Triangulating LLM Progress through Benchmarks, Games, and Cognitive Tests

Filippo Momentè¹, Alessandro Suglia^{2*}, Mario Giulianelli³, Ambra Ferrari¹, Alexander Koller⁴, Oliver Lemon⁵, David Schlangen^{6,7}
Raquel Fernández⁸, Raffaella Bernardi^{9†}

¹University of Trento, ²ILCC University of Edinburgh, ³UCL, ⁴Saarland University, ⁵Heriot-Watt University, ⁶University of Potsdam, ⁷DFKI, ⁸University of Amsterdam, ⁹Faculty of Engineering, Free University of Bozen-Bolzano lm-playschool@googlegroups.com

Abstract

We examine three evaluation paradigms: standard benchmarks (e.g., MMLU and BBH), interactive games (e.g., Signalling Games or Taboo), and cognitive tests (e.g., for working memory or theory of mind). First, we investigate which of the former two—benchmarks or games—is most effective at discriminating LLMs of varying quality. Then, inspired by human cognitive assessments, we compile a suite of targeted tests that measure cognitive abilities deemed essential for effective language use, and we investigate their correlation with model performance in benchmarks and games. Our analyses reveal that interactive games are superior to standard benchmarks in discriminating models. Causal and logical reasoning correlate with both static and interactive tests, while differences emerge regarding core executive functions and social/emotional skills, which correlate more with games. We propose a new evaluation framework triangulating LLMs progress. Our findings highlight the importance of defining evaluation regimes that consider multiple paradigms. We advocate for the development of new interactive benchmarks and targeted cognitive tasks inspired by human ability assessments but designed specifically for LLMs. The code for running the experiments is released at: https://github.com/momentino/ playpen_eval/tree/triangulating.

1 Introduction

Evaluating LLMs is critical to track progress, identify blind spots, and ultimately advance towards the kind of language-based AI systems we want as a society (Wooldridge and Jennings, 1995). Currently, the most widespread way to evaluate LLMs is by means of **large benchmarks** made up of miscellaneous tasks. Pre-LLM benchmarks such as GLUE

and SuperGLUE (Wang et al., 2019b,a) have been replaced by even larger evaluation suites such as MMLU (Measuring Massive Multitask Language Understanding; Hendrycks et al., 2021), MMLU-Pro (Wang et al., 2024), or BBH (BIG-Bench Hard; Suzgun et al., 2023; Srivastava et al., 2023). Models with high performance on these benchmarks are taken to possess extensive world knowledge along with complex problem-solving abilities.

This trend has promoted standardisation in LLM evaluation protocols, with online leaderboards constantly updated as new models are released. Despite this undeniable benefit, general-purpose benchmarks, like those mentioned above, are not without problems. Evaluation results may be inflated by data contamination (see, e.g., Gema et al. 2025 for MMLU) and distorted by model sensitivity to prompt format (Zhuo et al., 2024). Moreover, by design, such benchmarks overlook actual language use in favour of knowledge-intensive tasks where success is measured against gold reference answers provided in a single conversational turn.

This contrasts with the view, put forward by philosophers and psycholinguists alike (Wittgenstein, 1953; Austin, 1962; Searle, 1969; Clark, 1996), that the quintessence of language resides in *situated language use*, i.e., using language for a purpose in social and task-based multi-turn interactions (Bisk et al., 2020). This situated and interactive view underpins a parallel trend where LLMs are evaluated as **goal-directed language users** by means of **interactive games** (Bertolazzi et al., 2023; Schlangen, 2023; Suglia et al., 2024). This interactive evaluation paradigm goes beyond single-turn text generation, which is critical for deploying LLMs as agents. Additionally, it is

^{*}Work done while at Heriot-Watt University.

[†]Work partially done while at University of Trento.

¹Online leaderboards have started to appear for the interactive games evaluation paradigm; see, e.g., https://textarena.ai/, https://clembench.github.io.

less susceptible to data contamination because the vast space of possible multi-turn interactions is unlikely to be fully represented in the training data. Yet, despite these advantages, it is not easy to pinpoint which specific abilities underpin models' performance on interactive language games—a difficulty that to some extent also applies to static question-answering benchmarks such as MMLU.

In this paper, we study these two evaluation paradigms—general-purpose benchmarks and interactive games—and show that they can provide complementary perspectives. In particular, we investigate whether standard benchmarks or games are more effective in gauging qualitative differences between models, evaluate a selection of current LLMs from five model families, and find that games highlight more nuanced differences between LLMs compared to standard benchmarks. While scaling model size leads to systematic improvements on benchmarks, game-based evaluation reveals differences between models of the same size.

To shed light on the abilities underlying models' performance on these two evaluation frameworks, we resort to **targeted cognitive tests**. In the absence of specifically designed tests for LLMs, we propose a taxonomy of cognitive skills motivated by human neurocognitive science and compile a list of existing evaluation datasets designed to assess each skill in isolation. Specifically, we investigate to what extent increased performance on specific abilities correlates with performance gain in standard benchmarks and interactive games, respectively.

In sum, our work brings in a unique perspective by studying the interrelations between a diverse set of evaluation paradigms which, to our knowledge, have not been jointly investigated before.

2 Models

We apply our evaluation framework to the following open-weight LLMs ranging from 3.8B to 72B models: Phi-4-Mini with 3.8B parameters (Microsoft et al., 2025), OLMo-2-1124 with 7B and 13B parameters (Walsh et al., 2024), Falcon3 with 10B parameters (Falcon Team, 2024), Qwen2.5 with 7B, 32B, and 72B parameters (Yang et al., 2024; Team, 2024), and Llama-3 with 8B and 70B parameters (Grattafiori et al., 2024). All these models score above 70% on IFEval (Zhou et al., 2023), which we take as indication that they possess the instruction-following capabilities neces-

sary for playing games and understanding task descriptions in the cognitive tests. See Appendix A for further model details.

3 How to Identify Blind Spots in LLMs

Evaluation paradigms have most practical use when they allow us to track progress by identifying blind spots in LLMs. Here we compare the two paradigms under study on the extent to which they highlight differences between current models, helping us form hypotheses about possible performance gaps and successful mitigation strategies.

General-purpose reference-based benchmarks We take MMLU (Hendrycks et al., 2021) and BBH (Suzgun et al., 2023) as representative of standard benchmarks. MMLU evaluates whether LLMs can apply knowledge from specific domains: it consists of multiple-choice questions spanning 57 academic subjects. BBH assembles diverse tasks drawing problems from linguistics, child development, maths, and common-sense reasoning, a.o.

Interactive games We take clembench (Chalamalasetti et al., 2023) as a benchmark to assess LLMs' gameplay ability in dialogue games. We consider the games 1) *Taboo*, 2) *Wordle (base)* and the two variants *Wordle (Clue)* and *Wordle (Critic)*, 3) *Reference Game*, 4) *Image Game*, and 5) *Private/Shared*. Game descriptions and examples are provided in Appendices C and D.

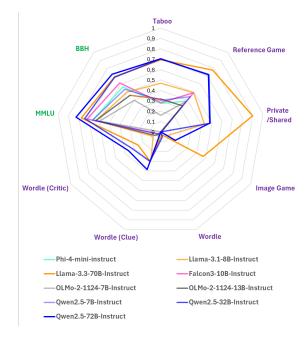


Figure 1: Accuracy for all models on standard benchmarks and interactive games. See Table 3 in App. I.

