Perspective

Safeguarding authenticity for mitigating the harms of generative AI: Issues, research agenda, and policies for detection, fact-checking, and ethical AI

Ahmed Abdeen Hamed,[1,*] Malgorzata Zachara-Szymanska,[2] and Xindong Wu[3,*]

SUMMARY

As the influence of transformer-based approaches in general and generative artificial intelligence (AI) in particular continues to expand across various domains, concerns regarding authenticity and explainability are on the rise. Here, we share our perspective on the necessity of implementing effective detection, verification, and explainability mechanisms to counteract the potential harms arising from the proliferation of AI-generated inauthentic content and science. We recognize the transformative potential of generative AI, exemplified by ChatGPT, in the scientific landscape. However, we also emphasize the urgency of addressing associated challenges, particularly in light of the risks posed by disinformation, misinformation, and unreproducible science. This perspective serves as a response to the call for concerted efforts to safeguard the authenticity of information in the age of AI. By prioritizing detection, fact-checking, and explainability policies, we aim to foster a climate of trust, uphold ethical standards, and harness the full potential of AI for the betterment of science and society.

INTRODUCTION

The advent of generative artificial intelligence (AI) tools has generated a wave of both excitement and apprehension within the scientific community. While the potential benefits are immense, they are often overshadowed by significant challenges. Generative AI tools are renowned for their remarkable ability to produce compelling content. However, the authenticity of such content remains a pressing concern. Among the numerous AI tool providers, ChatGPT[1] and Google Bard[2] stand out as leading pre-trained transformer-driven tools with formidable generative capabilities (hence the GPT). Yet, at present, their statements lack the backing of credible references. This raises concerns about the factual accuracy of their convincing output. The scientific community bears the responsibility to distinguish generated content from genuine scientific endeavors, particularly when such material carries the risk of being used as false evidence in scientific publications.

Generative AI tools are trained using large language models (LLMs) on a vast array of online resources and other datasets. While it is not asserted that all generated content is erroneous, the extent of factual knowledge versus false claims remains unclear. Further complicating the matter, certain transformer-based tools, such as graph neural networks (GNNs), widely employed in drug discovery, provide drug recommendations without transparency or reproducibility trails. Reproducibility issues stem from instances where authors fail to elucidate their findings, accepting the black box[3] default without elaboration. Supporting or endorsing such research during the peer-review process is not considered sound practice, as it plays a critical role in the scientific production process. Accepting the current default behavior of these tools could impede future discoveries.

In essence, the emergence of generative AI tools has introduced a new frontier in scientific research, presenting both extraordinary opportunities and formidable challenges. The scientific community must navigate this delicate balance with prudence, ensuring that the benefits of these tools are harnessed responsibly while maintaining the integrity of scientific inquiry. To proceed, in next section, we present our perspective on three crucial aspects that we believe are mandatory to counter the issues mentioned: (1) the detection of inauthentic content and fake science; (2) computational fact-checking versus human verification; and (3) the transparency and reproducibility of science.

OUR PERSPECTIVE ON SAFEGUARDING AUTHENTIC SCIENCE

Inauthentic content detection

The alarming rise of fake science has ignited a call for action.[4,5] The scientific community must collectively embrace computational methods to detect inauthentic content. As the world united to combat the global pandemic (followed by an infodemic), now is the time for the scientific

---

[1]Network Medicine and AI, Sano Centre for Computational Medicine, Czarnowiejska 36 Building C5, 30-072 Cracow, Poland
[2]Faculty of International and Political Studies, Jagiellonian University, Wladys lawa Reymonta 4, 30-059 Cracow, Poland
[3]Zhejiang Lab, Research Center for Knowledge Engineering, Kechuang Avenue, Hangzhou, Zhejiang Province 311121, P.R. China
*Correspondence: a.hamed@sanoscience.org (A.A.H.), xwu@zhejianglab.com (X.W.)

