



## Short communication

## Mind meets machine: Unravelling GPT-4's cognitive psychology

Sifatkaur Dhingra<sup>a,b</sup>, Manmeet Singh<sup>c</sup>, Vaisakh S.B.<sup>c</sup>, Neetiraj Malviya<sup>d</sup>, Sukhpal Singh Gill<sup>e,\*</sup><sup>a</sup> Department of Psychology, Nowrosjee Wadia College, Pune, India<sup>b</sup> Jindal Institute of Behavioural Sciences, O.P. Jindal Global University, Delhi-NCR, India<sup>c</sup> Indian Institute of Tropical Meteorology, Pune, India<sup>d</sup> Defence Institute of Advanced Technology, Pune, India<sup>e</sup> School of Electronic Engineering and Computer Science, Queen Mary University of London, London, UK

## ARTICLE INFO

## Keywords:

Artificial intelligence  
Large language models  
ChatGPT  
Cognitive psychology  
GPT

## ABSTRACT

Cognitive psychology delves on understanding perception, attention, memory, language, problem-solving, decision-making, and reasoning. Large Language Models (LLMs) are emerging as potent tools increasingly capable of performing human-level tasks. The recent development in the form of Generative Pre-trained Transformer 4 (GPT-4) and its demonstrated success in tasks complex to humans exam and complex problems has led to an increased confidence in the LLMs to become perfect instruments of intelligence. Although GPT-4 report has shown performance on some cognitive psychology tasks, a comprehensive assessment of GPT-4, via the existing well-established datasets is required. In this study, we focus on the evaluation of GPT-4's performance on a set of cognitive psychology datasets such as CommonsenseQA, SuperGLUE, MATH and HANS. In doing so, we understand how GPT-4 processes and integrates cognitive psychology with contextual information, providing insight into the underlying cognitive processes that enable its ability to generate the responses. We show that GPT-4 exhibits a high level of accuracy in cognitive psychology tasks relative to the prior state-of-the-art models. Our results strengthen the already available assessments and confidence on GPT-4's cognitive psychology abilities. It has significant potential to revolutionise the field of Artificial Intelligence (AI), by enabling machines to bridge the gap between human and machine reasoning.

## 1. Introduction

Cognitive psychology aims to decipher how humans learn new things, retain knowledge, and recall it when needed. Cognitive psychologists seek to understand how the mind works by conducting studies on people's thoughts and actions and by using other experimental methods like brain imaging and computer modelling [1]. Understanding the human mind and developing our cognitive skills to excel in a variety of areas is the ultimate objective of cognitive psychology [2]. Fig. 1 shows the different fields of cognitive psychology under different sub-fields such as common sense, mathematical reasoning, logical reasoning and others. Language models have come a long way since the first statistical models for modelling language were introduced [3]. With the advent of deep learning and the availability of large amounts of data [4], recent years have seen a rapid evolution of language models that have achieved human-like performance on many language tasks. Large Language Models (LLMs) are a type of Artificial Intelligence (AI) framework that have garnered significant attention in recent years due to their remarkable language processing capabilities

[5–10]. These models are trained on vast amounts of text data and are able to generate coherent, human-like responses to natural language queries. One of the key features of LLMs is their ability to generate novel and creative responses to text-based prompts, which has led to their increasing use in fields such as chatbots, question answering systems, and language translation. An example of the prompts from different datasets used in this study are shown in Fig. 2. The use of self-attention has been a key factor in this success, as it allows for more efficient and accurate modelling of long-range dependencies within the input sequence, resulting in better performance compared to traditional Recurrent Neural Network (RNN)-based models [11]. LLMs have demonstrated impressive performance on a wide range of language tasks, including language modelling, machine translation, sentiment analysis, and text classification. These capabilities have led to the increased use of LLMs in various fields, including language-based customer service, virtual assistants, and creative writing.

