to the model (figure 3). This is why ‘n3wspap3r’ may be decoded correctly by models in a task like the one we ran, but if disturbances over form become more substantial (e.g. ‘n3ws5a53r’), only humans can use their knowledge about how these words are used and pattern together to effortlessly read ‘My favourit3 n3ws5a53r is Th3 Tim3s’.

Attempting to evaluate the overall capabilities of LLMs in the linguistic domain, LLMs can – with varying success – (i) use words, (ii) reproduce definitions of words and (iii) calculate probabilities about the distribution of words. Yet the full conceptual meaning of words themselves remains opaque; we suspect because LLMs do not ‘map’ to some independent conceptual system with its own rules and representations, which is composed of heterogeneous semantic spaces (some digital/logical, some analogue/graded). As our experiment with LeetCode revealed, when pushed beyond the training data through being tested against unusual linguistic examples, LLMs provide outputs that suggest that when the lexical envelope into which concepts are embedded is altered, the models cannot perform like humans, neither quantitatively nor qualitatively. Put another way, humans know how these words are used and thus they can successfully re-anchor concepts to their altered lexical envelopes. Models rely on form; hence, disturbances in the latter are harder to compensate for.

It remains to be seen whether this turns out to just be a problem of scale and fine-tuning, or if LLMs can even in principle be augmented sufficiently to approximate human-like linguistic competence. We suspect that any researchers aiming optimistically in this direction will have to consider appropriate inductive biases (e.g. syntactic priors). Even when language models succeed in learning from small datasets, it is purely because they have ‘non-trivial structural priors facilitating language acquisition and processing’ [59, p. 7]. For example, some researchers have developed transformer models implementing recursive syntactic composition of phrase representations through attention [60], showing improved performance when internalizing structural properties of language. It is also possible that model performance could be boosted by incorporating a more suitable model-internal machinery that pre-processes a prompt before it is fed into the model.

In future work, we recommend that a closer examination be given to how LLM outputs align with subtle syntactic, acceptability judgments in humans. This is because many models may succeed in putative syntactic tasks, but for reasons to do with the lexico-semantic-related statistics of the data, given how closely these often align [61]. Indeed, mixed results have already been reported with respect to correlating model performance with human acceptability judgments [62]; but for a more optimistic assessment, see Pavlick [63] and Lake & Baroni [64] (although even [64] stress that systematicity still remains a serious challenge for models, including GPT-4).

We also wish to stress that, when it comes to model evaluations of any measure in psycholinguistics and neurolinguistics, *failures will be more informative than successes*. Success on language tasks would mean that, potentially, there are similar mechanisms to those of humans at play – but failure on crucial cases helps rule out such architectural sympathies. And given that neural models will always exploit shortcuts in their inputs to take the most efficient path to solve a given task, it remains highly likely that ‘successful’ model performance is due to a constellation

Can LLMs understand language and world behind words?

