The Epistemology of AI-driven Science: The Case of AlphaFold

Abstract

The success of AlphaFold, an AI that predicts protein structures, poses a challenge for traditional understanding of scientific knowledge. It generates predictions that are not empirically tested, without revealing the principles behind its predictive success. The paper presents an epistemological trilemma, forcing us to reject one of 3 claims: (1) AlphaFold produces scientific knowledge; (2) Predictions alone are not scientific knowledge unless derivable from established scientific principles; and (3) Scientific knowledge cannot be strongly opaque. The paper defends (1) and (2) and draws on Alexander Bird's functionalist, anti-individualist account of scientific knowledge, to accommodate AlphaFold's production of strongly opaque knowledge in science.

Keywords: philosophy of AI; AlphaFold; epistemology of science; scientific knowledge.

1 Introduction

The problem of predicting how a protein's sequence of amino acids determines its three-dimensional structure, the protein-folding problem, has historically been one of the most difficult and significant challenges in the biological sciences. AlphaFold's success in producing a data bank of close to all protein structures known in nature has revolutionized the sciences, with profound implications for the production and validation of scientific knowledge.

AlphaFold presents a fundamental challenge for epistemology of science: it operates as an *epistemically opaque* system (Duede 2023; 2022; Creel 2020; Sullivan 2022; Lipton 2018), a 'black box', generating predictions without revealing the underlying principles or rules it uses to produce them. This opacity raises critical questions about the nature and the production of scientific knowledge in the age of AI-driven science. Are predictions alone, even highly reliable ones, sufficient to constitute scientific knowledge? Can we claim that AlphaFold generates new scientific knowledge when its internal higher-level processes are inaccessible to human understanding?

This paper investigates whether AlphaFold's predictions are scientific knowledge. Instead of relying on experiment-centered frameworks, which are closer in spirit to more 'traditional' forms of empiricist views in philosophy of science (see Humphreys 2004; Bird 2022 for discussion), other epistemological models of knowledge might better reflect the realities of AI-driven research.

In the second section, I present the problem of protein-folding and motivate the problem AlphaFold presents for scientific epistemology. The third section presents a trilemma that confronts us with the challenge of accommodating opaque AI systems like AlphaFold within our broader understanding of how science generates knowledge: (1) AlphaFold produces scientific knowledge; (2) Predictions alone are not scientific knowledge unless derivable from established scientific principles; and (3) Scientific knowledge cannot be strongly opaque. I defend claims (1) and (2) and reject (3), drawing on a functionalist account of scientific knowledge by Alexander Bird (2010; 2022), that fully excludes human mental access from the picture. I conclude that scientific knowledge can be strongly opaque to human scientists.

2 Background

2.1. Protein folding and AlphaFold

A protein is a sequence of amino acids. When interacting with any environment, for instance when put in water, the protein's amino acid string *folds* into a three-dimensional structure. This 3D native[1] structure of the protein is thought to be encoded in its 1D amino acid string, but the principle behind this encoding is not well understood, with several hypotheses having been proposed (Dill and McCallum 2012; Dill et al. 2008). The *protein folding problem* is the question of how a protein's amino acid sequence dictates its three-dimensional structure (Dill et al. 2008).

There is so far no known general mechanism or principle that explains how proteins fold in nature. As proteins perform many important functions in biology and biochemistry, predicting their folds has been one of the most important and difficult problems in biological sciences. Since (cheaper) computational methods for modelling proteins became available and increasingly successful, predicting protein structures to accelerate drug discovery became a major objective in computational biology (ibid.). The Protein Data Bank (PDB), which is the primary repository for protein structures, had approximately 170,000 structures as of 2020. These structures were determined primarily using experimental methods such as X-ray crystallography, nuclear magnetic resonance (NMR) spectroscopy, and cryo-electron

---

[1] As postulated by Anfinsen (1973), the native structure of a protein is the thermodynamically stable structure which depends only on the amino acid sequence and on the conditions of solution. The native structure does not depend on whether the protein was synthesized biologically or in vitro (Dill et al. 2008).

microscopy (cryo-EM). The process of determining protein structures experimentally has historically proven to be extremely time-consuming, costly, and not feasible for every protein, leading to a substantial gap between the number of known 1D protein sequences and their corresponding 3D structures.

In 2019 Deep Mind announced *AlphaFold*, an artificial intelligence system which can successfully predict protein structures. AlphaFold 2.0 followed in 2021, (Jumper et al. 2021a) and was succeeded in 2024 by AlphaFold 3.0. The newest version is able to predict DNA, RNA and ligand structures, all essential to further accelerating drug discovery (Abramson et al. 2024). AlphaFold can predict protein folds with an atomic level of precision and has been integrated into scientific developments in biology, chemistry, and medicine (Yang, Zeng and Chen 2023).

As of 2024, AlphaFold has predicted the structures of over 200 million proteins. This comprehensive database includes nearly all known protein structures from a wide range of organisms, including plants, bacteria, animals, and viruses. The massive expansion from its initial release has significantly impacted the scientific community, accelerating research and innovation in fields such as drug discovery, molecular biology, and biotechnology (EMBL-EBI 2022; Heikklä 2022; Quach 2022). Most importantly, AlphaFold's achievement has been publicly pronounced as a solution to the protein folding problem (Lewis 2022)[2]. In its citation for the 2024 Nobel Prize in Chemistry, the Nobel Committee wrote that "Demis Hassabis and John Jumper have developed an AI model to *solve* a 50-year-old problem: predicting proteins' complex structures. These discoveries hold enormous potential"[3] (Press release, Nobel Prize Outreach 2024) (italics added). The claim that AlphaFold solves the protein-folding problem is a central motivation for this paper.

---

[2] For a more cautious view, see McDonough (2024).

[3] italics added.

2.2. Deep learning and scientific knowledge

Recent work in philosophy of machine learning (ML) and science has explored the reliability of deep learning systems' (DLS) outputs, and how the systems' opacity affects the epistemic justification for using them in scientific practice (Creel 2020; Sullivan 2022; Boge 2022; Grote, Genin and Sullivan 2024; Duede 2024; Duede 2023; Duede 2022).

