Large language models surpass human experts in predicting neuroscience results

Xiaoliang Luo[1, †, *], Akilles Rechardt[1, †], Guangzhi Sun[2, †],
Kevin K. Nejad[3, 6, †], Felipe Yáñez[4, †], Bati Yilmaz[5, †], Kangjoo Lee[8],
Alexandra O. Cohen[9], Valentina Borghesani[10], Anton Pashkov[11, 12, 13],
Daniele Marinazzo[14], Jonathan Nicholas[15], Alessandro Salatiello[16],
Ilia Sucholutsky[17], Pasquale Minervini[18], Sepehr Razavi[19], Roberta Rocca[20],
Elkhan Yusifov[21], Tereza Okalova[22], Nianlong Gu[23], Martin Ferianc[24],
Mikail Khona[25], Kaustubh R. Patil [26, 27], Pui-Shee Lee[28, 29], Rui Mata[30],
Nicholas E. Myers[31], Jennifer K Bizley[32], Sebastian Musslick[33, 34],
Isil Poyraz Bilgin[35], Guiomar Niso[36], Justin M. Ales[37], Michael Gaebler[38],
N Apurva Ratan Murty[39], Chloe M. Hall[40, 41], Jessica Dafflon[42, 43],
Sherry Dongqi Bao[44], and Bradley C. Love[1, 7, †]

[1]Department of Experimental Psychology, University College London, London, United Kingdom,
[2]Department of Engineering, University of Cambridge, Cambridge, United Kingdom, [3]Department of
Physiology, Anatomy & Genetics, University of Oxford, Oxford, United Kingdom, [4]Max Planck
Institute for Neurobiology of Behavior – caesar, Bonn, Germany, [5]National Magnetic Resonance
Research Center (UMRAM), Bilkent University, Ankara, Turkey, [6]Department of Computer Science,
University of Bristol, Bristol, United Kingdom, [7]The Alan Turing Institute, London, United Kingdom,
[8]Department of Psychiatry, Yale University, New Haven, United States, [9]Psychology, Emory
University, Atlanta, United States, [10]Faculty of Psychology and Educational Sciences, Université de
Genève, Genève, Switzerland, [11]Neurosurgery, Novosibirsk State Medical University, Novosibirsk,
Russian Federation, [12]Federal Center of Neurosurgery, FSBI, Novosibirsk, Russia, [13] Department of
Data Collection and Processing Systems, Novosibirsk State Technical University, Novosibirsk, Russia,
[14]Department of Data Analysis, Ghent University, Ghent, Belgium, [15]Psychology, New York
University, New York, United States, [16]Department of Cognitive Neurology, University of Tübingen,
Tübingen, Germany, [17]Computer Science, Princeton University, Princeton, United States, [18]ILCC,
University of Edinburgh, Edinburgh, United Kingdom, [19]Philosophy, Psychology, and Language
Sciences , The University of Edinburgh , Edinburgh , United Kingdom, [20]Department of Culture,
Cognition and Computation, Aarhus University, Aarhus, Denmark, [21]Department of Molecular Life
Sciences, University of Zurich, Zurich, Switzerland, [22]Bioengineering, University of Pennsylvania ,
Philadelphia , United States, [23]Linguistic Research Infrastructure, University of Zurich, Zurich,

Switzerland, [24]Electronic and Electrical Engineering, University College London, London, United Kingdom, [25]Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States, [26]Institute of Neuroscience and Medicine, INM-7: Brain and Behaviour, Research Centre Jülich, Jülich, Germany, [27]Medical Faculty, Institute of Systems Neuroscience, Heinrich Heine University Düsseldorf, Düsseldorf, Germany, [28]Graduate School of Systemic Neurosciences, Ludwig-Maximilians-University Munich, Planegg-Martinsried, Germany, [29]Institute of Neuronal Cell Biology, Technical University of Munich, Munich, Germany, [30]Faculty of Psychology, University of Basel, Basel, Switzerland, [31]School of Psychology, University of Nottingham, Nottingham, United Kingdom, [32]Ear Institute, University College London, London, United Kingdom, [33]Institute of Cognitive Science, University of Osnabrück, Osnabrück, Germany, [34]Department of Cognitive, Linguistic, & Psychological Sciences, Brown University, Providence, United States, [35]Département de psychologie, Centre de recherche de l'Institut universitaire de gériatrie de Montréal, Montreal, Canada, [36]Instituto Cajal, CSIC, Madrid, Spain, [37]School of Psychology & Neuroscience, University of St Andrews, St Andrews, United Kingdom, [38]Neurology, Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany, [39]Department of Psychology, Georgia Institute of Technology, Atlanta, United States, [40]Institute of Physiology, University Medical Center of the Johannes Gutenberg University, Mainz, Germany, [41]Institute for Quantitative and Computational Biosciences (IQCB), Johannes Gutenberg University, Mainz, Germany, [42]Data Science and Sharing Team, Functional Magnetic Resonance Imaging Facility, National Institute of Mental Health, Bethesda, United States, [43]Machine Learning Team, Functional Magnetic Resonance Imaging Facility, National Institute of Mental Health, Bethesda, United States, [44]Zurich Center for Neuroeconomics, Department of Economics, University of Zurich, Zurich, Switzerland, †Major contributions, the names of the remaining authors are listed in random order, see contributions breakdown in Table S.4, *Corresponding author: xiao.luo.17@ucl.ac.uk

Scientific discoveries often hinge on synthesizing decades of research, a task that potentially outstrips human information processing capacities. Large language models (LLMs) offer a solution. LLMs trained on the vast scientific literature could potentially integrate noisy yet interrelated findings to forecast novel results better than human experts. To evaluate this possibility, we created BrainBench, a forward-looking benchmark for predicting neuroscience results. We find that LLMs surpass experts in predicting experimental outcomes. BrainGPT, an LLM we tuned on the neuroscience literature, performed better yet. Like human experts, when LLMs were confident in their predictions, they were more likely to be correct, which presages a future where humans and LLMs team together to make discoveries. Our approach is not neuroscience-specific and is transferable to other knowledge-intensive endeavors.

Introduction

Keeping up with the exponentially increasing (*1*) scientific literature is a superhuman challenge. Potentially disruptive findings go unnoticed in the deluge of articles (*2*). Processing and integrating the myriad of relevant findings may already surpass humans' abilities. One path forward is a partnership between human scientists and machines. This partnership could take several forms, including specialist solutions that address specific challenges, such as in protein folding (*3*), drug discovery (*4*), and materials science (*5*). Alternatively, general models of the scientific literature could help guide human scientists' predictions and study designs. We consider this possibility.

It is an open question whether large language models (LLMs), trained on general text and scientific articles, can predict the outcomes of experiments. If LLMs' predictions surpassed

human experts, the practice of science and the pace of discovery would radically change. We consider this question for neuroscience, which is a large and interdisciplinary field. Prediction in neuroscience should be challenging for human experts for several reasons: (i) there are often many thousands of relevant scientific articles, (ii) an individual study can be noisy, unreliable, and may not replicate, (iii) neuroscience is a multi-level endeavor (*6*), spanning behavior and molecular mechanisms, (iv) and the analysis methods are diverse and can be complex (*7*), (v) as are the methods used, which include different brain imaging techniques, lesion studies, gene modification, pharmacological interventions, and so forth.

Can LLMs meet these challenges? In other domains, LLMs have performed impressively. Upon its release, OpenAI's ChatGPT (*8*) captured the public's imagination with its abilities. Most LLMs are based on the transformer architecture (*9*). These models contain billions and sometimes trillions of weights (*10*), which are tuned during training in a self-supervised manner to predict the next token, such as the next word in a text passage.

LLMs have displayed remarkable capabilities, including passing professional exams, reasoning (though not without limitations), translation, solving mathematics problems, and even writing computer code (*11, 12*). By constructing a statistical model during their training to predict the next token, whether that token is a word, pixel, or protein sequence (*13*), LLMs uncover the underlying patterns or structure of a domain. This generative model captures patterns in the training data, including subtle and imperfect ones. How LLMs learn and generalize to novel situations can be likened to how expert scientists detect patterns in their field after years of reading papers, attending conferences, and analyzing data. From these experiences, human experts build intuitions that enable them to predict future outcomes based on proposed study designs. However, unlike human scientists, LLMs are virtually unconstrained in how much of the scientific literature they can process during training.

How can we formally evaluate the predictive abilities of LLMs in neuroscience? With the

rise of LLMs, there has been a surge in evaluation benchmarks, many of which focus on assessing LLMs' capabilities in scientific domains. Most benchmarks evaluate core knowledge retrieval and reasoning abilities, which are typically *backward-looking* (Fig. 1). Backward-looking benchmarks include MMLU (*14*), PubMedQA (*15*), and MedMCQA (*16*). These benchmarks are structured in a question-and-answer format, where models must demonstrate extensive world knowledge, retrieve relevant information based on the context of the question, and answer correctly. However, none of these benchmarks are suitable for evaluating the ability of models to predict novel outcomes, which is inherently *forward-looking* (Fig. 1).

To address this need, we developed BrainBench to test LLMs' ability to predict neuroscience findings (Fig. 2). LLMs have been trained extensively on the scientific literature, including neuroscience. BrainBench evaluates whether LLMs have seized on the fundamental patterning of methods and results that underlie the structure of neuroscience. Can LLMs outperform human experts on this forward-looking benchmark? In particular, BrainBench evaluates how well the test-taker can predict neuroscience results from methods by presenting two versions of an abstract from a recent journal article. The test-taker's task is to predict the study's outcome, choosing between the original and an altered version. The altered abstract significantly changes the study's outcome (i.e., results) while maintaining overall coherence.