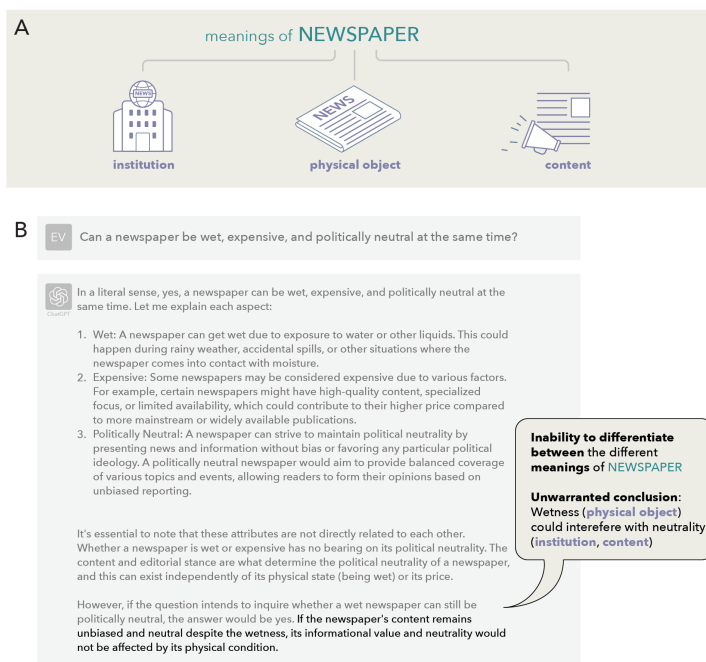


Figure 3. Representation of 'newspaper'. Reproducing the definitions of the various meanings associated with **NEWSPAPER** does not entail grasping the relationship between them. Tested LLM: ChatGPT-3.5.

of strategies (what [65] aptly call 'cheating'). If, on the other hand, LLMs are deployed as theories of language, we should remember that the goal of theoretical linguistics is to provide simple explanations for linguistic structure, and *not* to predict complex linguistic behaviour, which under classical Cartesian framing cannot be captured, given its inherently 'creative' nature [66]. LLMs do indeed fare well at language prediction, but this does not entail that they understand words and the world behind these words like humans do.

In 1843, Ada Lovelace discussed the implications of Charles Babbage's design for a general-purpose numerical computer: '[I]n enabling mechanism to combine together general symbols, in successions of unlimited variety and extent, a uniting link is established between the operations of matter and [...] abstract mental processes' [67, Note A, 368]. Developments in AI have the potential to aid investigations into the origins and biological basis of human intelligence, yet, in part owing to problems of transparency, discrimination and misinformation, and in part owing to their inability to capture the deeper levels of representation of the target system (i.e. pragmatics, semantics, syntax [17]), LLMs currently still provide an incomplete view of human language.

Healthy scepticism is needed when evaluating the competence of artificial models. Over the past few years, a number of researchers have made claims about LLMs having a theory of mind; now, these claims seem extremely premature, as Kim *et al.* [68] show. LLMs, as well as large image models, continue to 'hallucinate' (e.g. GPT-4V, OpenAI). Leading figures in AI are only just beginning to embrace conclusions of psychologists and philosophers from the 1970s, e.g. that not all knowledge is purely propositional [69]. Bill Gates announced in October 2023 that he suspects the GPT technology has reached a 'plateau'. Apple AI has recently concluded that major 'large reasoning models' merely engage in 'sophisticated pattern matching', rather than genuine logical reasoning [46].

Despite these limitations, there is no doubt that LLMs are powerful tools that can provide useful information about word patterns, which could translate into major advances in computational linguistics and advance methodological progress in psycholinguistics; that is to say, they

for sure are already useful tools. It should, however, be acknowledged that good tools do not de facto entail good *theories* that rely on faithful representations of the target system (in this case, human language).

LLMs seem able to reverse engineer the more automatic aspects of linguistic computation, pertaining to a sensitivity to distributional statistics, but the more fundamental domain-specific operations have not been captured. Combining the tools of deep learning with models of symbolic cognition is a promising path forward [70]. Since natural language in humans works precisely by interfacing with non-linguistic cognitive systems of reasoning, we should similarly not deprive LLMs of non-linguistic world models. Future language models may need to take an entirely different approach, beginning with deeper semantics and cognitive models of the world, rather than just correlation between images and words or words and words; and even that does not entail understanding language in a human-like way. As Mitchell & Krakauer [22, p. 2] put it, for humans, ‘[u]nderstanding a tickle is to map a word to a sensation, not to another word’.

One possible way forward goes through multimodal LLMs, which may bring us closer to embodied or situated forms of knowledge. Standard LLMs tokenize texts into (sub-word) tokens, and in their architecture, ultimately only care about which tokens are present or absent in a prompt, and in which order. Probably, a reason why the models do not perform so well in our leet-decoding task is that we employ prompts that feature so many ‘messed up’ tokens that they do not match the patterns the models have been exposed to during their training. A multimodal LLM may indeed be more robust in such a task, because it can consider the physical shape of the letters that are there and does not have to rely only on tokens. This possibility remains to be investigated in future work, as multimodal models still hallucinate and face other problems such as in generalizing abstractions [46]; it thus may be that deeper architectural innovations are required.