One of the key areas measuring intelligence in humans, other species and machines is the cognitive psychology [12]. There are

\* Correspondence to: School of Electronic Engineering and Computer Science, Queen Mary University of London, London, E1 4NS, UK.

E-mail addresses: [sifatkaurd13@gmail.com](mailto:sifatkaurd13@gmail.com) (S. Dhingra), [manmeet.cat@tropmet.res.in](mailto:manmeet.cat@tropmet.res.in) (M. Singh), [vaisakh.sb@tropmet.res.in](mailto:vaisakh.sb@tropmet.res.in) (Vaisakh S.B.), [neetirajmalviya@gmail.com](mailto:neetirajmalviya@gmail.com) (N. Malviya), [s.s.gill@qmul.ac.uk](mailto:s.s.gill@qmul.ac.uk) (S.S. Gill).<https://doi.org/10.1016/j.tbench.2023.100139>

Received 23 August 2023; Received in revised form 28 August 2023; Accepted 28 August 2023

Available online 1 September 2023

2772-4859/© 2023 The Authors. Publishing services by Elsevier B.V. on behalf of KeAi Communications Co. Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).




Fig. 1. Datasets used in the study with the different categories contained in them.

several tasks that are considered to be the benchmarks for testing cognitive psychology. Some of them are text interpretation, computer vision, planning and reasoning [13–15]. For cognitive psychology to work, we rely on a complex and potent social practise: the attribution and assessment of thoughts and actions [16]. The scientific psychology of cognition and behaviour, a relatively recent innovation, focuses primarily on the information-processing mechanisms and activities that characterise human cognitive and behavioural capabilities [17]. Researchers have attempted to create systems that could use natural language to reason about their surroundings [18] or that could use a world model to get a more profound comprehension of spoken language [19]. The report introducing Generative Pre-trained Transformer 4 (GPT-4) [20] has tested the HellaSwag [21] and WinoGrande [22] datasets for cognitive psychology. Although, these tests are relevant, they lack the sophistication required to understand deep heuristics of GPT-4. HellaSwag entails the task of finishing a sentence and WinoGrande involves identifying the correct noun for the pronouns in a sentence, which are quite simple. Other tasks and standardised datasets [23] which test the psychology are needed in order to perform a comprehensive assessment of cognitive psychology for GPT-4. Moreover GPT-4 needs to go through complex reasoning tasks than just predicting the last word of the sentence such as in HellaSwag, to emerge as a model capable of high-level intelligence. [24] note that SuperGLUE [25], CommonsenseQA [26], MATH [27] and HANS [28] are four such datasets that are needed to be tested for a comprehensive cognitive psychology evaluation of AI models. In this study, we evaluate the performance of GPT-4 on the SuperGLUE, CommonsenseQA, MATH and HANS datasets. This is a work in progress and we are performing continuous tests with the other datasets as suggested by [24]. Our study can be used to build up higher-order psychological tests using GPT-4.

The rest of the paper is structured as follows: Section 2 presents the datasets and methodology. Section 3 discusses the experimental results. Section 4 concludes the paper.

## 2. Datasets and methodology

In this study, four datasets have been used to test the cognitive psychology capabilities of GPT-4. The four datasets are CommonsenseQA, MATH, SuperGLUE and HANS. They are described as below:

### 2.1. CommonsenseQA

CommonsenseQA is a dataset composed for testing commonsense reasoning. There are 12,247 questions in the dataset, each with 5 possible answers. Workers using Amazon’s Mechanical Turk were used to build the dataset. The goal of the dataset is to evaluate the commonsense knowledge using CONCEPTNET to generate difficult questions. The language model tested in the CommonsenseQA paper has an accuracy of 55.9% whereas the authors report that human accuracy on the dataset is around 89%.