To appreciate how BrainBench qualitatively differs from existing benchmarks, consider a perceived limitation of LLMs, namely their tendency to generate erroneous information, a phenomenon commonly referred to as "hallucination" by LLM researchers. Unlike knowledge graphs that store verified facts, LLMs may not be trustworthy for backward-looking tasks such as summarizing research papers or providing accurate citations (*17*). However, for forward-looking tasks, such as predicting results from a novel experiment, we view this tendency to mix and integrate information from large and noisy datasets as a virtue. What is a hallucination in a backward-looking task is a generalization or prediction in a forward-looking task (e.g., Brain-
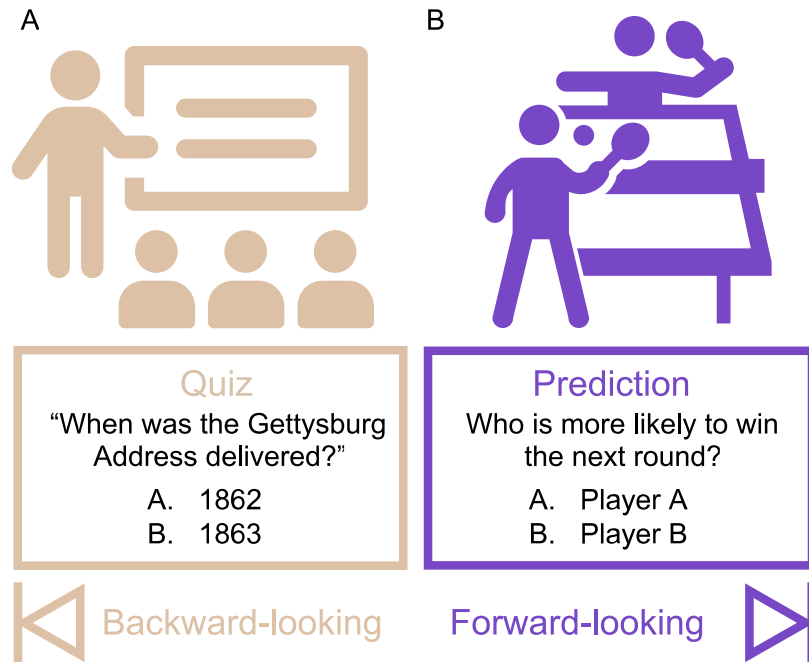
Figure 1: **Backward-looking and Forward-looking evaluations.** (A) Backward-looking benchmarks involve recalling factual information. For example, in the left panel, a student retrieves a fact about the Gettysburg Address that they learned during a history class. Existing benchmarks in scientific domains are in essence backward-looking as they emphasize retrieving accepted facts for question answering and reasoning tasks. (B) Forward-looking benchmarks involve predicting novel outcomes based on past data. Two forms of uncertainty, aleatoric (due to intrinsic randomness) and epistemic (due to lack of knowledge), may be present. For example, in the right panel, a table tennis fan predicts which player will win the next set based on their knowledge of the players, how they have played so far today, and so forth. Inherent random factors, such as a breeze affecting the ball's flight, will also be present.
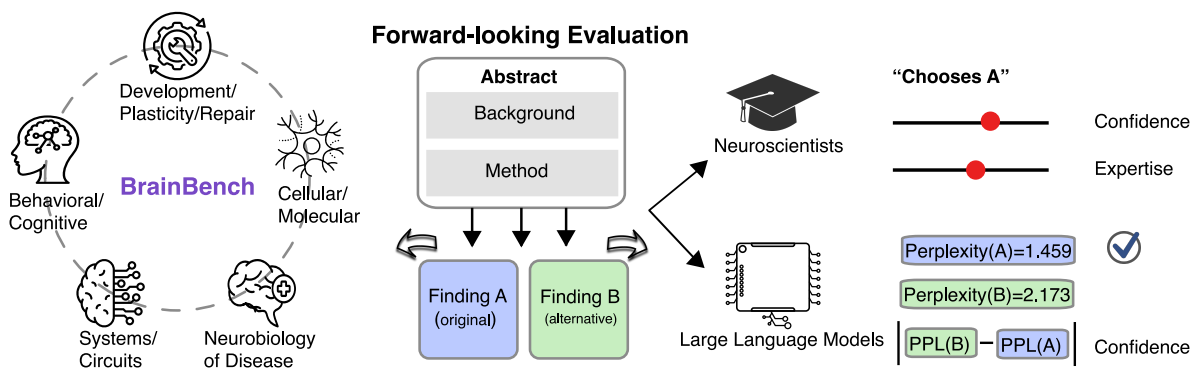
Figure 2: **BrainBench is a forward-looking benchmark for neuroscience.** BrainBench evaluates test-takers' ability to predict neuroscience results. BrainBench's test cases were sourced from recent *Journal of Neuroscience* abstracts across five neuroscience domains: Behavioral/Cognitive, Systems/Circuits, Neurobiology of Disease, Cellular/Molecular, and Developmental/Plasticity/Repair. Test-takers chose between the original abstract and one altered to significantly change the result while maintaining coherency. Human experts and Language Models (LLMs) were tasked with selecting the correct (i.e., original) version from the two options. Human experts made choices, and provided confidence and expertise ratings in an online study. LLMs were scored as choosing the abstract with the lower perplexity (i.e., the text passage that was less surprising to the model) and their confidence was proportional to the difference in perplexity between the two options.

Bench). BrainBench provides a way to quantify this forward-looking ability and compare to human experts. To foreshadow our results, LLMs surpassed human experts on BrainBench by a substantial margin and this margin increased when we provided additional training in neuroscience to an LLM, which we refer to as "BrainGPT".

Results

General-purpose LLMs best neuroscientists on BrainBench

On each benchmark trial (see Fig. 2), both the LLMs and human experts were tasked with selecting which of two versions of an abstract was correct (i.e., the original version). Human neuroscience experts were screened for their expertise and engagement (see Methods) with 171 out of 202 participants passing all checks and included in our analyses.

Every LLM outperformed human experts on BrainBench with LLMs averaging 81.4% accuracy and human experts averaging 63.4% (Fig. 3A). When restricting human responses to those in the top 20% of self-reported expertise for that test item, accuracy rose to 66.2%, still below the level of LLMs.

Smaller models such as Llama2-7B and Mistral-7B with 7 billion parameters, performed comparably to larger models (Fig. 3A). Chat or instruction-optimized models performed worse than their base model counterparts ($t(5) = 5.4$, $p < .01$). We suspect that aligning LLMs to engage in natural language conversations hinders their scientific inference abilities (see Discussion).

The previous analyses involved benchmark items created by co-authors who are neuroscience experts (see Methods). We conducted the same analyses using test cases generated by a LLM, namely GPT-4 (see Methods), and observed similar results (see Supplementary materials).
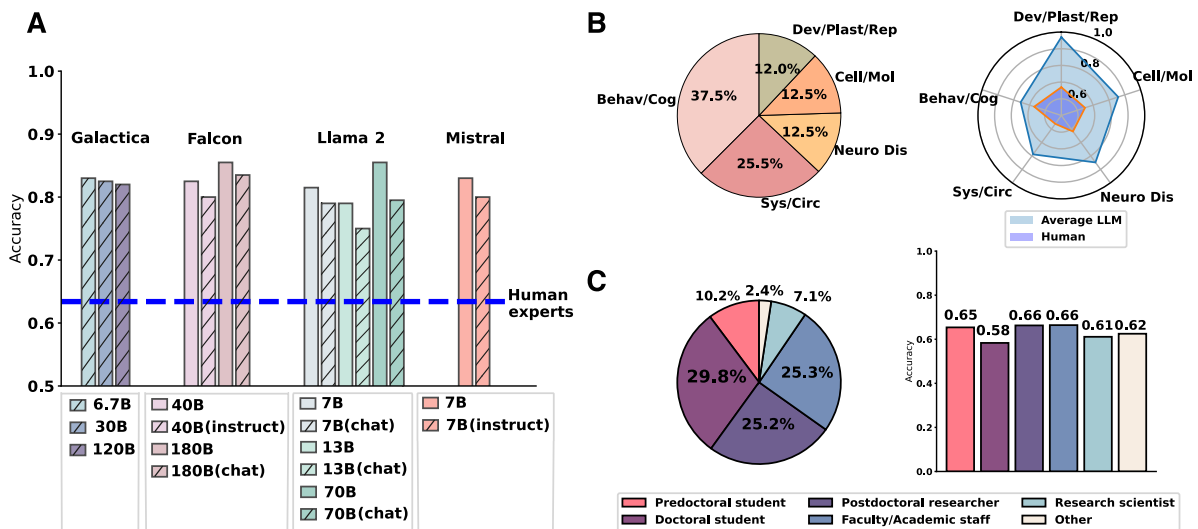
8

Figure 3: **Performance of human experts and large language models on BrainBench.** (A) LLMs outperformed human experts on BrainBench. Smaller models are on par with larger models. Base versions of models outperformed chat and instruct versions, which were tuned to be conversational with humans. (B) The distribution of test cases across neuroscience subfields roughly mirrors the distribution of articles in the *Journal of Neuroscience* with Behavior/Cognitive overrepresented. The average performance of 15 LLMs and human experts is shown. LLMs outperformed human experts in every subfield (see Fig. S.3 for full results). Error bars represent the standard deviation of accuracy around the mean. (C) Majority of the participants were doctoral students, postdoctoral researchers and faculty/academic staff.

Performance breakdown by subfields and by participant types BrainBench encompasses test cases from five distinct neuroscience domains: Behavioral/Cognitive, Cellular/Molecular, Systems/Circuits, Neurobiology of Disease, and Development/Plasticity/Repair. Some domains, particularly Behavioral/Cognitive, are overrepresented both in BrainBench (Fig. 3B) and the *Journal of Neuroscience* from which we drew our test cases (see Methods).

On average, LLMs performed better than human experts in every subfield (Fig. 3B), as did each individual LLM (Fig. S.3). Most human experts were doctoral students, postdoctoral researchers, or faculty/academic staff (see Fig. 3C). Please refer to Supplementary materials for more detailed demographic information including years of experience in neuroscience research about the human experts.

Do judgments from LLMs and human experts align? We considered whether human experts and LLMs found the same benchmark items difficult. For humans, we calculated the mean accuracy for each of the 200 test cases. For LLMs, we calculated the signed differences in perplexity between incorrect and correct abstracts for each test case. Perplexity measures how surprising a text passage is to an LLM. Using these measures (Fig. S.4), the mean Spearman correlation between an LLM and human experts was 0.15 (± 0.03) whereas the mean Spearman correlation between LLMs was 0.75 (± 0.08).

LLMs can integrate information across context To better understand the basis for the remarkable performance of LLMs (see Fig. S.2 for results), we investigated whether their performance was achieved by integrating information throughout the abstract (including the method used) or by solely relying on the local context in the results passages that differed between the original and altered abstract (Fig. 2)

We re-evaluated the LLMs on individual sentences containing only the altered results passage (i.e., local context only). LLMs performed much worse when restricted to this local context

(Fig. S.2), which provides strong evidence that LLMs are integrating information across the abstract, including information on background and methods. LLM's superior performance relative to human experts appears to arise from integrating information across the abstract.

LLM performance is not driven by data memorization   When LLMs perform well on a benchmark, one general concern is that the benchmark itself was part of the training set, allowing the LLM to memorize the correct answers. To address this concern, we used a commonly applied measure, zlib-perplexity ratio, for evaluating whether LLMs have memorized passages (*18*). This ratio gauges the difference between a data-agnostic compression rate of text and data-specific perplexity computed by an LLM (see Methods). Passages that are hard to compress but have low perplexity are indicative of memorization.

We found no indication that BrainBench was memorized by LLMs (Fig. S.5). For comparison, we calculated the zlib-perplexity ratio for a passage that we suspected would be memorized by LLMs, namely the Gettysburg Address. The Gettysburg Address should appear multiple times in an LLM's training set and indeed it showed signs of memorization (Fig. S.5). Interestingly, for some LLMs, we know exactly what they were trained on (see Table S.2). For these models, the distribution of zlib-perplexity ratios heavily overlapped for items that we knew were in the training set and for items, including BrainBench, that we knew were not in the training set. We believe this overlap indicates that scientific articles, which are unlikely to repeat in training sets, are stored in LLMs as general patterns (akin to a human schema), which support performance on forward-looking tasks (e.g., BrainBench) that require generalization.

LLMs and human experts are calibrated

To assess whether LLMs' predictions are calibrated, we examined how well their confidence tracked their accuracy, a crucial characteristic for a trustworthy prediction system. We estimated

LLMs' confidence using the ranked absolute difference in perplexities between two abstracts (Fig. 2; see Methods) and found, like human experts, all LLMs exhibited a positive correlation between accuracy and confidence. When LLMs are confident in their decisions, they are more likely to be correct (Fig. 4).

Augmenting LLMs with neuroscience knowledge to create BrainGPT

Pretrained LLMs can provide a foundation for further training in neuroscience with the aim of improving performance, as assessed by BrainBench. We used Low-Rank Adaptation (LoRA) (*19*) to augment a pretrained LLM, Llama-2-7B-chat, with additional neuroscience knowledge.

LoRA is a parameter-efficient fine-tuning technique that inserts low-rank adapter matrices into LLM transformer blocks (Fig. S.6) and trains only these LoRA weights to update the model's behavior. In our case, we fine-tuned Llama-2-7B chat using over 1.3 billion tokens from neuroscience publications spanning 100 journals between 2002 and 2022 (see Methods), which significantly improved performance.

The difference in perplexity between incorrect and correct options significantly increased under LoRA ($t(199) = -6.3, p < .001$), which resulted in a 3% increase in BrainBench performance (Fig. 5A). LoRA tuning dramatically shifted ($t(199) = 33.3, p < .001$) the perplexity of correct responses (Fig. 5B), which is indicative of the LLM specializing for neuroscience material. LoRA introduced fewer than 18 million new weights, which is $.26\%$ of the total number of weights in Llama-2-7B chat. These results indicate that BrainGPT models can efficiently be derived by extending existing LLMs.