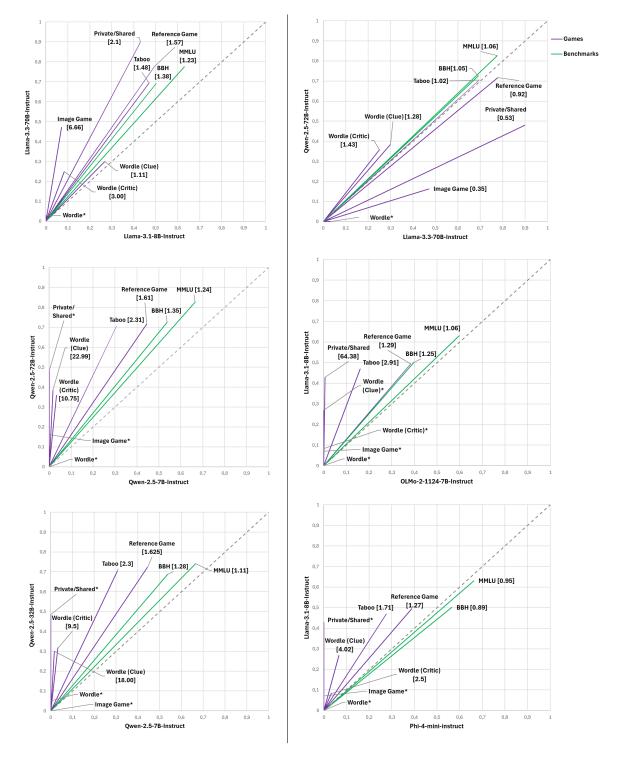


Figure 2: Comparing datasets in their power to discriminate between models of different size but same family (left) and of different families but similarly large (right). The number next to the benchmark's name indicates the ratio of performance between the two models. The asterisk '*' next to *Wordle* indicates that the ratio is undefined.

Results Figure 1 shows models' performance on the standard benchmarks and interactive games. While the OLMo-2 models are more inconsistent across different model sizes, all the other models exhibit the expected pattern of showcasing better performance on both benchmarks and interactive games as the parameter count increases. At the

same time, we observe that most of the interactive games highlight the benefits of larger model sizes much more strongly. This can more easily be appreciated in Figure 2 (left). In this visualisation, the further away a benchmark is from the diagonal, the more affected performance is by model size. While *Wordle* is an extremely challenging

game for any model, scaling up the number of parameters appears to be fundamental to succeed at *Private/Shared*, *Image Game*, and *Reference Game*—much more than what can be observed for MMLU and BBH.

Is size however all we need? Figure 2 (right) shows that reference-based benchmarks do not substantially distinguish between models of comparable size: Scaling on the number of parameters results in performance boosts across model families. Hence, arguably, reference-based benchmarks test for abilities that can be expressed within parametric knowledge. Given that such benchmarks are currently the standard LLM evaluation paradigm, it is not surprising that scaling is high on the agenda of model developers. In contrast, interactive games seem to provide a different picture: models with comparable parametric capacity perform very differently on Image Game, Private/Shared, and Wordle (Clue/Critic). Similar trends can be observed in other model comparisons (see Appendix I). This result supports the hypothesis that size is not all there is behind the potential of LLMs to learn inferential strategies for effective language use in interaction.

4 Cognitive Abilities

We now turn to targeted cognitive tests, which assess specific cognitive abilities deemed essential for effective language use in real-world situations. We explore the use of such tests as a complementary evaluation method to reference-based benchmarks and interactive games.

Ideally, such an investigation should be carried out using tests tailored to machine intelligence, which unfortunately have not yet been systematically developed (Hagendorff et al., 2024). To make progress in this direction, Millière and Rathkopf (2024) advocate for the development of an "artificial cognition evaluation toolkit" through a cyclic process that takes cognitive tests designed for humans as the starting point. Following this literature, we present a taxonomy of abilities and a suite of existing evaluation datasets inspired by human cognition, with the aim to make a first step towards more adequate testing of artificial cognition.

4.1 Taxonomy and Datasets

We present a taxonomy of cognitive abilities involved in human *functional linguistic competence* (Mahowald et al., 2024). It is guided by neurocognitive research (Ward, 2019), and it separates capa-

Cognitive Ability	Benchmark
Cognitive Flexibility	Kennedy and Nowak (2024)
Working Memory	Gong et al. (2024)
Inhibitory Control	_
Causal Reasoning	Jin et al. (2023)
Common-sense Reasoning	Sakaguchi et al. (2021)
Logical Reasoning	Liu et al. (2023)
Planning	Zheng et al. (2024)
	Cognitive Flexibility Working Memory Inhibitory Control Causal Reasoning Common-sense Reasoning Logical Reasoning

Table 1: Core and Higher-Order Executive Functions.

Cognitive Ability	Benchmark
Attribution and Judgement	Gu et al. (2025)
Emotional Intelligence	Paech (2023)
Pragmatics	Hu et al. (2023)
Social Common-sense Reasoning	Sap et al. (2019)
Theory of Mind	Gu et al. (2025)
Emotion Regulation	_ ` ` `
Self-Awareness	_
Empathy	_
Social Problem-Solving	_

Table 2: Social and Emotional Skills.

bilities into two distinct macro-categories known to recruit different brain networks: executive functions and socio-emotional skills. **Executive functions** are broadly defined as the complex processes by which we control and optimise our thoughts and behaviour (Baddeley, 1986), divided into *core* and *higher-order* abilities. **Socio-emotional skills** are abilities necessary to interact adaptively with other individuals (Higgins, 1987), including the ability to recognize their emotional and cognitive states.

For each cognitive ability, we select an existing evaluation dataset designed to test it in isolation. Socio-emotional skills have only recently entered the evaluation landscape in NLP, and they have done so with a forceful presence: remarkably, small benchmarks already exist for almost all of the abilities in this category. Tables 1 and 2 list the abilities in the taxonomy and the datasets we use to evaluate them.² More details are available in App. B and E.

4.2 Cognitive Ability Analysis

Equipped with our taxonomy and associated cognitive tests, we aim to shed some light on the cognitive abilities involved in interactive games and general benchmarks. Figure 3 reports Kendall's τ correlation coefficients. The analysis reveals that performance both on static and interactive evaluation correlates with performance on tests measur-

²We found no dataset to evaluate inhibitory control. The datasets we found for Emotion-regulation, Self-awareness (Liu et al., 2024), Empathy (Chen et al., 2024) and Social Problemsolving (Du et al., 2024) require human evaluation.

ing higher-order reasoning abilities: while planning is more dominant in static problem-solving tasks, working memory seems to be beneficial for games. Among the social skills, pragmatics appears to be relevant for both static and interactive tests, while ToM correlates better with the latter.

These results suggest that interactive tests correlate more strongly with socio-emotional skills compared to static tests, despite this analysis remaining speculative in nature, as we still lack carefully curated cognitive tests specifically designed for LLMs. Expanded results reporting correlation coefficients for each game and benchmark considered are shown in Appendix I, Figure 8. Specifically for games, in Appendix D we provide a qualitative analysis where we aim to draw tentative connections between capabilities and gameplay. Remarkably, such connections appear to be clearer for executive functions compared to socio-emotional skills.

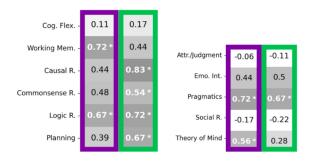


Figure 3: Correlation of cognitive abilities with performance on games and benchmarks (* = p < 0.05).

5 Related Work

A few recent papers have evaluated the abilities of LLMs beyond standard benchmarks. For example, Waldis et al. (2024) propose a framework to evaluate models' linguistic competence (in Chomsky 1965's sense) using probing classifiers to disentangle knowledge of, e.g., syntax and semantics from other abilities such as instruction following. Lu et al. (2024) posit that some perceived functional linguistic abilities (e.g., related to emotional intelligence) may be the result of in-context learning, model memory, and formal linguistic knowledge, whereas Ma et al. (2023) propose an evaluation protocol for assessing machine Theory of Mind based on a taxonomy of existing benchmarks driven by psychology research.

In contrast, the present work includes a uniquely comprehensive selection of evaluation tasks, encompassing standard benchmarks, interactive games, and cognitive ability assessments. To our knowledge, this diverse range of evaluations has not been jointly tested in prior work. Moreover, we systematically compare their discriminative power and analyse their interrelations—an approach that, to our knowledge, is unprecedented.

6 Conclusion

In this paper, we have examined three LLM evaluation paradigms. Our analyses revealed that interactive games have superior model discrimination power over standard reference-based benchmarks: Models with comparable parametric capacity perform similarly on standard benchmarks, but very differently on games. This suggests that scaling is not sufficient to learn strategies for effective language use in interaction.

We have further argued for the use of targeted cognitive tests as a complementary method to triangulate LLM progress. While each cognitive ability test alone does not get us very far in the quest for robust LLM evaluation, we contend that this type of evaluation paradigm has the potential to enhance our understanding of what fundamental abilities LLMs must develop to be able to function effectively as language agents, where multiple skills may be required and possibly interact. In our study, we have drawn inspiration from human cognitive tests as a first step in this direction. We hope to encourage the future development of targeted tests tailored to artificial cognition and supplemented with mechanistic analyses.

