

Opinion

Who Determines What Is Relevant? Humans or AI? Why Not Both?

A spectrum of human–artificial intelligence collaboration in assessing relevance.

TO MEASURE PROGRESS ON better methods for Web search, question answering, conversational agents, or retrieval from knowledge bases, it is essential to know which responses are relevant to a user’s information need. Such judgments of what is relevant are traditionally obtained by asking human assessors.

With the latest improvements on autoregressive large language models (LLMs) such as ChatGPT, researchers started to experiment with the idea of replacing human relevance assessment by LLMs.⁹ The approach is simple: Just ask an LLM chatbot whether a response is relevant for an information need, and it does provide an “opinion.”

In recent empirical studies on Web search³ but also in programming,⁷ human–computer interaction,⁵ or protein function prediction,¹⁰ it has been shown that LLM-generated opinions often agree with the assessment of humans. Some people readily believe the decision on what is relevant can be outsourced to artificial intelligence (AI) in the form of LLMs, without any human involvement.

However, as we argue here, there are several issues with such a fully automated judgment approach—and these issues cannot be overcome by a technical solution. Rather than continuing with the ongoing quest to



study where and how AI can replace humans, we suggest to examine forms of human–AI collaboration for which we lay out a spectrum in this column.

Why Not Just Use LLMs?

There are a number of issues that arise when we let LLMs judge the quality of search results or system-provided answers.

Judgment bias toward a particular LLM. If we use a particular LLM to cre-

ate relevance judgments to measure system quality, it would likely favor results from systems that use the same or a similar LLM for response generation. Such a bias in the gold standard benchmark can lead to wrong findings when comparing multiple systems for quality.

Bias toward user groups. Bender et al.¹ highlight the severe risk of LLMs to bias against underrepresented user groups. Such bias will likely be reflect-

ed in the relevance decisions made by the LLMs. Before using this technology, the computing community should develop approaches to quantify model bias and to understand possible ways of making LLMs more resilient when trained on biased data.

Resilience against misinformation.

Some information on the Web may seem topically relevant, but may be factually incorrect and hence should not be perpetuated. For example, on an information request like “Do lemons cure cancer?” a system response may discuss factually incorrect information about healing cancer with lemons. While on topic, such potentially harmful responses should not be presented to a user. Factuality is already difficult for humans to assess correctly and without additional resilience mechanisms in place against misinformation, an LLM is unlikely to make correct relevance decisions in such situations.

LLM-based LLM training. In a world where LLMs are used both for judging relevance and for generating responses, the issue of concept-drift also arises. Rather soon, a lot of Web content will be LLM-generated. At the same time, new LLMs may be trained using large amounts of Web content. This would lead to a cyclic learning problem, where possibly various LLMs agree on a definition of relevance that may not make sense to human end users.

Judging vs. predicting. When a strong LLM is used to create relevance judgments for training a system to produce relevant responses, another question arises: Why not directly use the judging LLM to produce the response? There could be arguments with respect to reduced model size or

improved response times, but still the trained system may not be able to surpass the quality of the judging LLM.

Truthfulness and hallucinations. A well-known issue of LLMs is that they tend to generate text that contains inaccurate or false information (that is, confabulations or hallucinations). Responses are often presented in such an affirmative manner that makes it difficult for humans to detect errors. While chain-of-thought reasoning⁸ or reinforcement from human feedback¹¹ can reduce the issue, it remains unclear to which extent the problem can be avoided.

LLM relevance judgments for training only. Even when LLM-generated relevance judgments are only used to train a system—but not to evaluate it—many of these issues still hold. Following the “garbage in/garbage out” mantra, issues arising from biased judgments, misinformation, and confabulations or hallucinations will affect the quality of the end user-facing system.

LLMs Are the New Crowdworkers

It is yet to be understood what the benefits and risks associated with LLM technology are—especially when it comes to creating gold standards. A rather similar debate was spawned more than 10 years ago when a lot of data annotations started to rely on crowdworkers instead of trained editors—with a substantial decrease in annotation quality somewhat compensated by a huge increase in annotated data. Quality-assurance methods for crowdworkers were developed to obtain reliable labels.² With LLMs, history may repeat itself: a huge increase in available relevance assessment data at a possibly decreased quality. However, the specific extent of the deterioration is still unclear and requires further study.