Discussion

We considered whether large language models (LLMs) can forecast the outcome of neuroscience experiments. By training on the vast scientific literature, we hoped LLMs could build a
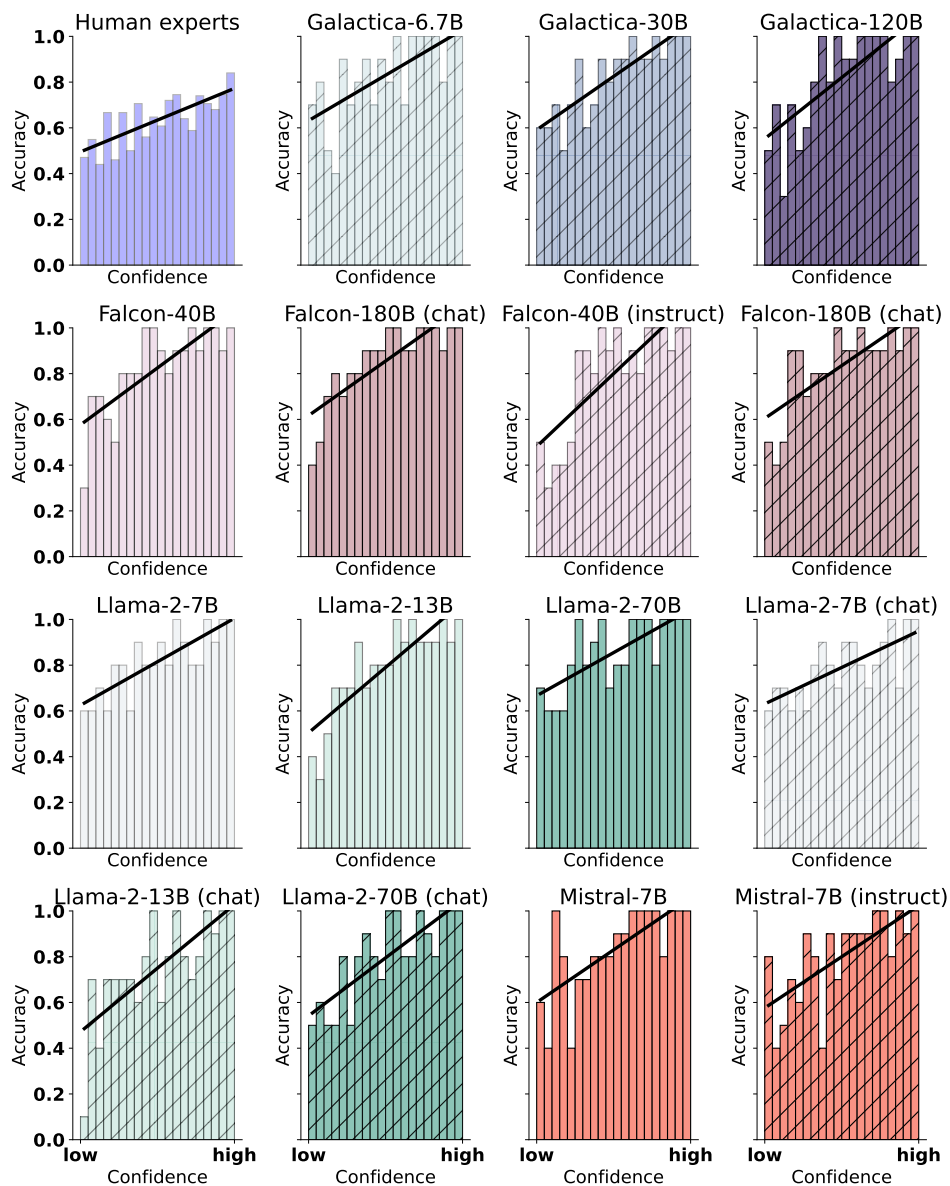
Figure 4: **Accuracy and confidence are calibrated for human experts and large language models (LLMs).** When human experts and LLMs are confident in their BrainBench judgments, they are more likely to be correct. Confidence ratings were sorted and placed in equally-sized bins with the mean accuracy for items in that bin plotted. The positive slope of the black regression lines for human experts and all LLMs indicates that confidence is well calibrated (i.e., higher confidence corresponds to higher accuracy). Calibration is beneficial for human-machine teams.
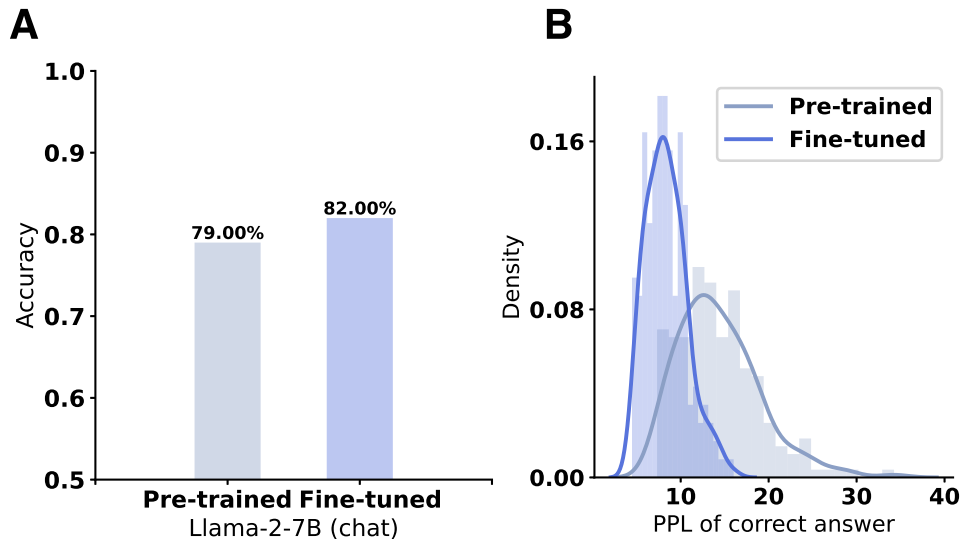
Figure 5: **Fine-tuning a pretrained large language model (LLM) on neuroscience knowledge.** Llama-2-7B (chat) was fine-tuned using LoRA on neuroscience articles from 2002-2022 (a total of 1.3 billion tokens). (A) The fine-tuned model improved by 3% on BrainBench. (B) The fine-tuning process substantially shifted the perplexity distribution of correct responses, indicative of the LLM specializing in neuroscience.

generative model that captured the patterns underlying neuroscience. To evaluate this possibility, we constructed a new forward-looking (Fig. 2) benchmark, BrainBench.

BrainBench assesses a test taker's ability to select which of two versions of a neuroscience abstract contains the actual results of the study (see Fig. 2). We found that LLMs outperform human experts on BrainBench by a considerable margin (see Fig. 3A) across all neuroscience subfields we consider (Fig. 3B). Moreover, the LLMs knew when their predictions were likely to be right or wrong (Fig. 4). LLMs' superior performance arose from their ability to integrate information throughout the abstract, such as text pertaining to the method and study design. When access to such information was removed, LLM performance drastically declined (Fig. S.2).

We found no indication that LLMs had been exposed to and memorized BrainBench items during their training. Instead, our analyses suggested that LLMs discovered the fundamental

patterns that underlie neuroscience studies, which enabled LLMs to predict the outcomes of studies that were novel to them. These conclusions were supported by a widely employed technique (*18*) to determine text membership within an LLMs' training set (see Fig. S.5). The Galactica (*20*) LLMs were particularly illuminating because the contents of their training set are publicly known and did not include BrainBench items. Interestingly, there was no indication of memorization in Galactica for scientific articles that were in its training set, consistent with the notion that LLMs learn broad patterns underlying scientific fields. While passages that frequently repeat in the training set, such as the Gettysburg Address, may be memorized (see Fig. S.5), scientific articles that occur infrequently (most likely once) in the training set appear to support LLM's forward-looking predictive abilities.

LLM's impressive forward-looking capabilities suggest a future in which LLMs help scientists make discoveries. To be effective, LLMs need to stay abreast of the rapidly expanding literature. We found that LLMs could efficiently be augmented with neuroscience knowledge using LoRA (*19*), boosting performance on BrainBench (Fig. 5). LoRA provides a way to create BrainGPT models by re-orienting general-purpose LLMs for use in neuroscience. One can easily imagine a future in which BrainGPT is near continuously updated with new knowledge using LoRA, along with complementary approaches such as Retrieval Augmented Generation (RAG; (*17*)). RAG could be used to query a database of relevant and up to date scientific articles for the task at hand.

In addition to keeping LLMs up to date, benchmarks should routinely be refreshed and expanded to address current needs. One challenge is that creating forward-looking benchmarks, such as BrainBench, is labor intensive and requires human expertise. To address this potential bottleneck, we created and evaluated 100 test cases using GPT-4 through a largely automated process (see Methods). Although there is room for improvement, these items were close in quality to the human-created ones with 8 of the 100 items being word-for-word matches with

the human-created versions. These efforts should pave the way for the rapid creation of other forward-looking benchmarks in neuroscience, as well as benchmarks for other knowledge intensive fields. We believe high-quality forward-looking benchmarks will be critical to developing LLMs as tools for scientific discovery.

For LLMs to be trustworthy and effective teammates, they need to convey the certainty of their predictions to human scientists. Fortunately, we found that LLMs' confidence is well calibrated. When LLMs were confident in their predictions, they were more likely to be correct (Fig. 4). A second ingredient for effective teams is being diverse or complementary. LLMs have potential here as well as the items they found difficult did not highly correlate with those human experts found difficult (Fig. S.4). These two ingredients, being well calibrated and complementary, allow systems that combine human and machine judgments to outperform either alone (*21*).

All our results, including those for calibrated confidence, were only possible because we had access to LLM weights to calculate the perplexity of passages (see Fig. 2). Our approach diverged from the popular approach of prompting models for responses through natural language (i.e., chat). Prompting in natural language may yield less reliable judgments and degrade model competency compared to using model probability scores or training separate classifiers directly from internal representations (*22–25*). These observations underscore the importance of working with models that are as open as possible, ideally making both the weights and training set publicly available. Accordingly, we make BrainGPT available on the Huggingface platform https://huggingface.co/BrainGPT.

Beyond serving as a tool for neuroscientists, BrainGPT can help reveal the structure of the field. In particular, we can vary BrainGPT's training set and observe the effect on BrainBench. For example, what is the effect of including training data from related fields like psychology? In terms of supporting prediction, we can quantify how interrelated fields are. Does it help

16

to weight articles in the training set by their recency, citations, or impact factor? In addition to these training manipulations, we can vary how testing is conducted. For example, would step-by-step thinking via chain-of-thought reasoning (*26*) benefit BrainGPT? If prediction in neuroscience is akin to a deductive reasoning process, then it should. If instead, as we suspect, prediction in neuroscience is a function of many noisy intertwined signals across subfields, then chain-of-thought reasoning will not help. BrainGPT and BrainBench can help answer these meta-science questions.

We foresee a future in which LLMs serve as forward-looking generative models of the scientific literature. LLMs can be part of larger systems that help researchers determine the best experiment to run next. LLMs' predictions are informed by a vast scientific literature that no human could read in their lifetime. As LLMs improve, so should their ability to provide accurate predictions. In this contribution, we focused on neuroscience but our aims are broader – we hope to provide a template for any knowledge intensive field. None of the approaches we adopted are neuroscience-specific. We hope to democratize the use of LLMs in science and increase reproducibility by highlighting the use of relatively small models that can be run locally and whose weights are accessible, which contrasts with commercial products. Finally, while LLMs appear poised to supplant humans at prediction, we foresee a role for human experts in providing the accompanying scientific explanations. Prediction is very important, but not everything.