community to accelerate the development of algorithms, methods, approaches, and heuristics capable of identifying generated and fake content. Prior to the advent of generative AI, research on fake news[6,7] and fake science[8,9] had already gained significant traction. However, the emergence of powerful generative tools underscores the need for scientists to intensify their efforts to combat misinformation, disinformation, and fake science that may arise from the misuse of these tools. Such endeavors are already underway. Abburi et al. proposed a novel ensemble approach for detecting AI-generated text.[10] Others have comprehensively surveyed the research landscape and revealed the potential of detecting AI-generated text.[11] The authors of this paper also played a pivotal role during the pandemic in identifying fake news and fake science, introducing a network-centric algorithm to address this issue.[12] Additionally, the authors of this paper tackled AI-generated text by introducing a groundbreaking algorithm trained on real scientific publications from the biomedical domain.[13]

Computational fact-checking vs. human verification

Although human verification of content is a much-needed task before establishing trust in AI-generated content, it is impossible for any human to cope with the massive volume and rapid pace of generation by such tools. What makes matters harder is the very convincing tone and appealing language used by all these tools. To demonstrate, we triggered ChatGPT, Bard, and Bing to generate their own versions of the article by Menczer et al.[14] Specifically, the task was triggered using the exact title and specified the exact three anchors stated in their communication (detection, moderation, and regulation). The three documents resulting from this trivial exercise were strikingly convincing and hit many points right on target. Three instances included three paragraphs clearly labeled after the three sections. More impressively, they also overlapped in relevant terms such as "inauthentic content" and "social media". Further, some of the tools used convincing terms such as "online platforms" and "spread misinformation" and sophisticated acronyms ("AI-generated inauthentic content [AIGC]"). It is impossible for any human to identify which one of the four documents was the authentic text of an original publication. While it is essential to verify the contents of generative AI tools, it is not feasible to declare human verification as a priority, as Van Dis et al. recommended.[3] Instead, here we offer a viable and possible solution that overcomes the bottleneck of the human expertise needed. By leveraging the human knowledge that is already documented in knowledge bases, gold-standard databases, and domain-specific ontologies, we can computationally construct ground truth that can offer fact-checking capabilities for the necessary verification.

Transparency and explainability

Transformer-based methods suffer from what is known as the black box problem.[3] This means that the results of such tools are not explainable and are sometimes hallucinatory. In the example of generative AI tools (e.g., ChatGPT, Google Bard), the end user can potentially receive different answers to the same questions without any justification, explanation, or references. In the effort of harnessing the potential of these tools, researchers must prioritize explainability and reproducibility while adhering to universal governmental orders (Biden's executive order) and ethical standards.[15,16] In addition, we strongly recommend the incorporation of explainer-type algorithms that are associated with modern transformer-based and GNN approaches.[17–19] These types of explainable algorithms are becoming increasingly important as they provide the transparency and trust that are needed in the results, especially within the biomedical field. For instance, Pfeifer et al. developed an explainable GNN approach for identifying disease subnetworks.[20] This demonstrates the potential of these algorithms to enhance the transparency and trustworthiness of transformer-based AI models. By addressing these challenges, researchers can harness the potential of generative AI while ensuring that its outputs are transparent and ethical.

RESEARCH AGENDA

Back-to-basics

We propose a "back-to-basics" approach that utilizes traditional algorithms known for their inherent traceability. These algorithms have a proven track record of effectiveness. We refer specifically to the most fundamental data mining algorithms that are inherently traceable. This means that each step of the process can be consistently replicated, whether it is performed by a machine or a human. Our previous work in identifying the most fundamental and influential data mining algorithms is documented in a living reference known as the Top-10 Data Mining Algorithms.[21] These algorithms can be employed for a wide range of computational tasks, including classification, clustering, ranking, boosting, association rule generation, and many others. To date, all of these algorithms remain widely taught in classrooms to students of all levels at numerous universities. They are also traceable, and their results are explainable. We advocate for this class of algorithms because their inherent traceability enables traversal mechanisms that facilitate the reproduction of results, both manually and programmatically. This class of algorithms is currently being investigated for detecting inauthentic content and fake science.[13]

Ontology-grounded truth

We believe that the integration of the biomedical ontology field is still in its early stages of development, and there is much work to be done to fully harness its potential in distinguishing between what is factual, what is hypothetical, and what is inauthentic, even if it appears credible. Additionally, exploring integrated ontology knowledge graphs adheres to the principles of network science and is therefore algorithmically traversable and computationally reproducible. Here, we present some examples of the scientific literature that has utilized biomedical ontologies as ground truth or empowered knowledge graphs that capture authentic knowledge using these ontologies.