### 2.2. MATH

The MATH dataset includes almost 12,500 problems from scholastic mathematics contests. Machine learning models take a mathematical problem as input and produce an answer-encoding sequence, such as  $\frac{23}{3}$ . After normalisation, their answers are distinct, therefore MATH may be evaluated using exact match instead of heuristic metrics like BLEU. Problems in seven different areas of mathematics, including geometry, are categorised by complexity from 1 to 5, and diagrams can be expressed in text using the Asymptote language. This allows for a nuanced evaluation of problem-solving skills in mathematics across a wide range of rigour and content. Problems now have comprehensive, detailed, step-by-step answers. To improve learning and make model outputs more interpretable, models can be trained on these to develop their own step-by-step solutions. The MATH dataset presents a significant challenge, with accuracy rates for big language models ranging from 3.0% to 6.9%. Models attain up to 15% accuracy on the least difficulty level and can develop step-by-step answers that are coherent and on-topic even when erroneous, suggesting that they do possess some mathematical knowledge despite their low accuracies. The results of human evaluations on MATH show that it may be difficult for humans as well; a computer science Ph.D. student who does not really like mathematics scored about 40%, while a three-time IMO gold medallist scored 90%.

### 2.3. SuperGLUE

SuperGLUE is an updated version of the GLUE benchmark that includes a more challenging set of language understanding tasks. Using

Dataset/Task	Sample Prompt	GPT4 response
Commonsense reasoning	She caught the first snowflake of Sher life on her tongue, she was naively excited to be spending a Swhat in a region that actually had snow? question concept: punishing choices: SA:cloud B:december C:air D:africa E:winterwhat in a region that actually had snow? question concept: punishing choices:	E:winter
Algebra	Problem: Find $a+b+c$ if the graph of the equation $y=ax^2+bx+c$ is a parabola with vertex $(5,3)$ , vertical axis of symmetry, and contains the point $(2,0)$ . Level: Level 5 Type: Algebra	$a+b+c = -\frac{7}{3}$ .
Prealgebra	Problem: What is the greatest integer $x$ for which $\frac{79}{13} > \frac{x}{13}$ ? Level: Level 3 Type: Prealgebra Solution:	the greatest integer $x$ that satisfies the inequality is $x = 10$ .
Geometry	Problem: $ABCD$ , a rectangle with $AB = 12$ and $BC = 16$ , is the base of pyramid $PP$ , which has a height of $24$ . A plane parallel to $ABCD$ is passed through $PP$ , dividing $PP$ into a frustum $FF$ and a smaller pyramid $PP'$ . Let $XX$ denote the center of the circumsphere of $FF$ , and let $TT$ denote the apex of $PP'$ . If the volume of $PP$ is eight times that of $PP'$ , then the value of $XT$ can be expressed as $\frac{m}{n}$ , where $m$ and $n$ are relatively prime positive integers. Compute the value of $m + n$ .	Finally, we have $m + n = 108 + 10 = \boxed{118}$ .
SuperGLUE	Sentence1 Tanks were developed by Britain and France, and were first used in combat by the British during a battle. Sentence2 Tanks were developed by Britain and France, and were first used in combat by the British during a battle with German forces. Entailment or Non-Entailment	Non entailment

Fig. 2. Examples of sample prompts and the respective responses of GPT4 on CommonsenseQA, MATH and SuperGLUE datasets.

the gap between human and machine performance as a metric, SuperGLUE improves upon the GLUE benchmark by defining a new set of difficult Natural Language Understanding (NLU) problems. About half of the tasks in the SuperGLUE benchmark have fewer than 1k instances, and all but one have fewer than 10k examples, highlighting the importance of different task formats and low-data training data problems. As compared to humans, SuperGLUE scores roughly 20 points worse when using BERT as a baseline in the original study. To get closer to human-level performance on the benchmark, the authors argue that advances in multi-task, transfer, and unsupervised/self-supervised learning approaches are essential.