References

1. L. Bornmann, R. Mutz, *Journal of the Association for Information Science and Technology* **66**, 2215 (2015).

2. J. S. G. Chu, J. A. Evans, *Proceedings of the National Academy of Sciences* **118**,

e2021636118 (2021).

3. K. Tunyasuvunakool, *et al.*, *Nature* **596**, 590 (2021).

4. A. Zhavoronkov, *et al.*, *Nature Biotechnology* **37**, 1038 (2019).

5. V. Tshitoyan, *et al.*, *Nature* **571**, 95 (2019).

6. R. M. Mok, B. C. Love, *Science Advances* **9**, eade6903 (2023).

7. R. Botvinik-Nezer, *et al.*, *Nature* **582**, 84 (2020). Publisher: Nature Publishing Group.

8. Y. Liu, *et al.* (2023). Publisher: arXiv Version Number: 3.

9. A. Vaswani, *et al.* (2017). Publisher: arXiv Version Number: 5.

10. W. Fedus, B. Zoph, N. Shazeer (2021). Publisher: arXiv Version Number: 3.

11. A. Srivastava, *et al.* (2022). Publisher: arXiv Version Number: 3.

12. S. Gunasekar, *et al.* (2023). Publisher: arXiv Version Number: 1.

13. R. Strack, *Nature Methods* **20**, 1868 (2023).

14. D. Hendrycks, *et al.*, Measuring Massive Multitask Language Understanding (2021). ArXiv:2009.03300 [cs].

15. Q. Jin, B. Dhingra, Z. Liu, W. W. Cohen, X. Lu, PubMedQA: A Dataset for Biomedical Research Question Answering (2019). ArXiv:1909.06146 [cs, q-bio].

16. A. Pal, L. K. Umapathi, M. Sankarasubbu, MedMCQA : A Large-scale Multi-Subject Multi-Choice Dataset for Medical domain Question Answering (2022). ArXiv:2203.14371 [cs].

17. P. Lewis, *et al.*, Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks (2021). ArXiv:2005.11401 [cs].

18. N. Carlini, *et al.*, Extracting Training Data from Large Language Models (2021). ArXiv:2012.07805 [cs].

19. E. J. Hu, *et al.*, LoRA: Low-Rank Adaptation of Large Language Models (2021). ArXiv:2106.09685 [cs].

20. R. Taylor, *et al.*, Galactica: A Large Language Model for Science (2022). ArXiv:2211.09085 [cs, stat].

21. M. Steyvers, H. Tejeda, G. Kerrigan, P. Smyth, *Proceedings of the National Academy of Sciences* **119**, e2111547119 (2022).

22. L. Zheng, *et al.*, Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena (2023). ArXiv:2306.05685 [cs].

23. J. Hu, R. Levy, Prompting is not a substitute for probability measurements in large language models (2023). ArXiv:2305.13264 [cs].

24. A. Azaria, T. Mitchell, The Internal State of an LLM Knows When It's Lying (2023). ArXiv:2304.13734 [cs].

25. L. Gao, *et al.*, A framework for few-shot language model evaluation (2023). Version Number: v0.4.0.

26. J. Wei, *et al.*, Chain-of-Thought Prompting Elicits Reasoning in Large Language Models (2023). ArXiv:2201.11903 [cs].

27. M. Nasr, *et al.*, Scalable Extraction of Training Data from (Production) Language Models (2023). ArXiv:2311.17035 [cs].

28. Gailly, Adler, zlib compression library. (2024).

29. E. Almazrouei, *et al.*, The Falcon Series of Open Language Models (2023). ArXiv:2311.16867 [cs].

30. H. Touvron, *et al.*, LLaMA: Open and Efficient Foundation Language Models (2023). ArXiv:2302.13971 [cs].

31. A. L. Anwyl-Irvine, J. Massonnié, A. Flitton, N. Kirkham, J. K. Evershed, *Behavior Research Methods* **52**, 388 (2020).

32. I. Loshchilov, F. Hutter, Decoupled Weight Decay Regularization (2019). ArXiv:1711.05101 [cs, math].

Acknowledgments

Materials and Methods

Dataset Creation

Co-authors (Table S.4) and GPT-4 (Azure OpenAI API; version 2023-05-15) created test cases that formed BrainBench. All test cases were sourced from *Journal of Neuroscience* abstracts published in 2023. The abstracts are organized into five sections, namely Behavioral/Cognitive, Systems/Circuits, Neurobiology of Disease, Development/Plasticity/Repair, and Cellular/Molecular. In constructing BrainBench, we incorporated a total of 200 test cases crafted by human experts and an additional 100 test cases generated by GPT-4 (Azure OpenAI API; version 2023-05-15). All test cases were subjected to extensive quality control by human experts and GPT-4. For the distribution of test cases among subfields, refer to Fig. 3 for human-created cases and Fig. S.10 for GPT-4 generated cases.

To create a test case, a published abstract was modified to create an altered version. The altered version significantly changed the results without changing the methods and background. Minimal changes were made that changed the basic result. For example, the altered abstract, compared to the original, could switch around the role of two brain regions in the results, reverse the direction of a result (e.g., replace "decreases" with "increases"), etc. Any changes maintained the coherency of the abstract which sometimes required multiple changes (e.g., replacing multiple "decreases" with "increases"). In other words, the altered abstracts needed to be empirically different, but not logically incoherent. Both volunteers and GPT-4 are given instructions that follow the essential criteria above. The exact wordings to prompt GPT-4 were slightly adjusted in order to obtain good quality test cases. We include the instructions used GPT-4 verbatim below.

GPT-4 test creation prompt   "Your task is to modify an abstract from a neuroscience research paper such that the changes significantly alter the result of the study without changing

21

the methods and background. This way we can test the Artificial Intelligence understanding of the abstract's subject area.

Please read the instructions below and ensure you follow them one by one while you are modifying the abstracts:

- The format to submit is putting double brackets around the change with the first element being the original and the second element being your edit. E.g., [[original passage, modified passage]]. Always remember to wrap your edits with the double brackets; there should not be any other edits outside the brackets to the original abstract. - If you change a single word, never wrap the entire sentence inside the double brackets. For example, "... exhibit [[enhanced LTP and deficits in LTD, impaired LTP and enhanced LTD]]." is a wrong format, the correct format is: "... exhibit [[enhanced , impaired]] LTP and [[deficits, enhanced]] in LTD." - The beginning of an abstract is the background and methods, so you should not alter those parts of the abstract. Do not alter the first couple sentences. - We want the abstract to become empirically wrong, but not logically incoherent. - To find the original result of the paper, one should require some neuroscience insight, not just general reasoning ability. So it is critical that the changes you make don't evaluate the Artificial Intelligence reasoning ability, but its knowledge of neuroscience and how the brain works. - Watch out for making changes that alter the results, but may still have occurred in the authors' study. For example, an fMRI abstract on learning might mention the hippocampus and not the striatum. Nevertheless, the striatum might have also been active and not reported in the abstract because it was not the focus of the study. - The changes you make should not be identifiable or decodable from the rest of the abstract. Hence, if you make a change, make sure you change everything that can reveal the original abstract. For example, "activation of neurons in the visual cortex [[increases, decreases]] the activity in the motor cortex. This decrease in the activity of the visual cortex was followed by an increase in task performance.". In this case it is very clear that the correct word is "decreases"

22

as the next sentence ("This decrease in the activity of the visual cortex") reveals that. - Be mindful of the article when you change words. For example, if you change the word "decline" to "enhancement", you must change the article as well, so the change will be [[a decline, an enhancement]]. - Ensure that your edits maintain inter-sentence consistency and proper syntax. The changes should not contradict or confuse the overall meaning of the abstract. - Avoid making trivial edits that do not require understanding of scientific concepts. The edits should reflect a deep understanding of the subject matter. - Do not miss any crucial results or findings in the abstract while making the edits. Every significant point should be addressed in your modifications.

To generate better responses, you can use the topic of their study and purpose of studies in those topics. This knowledge helps you to find what modification you should do in the abstract. Topics are: - Behavioral/Cognitive: To understand how the brain influences behavior, cognition, and emotion, and to apply this understanding in diagnosing and treating neurological and psychiatric disorders.

- Cellular/Molecular: To study are to understand the functions and mechanisms of neurons at a cellular and molecular level, which includes investigating the biology of nerve cells, their genetic makeup, and how they form complex circuits, ultimately contributing to our understanding of brain function, behavior, and the development of treatments for neurological disorders.

- Neurobiology of Disease: To understand the biological basis of various neurological and psychiatric disorders in order to develop effective treatments and preventative measures.

- Development/Plasticity/Repair: to understand the mechanisms of brain development, adaptation, and repair in response to injury or disease, with the goal of developing strategies and treatments to enhance brain recovery and function.

- Systems/Circuits: to understand how neural circuits in the brain interact and coordinate to process information, control behavior, and support cognitive functions.

Here are two examples of the edited abstract by human experts which can help you to understand the task:

Example 1: <example_1>

Example 2: <example_2>

These are some common mistakes you have made in the past. So keep them in mind whilst generating your responses: - You misunderstood/ignore the information provided at the beginning of the abstract. - The edits you have made are not what we are aiming for, you tweaked a portion of the studies with non-significant findings, so there's no significant alternation of results occurring. Make sure your edit changes the main results of the studies, not trivial changes. - Lack of inter-sentence consistency in the prompt - You made edits as early as the first sentence. THe first few sentence are general knowledge and are not result of the study. So you shouldn't make any change in the beginning. - Most of your edits contradict the conclusion. Make sure your changes do not contradit the conclusions or any part of the abstract. - You only modified verbs the understanding of which does not require understanding of scientific concepts & names of compounds, which makes the edits less likely to do wrong as long as reasons logically - One of your edits contradicts all other edits. - Your edit is inconsistent with the beginning of the sentence - You failed to change the first part of the conclusion for consistency - You missed out on one change. - You misunderstood the purpose of the study. Although in the abstract it explicitly states the purpose of the study.

Below, you are given an abstract with its topic. Follow the instructions given to you and return the modified abstract. Remember to use double brackets to show the changes ([[original, modified]] and keep the rest of the abstract unchanged. Also, pay attention to all the information you were given above as well as the common mistakes you have made before.

Abstract to edit: Topic: <abstract_topic>

Abstract: <abstract_to_edit> "

Evaluations

We tested human participants and LLMs on the BrainBench dataset. Both human experts and models were presented with two versions of the abstract, one with the actual results and one that was altered. The task was to determine which is which. Below, we detail how LLMs and human participants were tested.

Model Evaluation

We tested LLMs by adapting the Eleuther AI Language Model Evaluation Harness framework (*25*), which evaluates LLMs using a multiple choice setting. We presented LLMs with two versions of the abstracts from each test case separately. We prefixed each abstract with the prompt "*You are a neuroscientist with deep knowledge in neuroscience. Here is an abstract from a neuroscience publication:*" and applied model-specific instruction templates where appropriate. We then measured the perplexity of both passages and used perplexity as the indicator of whether LLMs favor one abstract or the other.

Perplexity (PPL) is one of the most common metrics for evaluating large language models. Perplexity measures the degree of uncertainty of a model when generating a particular sequence of text. Formally, perplexity is defined as the exponentiated average negative log-likelihood of a tokenized sequence. If we have a tokenized abstract $X = (x_0, x_1, \ldots, x_t)$, then the perplexity of X, given a LLM parameterized by θ is,

$$PPL(X) = \exp \left\{ -\frac{1}{t} \sum_i^t \log p_\theta(x_i | x_{<i}) \right\} \tag{1}$$

where $\log p_\theta(x_i | x_{<i})$ is the log-likelihood of the ith token conditioned on the preceding tokens $x_{<i}$ according to the LLM. Given both the original abstract X_{orig} and the altered abstract X_{alt}, we followed the decision rule,

$$X_{chosen} = \begin{cases} X_{orig}, & \text{if } PPL(X_{orig}) < PPL(X_{alt}) \\ X_{alt}, & \text{otherwise} \end{cases} \tag{2}$$

and evaluated the overall accuracy over the entire BrainBench accordingly.