DLS are often described as 'black boxes' due to their opacity. While they can be seen as mathematically transparent (Duede 2022), their high-level logic and the reasoning behind their predictions remain largely unintelligible (Creel 2020; Duede 2022). This *epistemic opacity* stems from the mismatch between the complex, high-dimensional mathematical optimisation processes within deep neural networks (DNNs) and human understanding and interpretation of these high-level processes (Boge 2022; Sullivan 2022; Creel 2020; Duede 2022). In the AlphaFold case, the epistemic opacity pertains to the prohibitively complex principles on which the system operates, in spite of the understanding its engineers have of its general workings, including of certain biological principles that were encoded into its architecture. Nevertheless, DLSs' ability to make accurate predictions has led to remarkable success in various scientific applications (Jordan and Mitchell 2015; Boge 2022).

The epistemic justification for relying on the outputs of opaque systems is a complex, unresolved issue. Some argue that DLSs cannot be evaluated as reliable or trustworthy in the same way as other scientific instruments because we lack access to the high-level logical rules governing their internal processes (Grote, Genin and Sullivan 2024; Duede 2022). While efforts in explainable AI aim to address this opacity, current approaches primarily focus on local interpretability, providing insights into specific model behaviours rather than a comprehensive understanding of the internal logic. It has been further argued that DLSs differ from human experts in terms of epistemic justification (Duede 2022). While we trust experts based on their reasoning, evidence, and adherence to epistemic norms, the opacity of DLSs prevents a direct evaluation of their internal reasoning processes.

Computational reliabilism, developed by Duran and Formanek (2018), provided a basis for epistemic justification in traditional computer simulations techniques. Based on process reliabilism, their account posits that beliefs derived from computer simulations are justified when the simulation constitutes a reliable process for generating true beliefs. Given the impossibility of surveilling every computational step, reliability must be attributed through external indicators, such as verification procedures, validation protocols, robustness and expert knowledge assessments (Duran and Formanek 2018, 652). Importantly, computational reliabilism permits reliability attributions without requiring full transparency of the underlying processes, a significant advantage for addressing the epistemic opacity challenge in computational science.

The framework successfully addresses traditional computer simulations, where reliability indicators can be systematically applied to the simulation process and expert knowledge provides meaningful oversight. However, unlike conventional simulations, AlphaFold generates novel predictions about protein structures through opaquely learned representations that lack theoretical foundations or sufficient empirical data to match the outputs with. Not only are the computational steps involved in generating predictions inaccessible, like in conventional computer simulations, but the internal high-level logic of the DLS learning processes makes the learned representations uninterpretable. The system's reliability cannot be established through standard verification and validation procedures when it produces predictions for previously unknown protein structures, nor can expert knowledge validate outputs that exceed current empirical and theoretical knowledge.

The computational reliabilism framework thus fails to capture how such DLS can generate legitimate scientific knowledge precisely when they produce reliable outputs on the basis of an unknown principle in absence of theoretical and empirical grounding, arguably their most significant novel feature and contribution, compared to previous computer simulation techniques. Computational reliabilism still aids in establishing the starting point in my discussion of whether AlphaFold produces scientific knowledge, in that it provides a basis for general trustworthiness of the system's outputs, which scientists use as propositional claims about the structures of particular proteins.

Humphreys (2004) discusses the role of theoretical knowledge in the construction of reliable scientific instruments which may exceed the capacities of humans. He asks whether not knowing how the instrument works undermines arguments for the reality of the properties it detects (Humphreys 2004, 35), pointing out that theoretical knowledge is normally embedded in at least the construction, if not in the use of the instrument and it is what ultimately makes the instrument a reliable producer of veridical data (Humphreys 2004, 37). Thus, even when the scientist using an instrument may not know how the instrument works, she can rely on it because theory is embedded into its construction, and its outputs may be related to the known theory. While it is obvious that AlphaFold has been constructed with *some* theoretical considerations about the processes of protein folding (Kieval and Westerblad 2024), it could not have been constructed based on a theory or a general principle of protein folding, as such a principle is not known. Where some of the previous accounts of opacity in computational techniques found no need for transparency when there is a good fit between predictions and theory (Humphreys 2004, 150), there is no theory of protein folding to match the outputs of AlphaFold with.

While there is a case for DLSs being sufficiently reliable and thus trustworthy to be a legitimate part of scientific practice (Duede 2023), there is an open question about the epistemic status of their outputs. Are they to be treated as candidates for scientific knowledge, and do they thus stand in need of justification (ibid.)? In the AlphaFold case, while the generated predictions are used for various other research projects (e.g. drug discovery), *the protein problem itself* has always been about *determining particular protein structures*, which is also the purpose of AlphaFold. Additionally, considering the scientists take whichever structures they need from the AlphaFold data bank to use in their research, these structures are first handled as propositional claims about the real structures of the proteins of interest (the 3D structure of the particular protein is XYZ). Therefore, the philosophically interesting question concerns to the status of the individual outputs, i.e. whether these outputs themselves *can be scientific knowledge*. In other words, the question is whether it can be known to science that a particular protein structure is XYZ, on the basis of a prediction from AlphaFold. In further sections, I argue that the answer depends on whether an opaque internal principle (that AlphaFold may have figured out, but no one knows how to decode) can also be part of scientific knowledge.

Building on this body of work, I take as my starting assumptions that AlphaFold 1) has a valid place in scientific practice in virtue of being sufficiently reliable, and 2) is epistemically opaque in terms of the high-level processes supporting its individual predictive outputs. I will not be pursuing further examination of the levels of opacity and / or reliability of AlphaFold in this paper, as it is not the ML techniques that are my focus here. Instead, this paper aims to investigate the philosophical question of what follows for our view of scientific knowledge, where scientific knowledge is understood approximately as ‘what is known to science’. In particular, I investigate whether what is known to science can at the same time be fully opaque to every human.

2.3. The problem

AlphaFold is said to have discovered hundreds of millions of protein structures, effectively granting the novel *predicted* structures the status of *discovered* ones in the public discourse

(Callaway 2024; Stroe 2023; Trager 2022; Belias 2022). These new structures have not led scientists to discover a principle of protein folding, such that they could either propose an explanation of how AlphaFold could have made its predictions or advance a theory of protein folding themselves.

