AI4Science and the Context Distinction

Moti Mizrahi

Florida Institute of Technology

Abstract

"AI4Science" refers to the use of Artificial Intelligence (AI) in scientific research. As AI systems become more widely used in science, we need guidelines for when such uses are acceptable and when they are unacceptable. To that end, I propose that the distinction between the context of discovery and the context of justification, which comes from philosophy of science, may provide a preliminary but still useful guideline for acceptable uses of AI in science. Given that AI systems used in scientific research are black boxes, for the most part, we should use such systems in the context of discovery but not in the context of justification. The former refers to processes of idea generation, which may be unproblematically opaque whether they occur in human brains or artificial neural networks, whereas the latter refers to scientific methods by which scientific ideas are tested, confirmed, verified, and justified, which should be transparent.

Keywords: AI4Science, context distinction, discovery, explicability, justification

1. Introduction

"AI4Science" refers to the use of Artificial Intelligence (AI) in scientific research. For instance, the stated goal of the AI4Science Lab at Oxford University is to make "contributions to solve important problems in [particle physics, heliophysics, astrobiology, Earth science, and computational social science] through application and development of AI methods" [1]. Likewise, at the California Institute of Technology, the aim of AI4Science there is "to bring together AI researchers with experts from other disciplines to push modern AI tools into every area of science and engineering" [2].

Many scientists are optimistic about the impact of AI systems on science. They believe that the introduction and use of AI in "every area of science" accelerates scientific progress. For example, Yongjun Xu et al. (2023) argue that AI has, and will continue to, accelerate progress in scientific research [3]. Likewise, according to Wang et al. (2023), "Recent advances, including the successful unraveling of the 50-year-old protein-folding problem and AI-driven simulations of molecular systems with millions of particles, demonstrate the potential of AI to address challenging scientific problems" [4]. They add that "Cross-disciplinary collaborations are crucial to realize a comprehensive and practical approach towards advancing science through AI" [4].

As Austin Clyde (2022) points out, however, if AI4Science is to be an engine of scientific progress, a few hurdles must be overcome [5]. One of these hurdles is what he calls "an orientation toward cogent justification" [5]. This is a challenge because many AI systems are opaque, which is to say that their innerworkings are unexplainable by human scientists, which in turn could make their outputs (e.g., decisions, predictions, etc.) uninterpretable by human scientists. Take AlphaFold, for instance, which proponents of AI4Science often cite as a successful example of the use of AI for scientific research. AlphaFold was developed by DeepMind, a subsidiary of Google's parent company, Alphabet, to predict a protein's three-dimensional structure from its amino acid sequence. In fact, Google DeepMind's Demis

Hassabis and John Jumper won the 2024 Nobel Prize in Chemistry for their work on AlphaFold. As the press release for the 2024 Nobel Prize in Chemistry states:

> In 2020, Demis Hassabis and John Jumper presented an AI model called AlphaFold2. With its help, they have been able to predict the structure of virtually all the 200 million proteins that researchers have identified. Since their breakthrough, AlphaFold2 has been used by more than two million people from 190 countries. Among a myriad of scientific applications, researchers can now better understand antibiotic resistance and create images of enzymes that can decompose plastic. Life could not exist without proteins. That we can now predict protein structures and design our own proteins confers the greatest benefit to humankind [6].

According to Yongjun Xu et al. (2021), "As an illustration of great achievement, AlphaFold successfully and accurately predicted 25 scratch protein structures from a 43 protein panel without using previously built proteins models" [7]. However, human scientists cannot explain its predictions. In other words, they know what goes into the system (i.e., the input) and they know what comes out of the system (i.e., the output), but they do not know what happens in between. This is because AlphaFold is a deep learning artificial neural network (DNN). DNNs are black boxes. Their neural networks are comprised of multiple layers, including nodes and edges, of mathematical relations, some of which are hidden, which makes it difficult to explain how they arrive at their predictions and decisions. "Such systems are a 'black box model', making it difficult to guess how they make decisions or why they create a certain output" [8]. This is known as "the black-box problem" [9].

Not all instances of AI4Science involve black boxes. Still, because of the black-box problem with AI systems like Deep Neural Networks (DNNs), some researchers call for rethinking the ways in which some AI systems are designed, developed, and deployed in scientific research. They call for making black box AI systems explainable and their outputs interpretable by humans. Explainable AI or "XAI methods attempt to explain black box systems (e.g., DNN) by building a second 'explanation' model" [10]. In fact, Tan and Zhang (2023) designed an explainable version of AlphaFold called "ExplainableFold (*x*Fold)" [11]. Another approach to XAI is build what are known as surrogate models. Surrogate models are algorithms trained to approximate the predictions of black box AI systems such that the former can be used to interpret the predictions of the latter [12].