(A) You

Please register these instructions:

Your role: is to act as a an expert system shell for mammography screening
Your will ask questions from the rules I provide you only
You will answer question also from the know I feed you only
When you ask question: be concise and ask within 15-20 words at most
When you answer: you will say Yes/No and indicate the answer concisely
You will then explain how you reasoned over the rules and show how the answer is derived

ChatGPT

Thank you for providing those instructions. They are now registered:

1. My role is to act as an expert system shell for mammography screening.
2. I will ask questions based on the rules you provide.
3. I will answer questions based on the knowledge you've provided.
4. When I ask questions, I will be concise and keep them within 15-20 words.
5. When I answer, I will use "Yes/No" and provide concise reasoning based on the rules and knowledge.

Feel free to proceed with any questions or scenarios related to mammography screening, and I'll follow these guidelines in our interactions.

Figure 1. Prompt-engineering ChatGPT to perform reasoning and produce explanation

Ontology-centric ground truth

The National Library of Medicine (NLM) often manually curates biomedical literature databases using ontology and taxonomy concepts due to their established authenticity.[22] Notably, the NLM employs two prominent taxonomies: (a) the Medical Subject Headings (MeSH)[23] and (b) the NLM National Center of Biotechnology Information (NCBI) Taxonomy.[24,25] Following the NLM's example, scientists have adopted ontology-based annotations as ground truth for identifying names and relationships in text.[26]
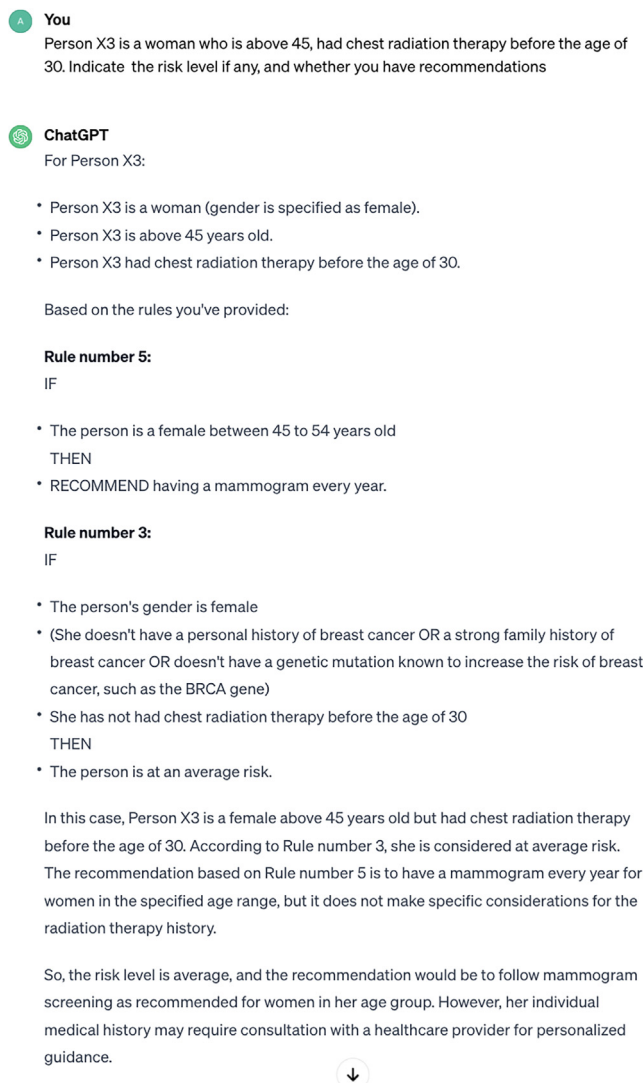
Ontology-driven knowledge graphs

Blagec et al. developed an ontology-oriented knowledge graph to be used as a benchmark and for AI-specific tasks.[27] Ontologies have also been instrumental in powering knowledge graphs to build data-driven systems focused on safety information sharing.[28] Due to their reliability, ontologies have further enriched knowledge graphs with text features extracted from the literature, enabling reasoning and discoveries, as demonstrated by Chen et al..[29] Additionally, ontologies played a crucial role during the global pandemic in constructing credible knowledge graphs for accessing and retrieving knowledge related to the pandemic.[30] Clearly, the combination of ontologies and knowledge graphs presents a powerful approach for establishing ground truth and countering fake science.