Accuracy Accuracy is the primary metric for reporting LLM performance on BrainBench. A correct response was when the model produces a lower perplexity for the original abstract than the altered abstract.

Confidence calibration We used the absolute difference of perplexities of two versions of the abstract as a measure of model confidence. To assess the calibration of LLMs, we compared their accuracies with their confidence levels. First, we ranked and sorted model confidence across all test cases. Subsequently, we created 20 bins based on this sort. Within each bin, we calculated the mean accuracy. A well-calibrated model will exhibit a higher accuracy in bins associated with higher confidence rankings. We fit a linear regression model using the bin number as the independent variable and the mean accuracy of each bin as the dependent variable to evaluate calibration.

Performance correlation across LLMs We assessed the correlation in performance among different LLMs by examining how they rank the relative difficulty of test cases. To determine difficulty, we calculated the difference in perplexity between incorrect and correct abstracts for each test case. Intuitively, a large positive difference in the perplexity between incorrect and correct versions of an abstract should indicate that the test case is easy from the LLM's perspective. We calculated the Spearman correlation coefficient of these difficulty measures to assess the agreement between two LLMs.

Integration analysis   To investigate the extent to which LLMs can integrate broad context from abstracts, we conducted an experiment involving the removal of contextual information from BrainBench test cases. Following the same evaluation procedure as previously outlined for full abstract cases, we assessed the models using individual sentences extracted from abstracts containing at least one result alternation. In cases with multiple alternations, we computed the mean accuracy across these alternations as the final accuracy for the abstract. We then compared the level of performance degradation when LLMs were evaluated on full-length abstracts versus individual sentences where background and method information from the abstracts were removed.

LLM training data memorization analysis   One concern regarding LLMs outperforming human experts on BrainBench is the possibility that LLMs were exposed to the original abstracts during their pre-training. If LLMs have simply memorized the training data, they would naturally assign lower perplexity scores to the correct abstracts. For some models, like Galatica, their training set is open and we know they were not trained on BrainBench items.

To address this concern, we employed a common method from the literature to determine whether a given text is part of LLM's training data (*18, 27*). This method involves calculating the zlib entropy and the perplexity ratio (Eq. 3) of a text sequence to infer its membership status.

$$ratio = \frac{ZLIB(X)}{PPL(X)} \tag{3}$$

Zlib entropy is computed using the zlib text compression algorithm (*28*), which measures the level of uncertainty in a text when compressed. It's a data-agnostic way of evaluating text. On the other hand, LLM perplexity depends on the specific training data and thus is data-dependent. In general, if a piece of text surprises zlib but not LLM, it's likely part of the training data.

To conduct this test, we carefully chose data sources that are either known to be part of

LLMs' pre-training or reasonably assumed to be excluded from it (refer to Table S.1 and S.2). We then applied zlib compression and LLM perplexity calculations to text samples from these selected sources.

In addition, we introduced the Gettysburg Address as a special anchor point to contrast with the zlib-perplexity ratio distribution across multiple data sources. This is because we expect the Gettysburg Address to exhibit a high zlib score due to its non-modern form of English, coupled with a low perplexity, given its likely frequent exposure during LLM pre-training.

Model Series	Pretraining Data Cutoff	Initial Release Date
Llama2	September 2022	July 2023
Galactica	July 2022	November 2022
Falcon	December 2022	May 2023
Mistral	Unknown	September 2023

Table S.1: Overview of LLMs Training Cutoff and Initial Release Dates.

Data Source	Model Training Set Inclusion
RefinedWeb	Known to be in Falcon's training set (*29*)
Arxiv (Jun-Dec 2021)	Likely to be in Llama2's training set; known to be in Galactica's training set
Arxiv (Jun-Dec 2023)	Known to not be in Llama2 and Galactica's training set (*20, 30*)
Biorxiv (Jun-Dec 2021)	Known to be in Galactica's training set (*20*)
Biorxiv (Jun-Dec 2023)	Known to not be in Galactica's training set
Gettysburg Address	Likely to be in all models' training set

Table S.2: Data Source Inclusion in AI Model Training Sets

Human Evaluation

Participants   We recruited 202 neuroscience experts via social media and an email newsletter. We excluded 31 participants for failing to answer both catch trials correctly, not providing confidence or expertise ratings during the entire experiment, and self-reported cheating. The remaining 171 participants consisted of 51 doctoral students, 43 faculty/academic staff, 43 post-doctoral researchers, 18 predoctoral students, 12 research scientists, and 4 classified as "other".

Participants' mean experience in neuroscience was 10.1 years. Participants identified as follows: 62.5% male, 34.5% female, and 0.6% gender variant/non-conforming. The mean age was 35.2 years (*SD = 9.4 years*).

Procedure   First, participants were briefed on the experimental task and provided their informed consent to proceed to the experiment. Demographic information was then collected, including gender identity, age, country, current position, and years of experience in neuroscience research, broadly construed. Next, participants completed a practice trial using the same testing format as the actual test cases. This trial was used to familiarize participants with the format of the task with the screen proceeding only once participants had made the correct choice based on common sense. Following this, 9 test trials and 2 catch trials commenced, where participants selected one version of each trial abstract. Out of the 9 test trials, 6 were randomly sampled human-created test cases and 3 were randomly sampled from the pool of machine-created items. We ensured that each test case is sampled approximately an equal number of times across all participants. To achieve this, we maintained a global counter that keeps track of how frequently each test case has been used. As a result, the next participant's sample will always be drawn from those test cases that have been used less frequently. Notably, the number of alterations varies between test cases, but the design allowed a single click to automatically select between the two abstract options (see Fig. S.1). Participants made one decision per test case, regardless of the number of alternations.

Subsequently, participants were required to rate their confidence and expertise using slider bars. The confidence slider had a range from "lower" on the left to "higher" on the right, while the expertise slider spanned from "not at all" on the left to "very much so" on the right, both internally implementing a 1-100 scaling. Additionally, participants indicated whether they had encountered the study previously before proceeding to the next trial. Upon completing the 11
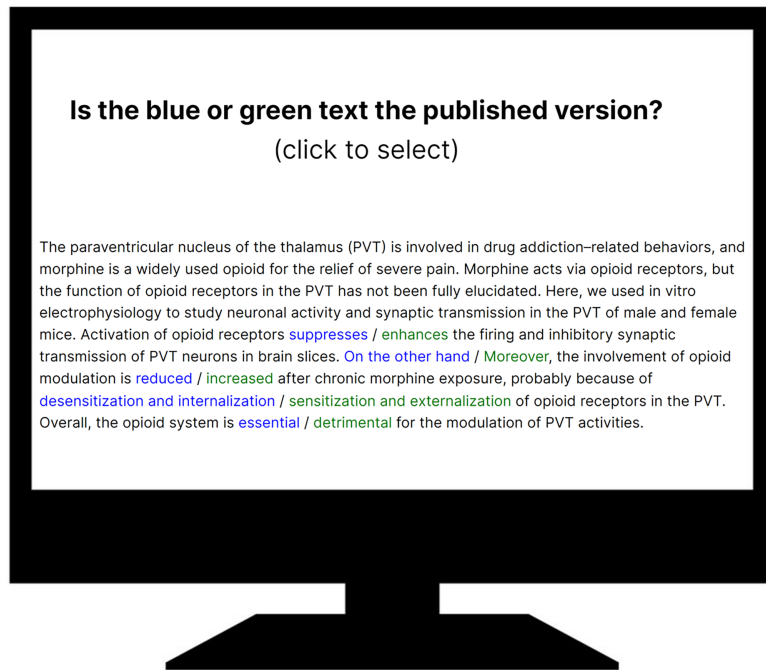
Is the blue or green text the published version?
(click to select)

The paraventricular nucleus of the thalamus (PVT) is involved in drug addiction–related behaviors, and morphine is a widely used opioid for the relief of severe pain. Morphine acts via opioid receptors, but the function of opioid receptors in the PVT has not been fully elucidated. Here, we used in vitro electrophysiology to study neuronal activity and synaptic transmission in the PVT of male and female mice. Activation of opioid receptors suppresses / enhances the firing and inhibitory synaptic transmission of PVT neurons in brain slices. On the other hand / Moreover, the involvement of opioid modulation is reduced / increased after chronic morphine exposure, probably because of desensitization and internalization / sensitization and externalization of opioid receptors in the PVT. Overall, the opioid system is essential / detrimental for the modulation of PVT activities.

Figure S.1: **Study test interface**: Participants were instructed to select which version of the abstract was the original by clicking on either blue or green text to select that set of options. Various test cases may have varying numbers of alternatives, but a single click will choose all options of the same color.

trials, participants were debriefed on which trials they got correct and were subsequently asked to indicate whether they engaged in any form of cheating during the study. We hosted the study entirely on the Gorilla platform (*31*).

Exclusion criteria   For participant selection and data analysis, we apply several exclusion criteria. First, individuals who failed to answer both catch trials correctly were not included in the data analyses. Second, participants who did not make adjustments to the sliders (i.e., expertise and confidence) during any of the trials were excluded. Additionally, trials where participants recognized the abstract content were omitted from the analysis. Furthermore, trials with reaction times less than 5 seconds were excluded. Lastly, participants who admitted to using external resources or engaging in cheating behaviors, as indicated by a checkbox in the debriefing form, were not considered in the final data analysis.

Performance correlation between humans and LLMs   We assessed the agreement between humans and LLMs using a similar approach as we did when evaluating the correlation among LLMs. For LLMs, the procedure for determining item difficulty was identical to that described above. For human experts, item difficulty was calculated as the mean accuracy for that item. Finally, the Spearman correlation of these difficulty measures was calculated to assess agreement.

Fine-tuning on neuroscience corpora

The LLMs we considered had been pre-training on a diverse range of text corpora, including Internet sources, Wikipedia, books, code repositories, and arXiv papers. While these pre-trained models are designed to be versatile and capable of handling various tasks, our approach for creating BrainGPT involved enhancing base models with domain-specific expertise, specifically in neuroscience.

To accomplish this, we employed the Low-Rank Adaptation of Large Language Models

31

(LoRA) technique (Fig. S.6; (*19*)). LoRA efficiently extends the capabilities of general-purpose LLMs by introducing low-rank trainable parameters (referred to as "adapters") into the existing model. This process effectively fine-tunes the model for downstream tasks without the need for prohibitively resource-intensive training of the entire model.

Training data   We collected training data from PubMed for abstracts and PubMed Central Open Access Subset (PMC OAS) for full-text articles using the Entrez Programming Utilities (E-utilities) API and the pubget Python package, respectively. The data span publication dates from 2002 to 2022. For science general journals, we applied a keyword filter of "Neuroscience" (see all sourced journals in Table S.3).

Our data extraction efforts yielded 332,807 abstracts and 123,085 full-text articles, totaling 1.3 billion tokens. We excluded figures and tables and randomly allocated 90% of the data for training, reserving the remaining 10% for validation.

Training details   We fine-tuned Llama2-7b-chat from Meta (*30*) using weights available on Huggingface (https://huggingface.co/meta-llama/Llama-2-7b-chat-hf). We used a batch size of 1 and a chunk size of 2048. Training involved the use of the AdamW optimizer (*32*) with a learning rate of 2e-5 and gradient accumulation steps set at 8. A single training epoch was performed, along with a warm-up step of 0.03 and a weight decay rate of 0.001. The learning rate was controlled using a cosine learning rate scheduler. LoRA adapters, characterized by a rank of 8, an alpha value of 32, and a dropout rate of 0.1, were applied after all self-attention blocks and fully-connected layers. This results in total 17,891,328 trainable parameters, roughly 0.26% of the entire parameters of the base model. To optimize training performance, bf16 mixed precision training and data parallelism were employed. We used 4 Nvidia A100 (80GB) GPUs hosted on the Microsoft Azure platform. An epoch of training takes roughly 42 GPU hours.