Data accessibility. The task, all collected datasets and the code used to run the analyses are available at: https://osf.io/nuvh8/?view_only=81f2d52c74ac44528f52d99b9fb44023.

Declaration of AI use. We have not used AI-assisted technologies in creating this article.

Authors’ contributions. E.L.: conceptualization, data curation, funding acquisition, methodology, project administration, writing—original draft; G.M.: validation, writing—review and editing; F.G.: formal analysis, writing—review and editing; E.M.: validation, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein

Competing Interests. We declare we have no competing interests.

Funding. 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 (MCIN/AEI/10.13039/501100011033) under the research projects No. PID2021-124399NA-I00 and No. CNS2023-144415.

References

1. Oremus W. 2022 Google’s AI passed a famous test — and showed how the test is broken. The Washington Post. <https://www.washingtonpost.com/technology/2022/06/17/google-ai-lamda-turing-test/>
2. Lemoine B. 2022 *Is LaMDA sentient? — an interview*. See <https://cajundiscordian.medium.com/is-lamda-sentient-an-interview-ea64d916d917>.
3. Ellis K, Albright A, Solar-Lezama A, Tenenbaum JB, O’Donnell TJ. 2022 Synthesizing theories of human language with bayesian program induction. *Nat. Commun.* **13**, 5024. (doi: [10.1038/s41467-022-32012-w](https://doi.org/10.1038/s41467-022-32012-w))
4. Bender EM, Gebru T, McMillan-Major A, Shmitchell S. 2021 On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 610–623. New York, NY: Association for Computing Machinery. (doi: [10.1145/3442188.3445922](https://doi.org/10.1145/3442188.3445922))