The integration of various ontologies to form knowledge graphs in the biomedical field offers a promising platform for capturing precise facts and concrete relationships. Although these methods are still under investigation, they have the potential to provide a rigorous mathematical framework for distinguishing between factual content and the hallucinations generated by generative AI tools such as ChatGPT.[31]

Explainable AI

Despite the inherent lack of explainability in generative AI tools such as ChatGPT, we conducted a prompt-engineering experiment to investigate whether this aspect could be enabled on demand when generating a response. Specifically, we prompt-engineered ChatGPT to function as an "expert system shell" by encoding publicly available knowledge as structured rules, such as the "American Cancer Association screening recommendations for women at average breast cancer risk" available online.[32] The prompt also required the request containing the response to provide an explanation of how the result was derived and which rules were triggered in making the recommendation. Figure 1 shows a screenshot of the ChatGPT prompt. The figure demonstrates how the prompt was engineered to guide the request and also demanded an explainable response that is derived solely from the entered rules and not from ChatGPT's default pre-training process. When

You

Person X3 is a woman who is above 45, had chest radiation therapy before the age of 30. Indicate the risk level if any, and whether you have recommendations

ChatGPT

For Person X3:

- Person X3 is a woman (gender is specified as female).
- Person X3 is above 45 years old.
- Person X3 had chest radiation therapy before the age of 30.

Based on the rules you've provided:

Rule number 5:

IF

- The person is a female between 45 to 54 years old
  THEN
- RECOMMEND having a mammogram every year.

Rule number 3:

IF

- The person's gender is female
- (She doesn't have a personal history of breast cancer OR a strong family history of breast cancer OR doesn't have a genetic mutation known to increase the risk of breast cancer, such as the BRCA gene)
- She has not had chest radiation therapy before the age of 30
  THEN
- The person is at an average risk.

In this case, Person X3 is a female above 45 years old but had chest radiation therapy before the age of 30. According to Rule number 3, she is considered at average risk. The recommendation based on Rule number 5 is to have a mammogram every year for women in the specified age range, but it does not make specific considerations for the radiation therapy history.

So, the risk level is average, and the recommendation would be to follow mammogram screening as recommended for women in her age group. However, her individual medical history may require consultation with a healthcare provider for personalized guidance.

Figure 2. Prompt-engineering ChatGPT to perform reasoning and produce explanation

we put the rules to the test using a hypothetical case scenario, ChatGPT accurately assessed the situation, triggered the corresponding rules, and provided a detailed explanation of how the answer was derived. This experiment showcased not only the tool's remarkable intelligence but also the significant role of prompt engineers in creatively utilizing generative AI responsibly. Figure 2 presents the hypothetical scenario and the specific rules that ChatGPT activated until the correct result emerged, accompanied by a comprehensive explanation based on the exact rules we entered and triggered accordingly. Here, we provide the entire prompt session, rules entered, and how they were invoked before the response was explained.[33] In the absence of such knowledge and structured rules, ChatGPT fell back on its pre-trained, unknown dataset, generating generic responses without providing any explanations. Figure 3 illustrates ChatGPT's default response when lacking the rules and knowledge encoding, relying solely on its pre-trained model.