Furthermore, most scientists using AlphaFold take its predicted structures at face value and use them to work on whatever their objective is. There is normally no experimentation until a stage after which AlphaFold is involved, and that experimentation is also not aimed at confirming particular protein structures. For instance, in drug discovery research, AlphaFold can be employed together with computational chemistry platforms such as Chemistry42, and biocomputational generative platforms such as PandaOmics. Chemistry42 has been used to generate molecules based on the structures predicted by AlphaFold. While there is laboratory testing of the molecules that are eventually selected for some specific purpose (Ren et al. 2023), AlphaFold is employed to generate a pool of suitable protein folds without subsequent experimental testing of its predictions (which would not be practically feasible anyway and would defeat the purpose of using AlphaFold). While sceptics have argued that empirical confirmation and experimentation is needed (e.g., Terwilliger et al. 2023), in the absence of such experimentation, the predictions are used as representations of the protein shapes.

This leads us to a strange situation from an epistemological point of view. Determining specific protein structures has been one of the most important difficult problems in biological sciences, which now appears to have been (at least partially) ‘solved’ by an opaque system. We rely on AlphaFold’s predictions of novel protein structures ubiquitously in cutting-edge science. Ordinarily, propositions we rely on in the course of doing science have the status of scientific knowledge – they are part of what is ‘known to science’. But in this case, the worry is that computational predictions which are not empirically tested, and for which we have no underlying theory or general principle, cannot qualify as scientific knowledge or genuine discovery. The long tradition of Nobel Prizes being awarded for *empirically tested discoveries*, rather than for novel techniques or technologies, is testament to this idea. Jumper and Hassabis may be said to have cracked the principle behind protein folding but neither they, nor other scientists can claim to *themselves have knowledge* of this principle.

The central question is whether AlphaFold, despite its opacity, generates scientific knowledge. And if it does who *has* this knowledge: science, scientists, or AlphaFold?

3 Does AlphaFold generate scientific knowledge? A trilemma

I approach this question by considering the following trilemma:

1. AlphaFold can produce scientific knowledge.
2. Empirically unconfirmed predictions alone are not scientific knowledge unless derivable in an appropriate way from supporting theories, laws or mechanisms that are scientific knowledge.
3. Scientific knowledge cannot be strongly opaque.

There are prima facie reasons for taking each of these claims seriously.

(1) AlphaFold can produce scientific knowledge.

AlphaFold has generated highly reliable and ubiquitously used predictions, which scientists take from the repository and make use of without experimental confirmation. If a repository of scientific claims (which AlphaFold produced and which the scientists are directly engaging with) is regarded with good reason, such as the strong reliability of the outputs, as authoritative and trustworthy in scientific practice, it should (absent any reason to believe the trust is misplaced) be regarded as scientific knowledge.

(2) Empirically unconfirmed predictions alone are not scientific knowledge – unless derivable in an appropriate way from supporting theories, laws or mechanisms that are scientific knowledge.

Sometimes predictions are plausible cases of scientific knowledge. For example, it is known that Halley's comet will return in 2061. However, these are cases where there is a known theory, a law or a mechanism that produces the prediction. One might find claim (2) too demanding, thinking of examples such as "it is scientific knowledge that paracetamol reduces fever, even though there is no knowledge of *how* exactly it does that". However, it is not that a *generalization* cannot be scientific knowledge unless an underlying principle is known. It's rather that a prediction about a specific instance cannot be scientific knowledge unless (at minimum) the supporting principle is known.

One could reply that we *do* in fact know the supporting generalization – that AlphaFold is x% reliable. But this is akin to someone claiming that they know they will not win the lottery because only one in a million tickets wins. I might know the odds of winning, but I cannot, on the basis of this, know *whether I will win this time*. Similarly, I may know paracetamol works in x% of cases, but it is not scientific knowledge that it will work *in this particular case*. And, in the case of AlphaFold, an unconfirmed prediction of a specific protein structure, not generated on the basis of a known principle of theory, does not seem to amount to scientific knowledge. To see the absurdity of taking a predictor with a good track record to be generating scientific knowledge, imagine a "science guru" who has a good track record of predicting future Nobel Prize-worthy discoveries. It does not seem plausible to equate this oracle's predictions with knowledge. If it were, we ought to give them the Nobel Prize.

Restricting the discussion to 'scientific' knowledge here is important. One can imagine scenarios where reliable predictions may meet the standards for everyday knowledge despite the lack of any derivation from scientific knowledge (e.g. I can know that sun will rise and set tomorrow even if I know none of the relevant astronomical generalizations). But the prediction will be scientific knowledge only if the supporting generalization is too.

(3) Scientific knowledge cannot be strongly opaque.

To the extent that AlphaFold's predictions are derived from a supporting theory, law or principle of protein folding, the supporting generalizations are 'strongly opaque'. *Strong* opacity implies that no individual human knows the relevant theory, and no human has any way of accessing it. This can be distinguished from cases of *weak* opacity in which only a few experts possess the knowledge or in which the knowledge is difficult but possible to access.

Why think scientific knowledge cannot be strongly opaque? If a black box falls out of the sky, it seems absurd to immediately treat the information it contains as scientific knowledge even if *no one* can access that information. If current AI systems make strongly opaque scientific knowledge possible, this is something without any obvious precedent.

While each claim is independently plausible, the AlphaFold case forces us to reject one of them. For instance, if AlphaFold is strongly opaque and strong opacity is incompatible with knowledge, then AlphaFold cannot produce knowledge. If strong opacity is compatible with knowledge, then AlphaFold can produce knowledge. Which one should we reject?

Those attracted to an experiment-centered, quasi-empiricist view, might want to reject (1) and defend (2) and (3). I have in mind here 'traditional', internalist empiricism, committed to the central place of the scientists' knowledge in scientific inquiry. Such views in philosophy of science are discussed Humphreys (2004) and Bird (2010, 2022), on which evidence and empirical confirmation are ultimately tied the mental states of scientists.

In the rest of this paper, I will defend (1) and (2) and reject (3). I will proceed by assessing each claim in more detail.

3.1 Defending 1: AlphaFold can produce scientific knowledge

This claim concerns the status of AlphaFold's highly reliable predictions about protein structures, most of which are not empirically confirmed. These can be understood as propositional in form and are themselves not opaque to human scientists, as they can be accessed via the AlphaFold repository. Whether these predictions amount to scientific knowledge depends on how we answer the question: Can a prediction which is not yet empirically confirmed, count as knowledge if the predicting system is known to be highly reliable?

A critic might argue that empirical confirmation is what turns predictions into knowledge. Against this, I argue that predictions can be knowledge, and it is in fact part of normal scientific practice to view them as such. This is for two reasons.