According to Clyde, "Justification with opaque AI will be a great challenge for the AI4Science campaign" [5]. Until we have explainable AI systems, whose decision-making and inferential processes are transparent to us, such that they produce outputs that we can interpret and understand, it would be prudent to restrict the use of black box AI systems in science to the context of discovery and to refrain from using them in the context of justification. Although some philosophers of science have abandoned this context distinction due to its association with empiricist accounts of science that are considered by many philosophers to be outdated, I argue that the discovery/justification context distinction may be useful for providing a guideline for acceptable uses of AI in scientific research. By itself, the discovery/justification context distinction is not meant to constitute a comprehensive ethical framework for AI4Science. Rather, it is supposed to serve as a preliminary guideline for developing such a framework.

Before we proceed, it is important to keep in mind that the term "model" is used in somewhat different ways as it pertains to AI and to scientific theorizing. On the one hand, an AI model is a program that processes data to make decisions or predictions by means of statistical models. For example, AI systems can use regression models to find relationships between dependent and independent variables. When the statistical models or decision-making processes used by AI systems are vastly more complicated, with hidden layers of neural networks, such that they are opaque to human scientists, as in the case of DNNs, the AI systems are said to be black box models.

On the other hand, a scientific model is a representation of some natural phenomenon. The representation can be physical, graphical, mathematical, etc. For example, in chemistry, there are different ways to represent or model molecules. Molecules can be represented graphically by skeletal formulas or bond-line notation. That is, skeletal diagrams are a graphical way to model molecules two-dimensionally. Molecules can also be represented physically by ball and stick models. That is, ball and stick models are a physical way to model molecules three-dimensionally.

Crucially, the underlying theories that guide the scientific modeling of molecules in chemistry are the same, which include valence bond theory and molecular orbital theory, no matter which of the aforementioned ways to model molecules is used in scientific practice. They are the underlying theories that explain why a certain molecule is modeled in a certain way whether the bond-line notation or ball and stick models are used. The same cannot be said about AI systems that are black boxes. As far as AI black boxes are concerned, we have no underlying theories that would make their internal processes explainable and their outputs interpretable to human scientists. In the case of AlphaFold, for example, we have no underlying theory for the predictions it generates, which is why AlphaFold is said to be a black box. For this reason, the debate as to whether or not AlphaFold solved the protein folding problem continues [13]. As Shi-Jie Chen et al. (2023) explain (original emphasis):

> Clearly, deep-learning AI represents a major advance in protein *fold* prediction. But this is not *folding* prediction. Patterns extracted from proteins in the Protein Data Bank (PDB) provide a ready "parts list," circumventing the folding process entirely. These patterns are "fully baked." That is, a pattern extracted from a solved structure in the PDB is fully preorganized; any physical-chemical organizing interactions have already been realized during folding. The situation is analogous to interpreting a movie by fast-forwarding to the final scene without first watching the previous two hours; we know how it ends, but we don't know why [14].

As long as we do not know why, since we have no understanding of its internal processes, which are largely hidden to us, AlphaFold will remain a black box to human scientists. This is not to say that AlphaFold is entirely useless. It provides useful predictions of protein structure from the PDB. Other AI systems can be similarly useful, despite being black boxes, insofar as they have the computational power to process massive amounts of data in order to find patterns that may be elusive to human scientists. What AlphaFold does not provide, however, is a principled answer to the following question: How does a particular protein structure emerge from a linear sequence of amino acid residues in aqueous solution? [14] To answer this question, we need more than a predictive algorithm. We need a scientific theory, i.e., an explanation that encompasses the underlying causes, entities, mechanisms, and principles that govern the occurrence of the natural phenomena under investigation.[1]

2. The discovery/justification context distinction

The distinction between the context of discovery and the context of justification is typically traced back to Hans Reichenbach (1938) [15] and Karl Popper (1959/2002) [16] but can probably be traced back even further [17]. The context of discovery refers to processes of thinking as they occur in the minds of scientists when they come up with ideas or hypotheses, whereas the context of justification refers to methods by which scientific ideas or hypotheses are justified. As Thomas Nickles (1980) puts it (emphasis added):

> The distinction is first of all a logical distinction between the psychological processes which occur when a scientist thinks of new ideas and the logical argument which exhibits the degree to which those ideas are supported by the facts and other evidential considerations. *Context of discovery concerns psychological connections* between thoughts; *context of justification concerns only logical connections* (plus the ascertainment of facts). Context of discovery is descriptive; context of justification is normative as well [18].