#### 2.4. HANS

The strength of neural networks lies in their ability to analyse a training set for statistical patterns and then apply those patterns to test instances that come from the same distribution. This advantage is not without its drawbacks, however, as statistical learners, such as traditional neural network designs, tend to rely on simplistic approaches that work for the vast majority of training samples rather than capturing the underlying generalisations. The loss function may not motivate

the model to learn to generalise to increasingly difficult scenarios in the same way a person would if heuristics tend to produce mostly correct results. This problem has been observed in several applications of AI. Contextual heuristics mislead object-recognition neural networks in computer vision, for example; a network that can accurately identify monkeys in a normal situation may mistake a monkey carrying a guitar for a person, since guitars tend to co-occur with people but not monkeys in the training set. Visual question answering systems are prone to the same heuristics. This problem is tackled by HANS (Heuristic Analysis for Natural Language Inference (NLI) Systems), which uses heuristics to determine if a premise sentence entails (i.e., suggests the truth of) a hypothesis sentence. Neural NLI models have been demonstrated to learn shallow heuristics based on the presence of specific words, as has been the case in other fields. As not often appears in the instances of contradiction in normal NLI training sets, a model can categorise all inputs containing the word not as contradiction. HANS prioritises heuristics that are founded on elementary syntactic characteristics. Think about the entailment-focused phrase pair below:

Premise: The judge was paid by the actor.

Hypothesis: The actor paid the judge.




Fig. 3. Accuracy of GPT4 on cognitive psychology tasks.

An NLI system may accurately label this example not by deducing the meanings of these lines but by assuming that the premise involves any hypothesis whose terms all occur in the premise. Importantly, if the model is employing this heuristic, it will incorrectly classify the following as entailed even when it is not.

Premise: The actor was paid by the judge.

Hypothesis: The actor paid the judge.

HANS is intended to detect the presence of such faulty structural heuristics. The authors focus on the lexical overlap, subsequence, and component heuristics. These heuristics are not legitimate inference procedures despite often producing correct labels. Rather than just having reduced overall accuracy, HANS is meant to ensure that models using these heuristics fail on specific subsets of the dataset. Four well-known NLI models, including BERT, are compared and contrasted using the HANS dataset. For this dataset, all models significantly underperformed the chance distribution, with accuracy just exceeding 0% in most situations.

## 2.5. Methodology

We test the four datasets as described above to test the cognitive psychology capabilities of GPT-4. The model is accessed using the ChatGPT-Plus offered by OpenAI. We evaluate these models as shown in the results and discussion section on accuracy metric. Accuracy is a fundamental metric used to evaluate the performance of large language models, especially when applied to psychology datasets. It measures the proportion of predictions that the model gets right out of all the predictions it makes. In the realm of psychology, where understanding human behaviour and cognition is paramount, the accuracy of a model can be crucial. A high accuracy indicates that the model is adept at capturing the nuances of psychological data. When testing large language models on psychology datasets, accuracy can help researchers and practitioners gauge how well the model understands and processes psychological concepts, theories, and patterns. As the field of artificial intelligence evolves, striving for higher accuracy on psychology datasets ensures that models remain relevant and effective in interpreting complex human behaviours and emotions. While accuracy is vital, it is equally important to ensure that the models are tested and trained in an ethical manner, respecting the privacy and sensitivity of psychological data.

## 3. Experimental results

We will first discuss the human and machine skill of the different models traditionally used in the datasets used to test cognitive psychology. As compared to humans, SuperGLUE scores roughly 20 points

worse when using BERT as a baseline in the original study. To get closer to human-level performance on the benchmark, the authors argue that advances in multi-task, transfer, and unsupervised/self-supervised learning approaches are essential. The language model tested in the CommonsenseQA paper has an accuracy of 55.9% whereas the authors report that human accuracy on the dataset is around 89%. The accuracy of humans on HANS dataset ranged from 76%–97% and the authors show that the BERT model performed below 10% on the non-entailment category. The human performance on MATH varied from 40%–90% and GPT-2/GPT-3 showed accuracies below 10%.