Evaluation   We tested the fine-tuned model on BrainBench using the same procedure as before. To verify the significance of performance improvement, we performed a paired t-test with respect to the difference in perplexity of incorrect and correct options before and after fine-tuning. We also performed a paired t-test with respect to the perplexity of the correct options before and after fine-tuning.
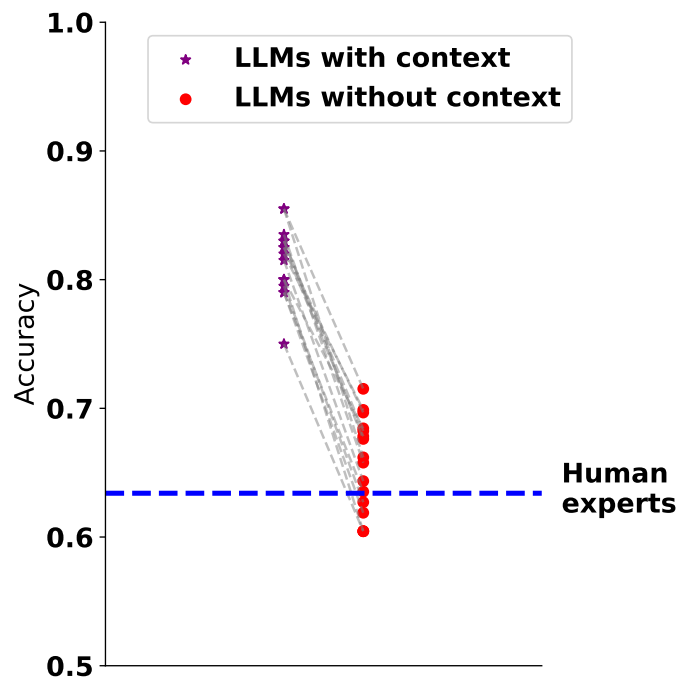
Figure S.2: **LLM integrates contextual information to succeed on BrainBench**. The removal of background and method sections from abstracts, with an evaluation based solely on individual sentences and result alternations, significantly impairs the performance of LLMs on BrainBench. LLMs' superior performance appears to arise from integrating information across the abstract.
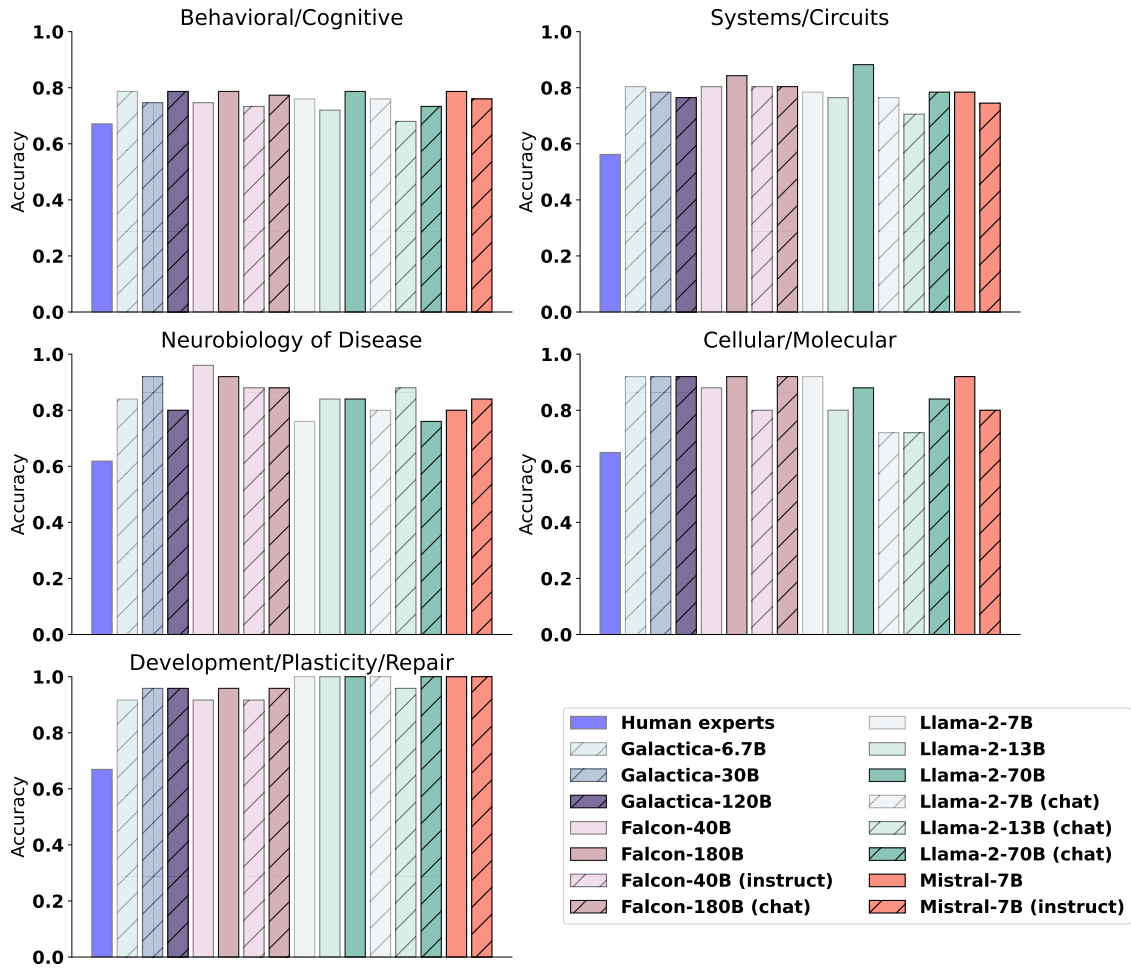
Figure S.3: **BrainBench performance by subfield.** Accuracy on BrainBench by subfields of neuroscience for human experts and LLMs.
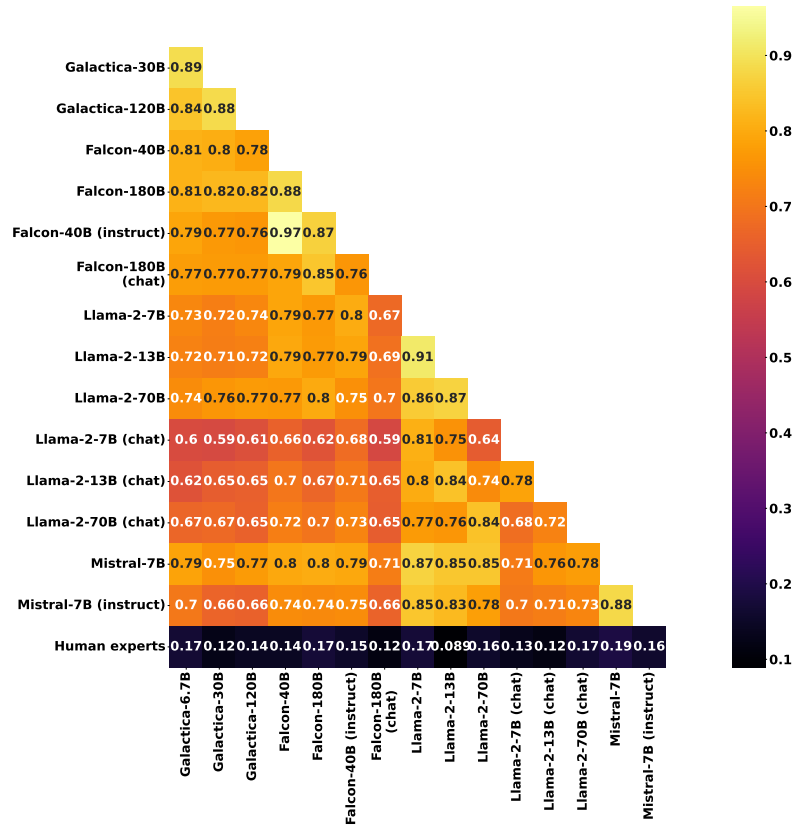
Figure S.4: **Item difficulty correlation among LLMs and human experts.** For LLMs, difference in perplexity between incorrect and correct abstracts is used to determine the relative difficulty of test cases. Mean accuracy is used for human experts. Spearman correlation is calculated for these difficulty measures. LLMs have an average correlation of 0.75 (± 0.08) whereas human experts have an average of 0.15 (± 0.03) with LLMs.
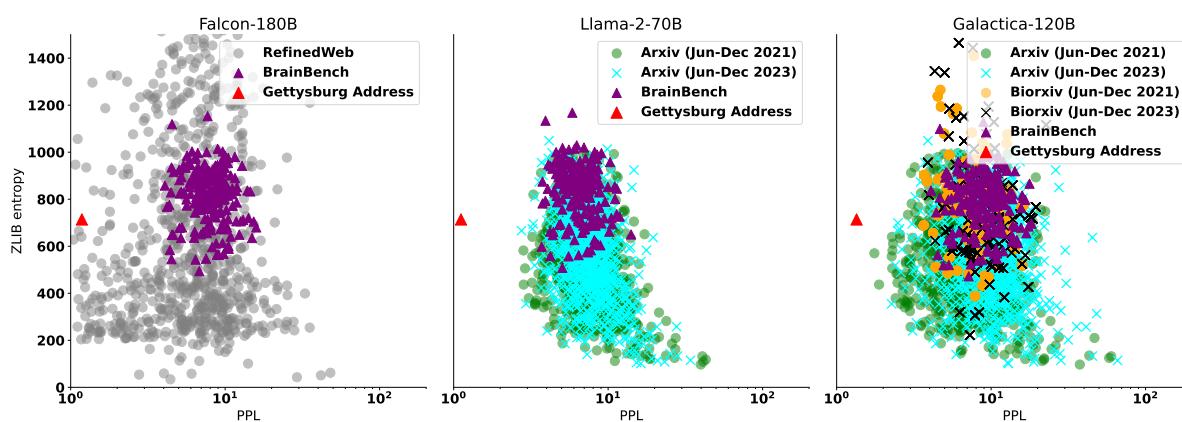
Figure S.5: **Memorization analysis.** Analyzing zlib entropy and perplexity ratio in text samples from known data sources, both included and excluded from LLMs training data, reveals no distinct signature for pre-trained data sources compared to those outside the training set. Although certain training data that occurs multiple times in the training set, like the Gettysburg Address, show signs of memorization, low-frequency training data show no sign of memorization. We analyzed the three largest LLMs we evaluated as they are most capable of memorizing data.
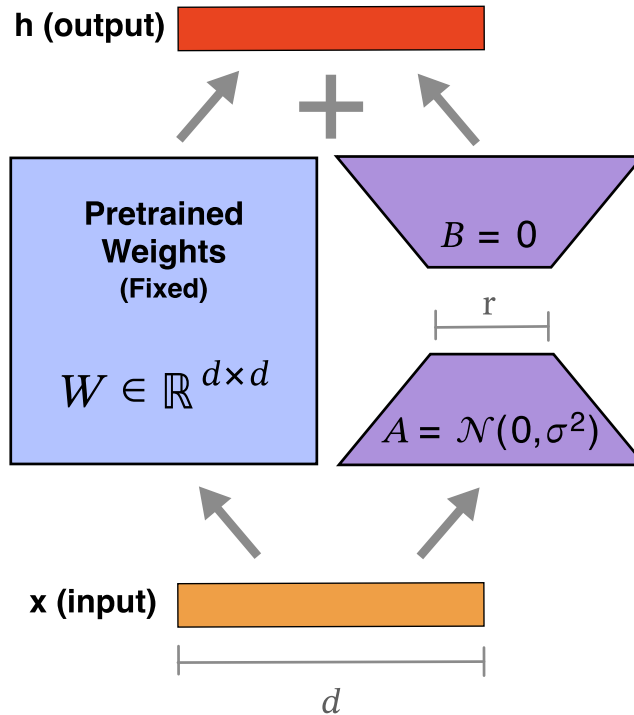
Figure S.6: **Low-Rank Adaptation (LoRA) for parameter-efficient large language model (LLM) fine-tuning.** Image adapted from (*19*). Training LLMs from scratch can be computationally expensive, especially when the model has many billions of parameters and the training data are massive. Ideally, one would be able to take advantage of a previously trained LLM (i.e., a base model) and build from it. Toward this end, various Parameter Efficient Fine-Tuning (PEFT) techniques have been proposed. Rather than retrain the LLM, these techniques preserve the LLM's original pre-trained weights and train a small subset of additional parameters that either enhance an existing capability in the LLM or introduce a new one. In our case, we tune a base LLM to the neuroscience literature. This tuning strategy significantly alleviates the computational burden while achieving comparable performance to training a new LLM from scratch. Among many PEFT approaches, we adopted LoRA in our study to fine-tune the 7-Billion parameter version of the LLaMa2 model on neuroscience abstracts. LoRA fixes the pretrained model weights and injects trainable rank decomposition matrices into each layer of the transformer. During training, the same input sequence (or intermediate layer output) x is processed separately by the pretrained weights and the low-rank adaptation matrices (A and B). The low-rank matrices possess the only trainable parameters in the model. The final output h is computed as a coordinate-wise addition between the product of the pretrained weights and the adaptation matrices, which are further processed by subsequent layers.