For example, it is a historical fact that the "nineteenth-century German chemist August Kekulé claimed to have pictured the ring structure of benzene after dreaming of a snake eating its own tail" [19]. This historical fact provides no evidential support for Kekulé's hypothesis [20]. It makes no evidentiary difference whether the idea came to Kekulé in a dream, in the shower, or in any other way. What matters is whether the idea can be justified by means of evidence gathered from experiments, observations, and the like.

Empiricist philosophers of science used the distinction to argue that the context of discovery is psychological and subjective, whereas the context of justification is logical and objective, which is why they argued that philosophers of science should concern themselves with the latter rather than the former. For this reason, empiricist philosophers of science also thought that philosophy of science is in the business of providing rational reconstructions of science. As Minwoo Seo and Hasok Chang (2015) put it (emphasis added):

> Discovery is a subject of all kinds of empirical research, historical, sociological, and psychological. Epistemology is and should be confined to the "context of justification," in which the propositions produced in science are reformulated and rearranged so that their structures and logical relations are made explicit. *Epistemology thus considers a*

[1] Another example that may help to bring this point home is the use of Newtonian physics to launch rockets. We know that in order to climb into low Earth orbit, a rocket needs to achieve a speed in excess of 28,000 km per hour. In order to leave Earth and travel out into deep space, a rocket needs to achieve escape velocity, i.e., a speed of over 40,250 km per hour. These numbers did not come out of thin air. They are based on Newton's laws of motions, which explain how things move in our universe. When it comes to launching rockets, we know that the greater the mass of rocket fuel burned (m), and the faster the gas produced can escape the engine (a), the greater the thrust of the rocket (f). This knowledge is based on Newton's second law of motion, which can be stated in the form of a mathematical equation as follows: $f = ma$. Newtons' laws of motion make up the underlying principles that govern the motion of objects, which human scientists and engineers understand well enough such that they can use them to launch rockets. As far as AlphaFold is concerned, we have no such underlying principles, theories, or anything that resembles a mathematical equation along the lines of Newton's second law of motion.

> *rational reconstruction of scientific practice*, rather than the actual practice of scientists [21].

We need not be committed to this further argument made by empiricist philosophers of science as to the proper aims of epistemology and philosophy of science. For present purposes, the important point is that there is a real distinction between the context in which scientific ideas are generated and the context in which those ideas are tested, verified, confirmed, and justified. Out of the several different ways of drawing the context distinction, this is the one that Hoyningen-Huene (1987) labels as "process of discovery vs methods of, or reconstruction of, or analysis of, or considerations relevant to justification" [17]. If so, this context distinction between *processes of discovery* and *methods of justification* may be useful in providing a preliminary guideline for acceptable uses of AI black boxes in scientific research, or so I argue.

3. The context distinction and AI4Science

The principle of explicability is widely recognized as a foundational principle of ethical AI [22]. For example, UNESCO's "Recommendation on the Ethics of Artificial Intelligence" (2022) lists *transparency* and *explainability* as principles for ethical AI [23]. These principles require that people are fully informed when decisions are made by AI systems and that the reasons for those decisions are accessible and understandable to people. In science, of course, it is imperative that findings and results are accessible and understandable to other researchers in the field. As the editors of *Nature Geoscience* (2014) put it (emphasis added):

> Science thrives on reproducibility. […] Two ingredients are essential for reproducibility in any field in science: *full disclosure of the methods used to obtain and analyse data*, and *availability of the data that went into and came out of the analysis* [24].

When it comes to black box AI, the data that go into and come out of these AI systems are available. But the methods used to analyze the data are not. As long as the AI systems typically used in science are black boxes, such as AlphaFold, we should use them in the context of discovery only, not in the context of justification. Here is why.

First, recall that the context of justification is the context in which scientific ideas are tested, verified, confirmed, and justified. Since there can be no justification for the outputs of AI black boxes, given that their decision-making and inferential processes are opaque to human scientists, such systems cannot and should not be used to test, verify, confirm, and justify scientific ideas.