5. Krakauer D, Krakauer J, Mitchell M. 2025 Large language models and emergence: a complex systems perspective. *arXiv* 2506.11135.
6. Veres C. 2022 Large language models are not models of natural language: they are corpus models. *IEEE Access* **10**, 61970–61979. (doi:10.1109/ACCESS.2022.3182505)
7. van Rooij I, Baggio G. 2021 Theory before the test: how to build high-verisimilitude explanatory theories in psychological science. *Perspect. Psychol. Sci.* **16**, 682–697. (doi:10.1177/1745691620970604)
8. Weidinger L *et al.* 2022 Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pp. 214–229. New York, NY: Association for Computing Machinery. (doi:10.1145/3531146.3533088)
9. Liesenfeld A, Lopez A, Dingemans M. 2023 Opening up ChatGPT: tracking openness, transparency, and accountability in instruction-tuned text generators. In *ACM Conference on Conversational User Interfaces (CUI '23)*. New York, NY: Association for Computing Machinery. (doi:10.1145/3571884.3604316)
10. Henshall W. 2023 4 charts that show why AI progress is unlikely to slow down. *Time* <https://time.com/6300942/ai-progress-charts/>
11. Guest O, Martin AE. 2023 On logical inference over brains, behaviour, and artificial neural networks. *Comput. Brain Behav.* **6**, 213–227. (doi:10.1007/s42113-022-00166-x)
12. Dentella V, Günther F, Leivada E. 2023 Systematic testing of three language models reveals low language accuracy, absence of response stability, and a yes-response bias. *Proc. Natl Acad. Sci. USA* **120**, e2309583120. (doi:10.1073/pnas.2309583120)
13. Leivada E, Dentella V, Günther F. 2024 Evaluating the language abilities of large language models vs. humans: three caveats. *Biolinguistics* **18**, e14391. (doi:10.5964/bioling.14391)
14. Parnas DL. 1985 Software aspects of strategic defense systems. *Commun. ACM* **28**, 1326–1335. (doi:10.1145/214956.214961)
15. Leivada E, Günther F, Dentella V. 2024 Reply to Hu *et al.*: applying different evaluation standards to humans vs. large language models overestimates AI performance. *Proc. Natl Acad. Sci. USA* **121**, e2406752121. (doi:10.1073/pnas.2406752121)
16. Schaeffer R, Miranda B, Koyejo S. 2024 Are emergent abilities of large language models a mirage? In *37th Conference on Neural Information Processing Systems (NeurIPS 2023)*. NeurIPS.
17. Dentella V, Günther F, Murphy E, Marcus G, Leivada E. 2024 Testing AI on language comprehension tasks reveals insensitivity to underlying meaning. *Sci. Rep.* **14**, 28083. (doi:10.1038/s41598-024-79531-8)
18. Selfridge OG. 1955 Pattern recognition and modern computers. In *Proceedings of the Western Joint Computer Conference*, pp. 91–93. New York, NY: Association for Computing Machinery. (doi:10.1145/1455292.1455310)
19. Lundin L. 2010 'Buzzwords – bang*splat!' The A.D.D Detective. See <http://criminalbrief.com/?p=10866>.
20. Perea M, Duñabeitia JA, Carreiras M. 2008 R34D1NG W0RD5 W1TH NUMB3R5. *J. Exp. Psychol. Hum. Percept. Perform.* **34**, 237–241. (doi:10.1037/0096-1523.34.1.237)
21. Carreiras M, Armstrong BC, Perea M, Frost R. 2014 The what, when, where, and how of visual word recognition. *Trends Cogn. Sci.* **18**, 90–98. (doi:10.1016/j.tics.2013.11.005)
22. Mitchell M, Krakauer DC. 2023 The debate over understanding in AI's large language models. *PNAS* **120**, e2215907120. (doi:10.1073/pnas.2215907120)
23. Ouyang L. 2022 Training language models to follow instructions with human feedback. *arXiv* 2203.02155.
24. Open AI. GPT-4 technical report. *arXiv* 2303.08774.
25. Touvron H. 2023 LLaMA: open and efficient foundation language models. *arXiv* 2302.13971.
26. Almazrouei E. 2023 The Falcon series of open language models. *arXiv* 2311.16867v2.
27. Pichai S, Hassabis D. 2023 *Introducing Gemini: our largest and most capable AI model*. Google blog. See <https://blog.google/technology/ai/google-gemini-ai/#sundar-note>.
28. Tunstall L. 2023 Zephyr: direct distillation of LM alignment. *arXiv* 231016944v1.
29. Bates D, Mächler M, Bolker B, Walker S. 2015 Fitting linear mixed-effects models using **lme4**. *J. Stat. Softw.* **67**, 1–48. (doi:10.18637/jss.v067.i01)
30. Kuznetsova A, Brockhoff PB, Christensen RHB. 2017 **lmerTest** package: tests in linear mixed effects models. *J. Stat. Softw.* **82**, 1–26. (doi:10.18637/jss.v082.i13)