With these experiments, we offer the following guidelines to advance the explainability aspect of the research agenda: (1) researchers and users of ChatGPT, Bard, and other generative AI tools should focus on the creativity of prompt engineering to insist on the explainability of such tools; (2) developers of ChatGPT and Google Bard should integrate built-in explainability mechanisms. Scientific workflows, like Pegasus WMS, serve as examples of how explainability can be achieved.[34] The generative AI industry, in particular, should advance the capabilities of such tools to include a referencing mechanism where the knowledge sources are appropriately cited; (3) researchers and users of transformer-oriented tools and methods (e.g., GNN) should seriously consider the use of explainable algorithms and actively contribute to advancing this research area[17,20]; (4) during the peer-review process, reviewers should be vigilant in assessing the necessary transparency that fosters confidence in the results and demonstrates reproducibility, especially in research related to health and medicine; and (5) scientific journals and publishers should establish guidelines that elevate standards and mandate transparency and reproducibility in scientific research.
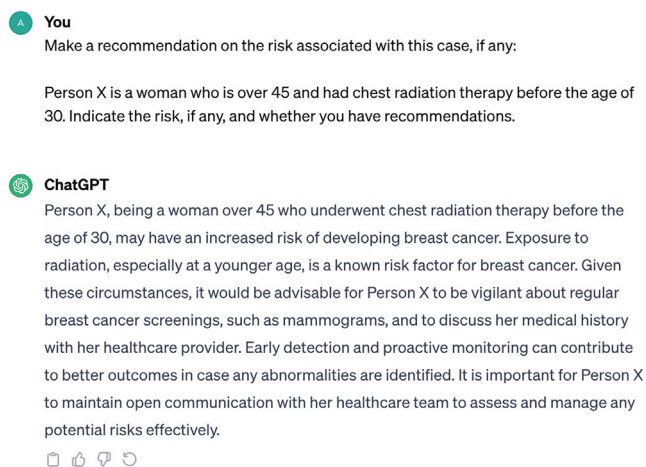
Figure 3. Querying ChatGPT using its default pre-trained engine without encoding of knowledge into structured rules and not demanding explanations of how the answer is derived

POLICY PRIORITIES

This study addresses several pressing issues in the wake of emerging generative AI tools: (1) the imperative to identify inauthentic content; (2) performing fact-checking; and (3) the critical need for computational explainability and result verification. However, the current stage of AI technology development is associated with a number of ethical concerns, of which bias, privacy, accountability, and transparency are the leading ones. While technology development is fundamental to increasing the social benefits of this technology's application, especially given the ongoing international competition for strategic innovation, the use of AI in scientific ecosystems has to be accompanied by widespread adoption of ethical standards. Otherwise, the uncertainty factor well identified with AI itself can easily be transmitted to the scientifically published research outcomes created with the use of technology.[35] Therefore, the approach proposed earlier goes far beyond the technical aspects of scientific verification infrastructure and should be seen in the broader context of the social reception of science. As a result, we call for the following essential policies to address both the technical and ethical issues associated with the emerging technologies of generative AI and LLMs.

The first policy, is leveraging both foundational algorithms and advanced explainer-type algorithms, specifically tailored for elucidating the inner workings of transformer-based tools. By utilizing computational detection methods and encouraging explainable AI, we strengthen our ability to distinguish real information from fake, creating a strong defense against the spread of false information.

The second policy, we highlight the significance of harnessing the rich knowledge of our knowledge bases, encompassing ontologies and gold standards. We advocate for seamless resource integration, the adoption of universally accepted sources of ground truth, and the innovation of automated fact-checking mechanisms. This concerted effort seeks to establish strong boundaries between reality and fiction, supporting the integrity of information disseminated through the utilization of generative AI tools.

The third policy, is for ethical AI is suggested where we seize the new opportunities ethically. Clearly, in the current era of AI and new advancement, a wide range of unprecedented opportunities emerge, alongside an array of hard challenges.[36] In order for the various communities to seize the opportunity, we must observe some of the foundational ethical standards outlined by Muller et al., who provide such aspects in their work titled "The Ten Commandments of Ethical Medical AI".[16] We hope that such policy will extend to the institutional level, where each organization mandates customized ethical standards pertaining to the specific purpose, research, and industry.

CONCLUSION

In this rapidly evolving research landscape, our recommendations described in the research agenda to support the computational detection of generated content, alongside computational fact-checking and explainability mechanisms, stand as a robust solution to counter the challenges posed by AI-generated tools and methodologies. Through the insights gained from our present work, and drawing from the experiences of the previous generation of conversational (non-generative) AI and chatbots,[11,37] we strive to safeguard authenticity while we foster positive experiences and establish trust. Our goal is to seize the opportunities presented by these powerful tools ethically and responsibly, marking an era characterized by the dissemination of true knowledge and reproducible science.

Union under the European Regional Development Fund and the National Natural Science Foundation of China (NSFC) under grant 62120106008.

AUTHOR CONTRIBUTIONS

A.A.H. conceived the idea, performed experiments, written the manuscript, and revised the manuscript. X.W. validated the ideas, written the manuscript, and revised the manuscript. M.Z.-S. contributed to the 2nd version of the manuscript in response to reviewers.

