

AI-assisted science needs a collective verification network

Luca Maria Del Bono

Artificial intelligence (AI) is already changing radically the way scientific research is carried out. Programming, numerical calculations, literature searches, data analysis, or the exploration of alternative hypotheses are increasingly being delegated to AI systems. This allows scientist to focus on more higher level and creative-based tasks. Additionally, this change has the potential to make science much more interdisciplinary. At variance with system-specific research that has been carried out in the last decades, a researcher can now use AI to explore questions that require knowledge or technical tools outside their main field of expertise. A physicist, for example, might use AI to construct a model involving biological mechanisms that they would previously have lacked the expertise or time to investigate.

However, the same development creates a fundamental problem: **generating scientific results may become much easier than verifying them**. Indeed, AI systems can hallucinate facts, produce incorrect code, overlook important assumptions, or apply valid approximations outside their domain of applicability. Moreover, a scientist working outside their own field may not have enough expertise to recognize these errors. (Niels Bohr allegedly said that "An expert is a person who has found out by his own painful experience all the mistakes that one can make in a very narrow field.")

While existing tools (such as Lean) can verify particular aspects of an AI-generated result, there is currently no general method able to guarantee the scientific validity of an arbitrary AI-generated conclusion.

Two complementary responses are therefore needed. First, AI systems for science should become more uncertainty-aware. They should not only return answers, but explicitly evaluate the level of confidence on a given answer (eventually refusing to give an answer altogether if they are not able to provide one). This ability should be considered a fundamental feature of scientific AI.

Second, critical scientific claims should remain subject to human expert verification. AI can help researchers cross disciplinary boundaries, but specialists in those disciplines are still needed to evaluate whether the resulting claims are scientifically meaningful.

I propose the development an **online collective verification network for AI-assisted science**. Researchers could openly share ideas, calculations, models, or hypotheses developed with AI and request verification of the parts extending beyond their own expertise. The platform could use AI itself to identify the relevant scientific fields and connect each contribution with researchers possessing the appropriate expertise. Instead of relying mainly on personal contacts, conferences, email, or word of mouth, interdisciplinary verification could become an explicit and scalable component of the scientific process.

Contributions should be structured so that assumptions, sources, approximations, code, data, and AI involvement are visible. Other researchers could validate specific steps, identify problems, propose improvements, or connect the work with results from their own fields. Importantly, these contributions should be recorded and credited. Scientific incentives currently reward the production of new results much more strongly than the verification of other people's work. A collective platform should recognize a researcher who detects a crucial error or validates an important assumption as making a genuine scientific contribution.

AI could also help prevent the platform itself from becoming overwhelmed: routine claims could undergo automated checks first, while uncertain or scientifically consequential steps would be routed to human specialists. Verification effort would therefore be concentrated where expertise is most valuable.

The result would be a new division of labor. **AI would expand the range of questions that each scientist can explore, while human collective intelligence would provide the distributed expertise needed to**

evaluate what AI produces. Rather than replacing scientific collaboration, AI could make a better-organized form of collaboration necessary. A scientific network designed around this principle could turn verification from a bottleneck into one of the central functions of collective intelligence.