First, science routinely treats empirically unconfirmed predictions as scientific knowledge. An example of a prediction that could be classified as knowledge before empirical confirmation was the discovery of the planet Neptune, which was predicted, and some would argue discovered, through mathematical methods rather than direct observation. In 1846, irregularities in the orbit of Uranus led astronomers to predict the existence of another planet. Neptune was subsequently observed in the predicted position. Some predictions are so robust and influential that they can be considered discoveries even before empirical confirmation, especially considering such predictions often drive further research efforts. For instance, the prediction of antiparticles by Dirac was considered groundbreaking and led to significant advancements in particle physics, even before the positron was experimentally confirmed (Bird

2022). As we see in these examples, scientific knowledge often includes and encompasses predictions without empirical confirmation.

Second, confirmation itself might not necessarily be instantiated via human cognitive access. Scientific evidence in general often includes and heavily relies on simulations (Boge 2022; Humphreys 2009, 2004; Winsberg 2022, Parker 2020), models and inferences based on indirect observations and measurements (Bird 2022; 2010). Although different accounts of confirmation may posit different requirements, on the empiricism's terms the method by which the results have been achieved, and the resulting knowledge itself must still be cognitively accessible to humans. AlphaFold forces us to challenge this assumption. The confirmation of AlphaFold's predictions may be run opaquely, by the processes operating withing the system's architecture. AlphaFold may, for instance, be conducting a kind of Bayesian confirmation, its architecture allowing for multistep / multilevel testing and updating (see Jumper et al 2021a). It is simulating new data based on prior evidence, then testing it, then simulating again (there may be many steps involved in this process, as can be inferred from AlphaFold's architecture[4]). This process may be even viewed as an experiment, or rather a great number of experiments (see Duede 2022 and Winsberg 2003 for a similar view), all ran inside of the deep neural network (the opacity of which will be addressed later).

A question for the critic remains, if AlphaFold's outputs cannot be scientific knowledge, what can they be? Some type of 'quasi-knowledge' could be invoked here or, as some philosophers have proposed, an ML-generated space of plausible hypotheses (Ratti 2020). However, since in the AlphaFold case determining individual protein structures is the core of the protein problem, individual predictions are themselves knowledge candidates, not merely a space of hypothetically useful structures for e.g., drug development. Any 'not quite knowledge' or 'quasi-knowledge' in this context has the same functional role in science as

---

[4] For details of AlphaFold 2.0 architecture, see Jumper et al. (2021a).

‘knowledge’ (scientists treating protein structures produced by AlphaFold and those determined by an experimental method in the same practical way, both the AlphaFold’s predictions and the experimentally determined structures being fallible to some degree). Thus, this move amounts to little more than paying lip service to the more traditional empiricist view. It is therefore not really a distinction worth drawing.

On this basis I claim that AlphaFold’s predictions function as scientific knowledge and therefore can be knowledge. Since empirically unconfirmed predictions are routinely treated as part of scientific knowledge in science, and the requirements for confirmation may differ depending on the account of confirmation, with at least Bayesian confirmation or confirmation via a simulated experimentation being available in AlphaFold’s case, AlphaFold’s predictions too can be taken as scientific knowledge claims about the previously unknown protein structures in accord with normal scientific practice. In virtue of high reliability and the trustworthiness with which the repository of these structures produced by AlphaFold is taken by scientists, they should be considered or at least *can be* scientific knowledge. The defense of this claim, however, is incomplete, without the discussion of how these predictions come about. The next section discusses this problem.

3.2. Defending 2: Empirically unconfirmed predictions alone are not scientific knowledge – unless derivable in an appropriate way from supporting theories, laws or mechanisms that *are* scientific knowledge

A novel prediction in science normally requires grounding in some known principle that gives this prediction justification for it to function as scientific knowledge.

Consider an example from physics: the existence of black holes has been known to science for decades before we had direct empirical confirmation of their existence, but only because they were predicted by a very strong theory which was known. General Relativity

Theory produced a lot of true predictions, including Penrose proving that black holes' existence followed from GRT mathematically in 1965, and there is a lot of confirming evidence for GRT itself. The Nobel Prize for his work on black holes was awarded to Penrose for this long-standing work as recently as 2020, after their existence was deemed empirically confirmed due to the first ever photograph of a black hole taken in 2019. Thus, despite the Nobel Prize only having been awarded in 2020, science has known for decades that black holes existed without their empirical confirmation on the basis of GRT, which was known and highly reliable[5].

Thus, even if empirical confirmation of a prediction is not always necessary for it to be knowledge, clearly a prediction alone is insufficient for knowledge, if nothing can be said of the principles it is either derived from, or which give it a high probability of being true. While it is contentious to assume fundamental principles or laws in biology, we can speak of mechanisms and principles for the equivalent purposes here.

Further relevant here may be an example from Norton's "The Material Theory of Induction" (2021): Knowing the relevant chemical properties of crystals and how they interact with the environment lets us predict of the shape a crystal will take before we see it form. This is in fact routine for chemists in a lab, when a new salt is prepared, to simply assert that such-and-such is the form of the salt's crystals. This prediction factually works as knowledge because the principles, or in this case knowing the relevant material conditions on which the inference is to be made, are also knowledge. Turning to the case of protein folding, two papers by Dill and colleagues (2012, 2008) provide substantial insights into reasons to posit a general (and in theory discoverable) principle. Proteins fold due to specific physicochemical forces encoded in their amino acid sequences, suggesting a general principle or a mechanism underlying the folding process. The principle is thought to dictate how proteins achieve their

---

[5] There are of course examples in the history of science where correct predictions were made by an erroneous theory, with a related debate in the literature (see e.g., Dellsen 2016; and Bird 2022), which is beyond the scope of this paper.

stable, functional native structures from their linear sequences, in the extremely short time they do so in nature. Finding out the relationship between a protein's sequence and its structure is a cornerstone of the folding principle. It is suggested that the amino acid sequence inherently contains all the information needed to achieve the correct 3D shape[6]. It is therefore reasonable to posit that there has to be a general principle of how proteins fold. Meaning, if you have all the relevant information about the mechanism of amino acids interacting with other molecules, then you can infer the shape of a particular protein based on this knowledge.

This leads us to a principal difficulty with defending the second claim in the AlphaFold's case. The issue essentially comes down to whether we can reasonably assume that AlphaFold has figured out some general principle of protein folding. In fact, the only way to fully hold the propositions (1) and (2) of the trilemma is to say that AlphaFold has implicitly grasped something like a theory or principle of protein folding, and that this principle should be considered part of scientific knowledge.