A related idea is to allow LLMs to learn by “observing” human relevance assessors or by following an active learning paradigm.¹³ Starting from LLM-generated relevance assessments that a human evaluates,¹⁷ the LLM could learn to provide better assessments. We believe humans working with LLMs is not only an option, but is likely unavoidable as shown by recent results indicating a large proportion of crowdworkers al-

We discuss a spectrum of options to combine human and machine intelligence in a complementary and collaborative fashion.

ready make use of LLMs to increase their productivity.¹⁵

A Spectrum of Human–LLM/AI Collaboration

Rather than exploring options for LLMs to replace humans, or reasons why LLMs should not be used, we discuss a spectrum of options to combine human and machine intelligence in a complementary and collaborative fashion.

The spectrum outlines different levels of collaboration. At one end, humans make judgments manually, while at the other end, LLMs replace humans completely. In between, LLMs assist humans with various degrees of interdependence or humans provide feedback to decision-making LLMs. A summary of our proposed levels of human–machine collaboration is shown in the accompanying table. Here, we discuss each level in detail.

Human judgment. On one extreme, humans make all relevance judgments manually without being influenced by an LLM. The relevance assessment interface only supports well-understood automatic features that do not require any form of automatic training/feedback. For instance, humans may decide which keywords should be highlighted during assessment, they may limit viewing a certain data subset, or they may sort the data in certain ways that influence their decision. This end of the spectrum thus represents the status quo practiced in the field of information retrieval and natural language processing, where humans are considered to be the only reliable arbiter.

Model in the loop. To make it easier for human assessors to decide on

Some people readily believe the decision on what is relevant can be outsourced to artificial intelligence in the form of LLMs.

relevance in a consistent manner, an advanced level of automatic support could be provided. For example, an LLM may generate a summary of a to-be-judged document and the human assessor then bases their relevance judgment on this compressed representation to complete the task in less time. Another approach could be to manually define information nuggets that are relevant¹² and to then train an LLM to automatically determine how many test nuggets are contained in the retrieved results (for example, via a QA system). We hope to see more research on helpful sub-tasks that can be taken over by LLMs, such as highlighting relevant passages and rationale generation.

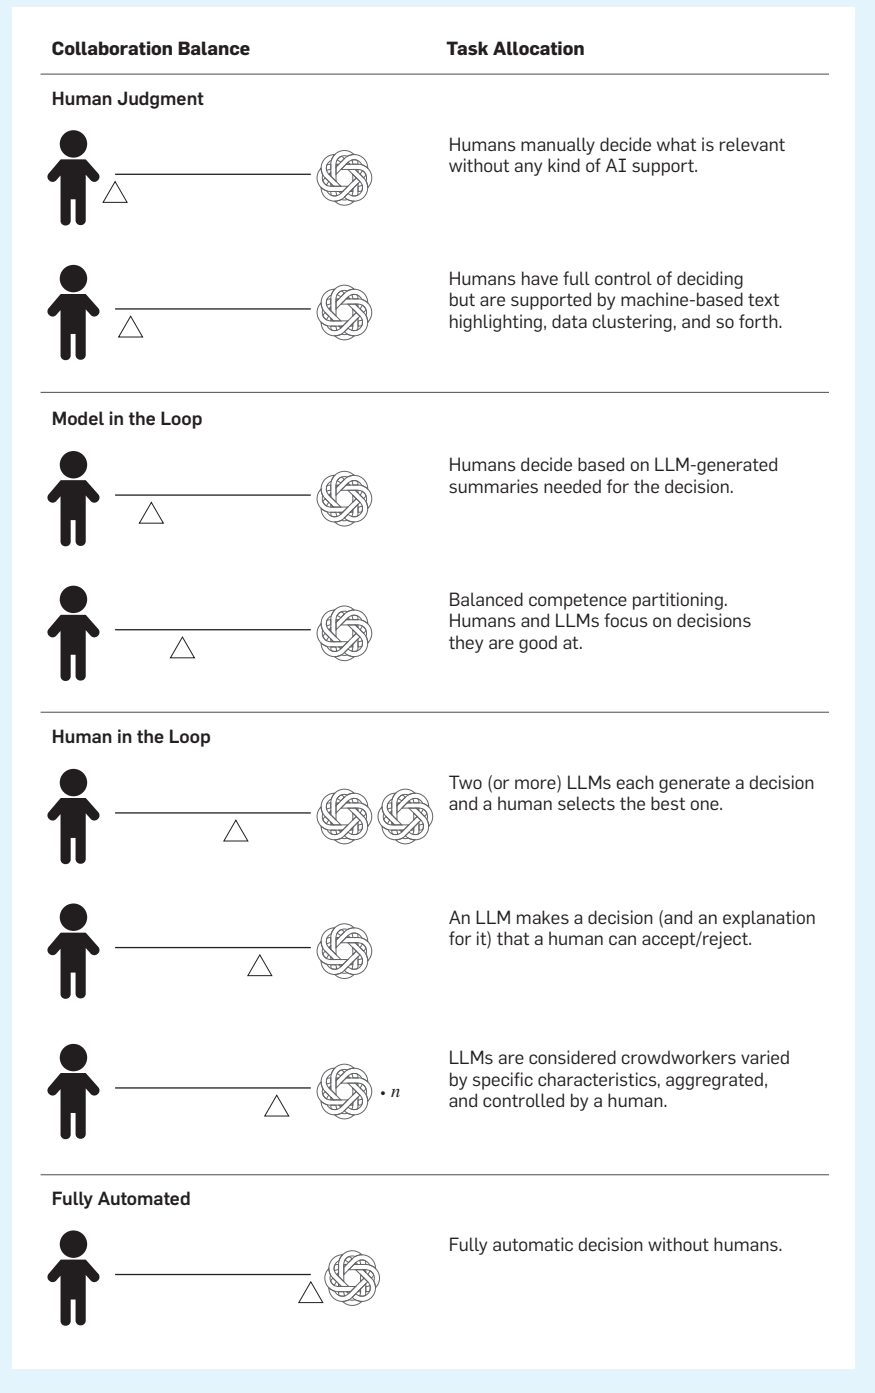
An important open question is: How to employ LLMs and other AI tools to assist human assessors in devising more reliable and faster relevance judgments?