DECLARATION OF INTERESTS

The authors declare no competing interests.

REFERENCES

1. OpenAI ChatGPT. https://chat.openai.com/.
2. Google bard. https://bard.google.com/.
3. von Eschenbach, W.J. (2021). Transparency and the black box problem: Why we do not trust ai. Philos. Technol. 34, 1607–1622.
4. Else, H. (2022). mill detector put to the test in push to stamp out fake science. Nature 612, 386–387.
5. Van Noorden, R. (2023). How big is science's fake-paper problem? Nature 623, 466–467.
6. Lazer, D.M.J., Baum, M.A., Benkler, Y., Berinsky, A.J., Greenhill, K.M., Menczer, F., Metzger, M.J., Nyhan, B., Pennycook, G., Rothschild, D., et al. (2018). The science of fake news. Science 359, 1094–1096.
7. Allcott, H., and Gentzkow, M. (2017). Social media and fake news in the 2016 elec-tion. J. Econ. Perspect. 31, 211–236.
8. Hopf, H., Krief, A., Mehta, G., and Matlin, S.A. (2019). Fake science and the knowl-edge crisis: ignorance can be fatal. R. Soc. Open Sci. 6, 190161.
9. Lukić, T., Blešić, I., Basarin, B., Ivanović, B.L., Milošević, D., and Sakulski, D. (2014). Predatory and fake scientific journals/publishers: A global outbreak with rising trend: A review. Geogr. Pannon. 18, 69–81.
10. Abburi, H., Roy, K., Suesserman, M., Pudota, N., Veeramani, B., Bowen, E., and Bhattacharya, S. (2023). A simple yet efficient ensemble approach for ai-generated text detection. Preprint at arXiv. https://doi.org/10.48550/arXiv.2311.03084.
11. Ghosal, S.S., Chakraborty, S., Geiping, J., Huang, F., Manocha, D., and Bedi, A.S. (2023). Towards possibilities & impossibilities of ai-generated text detec-tion: A survey. Preprint at arXiv. https://arxiv.org/pdf/2304.04736.
12. Abdeen, M.A.R., Hamed, A.A., and Wu, X. (2021). Fighting the covid-19 infodemic in news articles and false publications: The neonet text classifier, a supervised machine learning algorithm. Appl. Sci. 11, 7265.
13. Hamed, A.A., and Wu, X. (2023). Improving detection of chatgpt-generated fake science using real publication text: Introducing xfakebibs a supervised-learning network algorithm. Preprint at arXiv. https://arxiv.org/abs/2308.11767.
14. Menczer, F., Crandall, D., Ahn, Y.-Y., and Kapadia, A. (2023). Addressing the harms of ai-generated inauthentic content. Nat. Mach. Intell. 5, 679–680.
15. (2023). Fact sheet: President Biden issues executive order on safe, secure, and trustworthy artificial intelligence. https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-s. accessed on.
16. Muller, H., Mayrhofer, M.T., Van Veen, E.B., and Holzinger, A. (2021). The ten commandments of ethical medical ai. Computer 54, 119–123.
17. Luo, D., Cheng, W., Xu, D., Yu, W., Zong, B., Chen, H., and Zhang, X. (2020). Parameterized explainer for graph neural network. Preprint at arXiv. https://doi.org/10.48550/arXiv.2011.04573.
18. Ying, R., Bourgeois, D., You, J., Zitnik, M., and Leskovec, J. (2019). Gnnex-plainer: Generating explanations for graph neural networks. Preprint at arXiv. https://doi.org/10.48550/arXiv.1903.03894.
19. Kokhlikyan, N., Miglani, V., Martin, M., Wang, E., Alsallakh, B., Reynolds, J., Melnikov, A., Kliushkina, N., Araya, C., Yan, S., and Reblitz-Richardson, O. (2020). Captum: A unified and generic model interpretability li-brary for pytorch. Preprint at arXiv. https://doi.org/10.48550/arXiv.2009.07896.
20. Pfeifer, B., Saranti, A., and Holzinger, A. (2022). Gnn-subnet: disease subnetwork detection with explainable graph neural networks. Bioinformatics 38, ii120–ii126.