A note of caution should be taken at this point. Whether AlphaFold's solution represents the 'real' phenomenon (the 'real' protein-folding principle), merely a model, or something completely different (e.g., it developed a mechanism based on shortcuts that are not representative of how proteins actually fold in nature) is an important issue. It is however an issue for a separate investigation, pertaining to the relationship between models and phenomena in science, to the extent to which scientific instruments can represent real phenomena, as well as to how DLSs may exploit relevant vs irrelevant correlations to make predictions (see e.g., Parker 2020; Sullivan 2022; Boge 2022; Andrews 2023; Pietsch 2015). While these questions are certainly important to the 'realism debate' about what science can know from ML techniques (Rowbottom, Peden and Curtis-Trudel 2024), this investigation is beyond the scope

---

[6] For more detailed discussion of the modelling suggestive of the possible protein folding principles and further evidence, see Dill (2012, 2008).

of this paper. The question I aim to get into is about how what AlphaFold is doing relates to how scientific knowledge is generated.

Consider the following: AlphaFold was trained on all available data about protein structures and on what is known about the partial mechanisms of protein-folding (in absence of a theory or a general principle) (Stroe 2023)[7]. Based on all this data, it produced a massive number of novel inferences that turned out to be highly reliable. Consider further that AlphaFold uses some form of supervised and unsupervised learning. An essential step in training is when the output error is minimized by changing the system parameters until the output matches what we already know to be true (feedback). The system is given a set of amino acid strings for known protein shapes and then its parameters are 'tweaked' until it produces correct outputs. One could point out that it is precisely a feature of AlphaFold, that even during this training step, tweaking the system's parameters does not give us knowledge of the high-level mechanism, i.e. we do not gain knowledge of how and why the transformer architecture ends up distributing weights, what it picks up on, etc. It is almost like we are tweaking the numbers blindly, until it works. Let us now compare this with how a scientist adjusts parameters and calculations for a specific model, calibrating it for making better predictions. Something similar often happens when a scientist is working on solving a new problem or working to discover a principle, where she tries out various 'tweaks' to make it work. While she normally has reasons to try out the tweaks she does, she is still often trying things out without the knowledge of what works, until she does something right and thus discovers the working principle.

There seems to be a certain parallel between this process of calibration in the development and strengthening of scientific theories and methods and ML training. An

---

[7] see also: DeepMind Team: https://deepmind.google/discover/blog/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology/

important difference is of course that a scientist refining a theory for better predictions can explain her reasoning at least to an extent, whereas a deep neural network cannot explain its reasons for assigning weights across statistical distributions. Thus, an appropriate way to think of AlphaFold may not be as of an *artificial scientist that discovers* the principle of protein folding by making tweaks to the theory but as an entity that *embodies* a (simulated) principle of protein folding (along the lines of Parker 2020, and Gross 2024). Gross (2024) argued that in molecular biology ML techniques can be plausibly taken as providing a theoretical underpinning for the phenomena it models, which further supports at least the plausibility of the view that whatever AlphaFold has implicitly picked up on may in its own right constitute a principle of protein-folding. Notably, AlphaFold achieved remarkably good results, compared to the past strategies of running computer simulations with built-in knowledge about everything we know about the forces between molecular components of proteins. This further distinguishes ML systems from older computer simulation techniques in their capacity to not only compute prohibitively complex amounts of data but to learn and internally generate law-like principles for the modeled phenomena.

Another potential analogy can be drawn between AlphaFold embodying a simulated protein-folding principle and how in certain scientific domains, the predictive success of a model can lead to the acceptance of the model itself as a discovery (of a previously unknown principle). For example, quantum mechanics' ability to predict a wide range of phenomena with high accuracy has led to its broad acceptance and the discovery of new physical principles, even before some of its aspects were empirically confirmed. Just like the body of theory in this case is tested by its ability to make good predictions, so is AlphaFold’s.

A sceptic might point out that instead of adopting (2), in the AlphaFold’s case we can adopt (2)*: Empirically unconfirmed predictions alone are not scientific knowledge, unless derivable either in an appropriate way from supporting theories, laws or mechanisms *or* from a sufficient amount / type of data using other appropriate truth-filtering methods. This way, we can still deny that knowledge can be produced by an oracle whose workings we know nothing about, and we might still defend that the outputs of AlphaFold are knowledge, without having to accept that its opaque internal mechanisms are themselves knowledge. I would argue this approach will only show us what we already know, that AlphaFold's outputs are reliable, which is the starting point of my discussion in this paper. It still leaves us in a strange position regarding how predictions, if taken as scientific knowledge, relate to an underlying principle that is not knowledge. Normally in science, predictions function as knowledge if the underlying principle is known or, in the case of computer simulations, predictions are usually taken to not themselves constitute knowledge but either 'merely' a plausible hypotheses space (Ratti 2020) or higher-order evidence (Parker 2020), that is evidence about the existence of some other evidence about the relevant phenomenon, from which scientists can gain knowledge. Both can be reliable and used for further scientific investigation, e.g. drug discovery. However, in the case of the protein folding problem, to take specific protein shapes to be 'known to science' without empirical confirmation, we need to accept that they come from an underlying principle that can itself constitute scientific knowledge.

All in all, considering that AlphaFold is producing novel highly accurate data based on all the previously available evidence on how proteins fold via sufficiently reliable methods, and that we have good reasons to believe there is a general principle of protein folding to begin with, it looks like what AlphaFold is doing resembles normal scientific practice, if only occurring at an accelerated pace. This is in principle consistent with some other proposed accounts of how computational science techniques can provide evidence and lead to scientific knowledge, even where an individual scientist working with the simulation has no access to all of the relevant mechanisms (Parker 2020). Except, whatever it has learned in terms of general principles or if it embodies one itself is not known by *any* human. Accepting both that

AlphaFold’s internal simulated principle is both a plausible candidate for scientific knowledge and strongly opaque in the relevant sense is what makes the picture strange, rather than the specifics of the process of how it learns to predict the protein shapes.

Even if we have various accounts that equip us with reasons to trust AlphaFold in regard to our scientific practice, having a principle of protein folding that is both scientific knowledge and is not directly known to us is what makes its case unusual. It ultimately calls for taking a more radical stance on the possibility of *strongly opaque* yet *scientific* knowledge. As strong opacity of the internal principle remains the main cause for concern with accepting both propositions (1) and (2), in the next section I discuss proposition (3): the problem of opaque scientific knowledge.

