Harnessing the Power of Artificial Intelligence with Mathematics[1]

Gunnar Carlsson, Stanford University, and Karen Saxe, American Mathematical Society[2]

January 23, 2024

Overview

Artificial Intelligence (AI) is developing at an extraordinary pace. Rapid development has been going on for a long time, but the last year has seen a distinct acceleration, particularly in the area of large language models. Given the transformative power of AI and potential impact on all sectors of the U.S. economy, it is critical that we (a) harness the power of AI fully and (b) ensure that AI is used responsibly. Government, through its investments, should encourage increased research activity in all aspects of AI. Mathematics is very well positioned to contribute to AI, and should be an important component of AI initiatives. It is also the case that AI can contribute to the mathematics research effort.

How Mathematics can contribute to AI

One of the biggest hurdles facing AI is the enablement of productive human interaction with AI technology. The most powerful AI technologies, especially neural networks, are often regarded as "black boxes", in that they can obtain usable outputs but do so in a way which is incomprehensible to humans. This is a serious problem, because it means that we are unable to systematically and predictably design and improve the technology to perform as we would like. It also means that we cannot guarantee that the technology is not behaving in a way contrary to human value systems, in that it might introduce bias and inequities.

Example: As an illustration of how this lack of understanding creates problems, Tesla Autopilot constructed AI algorithms that were designed to recognize the digits in speed limit signs. Their algorithms were such that small changes in the number "35" on a speed limit (extending the middle "prong" on the digit "3" induced the algorithms to recognize "35" as "85", even though a human would easily recognize the difference. This kind of error could be catastrophic. Understanding when an AI trained in one context can reliably function in a new context is a fundamentally mathematical problem. This requires not only that we make an AI work but that we understand something about *how* it works.

Example: Amazon found that it had to scrap an AI-based recruiting tool because they found that it had introduced bias against women. If they had sufficient understanding of the internal workings of the tool, it might very well have been possible that adjustments could have been made so as to remove the bias.

Throughout history, mathematics has provided methods of improving "brute force" solutions, enabling technology to work smarter instead of harder. At the moment, many challenges in AI are addressed by deploying huge amounts of computational resources at a problem, trying many different approaches until an effective one is found. Mathematics enables optimal data architectures, more efficient algorithms, improved performance and reliability, and explainable decisions. In short, deeper mathematical understanding will enable AI development of the future to make the most effective choice first instead of eventually by chance.

How AI can contribute to Mathematics

New fundamental results in mathematics have been discovered with the assistance of machine learning. This subfield of AI is the process by which a machine learns from examples and thereby develops its

[1] This AMS white paper will be used to inform the "one-pager" on AI for Congress, found at https://www.ams.org/government/dc-prepare

[2] Written in consultation with Brendan Hassett, Brown University and ICERM Director; and Jerry McNerney, Pillsbury Winthrop Shaw Pittman LLP and retired from US Congress.

intelligence, demonstrating that AI can help mathematicians discover new conjectures and theorems, and more generally, that machines and humans can be genuine collaborators. It's becoming clear that AI, rather than replacing human scientific ingenuity, has the potential to enhance it in ways we are only beginning to understand.

For example, AI has contributed to solving important mathematical problems of imaging, such as denoising, edge detection, and inpainting. AI-driven systems already help mathematicians analyze the behavior of partial differential equations, which can model physical processes from ocean currents to nuclear explosions. In addition to discoveries with immediate real-world applications, AI has been used successfully to solve theoretical mathematical problems. As history has shown, this kind of development can be tremendously impactful in science and engineering, not just now but 10 or 100 years from now.

How to enable mathematicians to contribute

There are a number of specific measures that would simplify and accelerate mathematician's participation in AI.

- **Increased emphasis on experimental mathematics:** Much of mathematics has been driven by initial observation of empirical results, which are then formulated into theorems and frameworks for studying a given problem. Physics and games of chance are obvious examples, but other areas of mathematics also proceed by first performing computations which suggest the direction of theories. This initial experimentation is particularly important in working on AI, and experimentation guided by mathematical insight is potentially very effective.
- **Engage mathematicians with real world data and computational issues:** It is important that mathematicians engage directly with the data and the computations. Rather than dealing only with AI practitioners' ideas about how to build mathematical models and formalisms to solve problems in AI, we can take steps to understand directly the actual sources and computational methodologies and use the insights gained to formulate models. This will greatly increase the effectiveness of the contributions of mathematicians, and will enable them to see the problems in the most realistic light possible.
- **Involve mathematicians in the formulation stage of investigations:** Mathematicians prefer to deal with problems that are well formulated, so that they can apply their techniques to them. For many problems, however, the formulation step is the most important one. Given engagement with the actual empirical (data or computational) information, mathematicians are very well suited to the formulation of mathematical problems that will be of use for solution of a given problem in the AI application domain. Dealing exclusively with the formulation produced by those working directly with AI is likely an inefficient way to proceed.
- **Involve mathematicians in efforts to ensure responsible AI:** Mathematicians can inform the development of algorithms that permit the inclusion of ethical constraints, rather than trying to account for and mitigate ethical problems—including algorithmic bias—after the fact.
- **New funding sources:** The National Science Foundation has been the traditional source of much of the funding for mathematical research. The NSF has a number of excellent initiatives directed at the involvement of the scientific community in AI. Given the interest in AI demonstrated by many different agencies in AI with different foci, we should explore the possibilities for other funding sources as well.
- **Popularizing AI research within mathematics:** We should increase the activities that expose mathematicians to the issues and problems arising in AI. This could mean written documents or videos, but it could also mean activities such as workshops at the institutes. Much of this is already happening to a degree, and we believe it should be encouraged to accelerate.