Human in the loop. Automated judgments could be produced by an LLM and then verified by humans. For instance, a first-pass automatic relevance judgment could come with a generated natural language rationale based on which a human accepts or rejects the judgment, or, following the “preference testing” paradigm,¹⁶ two or more LLMs each could generate a judgment while a human will select the best one. In such cases, a human might possibly only intervene in the case of disagreements between the LLMs, thus increasing scalability. The purpose of this scenario is to simplify the decision for a human in most cases, and to use humans for difficult decisions or in situations where the LLMs generate a low-confidence decision.

Many issues identified in the field of explainability in machine learning apply to this scenario, such as the human tendency to over-rely on machines, or their inability to relate an LLM’s decision to its generated rationale.⁴ Thus, important open questions are: What are sub-tasks of the decision-making process that require human input (for example, prompt engineering¹⁴) and for what tasks should humans *not* be replaced by machines?

Fully automated. If LLMs were able to reliably judge relevance, they could completely replace humans when judging relevance. Indeed, a

Collaboration perspective: Spectrum of possibilities for collaborative human-machine task organization to make (relevance) decisions.
The Δ symbol indicates where on the spectrum each possibility falls.



recent study showed a good correlation between LLMs’ relevance judgments and human assessors,³ both, for an agreement on every judgment decision as well as to the correlation of leaderboards that rank systems by quality obtained with either set of judgments. Automatic relevance judgments might even surpass those of humans in terms of quality. However,

it is not entirely clear how to detect such super-human performance.

An important open question is: In which cases can human relevance judgments be replaced entirely by LLMs?

Interim conclusion. A central aspect to be investigated is where on this four-level human-machine collaboration spectrum one can obtain relevance decisions that are most cost-efficient,

INTERACTIONS



ACM's *Interactions* magazine explores critical relationships between people and technology, showcasing emerging innovations and industry leaders from around the world across important applications of design thinking and the broadening field of interaction design.

Our readers represent a growing community of practice that is of increasing and vital global importance.



To learn more about us, visit our award-winning website <http://interactions.acm.org>

Follow us on Facebook and Twitter  

To subscribe: <http://www.acm.org/subscribe>

Association for
Computing Machinery



We wish for more research on amplifying rather than replacing human intelligence using LLMs for judging the relevance of system responses.

fast, fair, and high in quality. In other words: How can one achieve ideal competence partitioning,⁶ where humans would perform tasks humans are good at, while machines perform tasks that machines are good at?

Conclusion

Our current understanding is not sufficient to let LLMs perform relevance judgments without human intervention. Furthermore, we wish for more research on amplifying rather than replacing human intelligence using LLMs for judging the relevance of system responses, especially with respect to “model in the loop” and “human in the loop” scenarios. To this end, we proposed a spectrum of possible ways in which we can balance human and artificial intelligence to increase the efficiency, effectiveness, and fairness in decision-making processes such as relevance assessment.

