https://doi.org/10.1038/s42254-022-00544-1

Defining physicists' relationship with AI

As physicists are increasingly reliant on artificial intelligence (AI) methods in their research, we ponder the role of human beings in future scientific discoveries. Will we be guides to AI, or be guided by it?

Throughout September and October, together with the Alan Turing Institute 'AI for science and government' programme, we organized a series of online seminars to explore the interface between machine learning and physics. The recordings are available on YouTube. We invited researchers from academia and industry, working in different areas, in theory, experiment and simulation. Despite the diversity of physics questions the speakers are trying to address in their research, they all agreed that AI is now playing a central role in trying to answer them.

In an Editorial earlier this year, we discussed Jim Gray's four paradigms of science, suggesting that a fifth is emerging in which "machines [are] no longer mere tools, but equal partners in scientific exploration, exchanging ideas, intuition and understanding with the human peers". But the fifth paradigm of science is yet to be formally defined. There are kindred ideas, for example, in a talk at the American Astronomical Society 2019 Meeting, astrophysicist Alexander Szalay defined it in the context of large astronomical surveys as when "computers decide objectively which experiments will yield the biggest gain in our knowledge". Earlier this year, in a blog post, Christopher Bishop, Director of Microsoft Research AI4Science, defined the fifth paradigm in the context of numerical solutions of scientific equations that can be tackled with machine learning methods which provide fast, robust emulators to replace some of the traditional numerical simulation methods. This view was also discussed in the final event of our seminar series.

No matter how the fifth paradigm is defined, it appears to be imminent, so we should start thinking about how to define our fast-evolving relationship with AI and what role the physicist is going to play in future discoveries. In a Perspective in this issue, Mario Krenn and colleagues overview how advanced computational systems, and AI in particular, can help humans reach new scientific understanding. They identify three dimensions of computer-assisted understanding: a 'computational microscope' as a tool to uncover new or deeper properties of a physical system, in ways not possible before; a source of inspiration that suggests new ideas and connections; and an agent of understanding. Whereas in the first two dimensions, the human scientist gains new understanding from the computer-aided insights and suggestions, in the latter the AI does. Krenn et al. propose a way to test whether AI truly understands something by requiring it to explain its understanding to someone else, a human scientist, for example. But is this always possible?

"We should start thinking about how to define our fast-evolving relationship with AI and what role the physicist is going to play in future discoveries."

In a Comment in this issue, Matthew Schwartz warns that the assumption that AI's understanding will always be transferable to humans will quickly become untenable because artificial intelligence evolves on dramatically faster timescales than biological intelligence does. Schwartz suggests that this inability to keep up with AI's understanding may not necessarily be a bad thing (read the Comment to find out why). As also discussed in a Feature in this issue, AI may develop very different representations and a completely alien understanding of the world. Perhaps instead of translating this understanding into human language, we should start learning AI's language, which may involve a change of mindset to be able to operate with concepts that do not come naturally in our perception of the physical world.

Such issues might now sound purely philosophical, but considering the fast-paced progress in AI for science and the number of breakthroughs in the past two years (for example, in protein folding, mathematics or density functional theory), they might become practical problems in a matter of years. Thinking more broadly, beyond science, related questions about the role of human creativity in art and design, for example, might be asked in the context of diffusion models, the machine learning models that generate images from text descriptions, such as DALL-E 2.

So it's neither too early, nor too speculative to ask these questions: Are you happy to take a back seat and enjoy AI's scientific endeavours or do you think humans will play a central role in driving the directions of future discoveries? Will humans be left behind or will we enhance our own cognitive abilities and be able to reach new levels of abstraction and sophisticated reasoning? We challenge our readers to give them some thought.

Published online: 1 December 2022