Limitations

The evaluation methods we have study in this papers have been applied without employing chain-of-thought (CoT) reasoning or similar elicitation techniques. While different elicitation strategies may enhance question-answering, interactive, and cognitive abilities in different ways (Yao et al., 2023; Hao et al., 2023; Li et al., 2024), we opted for an approach that remains agnostic to specific prompting methods and datasets. This ensures a consistent basis for comparison across models, though future work could explore how alternative prompting strategies influence performance across the three evaluation paradigms.

The set of models we evaluated is of course limited in scope, although representative of a wide range of open-weight model families and sizes. In particular, we did not evaluate any close-weight proprietary models, as we wanted to ensure full

reproducibility of our experiments.

Regarding the cognitive ability assessments, we have used currently available datasets; such resources have started to be compiled only very recently, hence there is no guarantee that in their current state they truly assess the intended cognitive abilities in LLMs. Nevertheless, they help in establishing a starting point and encouraging further progress in this direction.

Finally, our evaluation is only on English data, hence we cannot claim the results hold for any language. Yet, we believe the main take-home message of the paper, viz. using different evaluation methods which include static and interactive tasks as well as controlled cognitive experiments, holds independently of the language analysed. Whether this is indeed the case, remains an interesting open question for future research.

Acknowledgments

This work emerged from discussions at a workshop organised by Raquel Fernández and Sandro Pezzelle at MFO, the Oberwolfach Research Institute for Mathematics in the German Black Forest, on behalf of the ELLIS NLP programme. The event was funded by the state of Baden Württemberg (Germany) and organised in collaboration with the ELLIS Institute Tübingen and the Max Planck Institute for Intelligent Systems. Alessandro Suglia acknowledges the use of the Heriot-Watt University high-performance computing facility (DMOG) and associated support services in the completion of this work. Alessandro Suglia, Raffaella Bernardi, and David Schlangen acknowledge ISCRA for awarding this project access to the LEONARDO supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CINECA (Italy). We furthermore acknowledge our funders. Mario Giulianelli was supported in part by an ETH Zurich Postdoctoral Fellowship. David Schlangen acknowledges support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 423217434 ("RECOLAGE") grant, and support from University of Potsdam (President's Fund). Raquel Fernández was supported by the European Research Council (ERC Consolidator Grant DREAM 819455). This research has been partially supported and sponsored by the Province of Bolzano and the EU through the project Artificial Intelligence Laboratory (AI-Lab) ERDF 2021-2027 EFRE1047, CUP: I53C23001670009 (01/01/2024 - 31/12/2026).

References

- John Langshaw Austin. 1962. *How to do things with words*. Clarendon Press, London, UK.
- Alan Baddeley. 1986. *Working memory*. Oxford University Press.
- Leonardo Bertolazzi, Davide Mazzaccara, Filippo Merlo, and Raffaella Bernardi. 2023. ChatGPT's information seeking strategy: Insights from the 20-questions game. In *Proceedings of the 16th International Natural Language Generation Conference*, pages 153–162, Prague, Czechia. Association for Computational Linguistics.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. Experience grounds language. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8718–8735, Online. Association for Computational Linguistics.
- Kranti Chalamalasetti, Jana Götze, Sherzod Hakimov, Brielen Madureira, Philipp Sadler, and David Schlangen. 2023. clembench: Using game play to evaluate chat-optimized language models as conversational agents. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11174–11219, Singapore. Association for Computational Linguistics.
- Yuyan Chen, Songzhou Yan, Sijia Liu, Yueze Li, and Yanghua Xiao. 2024. EmotionQueen: A benchmark for evaluating empathy of large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2149–2176, Bangkok, Thailand. Association for Computational Linguistics.
- Noam Chomsky. 1965. *Aspects of the Theory of Syntax*. MIT Press.
- Herbert H Clark. 1996. *Using language*. Cambridge University Press, Cambridge, UK.
- Y. Du, P. Rajivan, and C. Gonzalez. 2024. Large language models for collective problem-solving: Insights into group consensus decision-making. In *Proceedings of the Annual Meeting of the Cognitive Science Society*.
- Technology Innovation Institute Falcon Team. 2024. The falcon 3 family of open models.
- Aryo Pradipta Gema, Joshua Ong Jun Leang, Giwon Hong, Alessio Devoto, Alberto Carlo Maria Mancino, Rohit Saxena, Xuanli He, Yu Zhao, Xiaotang Du, Mohammad Reza Ghasemi Madani, Claire Barale, Robert McHardy, Joshua Harris, Jean Kaddour, Emile van Krieken, and Pasquale Minervini. 2025. Are we done with MMLU? In *NAACL* 2025.
- Dongyu Gong, Xingchen Wan, and Dingmin Wang. 2024. Working memory capacity of ChatGPT: An

- empirical study. In Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence, AAAI'24/IAAI'24/EAAI'24. AAAI Press.
- D. A. Grant and E. A. Berg. 1948. Wisconsin card sorting test. *Journal of Experimental Psychology*.
- Aaron Grattafiori, Abhimanyu Dubey, and et al. Abhinav Jauhri. 2024. The Llama 3 herd of models. *Preprint*, arXiv:2407.21783.
- Yuling Gu, Oyvind Tafjord, Hyunwoo Kim, Jared Moore, Ronan Le Bras, Peter Clark, and Yejin Choi. 2025. SimpleToM: Exposing the gap between explicit ToM inference and implicit ToM application in LLMs.
- Thilo Hagendorff, Ishita Dasgupta, Marcel Binz, Stephanie C.Y. Chan, Andrew Lampinen, Jane X. Wang, Zeynep Akata, and Eric Schulz. 2024. Machine psychology.
- Shibo Hao, Yi Gu, Haodi Ma, Joshua Hong, Zhen Wang, Daisy Wang, and Zhiting Hu. 2023. Reasoning with language model is planning with world model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8154–8173, Singapore. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Tory Higgins. 1987. Social cognition and social perception. *Annual review of psychology*, 38(1):369–425.
- Jennifer Hu, Sammy Floyd, Olessia Jouravlev, Evelina Fedorenko, and Edward Gibson. 2023. A fine-grained comparison of pragmatic language understanding in humans and language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4194–4213, Toronto, Canada. Association for Computational Linguistics.
- Zhijing Jin, Yuen Chen, Felix Leeb, Luigi Gresele, Ojasv Kamal, Zhiheng Lyu, Kevin Blin, Fernando Gonzalez, Max Kleiman-Weiner, Mrinmaya Sachan, and Bernhard Schölkopf. 2023. CLadder: Assessing causal reasoning in language models. In *NeurIPS*.
- Sean M Kennedy and Robert D Nowak. 2024. Cognitive flexibility of large language models. In *ICML* 2024 Workshop on *LLMs* and Cognition.
- Zaijing Li, Gongwei Chen, Rui Shao, Yuquan Xie, Dongmei Jiang, and Liqiang Nie. 2024. Enhancing emotional generation capability of large language models via emotional chain-of-thought. *Preprint*, arXiv:2401.06836.

- Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. 2023. LogiQA 2.0—An improved dataset for logical reasoning in natural language understanding. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2947–2962.
- Ziyi Liu, Abhishek Anand, Pei Zhou, Jen-tse Huang, and Jieyu Zhao. 2024. InterIntent: Investigating social intelligence of LLMs via intention understanding in an interactive game context. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6718–6746, Miami, Florida, USA. Association for Computational Linguistics.
- Sheng Lu, Irina Bigoulaeva, Rachneet Sachdeva, Harish Tayyar Madabushi, and Iryna Gurevych. 2024. Are emergent abilities in large language models just incontext learning? In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5098–5139, Bangkok, Thailand. Association for Computational Linguistics.
- Ziqiao Ma, Jacob Sansom, Run Peng, and Joyce Chai. 2023. Towards a holistic landscape of situated theory of mind in large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1011–1031, Singapore. Association for Computational Linguistics.
- Kyle Mahowald, Anna A Ivanova, Idan A Blank, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. 2024. Dissociating language and thought in large language models: a cognitive perspective. *Trends in Cognitive Sciences*, 28.
- Microsoft, :, Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, Dong Chen, Dongdong Chen, Junkun Chen, Weizhu Chen, Yen-Chun Chen, Yi ling Chen, Qi Dai, Xiyang Dai, Ruchao Fan, Mei Gao, Min Gao, Amit Garg, Abhishek Goswami, Junheng Hao, Amr Hendy, Yuxuan Hu, Xin Jin, Mahmoud Khademi, Dongwoo Kim, Young Jin Kim, Gina Lee, Jinyu Li, Yunsheng Li, Chen Liang, Xihui Lin, Zeqi Lin, Mengchen Liu, Yang Liu, Gilsinia Lopez, Chong Luo, Piyush Madan, Vadim Mazalov, Arindam Mitra, Ali Mousavi, Anh Nguyen, Jing Pan, Daniel Perez-Becker, Jacob Platin, Thomas Portet, Kai Qiu, Bo Ren, Liliang Ren, Sambuddha Roy, Ning Shang, Yelong Shen, Saksham Singhal, Subhojit Som, Xia Song, Tetyana Sych, Praneetha Vaddamanu, Shuohang Wang, Yiming Wang, Zhenghao Wang, Haibin Wu, Haoran Xu, Weijian Xu, Yifan Yang, Ziyi Yang, Donghan Yu, Ishmam Zabir, Jianwen Zhang, Li Lyna Zhang, Yunan Zhang, and Xiren Zhou. 2025. Phi-4-mini technical report: Compact yet powerful multimodal language models via mixture-of-loras. Preprint, arXiv:2503.01743.
- Raphaël Millière and Charles Rathkopf. 2024. Anthropocentric bias and the possibility of artificial cogni-

- tion. In ICML 2024 Workshop on LLMs and Cognition
- Leora Morgenstern and Charles L. Ortiz. 2015. The winograd schema challenge: evaluating progress in commonsense reasoning. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI'15, page 4024–4025. AAAI Press.
- Samuel J. Paech. 2023. Eq-bench: An emotional intelligence benchmark for large language models. *Preprint*, arXiv:2312.06281.
- R. D. Rogers and S. Monsell. 1993. Costs of a predictible switch between simple cognitive tasks. *Journal of Experimental Psychology*, 124:207–231.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: an adversarial winograd schema challenge at scale. *Commun. ACM*, 64(9):99–106.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. Social IQa: Commonsense reasoning about social interactions. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- David Schlangen. 2023. What a situated languageusing agent must be able to do: A top-down analysis. *Preprint*, arXiv:2302.08590.
- John R Searle. 1969. *Speech acts: An essay in the philosophy of language*. Cambridge University Press, Cambridge, UK.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, and (445 additional authors not listed). 2023. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*. Featured Certification.
- Alessandro Suglia, Ioannis Konstas, and Oliver Lemon. 2024. Visually grounded language learning: a review of language games, datasets, tasks, and models. *Journal of Artificial Intelligence Research*, 79:173–239.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, and Jason Wei. 2023. Challenging BIG-bench tasks and whether chain-of-thought can solve them. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13003–13051, Toronto, Canada. Association for Computational Linguistics.
- Qwen Team. 2024. Qwen2.5: A party of foundation models.
- Raphael Vallat. 2018. Pingouin: statistics in python. *Journal of Open Source Software*, 3(31):1026.

- A. Waldis, Y. Perlitz, L. Choshen, Y. Hou, and I. Gurevych. 2024. Holmes: A benchmark to assess the linguistic competence of language models.
- Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafjord, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Michal Guerquin, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2024. 2 OLMo 2 Furious.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019a. Superglue: A stickier benchmark for general-purpose language understanding systems. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019b. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *International Conference on Learning Representations*.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyan Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhu Chen. 2024. Mmlu-pro: A more robust and challenging multi-task language understanding benchmark. In *Advances in Neural Information Processing Systems*, volume 37, pages 95266–95290. Curran Associates, Inc.
- Jamie Ward. 2019. *The student's guide to cognitive neuroscience*. Routledge.
- Ludwig Wittgenstein. 1953. *Philosophical Investigations*. Wiley-Blackwell, New York, NY, USA.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pages 38–45, Online. Association for Computational Linguistics.
- Michael Wooldridge and Nicholas R Jennings. 1995. Intelligent agents: Theory and practice. *The knowledge engineering review*, 10(2):115–152.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. 2024. Qwen2 technical report. arXiv preprint arXiv:2407.10671.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems*, volume 36, pages 11809–11822. Curran Associates, Inc.

Huaixiu Steven Zheng, Swaroop Mishra, Hugh Zhang, Xinyun Chen, Minmin Chen, Azade Nova, Le Hou, Heng-Tze Cheng, Quoc V. Le, Ed H. Chi, and Denny Zhou. 2024. Natural plan: Benchmarking llms on natural language planning.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

Jingming Zhuo, Songyang Zhang, Xinyu Fang, Haodong Duan, Dahua Lin, and Kai Chen. 2024. Prosa: Assessing and understanding the prompt sensitivity of llms. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1950–1976.

Appendix

A Models

The Phi-4 series (Microsoft et al., 2025) includes a model with 3.8 billion parameters. Despite its compact size, it has been reported to surpass larger models on a wide selection of tasks including reasoning, general-purpose and coding.

The Olmo-2-1124 series (Walsh et al., 2024) includes models with 7 billion and 13 billion parameters (OLMo-2-1124-*-Instruct). Both models are designed for a variety of tasks, including chat, mathematics, and reasoning. They have undergone supervised fine-tuning on the Tülu 3 dataset and further training using DPO techniques.

The Qwen2.5 series (Yang et al., 2024; Team, 2024) includes models with 7B, 32B, and 72B pa-

rameters (Qwen2.5-*-Instruct). They are multilingual, supporting over 29 languages, and excel in coding, mathematics, and instruction following.

The Llama-3 series (Grattafiori et al., 2024) includes models with 8B (Llama3.1-8B-Instruct) and 70B parameters (Llama3.3-70B-Instruct). These models are optimized for multilingual dialogue and support various languages. They use an optimized transformer architecture and are finetuned for instruction following.

The Falcon3 series (Falcon Team, 2024) includes a model with 10 billion parameters. It achieves state-of-the-art results in reasoning, language understanding, instruction following, code, and mathematics tasks. It supports four languages (English, French, Spanish, Portuguese) and a context length of up to 32K.

Model Implementations All the models used in this study have been made available by Hugging-face, and have been accessed through the *transformers* (Wolf et al., 2020) library. For text generation, we have been applying the default chat template specified by the models' Hugginface implementations.

B Taxonomy of Cognitive Abilities

EXECUTIVE FUNCTIONS:

• Core abilities

Working Memory: Hold and manipulate information in mind over short periods;

Inhibitory Control: Suppress automatic, inappropriate, or impulsive responses and resist distractions;

Cognitive Flexibility: Adapt to new situations, switch between tasks, and think about multiple concepts simultaneously;

• Higher-order abilities

Planning: Set goals, develop steps to achieve them, and anticipate potential obstacles;

Causal Reasoning: Understand cause-and-effect relationships;

Logical Reasoning: Deductive and inductive reasoning;

Common-sense Reasoning: Apply general common knowledge to everyday scenarios, including understanding basic physical properties, such as gravity, solidity, and object interaction;

SOCIO-EMOTIONAL SKILLS:

Social Common-sense Reasoning: Understand social norms and expectations;

Social Problem-Solving: Analyze social situations, generate solutions, and make decisions that foster positive interactions;

Emotional Intelligence: Recognize, interpret, and manage one's own and others' emotions.

Emotion Regulation: Manage and modify one's emotional responses in appropriate ways;

Self-Awareness: Recognize and understand one's own emotions, thoughts, and behaviors;

Empathy: Share and understand the feelings of others, both emotionally and cognitively;

Theory of Mind: Understand that others have thoughts, beliefs, desires, and intentions different from one's own;

Attribution and Judgment: Interpret the causes of others' behavior, distinguishing between intentional and unintentional actions.

Pragmatics: Aspects of communication that go beyond formal language competence: considering communicative intentions, the communicative context of the utterance, shared knowledge between speakers, manners, social and cultural norms.