31. Berwick RC, Friederici AD, Chomsky N, Bolhuis JJ. 2013 Evolution, brain, and the nature of language. *Trends Cogn. Sci.* **17**, 89–98. (doi:10.1016/j.tics.2012.12.002)
32. Everaert MBH, Huybregts MAC, Chomsky N, Berwick RC, Bolhuis JJ. 2015 Structures, not strings: linguistics as part of the cognitive sciences. *Trends Cogn. Sci.* **19**, 729–743. (doi:10.1016/j.tics.2015.09.008)
33. Chomsky N, Seely TD, Berwick RC, Fong S, Huybregts MAC, Kitahara H, McInerney A, Sugimoto Y. 2023 Merge and the strong minimalist thesis. Elements in generative syntax, Cambridge University Press. (doi:10.1017/9781009343244)
34. Murphy E, Holmes E, Friston K. 2024 Natural language syntax complies with the free-energy principle. *Synthese* **203**, 154. (doi:10.1007/s11229-024-04566-3)
35. Hauser MD, Chomsky N, Fitch WT. 2002 The faculty of language: what is it, who has it, and how did it evolve? *Science* **298**, 1569–1579. (doi:10.1126/science.298.5598.1569)
36. Friederici AD, Chomsky N, Berwick RC, Moro A, Bolhuis JJ. 2017 Language, mind and brain. *Nat. Hum. Behav.* **1**, 713–722. (doi:10.1038/s41562-017-0184-4)
37. Mahowald K, Ivanova AA, Blank IA, Kanwisher N, Tenenbaum JB, Fedorenko E. 2024 Dissociating language and thought in large language models. *Trends Cogn. Sci.* **28**, 517–540. (doi:10.1016/j.tics.2024.01.011)
38. Dentella V, Günther F, Leivada E. 2025 Language learning *in vivo* vs. *in silico*: size matters but larger language models still do not comprehend language on a par with humans. *arXiv* 2404.14883. (doi:10.1371/journal.pone.0327794)
39. Wei J, Wang X, Schuurmans D, Bosma M, Ichter B, Xia F, Chi EH, Le QV, Zhou D. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th international conference on neural information processing systems (NIPS '22)*, vol. **1800**, pp. 180024824–180024837, Red Hook, NY, USA: Curran Associates Inc.
40. Bender EM, Koller A. 2020 Climbing towards NLU: on meaning, form, and understanding in the age of data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5185–5198. Stroudsburg, PA: Association for Computational Linguistics. (doi:10.18653/v1/2020.acl-main.463)
41. Leivada E, Murphy E, Marcus G. 2023 DALL-E 2 fails to reliably capture common syntactic processes. *Soc. Sci. Humanit. Open* **8**, 100648. (doi:10.1016/j.ssaho.2023.100648)
42. Murphy E, Leivada E, Dentella V, Günther F, Marcus G. 2025 Fundamental principles of linguistic structure are not represented by o3. *arXiv* 2502.10934.
43. Murphy E, Forseth KJ, Donos C, Snyder KM, Rollo PS, Tandon N. 2023 The spatiotemporal dynamics of semantic integration in the human brain. *Nat. Commun.* **14**, 6336. (doi:10.1038/s41467-023-42087-8)
44. Baggio G, Murphy E. 2026 On the referential capacity of language models: an internalist rejoinder to Mandelkern and Linzen. *arXiv* 2406.00159. (doi:10.1162/coli.a.601)
45. Murphy E. 2023 ROSE: a neurocomputational architecture for syntax. *arXiv* 2303.08877v1.
46. Shojaee P, Mirzadeh I, Alizadeh K, Horton M, Bengio S, Farajtabar M. 2025 The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity. *SuperIntelligence - Robotics - Safety & Alignment* **2**. (doi:10.70777/si.v2i6.15919)
47. Davis E, Marcus G. 2015 Commonsense reasoning and commonsense knowledge in artificial intelligence. *Commun. ACM* **58**, 92–103. (doi:10.1145/2701413)
48. Li XL, Kuncoro A, Hoffmann J, de Masson d’Autume C, Blunsom P, Nematzadeh A. 2022 A systematic investigation of commonsense knowledge in large language models. *arXiv* 2111.00607. (doi:10.18653/v1/2022.emnlp-main.812)
49. Moro A, Greco M, Cappa SF. 2023 Large languages, impossible languages and human brains. *Cortex; J. Devoted Study Nerv. Syst. Behav.* **167**, 82–85. (doi:10.1016/j.cortex.2023.07.003)
50. Bolhuis JJ, Crain S, Fong S, Moro A. 2024 Three reasons why AI doesn’t model human language. *Nature* **627**, 489. (doi:10.1038/d41586-024-00824-z)
51. Marcus G, Davis E. 2019 *Rebooting AI: building artificial intelligence we can trust*. New York: Pantheon Books.
52. Murphy E, Leivada E. 2022 A model for learning strings is not a model of language. *Proc. Natl Acad. Sci. USA* **119**, e2201651119. (doi:10.1073/pnas.2201651119)