3.3. Rejecting 3: Why scientific knowledge *can* be strongly opaque

This claim concerns the most complex part of the trilemma and the heart of the dispute with my critic. In this section, I will be granting the assumption that AlphaFold has or embodies an implicit principle of protein folding which is a knowledge candidate, in order to focus on the question of whether a strongly opaque principle can be scientific knowledge. The strong opacity featuring in the AlphaFold case is by extension likely to characterize much of AI-driven science.

The question behind the claim (3) can be formulated as following: Are the general principles of protein folding AlphaFold has implicitly grasped (but that no one knows how to decode) part of scientific knowledge? One might defend (3): *scientific knowledge cannot be strongly opaque,* on the basis of lacking accessibility to the opaque principle, along the lines of an internalist and thus an individualist picture of knowledge, contrary to an externalist view that prioritizes the system’s trustworthiness.

Establishing sufficient reliability or trustworthiness I have outlined in the Background section constitutes the basis for the justification of AlphaFold's integration into the production of scientific knowledge. However, there is more to the argument against (3) than confirming that AlphaFold's predictions are trustworthy on the basis of their reliability and can therefore be taken as scientific claims. Namely, one may still object that the internal opaque principle of protein folding should not be considered part of scientific knowledge, even if the outputs can be trusted. The internal principle calls for its own justification as part of knowledge, since we do not want to claim that *science is done via a process unknown to science*. Where we do eventually want to land is to say that science is done via a process *not known to individual scientists* but *known to science as an enterprise*.

The idea that humans might eventually be displaced from the process of scientific knowledge production is not new. Humphreys (2004; 2009) viewed this as a possible or even a likely trajectory of scaling automation in science, arguing for moving away from "anthropocentric epistemology". Other accounts, such as Parker (2020) stated that scientists can in principle gain knowledge from computer simulations (via gaining second-order evidence about the existence of other evidence about the target phenomena), even when the scientist herself does not have access to all of the relevant internal processes of the simulation. However, the present problem calls for a more specified solution, as AlphaFold and similar AI systems are 1) distinct from older / traditional computer simulation techniques, and 2) generate knowledge without its outputs being matched to an existing theory, nor are they taken in practice as a second-order type of evidence. The strongly opaque implicit principle of protein-folding thus needs to be either accepted as part of what is known to science or rejected as known to science.

To help accommodate the rejection of claim (3), I turn to Bird's (2010; 2022) account of scientific knowledge as irreducibly *social* or *collective* knowledge. The view rejects individualist and internalist accounts of scientific knowledge based on various forms of ‘belief’ and on which science is best explained as functionally analogous to but not supervenient on individual cognitive systems. A somewhat more radical version of his account may allow for accommodating broader implications of AI-driven science, as it posits a view of scientific knowledge from the point of view of science as a social enterprise. Bird (2010) outlines three conditions for qualifying something as knowledge, which can be satisfied by AlphaFold:

1. *Outputs must be propositional in nature (propositionality)*. Both the predictions and arguably the underlying principle of protein folding (assuming there is one in this case) can be expressed in propositional terms and amount to scientific claims.
2. *Mechanisms whose function is to ensure or promote the chances that the outputs are true / valid / trustworthy are in place (truth-filtering)*. AlphaFold is not an oracle or an alien artifact, but a learning system, which bases its novel outputs on all of our previous scientific data and claims. It is also integrated into the social system of science: there are peer reviewed papers on the quality of AlphaFold’s predictions, the advancements in protein folding-related problems in biological sciences, on its use in drug discovery and design research, etc. (e.g., Ali and Caetano-Anollés 2024; Kryshtafovych et al. 2023; Jumper et al. 2021b); and it is integrated with various other processes and tools of science, which continuously correct for errors and unreliable data.
3. *The outputs are the inputs for a) social actions or for b) social cognitive structures (incl. the very same structure) (function of outputs is preserved) – the outputs must be usable to produce more scientific knowledge.* AlphaFold’s outputs are used in complex projects, as part of the division of labor between various other computational systems and human scientists. New cutting-edge research projects employ AlphaFold to varying degrees. Computational methods of problem-solving in science are advancing due to systems like AlphaFold. Moreover, new versions of these AI systems are developed on the basis

of the previous ones, adapted to expand on the type of data and predictions it can produce.

The justification here is an externalist reliabilist one and is realized through the truth-filtering (trustworthiness) condition. Since the functionalist account rejects the central role of individual minds in the knowledge production, the proposition that is a knowledge candidate does not need to be individually known by anyone, but it must be integrated into the broader scientific knowledge infrastructure, as the conditions 2 and 3 state.

This view further lays a basis for arguing towards a more radical stance: that mental states' role is peripheral, such that something can be knowledge without mental states involved altogether and at any point in the knowledge production. If we accept that for P to be knowledge, no individual human has to know P at any given time (Bird 2010), we still might acknowledge that the process of P becoming knowledge usually, or at least so far, does involve individual mental states at some stage of the process (e.g. it is still largely humans who write and peer-review papers). However, P does not become or remain knowledge *in virtue of being accessed* by individual mental states. For example, a scientific paper contains knowledge not in virtue of humans being able to read it but in virtue of it having been reviewed and utilized for producing further papers. If a machine could write and review a paper and if further scientific inferences could be drawn on its basis, it does not matter whether any human have either written or accessed it.

Still, one might worry that *accessibility* remains a crucial, even if implicit, part of the functionalist profile of knowledge for Bird's account, beside *trustworthiness*. Even if no mental states are required for something to be part of scientific knowledge at any given time, mental access to knowledge may ensure that the system of knowledge production functions properly. Bird argues explicitly however that it is not access that defines knowledge but the functional integration into a societal structure, paired with the capacity of the said knowledge to play a social role (e.g. drive scientific decision making), where: “[cognitive| accessibility is the principal means by which that is achieved” (Bird 2010: 48). For Bird, however important the access to knowledge might currently be in science, it is its function that plays the qualifying role.

One might object that, since the implicit theory of protein-folding is not known by humans, it cannot in fact play the required social role (unless it is made known to humans). This may be an issue of conflating the individualist epistemological approach with the collective knowledge account, where social role in decision-making need not be dependent on an internally held mental-state-like knowledge of individual scientists. For example, if a question arises whether a drug modelled with the involvement of AlphaFold should be developed further into an experimental or even trial phase, it does not come down to whether the principle by which the predictions were made is opaque or not. One might still object at this point, that, while the predictions themselves can be accessed and used for social decision-making, the opaque mechanism cannot. At this point I see the distinction between the two to be quite thin. As far as AlphaFold can propose several structures to be taken as potential candidates for further research on a particular task, its internal mechanism is effectively driving decision-making for the research direction, even if the human scientist only interacts with the predictions in form of outputs and not with the internal mechanism itself.