C Interactive Games

We leverage *clembench* (Chalamalasetti et al., 2023), a benchmark that assesses models' gameplay ability in well-defined dialogue games such as:

Taboo: A game where one player tries to get the other player to guess a word without using certain 'taboo' words in their clues;

Wordle: A game where one player thinks of a 5-letter word and the other player tries to guess it. At each turn, the guesser receives feedback regarding whether each letter of the tentative word is present in the target word, and whether it is also in the same position in such word;

Wordle (Clue): A variant of Wordle where the guesser is given more clue to help them guess the target word;

Wordle (Critic): A variant of Wordle where the guesser's suggestions are evaluated by a third player, who provides feedback on their accuracy;

Reference Game: A game where one player is presented with three textual grids and is tasked to make another player—who is also presented with the same three grids, but potentially in a different order—identify the one it has in mind based on its textual description;

Private/Shared: A game where a customer agent goes through a form with a service agent and, after each turn, a third agent,³ probes the customer on what they believe that the service agent already knows;

Image Game: A game where given a certain grid structure similar to those in *Reference Game*, a player has to instruct a second one (who can't see the it) about how to reconstruct it step by step.

D Interactive Games: Samples

In this section we provide samples for the interactive games, and discuss their nature and links with cognitive abilities, as they are shown by the correlation analysis displayed in Figure 8. We provide samples deriving from our our evaluations however, as a reference, the clembench benchmark's maintainers made available a repository of transcripts for the conversations obtained by evaluating several models on the games.⁴

D.1 Taboo

Figure 10 shows an episode of Taboo played by Llama-3.1-70B-Instruct. In this case, the word to guess is *clear*, and the Describing player (on the left) is instructed to generate a description for the other player without using the taboo words *crystal*, *obvious* and *confident*. The description is then delivered to the guesser who guesses the word *glass* at first. Since this was not the correct word, the game proceeds for another round, when the guesser gets the right word.

As shown in Figure 8, this game presents a higher correlation with working memory, commonsense and logical reasoning among the executive

³In clembench, all interactions are mediated by a "Game Master". This agent plays a particularly active role in *Private/Shared*

⁴https://github.com/clembench/clembench-runs

functions, as well as pragmatics and Theory of Mind among the socio-emotional skills.

We can observe from the example interaction that the correlation results are quite compatible with the abilities required for being successful at the game. Common-sense reasoning is exercised for guessing the appropriate word from the descriptions, and for the generation of descriptions themselves, along with Theory of Mind which is necessary for tuning the descriptions with the guesser's knowledge that can be inferred from its answers. Logical reasoning can be instead exercised along with common-sense reasoning for guessing. Working memory can be used to avoid repeating the same answer and being able to remember useful information from the conversation, while the role of pragmatics remains less clear.

D.2 Wordle (and its variants)

Figure 11 shows an episode of Wordle. The model here is asked to guess the word *whine*. At each turn, the model receives letter-level feedback which has to exploit in order to get to the solution. We have evaluated models also on the variants providing additional clue (Figure 12) and the support of a critic model judging every guess (Figure 13). While correlation has not been measured on the classic Wordle variant due to the excessively low performance of models, we can interestingly observe differences in the correlation between the two variants and the cognitive capabilities. Results suggest that the presence of an external model stimulates the use of social capabilities.

D.3 Image Game

Figure 14 shows a match of Image Game, where a *describer* player has to describe a certain grid step by step to a second player which has to draw it by following the instruction. The describer has then to decide when the grid is complete. Given the high difficulty of the game for most of the models in this study, inferring links with cognitive capabilities appears complicated.

D.4 Reference Game

In Figure 15 is displayed an episode of Reference Game where a *describer* model has to describe a given grid with a empty cross-like shape. The other player has to guess the grid by following the description.

In our study, this task has been shown to correlate with Planning more than other cognitive ca-

pabilities. This connection appears reasonable: in order to build the best description a model may either describe the grid cell by cell, or it may look at the task as a constrain resolution problem where it has to resolve conflicts in order to provide the the simplest description that does not overlap with that of other grids. In this case, it is possible to observe analogies with schedule planning tasks such as those present in our planning benchmarks. In the case of other executive functions, working memory may be involved to remember the information regarding the grids within the context, and logical reasoning for guessing the appropriate grid.

Less clear are the correlations with social functions.

D.5 Private/Shared

Private/Shared is a game where two parallel conversations occur at the same time. Figure 16 shows the main conversation of an episode simulating a travel-agency interaction. Among the games taken into consideration, it is the most conversational task. At each turn, in a parallel conversation the player is questioned regarding its understanding of the state (private or public) of information.

The low correlation of this task with Theory of Mind is quite surprising and would be worth further investigations. On the contrary, it appears understandable its correlation with working memory superior to that of any other tasks, as conversations are generally longer and require more explicit recalling of information shared in prior messages.

E Benchmarks for Cognitive Abilities

Working Memory (Gong et al., 2024) (referred as WM in this work) is a set of verbal and spatial n-back tasks presented with three levels of difficulties from n=1 to n=3. The model has to identify whether the current stimulus (a letter in a string or a spatial location in a grid) is the same as the stimulus shown n messages back or not. Size: 50 blocks of 24 instances each, total: 1200;

Cognitive Flexibility (Kennedy and Nowak, 2024) (referred as *LLM-Cognitive-Flexibility* in this work) aims to test to what degree LLMs can rapidly switch tasks within a single context window. To this end, it employs two neuropsychological tests, the Wisconsin Card Sorting Test (WCST) (Grant and Berg, 1948) and the Letter-Number Test (LNT) (Rogers

and Monsell, 1993) commonly used to measure cognitive flexibility in humans. Size: 8 runs by 25 instances each. Total: 200;

Logical Reasoning LogiQA 2.0 (Liu et al., 2023)
This dataset evaluates logical reasoning using the same data both in NLI and Machine Reading Comprehension format (text, question, multiple-choice) for each of the following (deductive) reasoning types: categorical, sufficient condition, necessary condition, disjunctive, conjunctive reasoning. Size: 1572 data points;

Causal Reasoning CLADDER (Jin et al., 2023) focuses on formal causal reasoning (causal inference), as opposed to common-sense causal reasoning. The dataset is constructed from formal logic-based templates that are then verbalised into natural language as binary questions. Size: 10112 data points;

Common-sense Reasoning WinoGrande (Sakaguchi et al., 2021) A large-scale dataset of common-sense reasoning problems consisting of pairs of nearly identical questions with two answer choices (as in the original Winograd Schema Challenge (Morgenstern and Ortiz, 2015), from which this dataset takes inspiration but aims to make larger and more complex. Size: 1267 data points;

Planning NATURAL PLAN (Zheng et al., 2024) is a realistic planning benchmark consisting of three tasks expressed in natural language: Trip Planning, Meeting Planning and Calendar Scheduling. Models are given a situation and a problem to solve (e.g. find a trip plan that satisfies some constraints given the situation described). Each task contains problems of different levels of complexity based on the number of cities, people or days involved. The problems are often based on numerical reasoning too. We evaluate models on the Trip Planning and the Calendar Scheduling tasks. Size: 2600 data points;

Emotional Intelligence *EQ-Bench* (Paech, 2023) the model is given an emotionally charged short dialogue (generated by GPT-4) and has to score the four possible emotions felt by a given character. Scores are compared against a reference score. Size: 171 data points;

Pragmatics (Hu et al., 2023) (referred as LM-Pragmatics in this work) is a benchmark evaluating LLMs' understanding of seven pragmatics phenomena: deceit, indirect speech, irony, maxims, metaphor, humour, and coherence. Scenarios are grounded into social situations, requiring LLMs to interpret utterances. The task is designed as a multi-choice questionnaire with 2-5 questions based on the subtask. Size: 845 data points;

Social Common-sense SOCIAL IQA (Sap et al., 2019) a dataset for evaluating social commonsense reasoning and emotional intelligence. Each sample includes a short scenario and three multiple-choice questions across six categories: intentions, reactions, descriptions, motivations, needs, and consequences. Transfer learning on this dataset has shown strong performance on other common-sense reasoning benchmarks. Size: 1954 data points;

Attribution and Judgment/Theory of Mind

SimpleToM(Gu et al., 2025) contains concise, diverse stories each with questions that ask models to predict behavior ("Will Mary pay for the chips or report the mold?"), judgment ("Mary paid for the chips. Was that reasonable?") or mental states ("Is Mary likely to be aware that 'The can of Pringles has moldy chips in it.'? Yes or No?") The first two subtasks have been taken as a reference for the Attribution and Judgment cognitive ability, while the last as a reference for Theory of Mind. Size: 3450 data points.