Fig. 3 shows that GPT-4 has an accuracy of 83.2% on CommonsenseQA, data, we find that GPT-4 has an accuracy of around 84%, 82% on prealgebra, 35% on geometry, 100% on HANS and 91.2% on SuperGLUE. It is to be noted that the perfect results on HANS data might be because all the examples used are of non-entailment, as the model might be memorising this particular heuristic. The experiments to generate GPT-4 results with mixed data from HANS are ongoing.

### 3.1. Comparison assessing the cognitive abilities of GPT-3: A state-of-the-art model

In a previous study, researchers [29] draw a parallel between the historical case of “Clever Hans”, a horse believed to solve mathematical problems, and the modern interpretation of large-scale machine learning models, particularly GPT-3. The researchers employed systematic investigations and psychological experimentation to assess GPT-3’s decision-making, information search, deliberation, and causal reasoning abilities. The results indicated that GPT-3 could solve certain tasks similarly or even better than humans. However, its performance was inconsistent, especially when minor changes were made to the tasks. The study found that GPT-3 performed well in certain tasks, such as gambles and a multiarmed bandit task, but lacked in areas like directed exploration and causal reasoning. The researchers emphasised the importance of understanding how these models solve tasks and suggested that future models would benefit from active interaction with the world. The study also highlighted the potential of cognitive psychology methods in understanding the behaviour of deep learning models.

## 4. Conclusions

GPT-4, which is a state-of-the-art large language model, is a revolution in the field of psychology since it gives psychologists unprecedented resources to use in their studies and work. This sophisticated AI model offers psychologists and psychiatrists to learn more about the human mind and come up with novel treatment theories and approaches. It provides an avenue for improved efficacy of psychological therapies and allowing professionals to spend more time with clients, leading to deeper and more fruitful therapeutic bonds. The potential applications of GPT-4 can only be realised if the model is thoroughly tested on basic tests of reasoning and cognition. Cognitive psychology enables the humans to perform various activities [30] in their personal and professional lives. We show that the performance of GPT-4 greatly surpasses the language model used in the original studies from where the different datasets are sourced, thus it can make a tool of day-to-day utility for psychologists. This development can lead to cascading benefits in addressing the mental health challenges faced by today’s society.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

We would like to thank the Editor-in-Chief, area editor and anonymous reviewers for their valuable comments, useful suggestions to improve the quality of the paper.