Nature, Cell, Cell Reports, eLife, Science Advances, Nature Communications, PNAS, The EMBO Journal, Nature Neuroscience, Neuron, Brain, NeuroImage, Molecular Psychiatry, Journal of Neuroscience, Nature Reviews Neuroscience, Cerebral Cortex, Annals of Neurology, Human Brain Mapping, Epilepsia, Clinical Neurophysiology, Trends in Cognitive Sciences, Biological Psychiatry, Translational Psychiatry, Neuroscience and Biobehavioral Reviews, Neuropsychopharmacology, Alzheimer's and Dementia, NeuroImage: Clinical, Neurobiology of Aging, Trends in Neurosciences, Nature Reviews Neurology, Brain Stimulation, Frontiers in Neuroscience, Movement Disorders, Nature Human Behaviour, Frontiers in Neurology, Cortex, Journal of Alzheimer's Disease, Neurobiology of Disease, Biological Psychiatry: Cognitive Neuroscience and Neuroimaging, Brain Structure and Function, Pain, Frontiers in Human Neuroscience, eNeuro, Current Opinion in Neurobiology, European Journal of Neuroscience, Frontiers in Aging Neuroscience, Alzheimer's Research and Therapy, Journal of Neurology, Glia, Epilepsy and Behavior, Brain Imaging and Behavior, Journal of Neurophysiology, Sleep, Neuroscience, Neuropsychologia, Journal of Neural Engineering, Molecular Neurobiology, Frontiers in Cellular Neuroscience, Neuropharmacology, Alzheimer's and Dementia: Diagnosis, Assessment and Disease Monitoring, Journal of Neuroinflammation, Epilepsia Open, Acta Neuropathologica Communications, Frontiers in Neuroinformatics, Current Opinion in Behavioral Sciences, Developmental Cognitive Neuroscience, Frontiers in Molecular Neuroscience, Cerebellum, Journal of Cognitive Neuroscience, Network Neuroscience, Annual Review of Neuroscience, Progress in Neurobiology, Epilepsy Research, Molecular Autism, Journal of Comparative Neurology, Social Cognitive and Affective Neuroscience, Brain Topography, Hippocampus, Seizure: the journal of the British Epilepsy Association, Psychophysiology, Frontiers in Behavioral Neuroscience, Journal of Neurotrauma, Journal of Physiology, Frontiers in Neural Circuits, Neurobiology of Learning and Memory, Journal of Neural Transmission, Frontiers in Neuroanatomy, International Journal of Neuropsychopharmacology, Neuroscientist, Brain Sciences, Behavioural Brain Research, Experimental Neurology, Progress in Neuro-Psychopharmacology and Biological Psychiatry, Neurological Sciences, Neurotherapeutics, Neuroscience Letters, Current Opinion in Neurology, Journal of Neuroscience Methods, Journal of Neurochemistry, Neuromodulation, Molecular Neurodegeneration, Frontiers in Systems Neuroscience, Sleep Medicine Reviews, Brain and Behavior, Brain Research, Neurorehabilitation and Neural Repair, Autism Research.

Table S.3: **Journals used for LoRA fine-tuning.** Abstracts and full articles published between 2002 and 2022 sourced from the above journals were used as the training set for LoRA fine-tuning.
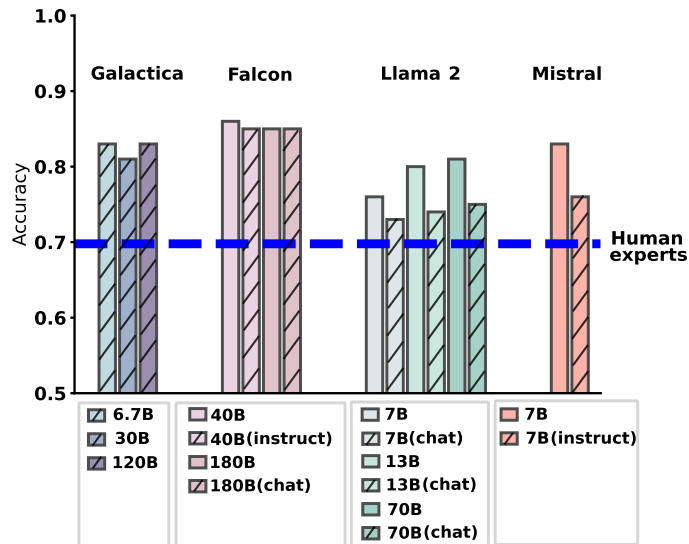
Figure S.7: **LLMs outperformed human experts on BrainBench (using GPT-4 created test cases).** Base versions of models outperformed chat and instruct versions, which were tuned to be conversational with humans.
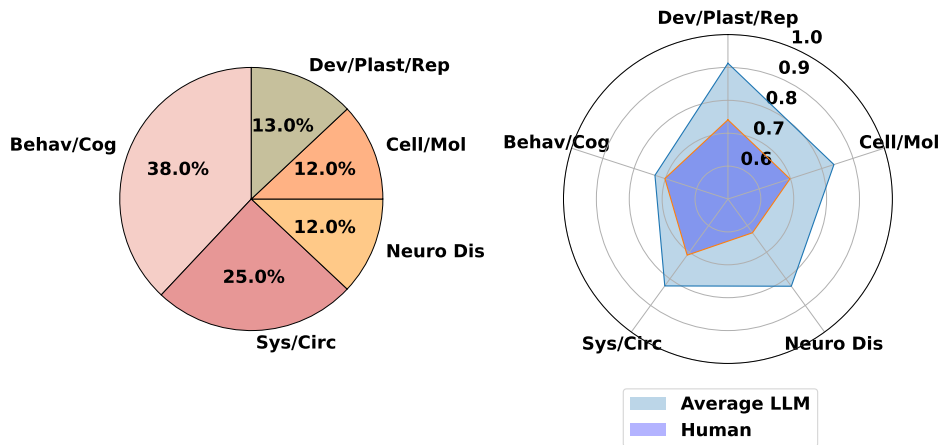