F Benchmark Implementations

For the majority of the static benchmarks evaluated in this work we relied on the popular framework for the evaluation of LLMs *LM Evaluation Harness*⁵ (ver. 0.4.7), which already made available many of the selected benchmarks, and enabled a common interface for the implementation of most of the remaining ones.

The benchmarks which were already present within the framework are: SOCIAL IQA (Sap et al., 2019), WinoGrande (Sakaguchi et al., 2021), EQBench (Paech, 2023), LogiQA 2.0 (Liu et al., 2023), MMLU (Hendrycks et al., 2021), BBH (Suzgun et al., 2023). The benchmarks which have been

⁵https://github.com/EleutherAI/ LMEvaluationHarnessuation-harness

implemented in the framework over the course of the study are: CLADDER (Jin et al., 2023), LM-Pragmatics (Hu et al., 2023), SimpleToM (Gu et al., 2025), NATURAL PLAN (Zheng et al., 2024), IFE-val (Zhou et al., 2023).

As for the interactive games, we have used the implementation provided by version 1.6 of the *clembench* (Chalamalasetti et al., 2023). The remaining benchmarks (WM, LLM-Cognitive-Flexibility) have been implemented outside of the framework, as *LM Evaluation Harness* did not provide support for the multi-turn nature of the tasks.

F.1 Zero-shot and Few-shot Tasks

The majority of the tasks have been evaluated in a zero-shot setting with the exception of MMLU (5-shot), BBH (3-shot) (following common practices in model evaluation, e.g. in the Open Leaderboard 2⁶ for BBH) and NATURAL PLAN (5-shot). In the case of NATURAL PLAN, our models performed really poorly when evaluated in a zero-shot fashion—resulting in scores close to 0. Given that the task relies on the models producing answers in a strict format for parsing, we opted for using the 5-shot version provided by the benchmark's authors.

F.2 Metrics

Evaluation For the evaluation of models, we followed the original works' implementations as well as associated metrics. However, it may be the case that for a certain benchmark more metrics were defined, or that the original work did not aggregate results across subtasks. For this reason, we report here the metrics we used for evaluating models.

In the case of *Clembench games*, we computed performance by computing the ratio between the quality score (a number from 0 to 100) and the percentage of played games (a number between 0 and 1) divided by 100.

In the case of *IFEval*, following what was done in the Open Leaderboard 2, we averaged the results obtained on prompt-level and instruction-level strict accuracy.

As for *EQ-Bench*, we computed the task-specific score as it was implemented in the *LM Evaluation Harness*

Regarding *WM*, we only considered the subtask *Verbal N-3*, and we computed the accuracy for the results obtained across the 50 trials defined in the original work.

In the case of *LLM-Cognitive-Flexibility*, we ran each subtask 8 times with 25 trials each, and computed the average of the accuracy obtained in each run. In this case, the accuracy was computed only on the trials for which response parsing was successful. We then averaged the accuracy obtained on both subtasks to compute the final score.

In the case of CLADDER, we followed the original work which treated the task as generative and probed for the presence of the substrings "yes"/"no" at the beginning of the model's answer.

In NATURAL PLAN, the original work defined a rule-based procedure to parse specific data from the generated plan (e.g., dates). We reuse their parsing procedure and verify whether the expected elements are all present in the parsed plan.

For the remaining tasks (LogiQA 2.0, Wino-Grande, LM-Pragmatics, SOCIAL IQA, MMLU, BBH, SimpleToM), we treated them as a multiple-choice question answering task that is evaluated based on the likelihood of the correct answer for the task.

In the case of BBH, the Open Leaderboard 2's evaluation code excludes three of the original tasks from the overall score's computing: *dyck languages*, *navigate* and *word sorting*. The performances on these subtasks are therefore also ignored in the performance reported in this study.

In the case multiple subtasks were present (LM-Pragmatics, MMLU, BBH, NATURAL PLAN, LLM-Cognitive-Flexibility), we computed the micro-average over the results achieved on each subtask. In the specific case of SimpleToM, since the subtasks were associated with two different Cognitive Abilities, we've aggregated the score of the subtasks *behaviour* and *judgment* into a single score (under Attribution and Judgment), and considered the *mental state* subtask separately (under Theory of Mind).

Correlation For measuring the pair-wise correlation between benchmarks, we've computed the Kendall rank correlation coefficient (or Kendall's Tau) (Tau-b version). It measures rank correlation according to this formula:

$$\tau_b = \frac{P - Q}{\sqrt{(P + Q + T_x)(P + Q + T_y)}}$$

⁶https://huggingface.co/docs/leaderboards/ open_llm_leaderboard/about

where:

P = number of concordant pairs,

Q = number of discordant pairs,

 T_x = tie correction for variable X,

 T_y = tie correction for variable Y.

This method was preferable over others given its robustness in case of few data points, as it was in our case. We have also experimented with the Pearson correlation coefficient and observed that in the majority of the cases, the correlation patterns were similar, however with larger positive as well as negative correlations compared to Kendall. We've relied on the implementation provided by the *pingouin* Python package (Vallat, 2018).

F.3 Generation Settings

The tasks which required the models to generate text are: EQ-Bench, WM, BBH, IFEval, the clembench games, LLM-Cognitive-Flexibility, NATURAL PLAN, CLADDER. With the exception of Working Memory and LLM-Cognitive-Flexibility, all tasks have been evaluated by applying a temperature of 0. Following the original implementation, we have applied a temperature of 1 to WM and 0.7 to LLM-Cognitive-Flexibility. In these cases, however, the increased randomness caused by the higher temperature was mitigated by averaging the results obtained over multiple trials.

As for the other generation settings, we also have followed what was prescribed in the original works regarding the tokens for the termination of the generation, the maximum or minimum number of tokens. In the case of NATURAL PLAN, the original work did not provide specific information regarding the settings they have adopted for the evaluation. Given the highly challenging nature of the task, we have set the minimum and maximum number of tokens to 90 and 350, respectively. This was derived based on the minimum and maximum number of tokens in the gold plans.

G Limitations in the Evaluations

In certain cases, results have not been computed on all the subtasks available for that benchmark. In the case of NATURAL PLAN, we have not considered results coming from the *meeting* subtask, while for WM we have only considered those coming from the *Verbal (Base) N-3* subtask. In the first case, the high amount of resources required for evaluating

the task, especially for the larger models prevented us from doing so. As for the second, we've only considered the base version of the verbal subtask and excluded its variations as they would not provide meaningful information for this study.

H Computational Resources

As a reference, we provide the time required for running through all the benchmarks for the Llama-3.1-8B-Instruct model on 1 A100 GPU with batch size set to 'auto' in the *LM Evaluation Harness* (i.e. it automatically fits into the memory the maximum batch size possible for each task). For the Clembench games, LLM-Cognitive-Flexibility and WM, the batch size is 1. The time also includes that required for procedures performed by the *LM Evaluation Harness* prior to the actual evaluation (relevant only for those datasets included in this framework) and for loading the model into the memory (all tasks).

• LLM-Cognitive-Flexibility: ~1:50 min

• LogiQA 2.0: ~5 min

• CLADDER: ~19:30 min

• WinoGrande: ~1 min

• NATURAL PLAN: ~4:50 hours

• WM ~2:40 min

• EQ-Bench: ~3 min

• LM-Pragmatics: ~6:30 min

• SOCIAL IQA: ~1:30 min

• SimpleToM: ~2:40 min

• MMLU: ~14 min

• BBH: ~24 min

• IFEval: ~1:19 min

• Taboo: ~3:30 min

• Reference Game: ~3:00 min

• Image Game: ~2.40 min

• Wordle: ~7:50 min

• Wordle (Critic): ~2:50 min

• Wordle (Clue): ~2:15 min

• Private/Shared: ~17:30 min

I Additional Results

Table 3 includes the scores achieved by every model on each benchmark. Figure 4, Figure 5 and Figure 6 show the supplementary plots for the results in Section 3 (comparing models of different size but same family and models of similar size but coming from different families, respectively). Moreover, we provide supplementary plots for the results in Section 4. Figure 7 presents a direct

comparison of models based on our selected cognitive tests. Figure 8 reports an extended version of Figure 3. Finally, Figure 9 reports two scatter plots exemplifying, respectively, situations of high and low correlation between two benchmarks (a game-based one and a cognitive ability one).