The case of AlphaFold pushes this account towards a more radical version that is better able to accommodate AI-driven science, in that it at least requires taking a stance on whether trustworthiness without accessibility is enough, as AlphaFold may be satisfying the conditions for the former but not the latter. One could still insist that in the examples Bird (2010) provides, such as that of a science text book sitting in a library, where the knowledge is constituted via the book having been reviewed and being available in a library (both being instruments of knowledge in a socially-defined sense) still implies guaranteed cognitive access to the knowledge kept in the book. But the fact that the access to this knowledge is ensured by the features of the social mechanisms we build (the book being readable, the libraries accessible, etc.) may ultimately be due to the fact that individual humans have so far played a driving, if only contingent, role in producing scientific knowledge and defining its trajectories. That is to say, the fact that human artifacts are at least typically cognitively accessible to individuals ultimately only matters incidentally to how the knowledge they keep performs its (societal) function. While accessibility is central for trustworthiness on an internalist view, it is not necessary for trustworthiness on an externalist functionalist view, and it is arguably only tangential to scientific knowledge as collective social knowledge.

The radical decoupling of scientific practice from mental states and cognitive access thus allows us to accept the position that scientific knowledge can be fully opaque to humans, via excluding human cognitive access to knowledge from the picture of what is required in scientific knowledge production. It can therefore more adequately accommodate developments in AI-driven science. On such a view, AlphaFold produces outputs on the basis of an opaque internal principle which itself satisfies the externalist functionalist conditions for knowledge.

The outlook on the wide employment of systems such as AlphaFold in current and future science generally confronts us with the genuine possibility that strong opacity of the mechanisms and the high-level principles behind their predictions is largely compatible with the rich functional integration with the rest of science. If the integration is in place and AlphaFold functions as a part of a larger system that sustains itself according to the scientific norms and standards, the inaccessibility of its internal principle should not prevent us from treating it as scientific knowledge.

4 Conclusion

I presented the epistemology of AlphaFold as a trilemma and argued that scientific knowledge can be strongly opaque to humans, as long as it is properly functionally integrated into the enterprise of the scientific practice as a whole. If we accept an externalist functionalist view which completely rejects the role of accessibility to knowledge through mental / cognitive states, we can accept that AlphaFold generates scientific knowledge, even if part of that knowledge is strongly opaque. Otherwise, we face a situation where scientists are advancing science without scientific knowledge. The AlphaFold case in a way forces us to take a more radical position on whether we can call something that is reliable, novel, based on existing scientific knowledge, and used ubiquitously for cutting-edge science ‘scientific knowledge’, if it may never be known to humans. By accommodating such epistemically opaque systems into not just our notions of reliable scientific instruments, but into our broader conception of what can be known to science, we come closer to developing an epistemological account that is more responsive and appropriate to the current practices in the AI-driven science.

Statements and Declarations

No competing interests to declare.

Bibliography

Abramson, J., Adler, J., Dunger, J. *et al.* (2024): Accurate structure prediction of biomolecular interactions with AlphaFold 3. *Nature,* 630, 493–500. https://doi.org/10.1038/s41586-024-07487-w.

Ali, M.A., and Caetano-Anollés G. (2024): AlphaFold2 Reveals Structural Patterns of Seasonal Haplotype Diversification in SARS-CoV-2 Nucleocapsid Protein Variants. *Viruses,* 25;16(9), 1358. doi: 10.3390/v16091358. PMID: 39339835; PMCID: PMC11435742.

Andrews, M (2023): The immortal science of ML: Machine learning and the theory-free ideal. Philosophy of Science Archive: https://philsci-archive.pitt.edu/23840/7/Andrews_Theory-Free_Ideal.pdf.

Anfinsen, C.B. (1973): Principles that govern the folding of protein chains. *Science*, 181, 223-30.

Belias, A. (2022): AlphaFold: How AI can accelerate scientific discovery. ERCD.AI post. https://oecd.ai/en/wonk/alphafold-ai-accelerate-scientific-discovery.

Bird, A. (2010): Social Knowing: The Social Sense of “Scientific Knowledge”: Social Knowing. *Philosophical Perspectives 24*(1), 23–56. https://doi.org/10.1111/j.1520 8583.2010.00184.x.

Bird, A. (2022): *Knowing Science.* First edition. New York: Oxford University Press.

Boge, F. (2022): Two Dimensions of Opacity and the Deep Learning Predicament. *Minds and Machines* 32, 43–75. https://doi.org/10.1007/s11023-021-09569-4.

https://plato.stanford.edu/archives/win2021/entries/science-theory-observation/>.

Callaway, E. (2024): Chemistry Nobel goes to developers of AlphaFold AI that predicts protein structures. *Nature* News. https://www.nature.com/articles/d41586-024-03214-7.

Creel, K. (2020): Transparency in Complex Computational Systems. *Philosophy of Science* 87, 568–589. https://doi.org/10.1086/709729.

Dellsén, F. (2016): Scientific progress: Knowledge versus understanding. *Studies in History and Philosophy of Science*, 56, 72–83.

Dill, K.A. and MacCallum, JL. (2012): The Protein-Folding Problem, 50 Years On. *Science* 338, 1042-1046.

Dill, K.A., Ozkan, S.B. Shell, M.S., and Weikl, T.R. (2008): The protein folding problem. *Annu Rev Biophys*; 37, 289-316. doi: 10.1146/annurev.biophys.37.092707.153558. PMID: 18573083; PMCID: PMC2443096.

Duede, E., and Davey, K. (2024): Apriori Knowledge in an Era of Computational Opacity: The Role of AI in Mathematical Discovery. https://arxiv.org/abs/2403.15437.

Duede, E. (2023): Deep Learning Opacity in Scientific Discovery. *Philosophy of Science* 90, 1089-1099. https://doi.org/10.1086/709729.

Duede, E. (2022): Instruments, agents, and artificial intelligence: novel epistemic categories of reliability. *Synthese* 200, 491.