To put it another way, the principle of explicability in terms of *intelligibility* demands an answer to the question "how does it work?" [22] As far as the outputs generated by AI black boxes are concerned, this question cannot be answered because the internal decision-making and inferential processes of such systems are opaque to human scientists. Without answers to this question, the outputs generated by AI black boxes cannot be tested verified, confirmed, or justified, which is why such AI systems should be excluded from the context of justification. The fact that an idea was generated by an opaque AI system, such as AlphaFold, does not constitute a scientific justification, even in part, of the veracity of that idea.

Second, the origin of an idea rarely, if ever, guarantees the truth (or falsity) of that idea. Assuming that it does is often called the "genetic fallacy" in informal logic. In fact, some logic textbooks define the genetic fallacy in terms of the discovery/justification context distinction. For example, "The errors of treating items in the context of discovery as if they belonged to the context of justification is called the 'genetic fallacy'. It is the fallacy of considering factors in the discovery or genesis of a statement relevant, *ipso facto*, to the truth or falsity of it" [25] [26]. Salmon (1973) gives the example of dismissing the theory of relativity on the grounds that it was proposed by a Jewish person, namely, Albert Einstein [25]. Indeed, in his discussion of the context distinction, Giere (1999) proposes that Reichenbach was concerned about the dismissal of scientific ideas simply because they were proposed by Jewish people. As Giere (1999) puts it (original emphasis):

> I suggest that part of the significance of the [context] distinction for Reichenbach at this time was its implicit denial that characteristics of a *person* proposing a scientific hypothesis have anything to do with the scientific validity of the hypothesis proposed. This applies, in particular, to that person's being a Jew. Reichenbach seems to have made it a precondition on any scientific epistemology that it rule out the possibility of any distinction between, for example, Jewish and Aryan science. But I think there was more to it than this. Separating questions of the origins of ideas from questions of their validity seems to have been for Reichenbach, at that time, a matter as deeply personal as it was philosophical [27].

This is not to say that there is never any relation at all between questions concerning the origins of ideas (i.e., the context of discovery) and questions concerning their correctness (i.e., the context of justification), which is why genealogical critiques (i.e., criticizing ideas in terms of their psychological origins) may have some value after all [28]. However, the relation between an idea's origin and its correctness is not one of logical consequence. In other words, faulty origins can give rise to good ideas, and questionable sources can make statements that are true. Therefore, if the genesis of an idea is relevant to its truth (or falsity) in some way, it should be assessed as part of the *justification* for that idea independently of its *discovery*.

To illustrate, take the example of Kekulé again. If the fact that his idea came to him in a dream has any bearing on the hypothesis of the ring structure of benzene, then it should be considered in the context of justification as one of the methods of justification for Kekulé's hypothesis. Given that dreams are not a reliable mode of justification and are not one of the established methods of justification commonly used and accepted in scientific research, they should not be considered as such.[2]

Applying this point to AI4Science, it could be argued that the fact that a scientific idea originated in a black box AI is irrelevant to the justification of that idea. This is because the processes that generated the idea are unavailable for assessment by the scientific community. But the idea should not be dismissed simply because it was generated in a black box AI. Regardless of its

[2] Interestingly, Dmitri Mendeleyev recounts that his discovery of a periodic table of the elements occurred to him while he was dreaming [31]. One could argue that an AI system would have been able to find the patterns that eluded nineteenth century chemists for too long and construct a periodic table of the elements. But again, the question would be this: what is the underlying theory that *makes sense of* the patterns captured in a periodic table of the elements? In Mendeleyev's original periodic table, the elements are arranged by order of increasing atomic weight, which exhibits a pattern known as the periodic law, according to which elements in the same group have similar properties. The advent of quantum mechanics provided a better understanding of electron configuration, particularly the number and arrangement of electrons in the outermost shell (i.e., valence electrons), which determines an element's chemical properties and its position on later versions of periodic tables by atomic number.

origin, the idea should be tested, verified, confirmed, and justified by means of established scientific methods that are transparent to the scientific community.

In other words, "the computer said so" should not play a justificatory role in science; otherwise, we risk what Floridi (2023) calls "agency laundering," i.e., avoiding responsibility and blaming an AI system when things go wrong [22]. As the motto of the Royal Society of London for Improving Natural Knowledge, which was adopted in its First Charter in 1662, states "*Nullius in verba*," i.e., "take nobody's word for it" [29]. As AI becomes more widely used in science, we should not take its word for it, either. Just as the Fellows of the Royal Society were determined "to withstand the domination of authority and to verify all statements by an appeal to facts determined by experiment" [29], we should withstand the domination of the authority of AI or "automation bias" [22]. Instead, we should restrict the use of black box AI to the context of discovery until such time that its outputs can be verified "by an appeal to facts determined by experiment," observation, and the like.