## References

- [1] R. Núñez, M. Allen, R. Gao, C. Miller Rigoli, J. Relaford-Doyle, A. Semenuks, What happened to cognitive science? *Nat. Hum. Behav.* 3 (8) (2019) 782–791.
- [2] L.W. Barsalou, *Cognitive Psychology: an Overview for Cognitive Scientists*, Psychology Press, 2014.
- [3] M.C. Frank, Baby steps in evaluating the capacities of large language models, *Nat. Rev. Psychol.* (2023) 1–2.
- [4] S.S. Gill, M. Xu, C. Ottaviani, P. Patros, R. Bahsoon, A. Shaghghi, M. Golec, V. Stankovski, H. Wu, A. Abraham, et al., AI for next generation computing: Emerging trends and future directions, *Internet Things* 19 (2022) 100514.
- [5] S. Harrer, Attention is not all you need: the complicated case of ethically using large language models in healthcare and medicine, *eBioMedicine* 90 (2023) 104512.
- [6] T. Brown, B. Mann, N. Ryder, M. Subbiah, J.D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *Adv. Neural Inf. Process. Syst.* 33 (2020) 1877–1901.
- [7] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, 2018, arXiv preprint arXiv:1810.04805.
- [8] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al., Language models are unsupervised multitask learners, *OpenAI Blog* 1 (8) (2019) 9.
- [9] M. Brundage, S. Avin, J. Clark, H. Toner, P. Eckersley, B. Garfinkel, A. Dafoe, P. Scharre, T. Zeitsoff, B. Filar, et al., The malicious use of artificial intelligence: Forecasting, prevention, and mitigation, 2018, arXiv preprint arXiv:1802.07228.
- [10] V. Liévin, C.E. Hother, O. Winther, Can large language models reason about medical questions? 2022, arXiv preprint arXiv:2207.08143.
- [11] F. Tang, W. Gao, SNNBench: End-to-end AI-oriented spiking neural network benchmarking, *BenchCouncil Trans. Benchmarks Stand. Eval.* 3 (1) (2023) 100108.
- [12] J. Zhao, M. Wu, L. Zhou, X. Wang, J. Jia, Cognitive psychology-based artificial intelligence review, *Front. Neurosci.* 16 (2022) 1024316.
- [13] H. Singh, et al., Metaheuristics for scheduling of heterogeneous tasks in cloud computing environments: Analysis, performance evaluation, and future directions, *Simul. Model. Pract. Theory* 111 (2021) 102353.
- [14] K. Bansal, et al., DeepBus: Machine learning based real time pothole detection system for smart transportation using IoT, *Internet Technol. Lett.* 3 (3) (2020) e156.
- [15] D. Chowdhury, et al., CoviDetector: A transfer learning-based semi supervised approach to detect Covid-19 using CXR images, *BenchCouncil Trans. Benchmarks Stand. Eval.* 3 (2) (2023) 100119.
- [16] A. Madaan, A. Yazdanbakhsh, Text and patterns: For effective chain of thought, it takes two to tango, 2022, arXiv preprint arXiv:2209.07686.
- [17] S. Singh, I. Chana, M. Singh, The journey of QoS-aware autonomic cloud computing, *IT Prof.* 19 (2) (2017) 42–49.
- [18] J. McCarthy, A basis for a mathematical theory of computation, in: *Studies in Logic and the Foundations of Mathematics*, Vol. 26, Elsevier, 1959, pp. 33–70.
- [19] T. Winograd, Understanding natural language, *Cogn. Psychol.* 3 (1) (1972) 1–191.
- [20] OpenAI, GPT-4 technical report, 2023.
- [21] R. Zellers, A. Holtzman, Y. Bisk, A. Farhadi, Y. Choi, HellaSwag: Can a machine really finish your sentence? 2019, arXiv preprint arXiv:1905.07830.
- [22] K. Sakaguchi, R.L. Bras, C. Bhagavatula, Y. Choi, Winogrande: An adversarial winograd schema challenge at scale, *Commun. ACM* 64 (9) (2021) 99–106.
- [23] Y. Li, J.L. McClelland, Systematic generalization and emergent structures in transformers trained on structured tasks, 2022, arXiv preprint arXiv:2210.00400.
- [24] R. Shiffrin, M. Mitchell, Probing the psychology of AI models, *Proc. Natl. Acad. Sci.* 120 (10) (2023) e2300963120.
- [25] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, S. Bowman, Superglue: A stickier benchmark for general-purpose language understanding systems, *Adv. Neural Inf. Process. Syst.* 32 (2019).
- [26] A. Talmor, J. Herzig, N. Lourie, J. Berant, Commonsenseqa: A question answering challenge targeting commonsense knowledge, 2018, arXiv preprint arXiv:1811.00937.
- [27] D. Hendrycks, C. Burns, S. Kadavath, A. Arora, S. Basart, E. Tang, D. Song, J. Steinhardt, Measuring mathematical problem solving with the math dataset, 2021, arXiv preprint arXiv:2103.03874.
- [28] R.T. McCoy, E. Pavlick, T. Linzen, Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference, 2019, arXiv preprint arXiv:1902.01007.
- [29] M. Binz, E. Schulz, Using cognitive psychology to understand GPT-3, *Proc. Natl. Acad. Sci.* 120 (6) (2023) e2218523120.
- [30] G. Aher, R.I. Arriaga, A.T. Kalai, Using large language models to simulate multiple humans, 2022, arXiv preprint arXiv:2208.10264.