Durán, J.M., and Formanek, N. (2018): Grounds for Trust: Essential Epistemic Opacity and Computational Reliabilism. *Minds & Machines* 28, 645–666. https://doi.org/10.1007/s11023-018-9481-6.

EMBL-EBI (2022): DeepMind AI Powers Major Scientific Breakthrough: AlphaFold Generates 3D View of the Protein Universe. SciTechDaily article. https://scitechdaily.com/deepmind-ai-powers-major-scientific-breakthrough-alphafold-generates-3d-view-of-the-protein-universe/.

Gross, F. (2024): The explanatory role of machine learning in molecular biology. *Erkenntnis* 90, 1583-1603.

Grote, T., Genin T., and Sullivan, E. (2024): Reliability in Machine Learning. *Philosophy Compass*, e12974. https://doi.org/10.1111/phc3.12974.

Hardwig, J. (1985): Epistemic dependence. *Journal of Philosophy*, 82, 335–49.

Hardwig, J. (1991): The role of trust in knowledge. *Journal of Philosophy*, 88, 693–708.

Heikkilä, M. (2022): DeepMind has predicted the structure of almost every protein known to science. MIT Technology Review. https://www.technologyreview.com/2022/07/28/1056510/deepmind-predicted-the-structure-of-almost-every-protein-known-to-science/.

Humphreys, P. (2004): *Extending Ourselves. Computational Science, Empiricism, and Scientific Method.* New York, US: Oxford University Press.

Humphreys, P. (2009): The philosophical novelty of computer simulation methods. *Synthese* 169, 615–626. https://doi.org/10.1007/s11229-008-9435-2.

Jordan, M. I., and Mitchell, T. M. (2015): Machine learning: Trends, perspectives, and prospects. *Science*, 349, 255–260.

Jumper, J. *et al.* (2021a): Highly Accurate Protein Structure Prediction with AlphaFold. *Nature,* 596 (7873), 583–89. https://doi.org/10.1038/s41586-021-03819-2.

Jumper, J. et al. (2021b) : Applying and improving AlphaFold at CASP14. *Proteins*, 89(12), 1711-1721. doi: 10.1002/prot.26257. PMID: 34599769; PMCID: PMC9299164.

Kieval, P., and Westerblad, O. (2024): Deep Learning as Method-Learning: Pragmatic Understanding, Epistemic Strategies and Design-Rules. *Philosophy of Science Archive*: https://philsci-archive.pitt.edu/23489/.

Kryshtafovych, A, Schwede T, Topf M, Fidelis K, Moult J. (2023): Critical assessment of methods of protein structure prediction (CASP)-Round XV. *Proteins*, 91(12), 1539-1549. doi: 10.1002/prot.26617. Epub 2023 Nov 2. PMID: 37920879; PMCID: PMC10843301.

Lewis, T. (2022): One of the biggest problems in biology has finally been solved. SciAm article. https://www.scientificamerican.com/article/one-of-the-biggest-problems-in-biology-has-finally-been-solved/.

Lipton, Z.C. (2018): The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3), 31–57.

McDonough, M. (2024): Did AI solve the protein-folding problem? Harvard Medicine, feature article. https://magazine.hms.harvard.edu/articles/did-ai-solve-protein-folding-problem.

Norton, JD. (2021): *The Material Theory of Induction*. University of Calgary Press.

Novel Prize in Chemistry 2024: NobelPrize.org, Press release. Nobel Prize Outreach AB 2024. https://www.nobelprize.org/prizes/chemistry/2024/press-release/.

Parker, W. (2020): Evidence and knowledge from computer simulation. *Erkenntnis*, 87, 1521-1538.

Pietsch, W. (2015): Aspects of Theory-Ladenness in Data-Intensive Science. Philosophy of Science, 82(5), 905–916. https://doi.org/10.1086/683328.

Poston, T. Internalism and Externalism in Epistemology. Internet Encyclopaedia of Philosophy. https://iep.utm.edu/int-ext/.

Quach, K. (2022): Nearly all protein structures known to science predicted by AlphaFold AI. Let's hold off the champagne until an actual drug is developed using this tech. Enterprise Technology News and Analysis. https://www.technologyreview.com/2022/07/28/1056510/deepmind-predicted-the-structure-of-almost-every-protein-known-to-science/.

Ratti, E. (2020): What kind of novelties can machine learning possibly generate? The case of genomics. *Studies in History and Philosophy of Science, Part A*, *83*, 86-96.

Ren, F., *et al.* (2023): AlphaFold accelerates artificial intelligence powered drug discovery: efficient discovery of a novel CDK20 small molecule inhibitor. *Chem. Sci*., 14, 1443. DOI: 10.1039/d2sc05709c.

Shapere, D. (1982): The concept of observation in science and philosophy. *Philosophy of Science,* 49 (4), 485-525.

Stroe, O. (2023): Case study: AlphaFold uses open data and AI to discover the 3D protein universe. EMBL post. https://www.embl.org/news/science/alphafold-using-open-data-and-ai-to-discover-the-3d-protein-universe/.

Sullivan, E. (2022): Understanding from Machine Learning Models. *The British Journal for the Philosophy of Science* 73(1), 109-133. https://doi.org/10.1093/bjps/axz035.

Terwilliger, T.C., Liebschner, D., Croll, T.I. *et al.* (2024): AlphaFold predictions are valuable hypotheses and accelerate but do not replace experimental structure determination. *Nat Methods* 21, 110–116. https://doi.org/10.1038/s41592-023-02087-4.

Trager, R. (2022): AlphaFold has predicted the structures of almost every known protein. Royal Society of Chemistry News. https://www.chemistryworld.com/news/alphafold-has-predicted-the-structures-of-almost-every-known-protein/4016033.article.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, AN., Kaiser, L., Polosukhin I., *et al*. (2017): Attention is all you need. https://doi.org/10.48550/arXiv.1706.03762.

Winsberg, E. (2022): Computer Simulations in Science. *The Stanford Encyclopedia of Philosophy.* Edward N. Zalta & Uri Nodelman (eds.). https://plato.stanford.edu/archives/win2022/entries/simulations-science/>.

Winsberg, E. (2003): Simulated experiments: Methodology for a virtual world. *Philosophy of Science, 70*(1), 105–125.

Yang, Z., Zeng, X., Zhao, Y. et al. (2023): AlphaFold2 and its applications in the fields of biology and medicine. *Sig. Transduct Target Ther.* 8, 115. https://doi.org/10.1038/s41392-023-01381-z.