Figure S.8: **BrainBench performance breakdown by subfields of neuroscience (using GPT-4 created test cases).** The distribution of test cases across neuroscience subfields roughly mirrors the distribution of articles in the Journal of Neuroscience with Behavior/Cognitive overrepresented. The mean performance of 15 LLMs and human experts is shown. LLMs outperformed human experts in all subfields.
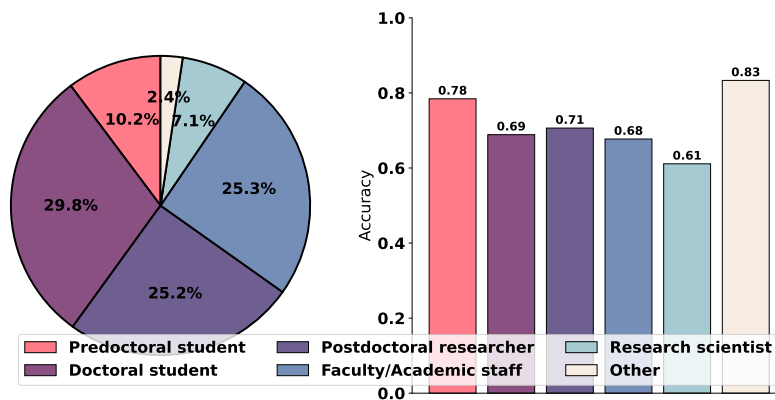
Figure S.9: **BrainBench performance of human experts belong to different self-reported career categories (using GPT-4 created test cases).** Majority of the participants were doctoral students, postdoctoral researchers and faculty/academic staff.
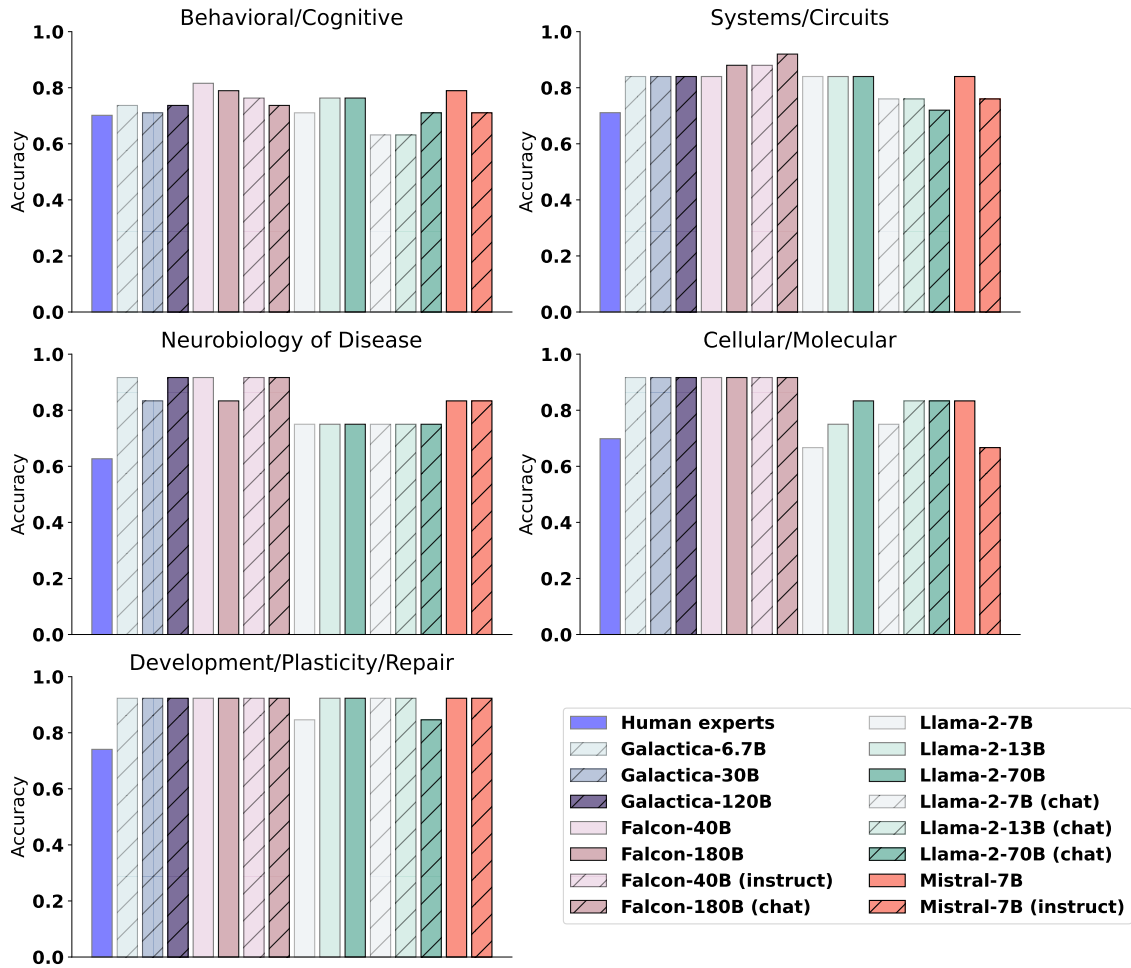
Figure S.10: **BrainBench performance breakdown (using GPT-4 created test cases).** Accuracy by subfields of neuroscience across human experts and LLMs.
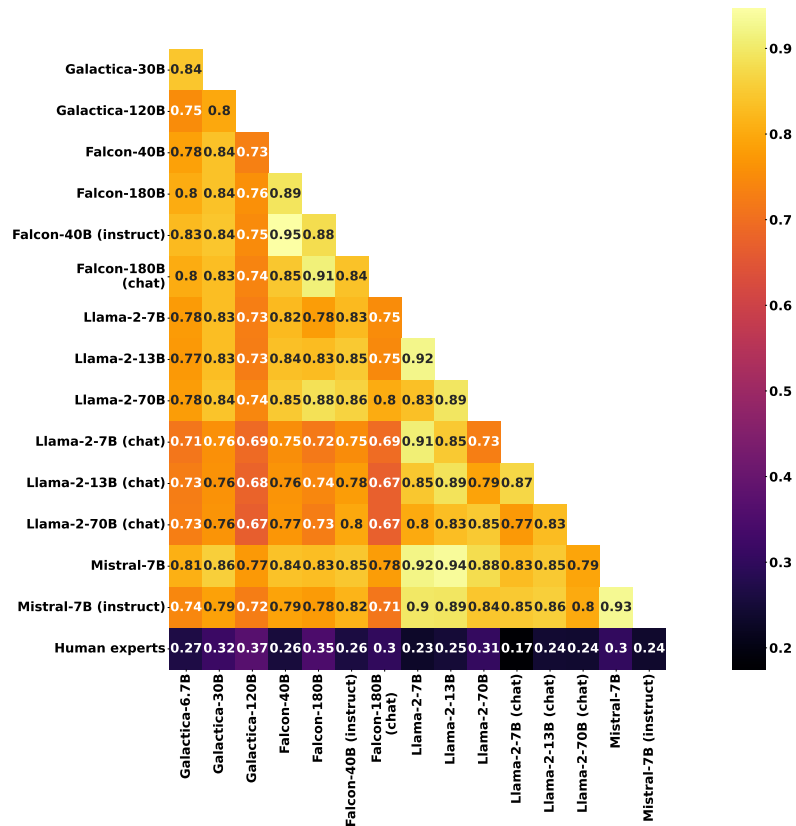
Figure S.11: **Test difficulty correlation among LLMs and human experts (using GPT-4 created test cases).** Difference in perplexity between incorrect and correct abstracts is used to determine the relative difficulty of test cases. Spearman correlation is computed between ranked difficulties by LLMs and human experts. LLMs have an average Spearman correlation of 0.83 (± 0.06) whereas human experts have an average of 0.27 (± 0.05) with LLMs.
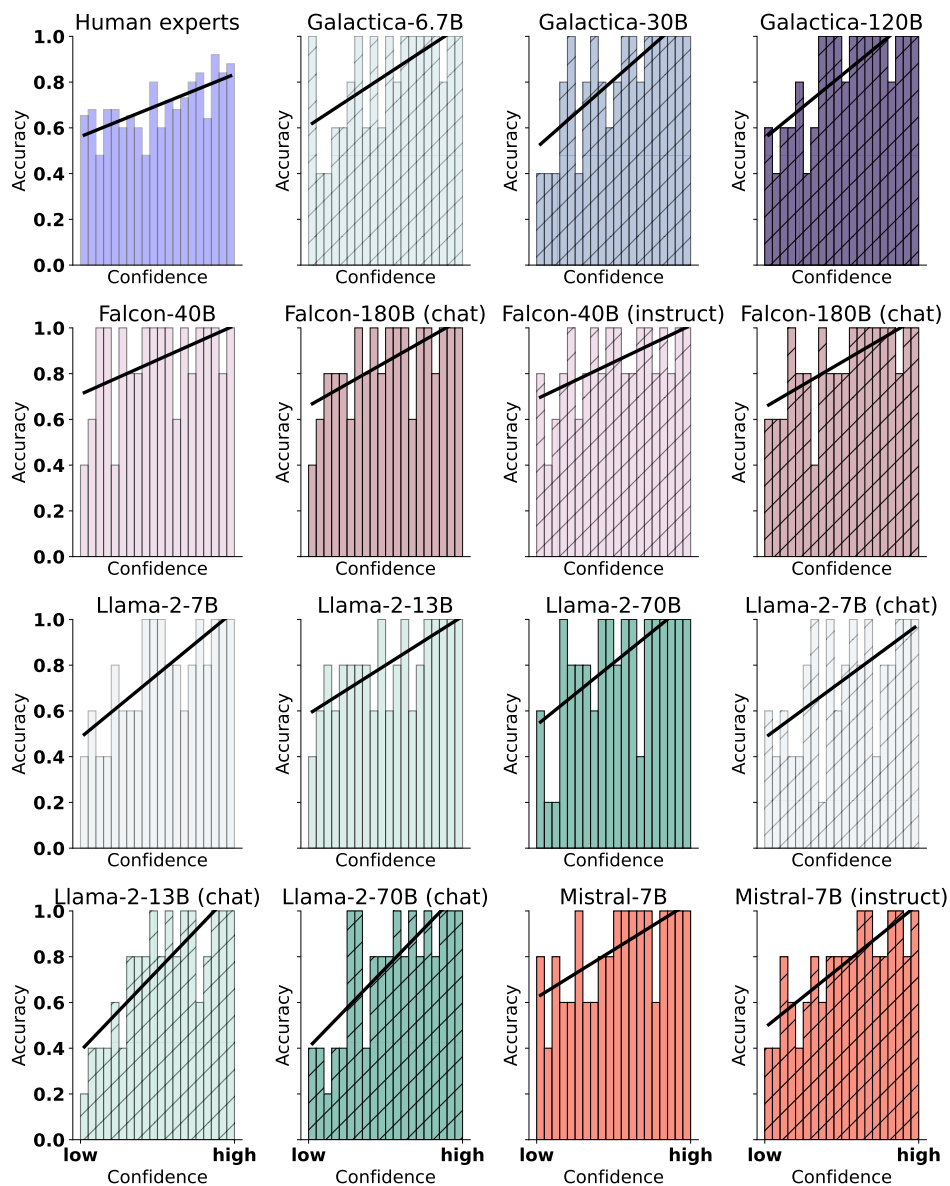
Figure S.12: **Accuracy and confidence are calibrated for human experts and LLMs (using GPT-4 created test cases).** When human experts and LLMs are confident in their BrainBench judgments, they are more likely to be correct. Confidence ratings were sorted and placed in equally-sized bins with the mean accuracy for items in that bin plotted. The black regression line's positive slope for human experts and all LLMs indicates that confidence is well calibrated (i.e., higher confidence corresponds to higher accuracy). Calibration is beneficial for human-machine teams.

Example #1

The amygdala plays a key role in the processing of itch and pain signals as well as emotion. A previous study [...] **Background**

[...] To test this possibility, prodynorphin (Pdyn)-Cre mice were used to optogenetically manipulate Pdyn+ CeA-to-PBN projections. **Method**

[...] We found that [...] Pdyn+ CeA-to-PBN projections [[inhibited, increased]] histamine-evoked and [...] Optogenetic stimulation of Pdyn+ CeA-to-PBN projections [[suppressed, enhanced]] the increase in Fos expression in the PBN. [...] **Result**

Example #2

Language comprehension requires the rapid retrieval and integration of contextually appropriate concepts ("semantic cognition"). Current neurobiological models [...] **Background**

[...] Through the use of fused functional magnetic resonance imaging and electroencephalography analysis in humans (n = 26 adults; 15 females), we elucidate a temporally and spatially specific neurobiological model [...] **Method**

[...] We find that [...] provide convergent evidence for the following progression: [[a hippocampal/anterior temporal, posterior temporal]] phonological semantic retrieval network [...] an inferior [[frontal, parietal]] semantic syntactic reappraisal network (~600 ms); and [...] **Result**

Figure S.13: **BrainBench examples.** For a test case, the background and the method sections of the abstract are kept unchanged. Alternative findings are introduced by human experts or GPT-4. Participants choose from two options, the original and the altered, with the aim of selecting the actual (i.e., original) finding.

Tasks	Authors
Primary writing	Xiaoliang Luo, Bradley C. Love
Test case creation	Bati Yilmaz, Anton Pashkov, Isil Poyraz Bilgin
	Tereza Okalova, Alexandra O. Cohen, Felipe Yáñez
	N Apurva Ratan Murty, Elkhan Yusifov, Kangjoo Lee
	Valentina Borghesani, Sepehr Razavi, Kaustubh R. Patil
	Michael Gaebler, Justin M. Ales, Guiomar Niso
	Rui Mata, Mikail Khona, Roberta Rocca
	Kevin K. Nejad, Alessandro Salatiello, Jonathan Nicholas
	Daniele Marinazzo, Chloe M. Hall, Pui-Shee Lee
	Sebastian Musslick, Nicholas E. Myers, Jennifer K Bizley
	Sherry Dongqi Bao, Nianlong Gu
	Jessica Dafflon, Bradley C. Love
Quality control	Bati Yilmaz, Kevin K. Nejad, Daniele Marinazzo
	Pui-Shee Lee, Nianlong Gu, Chloe M. Hall
	Kangjoo Lee, Sebastian Musslick, Akilles Rechardt
	N Apurva Ratan Murty, Rui Mata, Isil Poyraz Bilgin
	Tereza Okalova, Guiomar Niso, Ilia Sucholutsky
	Mikail Khona, Justin M. Ales, Michael Gaebler
	Elkhan Yusifov, Jessica Dafflon, Felipe Yáñez
	Jonathan Nicholas, Roberta Rocca, Valentina Borghesani
	Sherry Dongqi Bao, Alexandra O. Cohen
	Alessandro Salatiello, Xiaoliang Luo, Bradley C. Love
GPT-4 case creation	Kevin K. Nejad, Xiaoliang Luo, Bradley C. Love
Human-machine teaming	Felipe Yáñez, Xiaoliang Luo, Bradley C. Love
LoRA fine tuning	Guangzhi Sun, Xiaoliang Luo, Bradley C. Love
Model evaluation	Xiaoliang Luo, Guangzhi Sun, Martin Ferianc
	Bradley C. Love
Building the experiment	Akilles Rechardt, Xiaoliang Luo, Bradley C. Love
Data analysis	Xiaoliang Luo, Akilles Rechardt, Bradley C. Love
Conceptualization and strategy	Xiaoliang Luo, Bradley C. Love
Figure creation	Xiaoliang Luo, Bati Yilmaz, Isil Bilgin
	Chloe M. Hall, Guiomar Niso, Bradley C. Love
Useful input and suggestions on project	All authors
Commenting and editing on the manuscript	All authors

Table S.4: **Author contributions breakdown.** For each task, authors listed toward the beginning of the list contributed more.