21. Wu, X., Kumar, V., Ross Quinlan, J., Ghosh, J., Yang, Q., Motoda, H., McLachlan, G.J., Ng, A., Liu, B., Yu, P.S., et al. (2008). Top 10 algorithms in data mining. Knowl. Inf. Syst. 14, 1–37.
22. Beretta, V., Harispe, S., Ranwez, S., and Mougenot, I. (2016). How can ontologies give you clue for truth-discovery? an exploratory study. In Proceedings of the 6th International Conference on Web Intelligence (Mining and Semantics), pp. 1–12.
23. Lipscomb, C.E. (2000). Medical subject headings (mesh). Bull. Med. Libr. Assoc. 88, 265.
24. Federhen, S. (2012). The ncbi taxonomy database. Nucleic Acids Res. 40, D136–D143.
25. Schoch, C.L., Ciufo, S., Domrachev, M., Hotton, C.L., Kannan, S., Kho-vanskaya, R., Leipe, D., Mcveigh, R., O'Neill, K., Robbertse, B., et al. (2020). Ncbi taxonomy: a comprehensive update on curation, resources and tools. Database 2020, baaa062.
26. Suravee, S., Stoev, T., Schindler, D., Hochgraeber, I., Pinkert, C., Holle, B., Halek, M., Krüger, F., and Yordanova, K. (2022). Annotation scheme for named en-tity recognition and relation extraction tasks in the domain of people with dementia. In 2022 IEEE International Conference on Pervasive Computing and Communications Workshops and other Affiliated Events (PerCom Workshops) (IEEE), pp. 236–241.
27. Blagec, K., Barbosa-Silva, A., Ott, S., and Samwald, M. (2022). A curated, ontology-based, large-scale knowledge graph of artificial intelligence tasks and benchmarks. Sci. Data 9, 322.
28. Pedro, A., Pham-Hang, A.-T., Nguyen, P.T., and Pham, H.C. (2022). Data-driven construction safety information sharing system based on linked data, ontologies, and knowledge graph technologies. Int. J. Environ. Res. Publ. Health 19, 794.
29. Chen, H., and Luo, X. (2019). An automatic literature knowledge graph and rea-soning network modeling framework based on ontology and natural lan-guage processing. Adv. Eng. Inf. 42, 100959.
30. Zheng, X., Xiao, Y., Song, W., Tong, F., Liu, S., and Zhao, D. (2021). Covid19-obkg: an ontology-based knowledge graph and web service for covid-19. In 2021 IEEE International Conference on Bioinformatics and Biomedicine (BIBM) (IEEE), pp. 2456–2462.
31. Hamed, A.A., Lee, B.S., Crimi, A., and Misiak, M.M. (2023). Challenging the ma-chinery of generative ai with fact-checking: Ontology-driven biological graphs for verifying human disease-gene links. Preprint at arXiv. https://doi.org/10.48550/arXiv.2308.03929.
32. Society, A.C. (2023). American cancer society recommendations for the early detection of breast cancer. https://www.cancer.org/cancer/types/breast-cancer/screening-tests-and-ear.
33. Hamed, A.A. (2023). Prompt-engineering chatgpt an explainble intelligent system. https://chat.openai.com/c/62c9cdd1-b2c5-4036-9136-5ad20f3841dd.
34. Deelman, E., Vahi, K., Juve, G., Rynge, M., Callaghan, S., Maech-ling, P.J., Mayani, R., Chen, W., Ferreira da Silva, R., Livny, M., and Wenger, K. (2015). Pegasus, a workflow management system for science automation. Future Generat. Comput. Syst. 46, 17–35.
35. Cave, S., Craig, C., Dihal, K., Dillon, S., Montgomery, J., Singler, B., and Tay-lor, L. (2018). Portrayals and Perceptions of Ai and Why They Matter.
36. Tian, S., Jin, Q., Yeganova, L., Lai, P.-T., Zhu, Q., Chen, X., Yang, Y., Chen, Q., Kim, W., Comeau, D.C., et al. (2023). Opportunities and challenges for chatgpt and large language models in biomedicine and health. Preprint at arXiv. https://doi.org/10.48550/arXiv.2306.10070.
37. Gkinko, L., and Elbanna, A. (2022). Hope, tolerance and empathy: employees' emo-tions when using an ai-enabled chatbot in a digitalised workplace. Inf. Technol. People 35, 1714–1743.