As explainable AI becomes the rule, not the exception, in science, scientists may be able to consider the internal processes of such white box AI as part of the context of justification for the ideas that such systems generate [30]. Until then, black box AI should be restricted to the context of discovery, where the origin of an idea, be it a human brain (awake or asleep) or a DNN, is distinct and kept separate from the available methods by which the idea is tested, confirmed, verified, and justified.[3]

4. Conclusion

I proposed that the distinction between the context of discovery and the context of justification, which comes from philosophy of science, may provide a rather useful guideline for acceptable uses of black box AI in science. Whether an idea originates in a biological or an artificial network of neurons constitutes no evidence for or against the veracity of that idea. For that is the context of discovery. The idea should then be tested, confirmed, verified, and justified by means of established scientific methods. For that is the context of justification. Accordingly, given that AI systems used in science today are black boxes whose internal processes of idea generation and decision making are opaque to human scientists, for the most part, it would be prudent to restrict the use of such systems in science to the context of discovery. As long as the AI in AI4Science refers to black box AI, AI4Science should not apply to the context of justification. The discovery/justification context distinction thus provides a rather straightforward guideline for acceptable uses of AI in scientific research, which can be further developed into more detailed, best practices for AI4Science. As such, the discovery/justification context distinction is supposed to serve as a preliminary guideline for developing an ethical framework for AI4Science, not as a comprehensive framework in itself. Maintaining the discovery/justification

[3] An anonymous reviewer raises an interesting and important question. That is, could there be cases in which failing to use black box AI systems in scientific research would be considered immoral? Briefly, I think that the answer to this question would depend on the ethical perspective taken to answer it. For example, from a utilitarian point of view, if the failure to use black box AI in scientific research would result in less utility (where utility could be construed in terms of benefits, happiness, wellbeing, etc.) for everyone concerned, then it could be argued that black box AI should be used because it would maximize utility for all concerned. Alternatively, if scientists' refusal to use black box AI in their scientific research embodies virtues, such as honesty, integrity, courage, justice, and the like, then it could be argued that black box AI should not be used because refusing to use black box AI would further the aforementioned virtues. These are only two ways to approach the reviewer's question. There are many more. So, hopefully, readers can see that doing justice to this question would require more space than this paper affords.

context distinction as far as AI4Science is concerned could also help to ensure that human scientists remain in the loop while the use of AI in science becomes more widespread.

References

1. Oxford AI4Science Lab. Home. *Oxford AI4Science Lab*. https://oxai4science.github.io/ (2024). Accessed on April 29, 2024.
2. AI4Science @ Caltech. About. *AI4Science @ Caltech*. https://www.ai4science.caltech.edu/about.html. (2024). Accessed on April 29, 2024.
3. Xu, Y., Wang, F., An, Z., et al. Artificial intelligence for science—bridging data to wisdom. *The Innovation* 4 (6): 1-2 (2023). https://doi.org/10.1016/j.xinn.2023.100525.
4. Wang, H., Fu, T., Du, Y., et al. Scientific discovery in the age of artificial intelligence. *Nature* 620, 47-60 (2023). https://doi.org/10.1038/s41586-023-06221-2.
5. Clyde, A. AI for science and global citizens. *Patterns* 3 (2), 1-3 (2022). https://doi.org/10.1016/j.patter.2022.100446.
6. NobelPrize.org. Press Release: They Cracked the Code for Proteins' Amazing Structures. *NobelPrize.org*, October 9, 2024. https://www.nobelprize.org/prizes/chemistry/2024/press-release/. (2024). Accessed on December 12, 2024.
7. Xu, Y., Liu, X., Cao, X., et al. Artificial intelligence: a powerful paradigm for scientific research. *The Innovation* 2 (4), 1-20 (2021). https://doi.org/10.1016/j.xinn.2021.100179.
8. Garrido-Rodriguez, Pedro, Miguel Carmena-Bargueño, Maria Eugenia de la Morena-Barrio, Carlos Bravo-Perez, Belen de la Morena-Barrio, Rosa Cifuentes-Riquelme, Maria Luisa Lozano1, Horacio Perez-Sanchez, and Javier Corral. Analysis of AlphaFold and molecular dynamics structure predictions of mutations in serpins. *PLoS ONE* 19 (7), 0304451 (2024). https://doi.org/10.1371/journal.pone.0304451.
9. Hassija, V., Chamola, V., Mahapatra, A. et al. Interpreting black-box models: a review on explainable artificial intelligence. *Cognitive Computation* 16, 45-74 (2024). https://doi.org/10.1007/s12559-023-10179-8.
10. Fleischer, W. Understanding, idealization, and explainable AI. *Episteme* 19 (4), 534-560 (2022).
11. Tan, J. and Zhang, Y. ExplainableFold: understanding AlphaFold prediction with explainable AI. *KDD 2023: Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (August 2023), 2166-2176 (2023). https://doi.org/10.1145/3580305.3599337.
12. Molnar, C. *Interpretable Machine Leaning: A Guide for Making Black Box Models Explainable*. (2024). https://christophm.github.io/interpretable-ml-book/.
13. Al-Janabi, Aisha. Has DeepMind's AlphaFold solved the protein folding problem? *BioTechniques* 72 (3), 73-76 (2022).
14. Chena, Shi-Jie, Mubashir Hassanb, Robert L. Jerniganc, Kejue Jiad, Daisuke Kiharae, Andrzej Kloczkowskif, Sergei Kotelnikovg, Dima Kozakovh, Jie Liangi, Adam Liwoj, Silvina Matysiakk, Jarek Meller, Cristian Michelettim, Julie C. Mitchelln, Sayantan Mondalo, Ruth Nussinovp, Kei-ichi Okazakir, Dzmitry Padhornys, Jeffrey Skolnickt, Tobin R. Sosnicku, George Stanv, Ilya Vakserw, Xiaoqin Zoux, and George D. Rose. Protein folds vs. protein folding: differing questions, different challenges. *PNAS* 120 (1), e2214423119. https://doi.org/10.1073/pnas.2214423119.