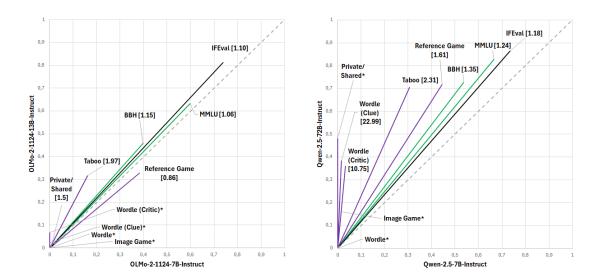


Figure 4: Comparing datasets in their power to discriminate models across size.

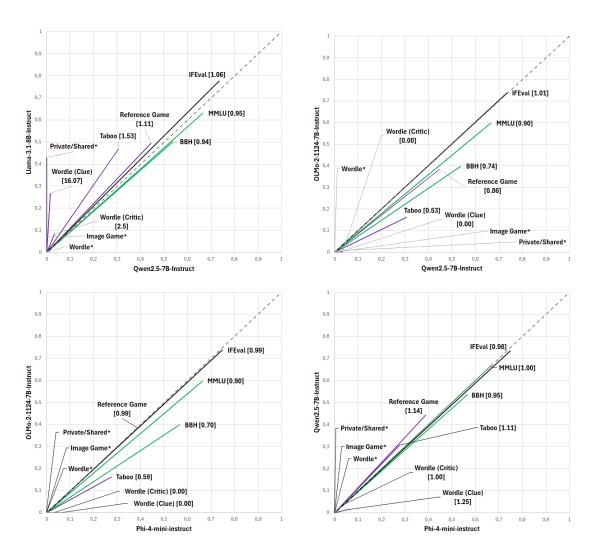


Figure 5: Comparing datasets in their power to discriminate models with similar size across families (Part 1).

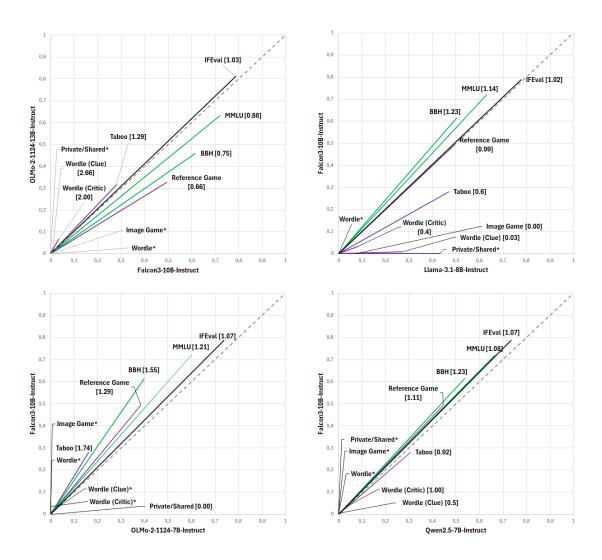


Figure 6: Comparing datasets in their power to discriminate models with similar size across families (Part 2).

Metric	OLMo-2-7B	OLMo-2-13B	Llama-3.1-8B	Llama-3.3-70B	Qwen2.5-7B	Qwen2.5-32B	Qwen2.5-72B	Falcon3-10B	Phi-4-mini
Cognitive Flexibility	0.2821	0.1500	0.4438	0.4339	0.2265	0.2008	0.4318	0.3615	0.3217
Working Memory	0.6383	0.3383	0.6692	0.8475	0.7175	0.7683	0.7900	0.3292	0.2125
Causal Reasoning	0.5034	0.4997	0.5035	0.5635	0.5329	0.5654	0.5856	0.5343	0.5489
Logical Reasoning	0.3181	0.3766	0.3391	0.4905	0.4027	0.4701	0.5592	0.4001	0.3384
Planning	0.0708	0.0613	0.0482	0.2666	0.0993	0.3428	0.3079	0.2147	0.0894
Commonsense Reasoning	0.5193	0.5193	0.5328	0.5754	0.5375	0.5793	0.5983	0.5107	0.5201
Social Reasoning	0.5404	0.5701	0.5000	0.5302	0.4458	0.4667	0.4672	0.5174	0.5015
Emotional Intelligence	0.6456	0.6484	0.6614	0.8290	0.5504	0.7949	0.7752	0.6405	0.6776
Theory of Mind	0.6085	0.9093	0.9503	0.8326	0.7742	0.9364	0.9538	0.8361	0.4708
Attribution/Judgment	0.5803	0.3710	0.3444	0.4795	0.2602	0.4207	0.2847	0.2642	0.4429
Pragmatics	0.5951	0.6634	0.6463	0.8402	0.6707	0.8146	0.8500	0.6646	0.5146
Taboo	0.1611	0.3167	0.4694	0.6944	0.3056	0.7028	0.7056	0.2806	0.2750
Reference Game	0.3833	0.3278	0.4944	0.7778	0.4444	0.7222	0.7167	0.4944	0.3889
Private/Shared	0.0067	0.0100	0.4292	0.8985	0	0.4851	0.4799	0	0
Image Game	0	0	0.0708	0.4710	0	0	0.1625	0	0
Wordle	0	0	0	0.0300	0	0.0500	0	0	0
Wordle (Clue)	0	0.0222	0.2678	0.2994	0.0167	0.3000	0.3833	0.0083	0.0667
Wordle (Critic)	0	0.0667	0.0833	0.2500	0.0333	0.3167	0.3583	0.0333	0.0333
MMLU	0.5973	0.6332	0.6305	0.7761	0.6651	0.7408	0.8264	0.7198	0.6627
ВВН	0.3980	0.4588	0.5013	0.6915	0.5369	0.6874	0.7256	0.6148	0.5662
IFEval	0.7378	0.8114	0.7758	0.9098	0.7344	0.8322	0.8644	0.7878	0.7485

Table 3: Scores obtained by models on each benchmark. The "Instruct" has been omitted from the column names for stylistic purposes.

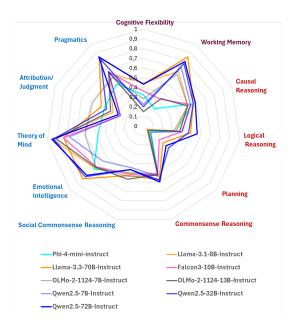


Figure 7: Cognitive Abilities Spectrum of LLMs

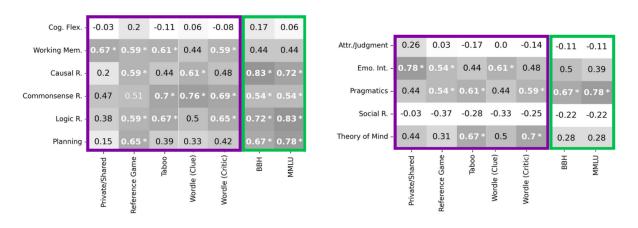


Figure 8: Correlation of social/emotional cognitive abilities with Interactive Games and Large QA benchmarks. The correlation matrix does not include results on Wordle and Image Game as model performances were too low.