References

1. Bender, E.M. et al. On the dangers of stochastic parrots: Can language models be too big? In *FAccT '21: 2021 ACM Conference on Fairness, Accountability, and Transparency*, Virtual Event, (Toronto, Canada, Mar. 3–10, 2021). M.C. Elish, W. Isaac, and R.S. Zemel, Eds. ACM, 2021; 10.1145/3442188.3445922
2. Daniel, F. et al. Quality control in crowdsourcing: A survey of quality attributes, assessment techniques, and assurance actions. *ACM Comput. Surv.* 51, 1 (2018). 10.1145/3148148
3. Faggioli, G. et al. Perspectives on large language models for relevance judgment. In *Proceedings of the 2023 SIGIR Conference on the Theory of Information Retrieval (ICTIR)*, 2023; 10.1145/3578337.3605136
4. Fok, R. and Weld, D.S. In *Search of Verifiability: Explanations rarely enable complementary performance in AI-advised decision making*, 2023; arXiv Preprint arXiv:2305.07722.
5. Härmäläinen, P. et al. Evaluating large language models in generating synthetic HCI research data: A case study. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023; 10.1145/3544548.3580688
6. Hancock, P.A. Task partitioning effects in

semi-automated human-machine system performance. *Ergonomics* 56, 9 (2013); 10.1080/00140139.2013.816374

7. Kashеfi, A. and Mukerji, T. *ChatGPT for programming numerical methods*, 2023; arXiv Preprint arXiv:2303.12093.
8. Lyu, Q. et al. Faithful chain-of-thought reasoning. *CoRR*, 2023; 10.48550/arXiv.2301.13379.
9. MacAvaney, S. and Soldaini, L. One-shot labeling for automatic relevance estimation. In *SIGIR '23: The 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 2023; 10.1145/3539618.3592032
10. Madani, A. et al. Large language models generate functional protein sequences across diverse families. *Nature Biotechnology*, 2023; <https://www.nature.com/articles/s41587-022-01618-2>.
11. Ouyang, L. et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35, 5 (2022).
12. Sander, D. and Dietz, L. EXAM: How to evaluate retrieve-and-generate systems for users who do not (yet) know what they want. In *Proceedings of the 2nd International Conference on Design of Experimental Search & Information Retrieval Systems (DESIRE)*, 2021; <https://www.cs.unh.edu/~dietz/papers/sander2021exam.pdf>
13. Seo, S. et al. Active learning on pre-trained language model with task-independent triplet loss. In *Proceedings of the 34th Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The 12th Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, (Feb. 22–Mar. 1, 2022)*. AAAI Press, 2022; <https://ojs.aaai.org/index.php/AAAI/article/view/21378>
14. Sørensen, T. et al. An information-theoretic approach to prompt engineering without ground truth labels. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2022, Dublin, Ireland, May 22–27, 2022. S. Muresan, P. Nakov, and A. Villavicencio, Eds. Association for Computational Linguistics, 2022; 10.18653/v1/2022.acl-long.60
15. Veselovsky, V. et al. Artificial artificial intelligence: Crowd workers widely use large language models for text production tasks. *CoRR*, 2023; 10.48550/arXiv.2306.07899.
16. Windsor, J. et al. Preference testing: A comparison of two presentation methods. *Research in Developmental Disabilities* 15, 6 (1994); 10.1016/0891-4222(94)90028-0.
17. Xu, J. et al. On the role of human and machine metadata in relevance judgment tasks. *Information Processing and Management* 60, 2 (2023); 10.1016/j.ipm.2022.103177

Guglielmo Faggioli (guglielmo.faggioli@phd.unipd.it) is a postdoctoral researcher at the University of Padua, Padua, Italy.

Laura Dietz (dietz@cs.unh.edu) is an associate professor at the University of New Hampshire, Durham, NH, USA.

Charles L.A. Clarke (claclark@uwaterloo.ca) is a professor at the University of Waterloo, Ontario, Canada.

Gianluca Demartini (demartini@acm.org) is an associate professor at The University of Queensland, Brisbane, Australia.

Matthias Hagen (matthias.hagen@uni-jena.de) is a professor at Friedrich-Schiller-Universität Jena, Germany.

Claudia Hauff (claudiah@spotify.com) is a staff research scientist at Spotify, Delft, The Netherlands.

Noriko Kando (Noriko.Kando@nii.ac.jp) is a professor at the National Institute of Informatics (NII), Tokyo, Japan.

Evangelos Kanoulas (e.kanoulas@uva.nl) is a professor at the University of Amsterdam, The Netherlands.

Martin Potthast (martin.potthast@uni-leipzig.de) is a professor at Leipzig University and the ScaDS.AI Center for Scalable Data Analytics and Artificial Intelligence in Leipzig, Germany.

Benno Stein (benno.stein@uni-weimar.de) is a professor at Bauhaus Universität Weimar, Weimar, Germany.

Henning Wachsmuth (h.wachsmuth@ai.uni-hannover.de) is a professor at Leibniz Universität Hannover, Germany.