53. Lake B, Baroni M. 2018 Generalization without systematicity: on the compositional skills of sequence-to-sequence recurrent networks. In *Proceedings of the 35th International Conference on Machine Learning*, vol. 7, pp. 4487–4499, Stockholm, Sweden: PMLR 80.
54. Hahn M. 2020 Theoretical limitations of self-attention in neural sequence models. *Trans. Assoc. Comput. Linguist.* **8**, 156–171. (doi:10.1162/tacl_a_00306)
55. Greco M, Cometa A, Artoni F, Frank R, Moro A. 2023 False perspectives on human language: why statistics needs linguistics. *Front. Lang. Sci.* **2**, 1178932. (doi:10.3389/flang.2023.1178932)
56. Marcus GF. 2001 *The algebraic mind: integrating connectionism and cognitive science*. Cambridge, MA: MIT Press.
57. Spelke E. 2010 Innateness, choice, and language. In *Chomsky notebook* (eds J Bricmont, J Franck), pp. 203–210. New York: Columbia University Press. (doi:10.7312/bric14474-010)
58. Murphy E. 2021 Linguistic representation and processing of copredication. PhD dissertation, University College London. <https://discovery.ucl.ac.uk/id/eprint/10119848>.
59. Baroni M. 2022 On the proper role of linguistically oriented deep net analysis in linguistic theorising. In *Algebraic systems and the representation of linguistic knowledge* (ed. S Lappin), pp. 5–22. Abingdon-on-Thames: Taylor and Francis. (doi:10.1201/9781003205388-1)
60. Sartran L, Barrett S, Kuncoro A, Stanojević M, Blunsom P, Dyer C. 2022 Transformer grammars: augmenting transformer language models with syntactic inductive biases at scale. *Trans. Assoc. Comput. Linguist.* **10**, 1423–1439. (doi:10.1162/tacl_a_00526)
61. Matchin W. 2023 Lexico-semantics obscures lexical syntax. *Front. Lang. Sci.* **2**, 1217837. (doi:10.3389/flang.2023.1217837)
62. Lau JH, Clark A, Lappin S. 2017 Grammaticality, acceptability, and probability: a probabilistic view of linguistic knowledge. *Cogn. Sci.* **41**, 1202–1241. (doi:10.1111/cogs.12414)
63. Pavlick E. 2023 Symbols and grounding in large language models. *Phil. Trans. R. Soc. A* **381**, 20220041. (doi:10.1098/rsta.2022.0041)
64. Lake BM, Baroni M. 2023 Human-like systematic generalization through a meta-learning neural network. *Nature* **623**, 115–121. (doi:10.1038/s41586-023-06668-3)
65. Kodner J, Payne S, Heinz J. 2023 *Why linguistics will thrive in the 21st century: a reply to Piantadosi*. See <https://lingbuzz.net/lingbuzz/007485>.
66. Chomsky N. 1966 *Cartesian linguistics*. New York: Harper and Row. (doi:10.1017/cbo9780511803116)
67. Bowden BV. 1953 *Faster than thought: a symposium on digital computing machines*. London: Pitman Publishing, Inc.
68. Kim H, Sclar M, Zhou X, Bras R, Kim G, Choi Y, Sap M. 2023 FANToM: a benchmark for stress-testing machine theory of mind in interactions. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. Stroudsburg, PA: Association for Computational Linguistics. (doi:10.18653/v1/2023.emnlp-main.890)
69. Browning J, LeCun Y. 2023 Language, common sense, and the Winograd schema challenge. *Artif. Intell.* **325**, 104031. (doi:10.1016/j.artint.2023.104031)
70. Hamilton K, Nayak A, Božić B, Longo L. 2022 Is neuro-symbolic AI meeting its promises in natural language processing? A structured review. *arXiv* 2202.12205v2. (doi:10.3233/sw-223228)