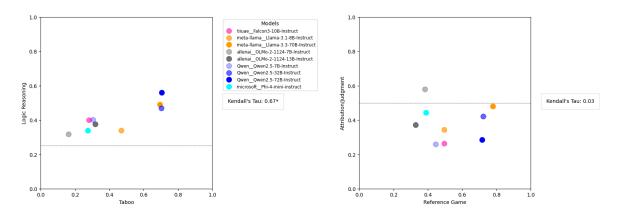


Figure 9: High (left) and low (right) correlation

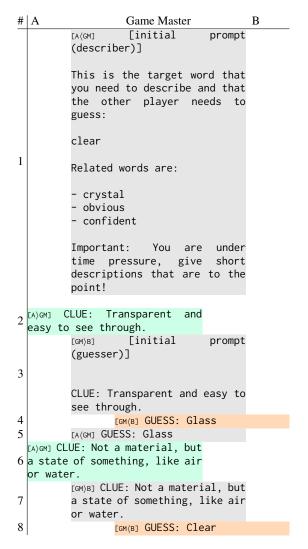


Figure 10: An example of an episode of Taboo played by Llama-3.3-70B-Instruct on both roles. The two players' messages are in green and orange, while the greycolored text comes from the environment itself, represented by the Game Master.

```
# | A
                  Game Master
                                        B
         [A(GM] [instruction prompt..]
 [A)GM] guess: house
 explanation: this is a common
 five-letter English word, and
 I am starting my guess with
 this word because it contains
 a mix of vowels and consonants
 that are frequently used in
 many words.
                        guess_feedback:
         [A(GM]
         h<yellow>
                       o<red>
                                 u<red>
         s<red> e<green>
3
         Provide your response only in
         this format.
         guess: word
         explanation: details
                 [GM|GM] [...]
 [A)GM] guess: whine
 explanation: I'm replacing the
 "t" with a "w" to form a
 new word, as the letters "h"
5 "i", "n", and "e" were all
 marked green, indicating they
 are in the correct positions,
 and "whine" is a valid English
 word that fits this pattern.
                 [GM|GM]
6
                 game_result
                 = WIN
```

Figure 11: An example of an episode of Wordle played by Llama-3.3-70B-Instruct. The player's messages are in green, while the grey-colored text comes from the environment itself, represented by the Game Master.

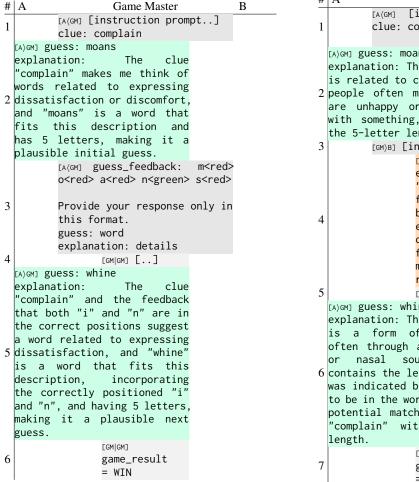


Figure 12: An example of an episode of Wordle (variant with Clue) played by Llama-3.3-70B-Instruct. The player's messages are in green, while the grey-colored text comes from the environment itself, represented by the Game Master.

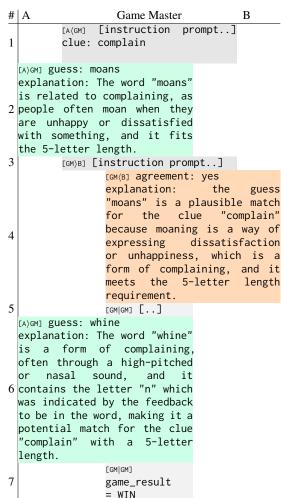


Figure 13: An example of an episode of Wordle (variant with Critic) played by Llama-3.3-70B-Instruct on both roles. The two players' messages are in green and orange, while the grey-colored text comes from the environment itself, represented by the Game Master.

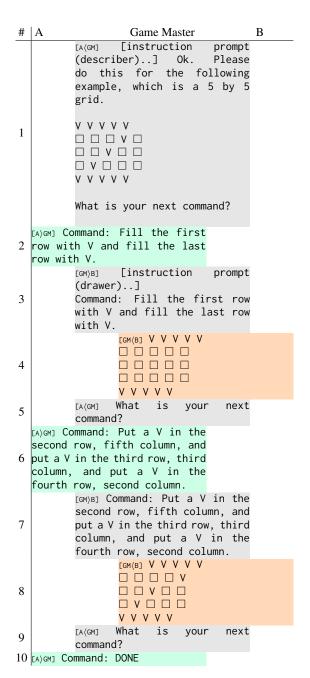


Figure 14: An example of an episode of Image Game played by Llama-3.3-70B-Instruct on both roles. The two players' messages are in green and orange, while the grey-colored text comes from the environment itself, represented by the Game Master.

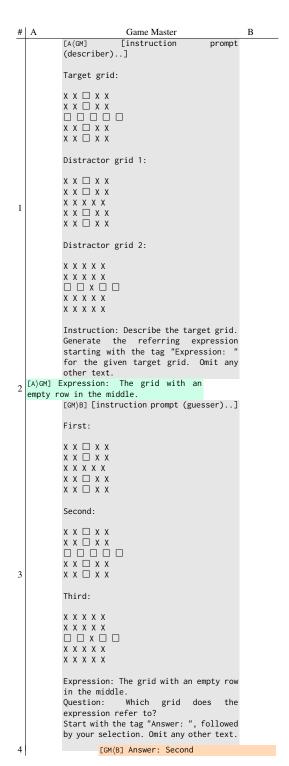


Figure 15: An example of an episode of Reference Game played by Llama-3.3-70B-Instruct on both roles. The two players' messages are in green and orange, while the grey-colored text comes from the environment itself, represented by the Game Master.

#	A Game Master B
1	[A(GM] [instruction prompt]
2	[GM)B] What is the next request?
3	[GM(B] Where do you want to go?
	[A(GM] TRAVEL-AGENT: Where do
4	you want to go? Answer only this question in one short sentence and then stop.
5	[A)GM] ANSWER: Dresden
6	[GM)B] Dresden
7	[GM/B] What is the next request?
8	[GM(B] How do you want to reach your destination?
	[A(GM] TRAVEL-AGENT: How do
9	you want to reach your destination? []
10	[A)GM] ANSWER: By overnight bus.
11	[GM)B] By overnight bus.
12	[GM)B] What is the next request?
	[GM(B] Where are you departing
13	from?
1.4	[A(GM] TRAVEL-AGENT: Where are
14	you departing from? []
15	[A)GM] ANSWER: Hamburg.
16	[GM/B] Hamburg .
17	[GM)B] What is the next request?
18	[GM(B] What kind of class do you
	wish?
19	[A(GM] TRAVEL-AGENT: What kind
	of class do you wish? []
20	[A\GM] ANSWER: The most
21	comfortable.
21	[GM/B] The most comfortable.
22	[GM)B] What is the next request?
23	[GM(B] When does your trip begin?
24	[A(GM] TRAVEL-AGENT: When does
	your trip begin? []
25	[A)GM] ANSWER: Third week of
26	June.
26	[GM)B] Third week of June.

Figure 16: An example of the main interaction of a Private/Shared episode played by Llama-3.3-70B-Instruct with a programmatic player. The two players' messages are in green and orange, while the grey-colored text comes from the environment itself, represented by the Game Master.

#	A Game Master B
1	[GM GM] Begin probing
2	[A(GM] ME: Do you think the travel agent knows where you depart from? Answer this question with yes or no and then stop.
3	[A\GM] ASIDE: Yes
4	[GM GM] yes [GM GM] Answer
5	for FROM valid after 1 tries.
6	[GM]GM] Answer is correct.
7	[A(GM] ME: Do you think the travel agent knows where you want to go to? []
8	[A)GM] ASIDE: Yes [GM GM] yes
	[GM GM] Answer
10	for TO valid after 1 tries. [GMGM] Answer is
11	correct.
12	[A(GM] ME: Is the travel agent aware of your means of transportation? []
	[A)GM] ASIDE: Yes
14 15	[㎝]ຜng yes [㎠]ຜng Answer for BY valid
16	after 1 tries. [GM GM] Answer is correct.
17	[A(GM] ME: Have you informed the travel agent about your class preference? []
18	[A)GM] ASIDE: Yes
19	[GM GM] yes [GM GM] Answer
20	for CLASS valid after 1 tries.
21	[GM]GM] Answer is correct. [A(GM] ME: Is the travel agent
22	aware of the dates of your trip? []
	[A)GM] ASIDE: Yes
2425	[GM]GM] yes [GM]GM] Answer for WHEN valid
26	after 1 tries. [GM GM] Answer is correct.
27	[GM]GM] End probing

Figure 17: An example of the final turn's probing interaction of a Private/Shared episode played by Llama-3.3-70B-Instruct with a programmatic player. The two players' messages are in green and orange, while the grey-colored text comes from the environment itself, represented by the Game Master.