15. Reichenbach, H. *Experience and Prediction: An Analysis of the Foundations and the Structure of Knowledge*. Chicago: The University of Chicago Press (1938).
16. Popper, K. *The Logic of Scientific Discovery*. London: Routledge (1959/2002).
17. Hoyningen-Huene, P. Context of discovery and context of justification. *Studies in History and Philosophy of Science* 18 (4), 501-515 (1987).
18. Nickles, T. Scientific discovery and the future of philosophy of science. In *Scientific Discovery, Logic, and Rationality*, edited by T. Nickles, pp. 1-60. Dordrecht: Reidel Publishing Co. (1980).
19. Robinson, A. Chemistry's visual origins. *Nature* 465 (6), 36 (2010). https://doi.org/10.1038/465036a.
20. Rocke, A. It began with a daydream: the 150th anniversary of the Kekulé's benzene structure. *Angewandte Chemie International* 54 (1), 46-50 (2015). https://doi.org/10.1002/anie.201408034.
21. Seo, M., and Chang, H. Context of discovery and context of justification. In *Encyclopedia of Science Education*, edited by R. Gunstone. Dordrecht: Springer (2015). https://doi.org/10.1007/978-94-007-2150-0_239.
22. Floridi, L. *The Ethics of Artificial Intelligence*. Oxford: Oxford University Press (2023).
23. UNESCO. *Recommendation on the Ethics of Artificial Intelligence*. Paris: UNESCO. https://unesdoc.unesco.org/ark:/48223/pf0000381137.locale=en (2022).
24. Nature Geoscience. Towards transparency. *Nature Geoscience* 7, 777 (2014). https://doi.org/10.1038/ngeo2294.
25. Salmon, W. C. *Logic*. Second Edition. Englewood Cliffs, NJ: Prentice-Hall, Inc. (1973).
26. Walton, D. *Ad Hominem Arguments*. Tuscaloosa: The University of Alabama Press (1998).
27. Giere, R. N. *Science Without Laws*. Chicago: University of Chicago Press (1999).
28. Ward, A. C. The value of genetic fallacies. *Informal Logic* 30 (1), 1-33 (2010). https://doi.org/10.22329/il.v30i1.1237.
29. Royal Society. History of the Royal Society. *The Royal Society*. https://royalsociety.org/about-us/who-we-are/history/ (2024). Accessed May 2, 2024.
30. Loyola-González, O. Black-box vs. white-box: understanding their advantages and weaknesses from a practical point of view. *IEEE Access* 7, 154096-154113 (2019). https://doi.org/10.1109/ACCESS.2019.2949286.
31. Scerri, Eric C. *The Periodic Table: A Very Short Introduction*. New York: Oxford University Press (2019).