Deep
Learning
+
Deep
Thinking
=
Deeper
Understanding

Infusing physics intelligence into artificial intelligence

by Mike Williams and Jesse Thaler

Artificial intelligence (AI) is transforming many aspects of society, including the ways that scientists are pursuing groundbreaking discoveries. For many years, physicists have been at the forefront of applying AI methods to investigate fundamental questions about the universe. Building on these successes, we founded the Institute for Artificial Intelligence and Fundamental Interactions ("IAIFI," rhymes with "WiFi")[1], a collaboration of MIT, Harvard, Northeastern and Tufts universities. The IAIFI is one of the inaugural National Science Foundation AI research institutes[2], with 11 of our 26 senior members coming from the MIT Physics Department.

The primary goal of the IAIFI is to develop and deploy the next generation of AI technologies, based on the transformative idea that artificial intelligence can directly incorporate physics intelligence. IAIFI researchers are using these new AI technologies to tackle some of the most challenging problems in physics, from precision calculations of the structure of matter to gravitational wave detection of merging black holes. In addition, this interdisciplinary research leverages first principles from physics to drive AI innovations with widespread applications in both fundamental physics research and AI studies.

By fusing the deep learning revolution with the time-tested strategies of deep thinking in physics, we are gaining a deeper understanding of our universe—from the smallest building blocks of nature to the largest structures in the universe—and of the principles underlying intelligence itself.

Research in AI and fundamental interactions

Fundamental interactions are described by two pillars of modern physics: at short distances by the Standard Model of particle physics, and at long distances by the Lambda Cold Dark Matter model of Big Bang cosmology. Both models are based on physical first principles such as causality and space-time symmetries. An abundance of experimental evidence supports these theories, but also exposes where they are incomplete—most pressingly, that the Standard Model does not explain the nature of dark matter, which plays an essential role in cosmology. AI has the potential to help us solve this puzzle and many others in physics.

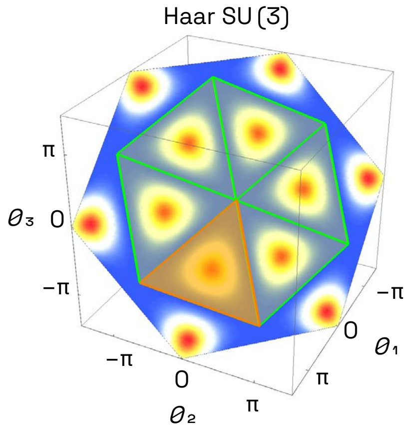

FIGURE 1:
Visualization of the complex symmetry space of variables in first-principles theoretical physics calculations, which have been built into machine learning models. Credit: Shanahan et al: *arxiv.org/abs/2008.05456*

IAIFI researchers, including MIT physics faculty Tracy Slatyer and Lina Necib, are using AI to study dark matter. Prior to joining the MIT faculty, Necib used AI to discover a massive unknown stellar structure that may have helped shape the Milky Way, for which she was awarded the 2023 Valley Prize by the American Physical Society. Now as an IAIFI researcher, Necib is developing new, groundbreaking AI methods for studying the dark matter halo and growth history of our galaxy [3].

For many physics problems, the governing equations that encode the fundamental physical laws are known. However, undertaking key calculations within these frameworks—essential to test our understanding of the universe and guide physics discovery—can be computationally demanding or even intractable. IAIFI researchers, including MIT physics faculty Phiala Shanahan, Will Detmold and Jesse Thaler, are developing AI for such first-principles theory studies, which naturally require AI approaches that rigorously encode physics knowledge.

One such application is focused on developing AI methods for performing calculations involving the strong nuclear force that binds quarks into protons and neutrons, and protons and neutrons into nuclei. Phiala Shanahan and her group in collaboration with colleagues at Google DeepMind have developed machine learning architectures that include the physical symmetries of the theory of quantum chromodynamics, incorporating guarantees of exactness into the novel AI algorithms [4]. If these results can be successfully scaled up to current state-of-the-art applications, they will enable novel first-principles studies of nuclear and particle physics. Beyond the domain of physics, this work can be used in robotics or for artificial limbs, where exact rotational symmetries inherently arise in joints.

Incorporating physics principles into AI is also having a major impact on many experimental applications, such as designing AI methods that are more robust and more easily verifiable. IAIFI researchers are working to enhance the scientific potential of various facilities, including groups led by MIT physics faculty members Phil Harris and Mike Williams at the Large Hadron Collider (LHC), working on the CMS and LHCb experiments, respectively, and by MIT Senior Research Scientist Lisa Barsotti at the Laser Interferometer Gravity Wave Observatory (LIGO).

The IAIFI LHCb group has developed a novel type of neural network that guarantees the interpretability and robustness required for use in real-time data processing at the LHC—applications with some of the largest data rates in the world. Since verification and interpretability of AI solutions are also important in other AI application domains, it is not surprising that these novel neural networks developed for the LHC have also been shown to beat state-of-the-art models in various problems in other domains, including criminal justice, medicine and finance. A related effort led by the IAIFI CMS group is in the area of ultra-low-latency AI inference, where neural networks make decisions in under a microsecond, motivated by the relentless pressure of the LHC's 40 MHz proton-bunch collision rate. IAIFI technology here is also being applied to other domains where real-time decision-making is critical [5].

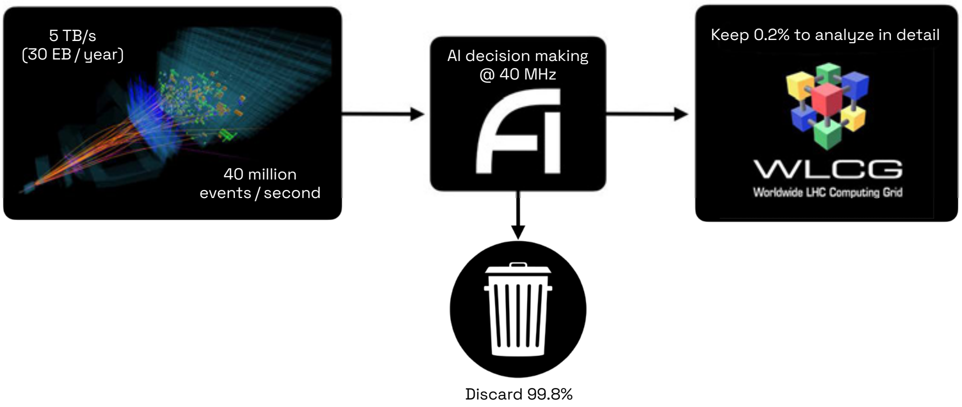

FIGURE 2:

Illustrating the data flow at the LHCb experiment at the Large Hadron Collider. (*From left*) Proton bunches collide at 40 MHz creating 5 TB of data per second. These events are analyzed in real time, with AI making the majority of decisions regarding which data to keep and which to discard. The vast majority of LHCb data must be discarded without ever being examined by human physicists.

The unique features of these physics applications also offer compelling research opportunities in AI more broadly, which is the primary focus of IAIFI and MIT physics faculty members Max Tegmark, Marin Soljačić and Isaac Chuang. A common AI application involves the denoising of images, for which a leading method is the diffusion approach inspired by thermodynamics. Tegmark and colleagues have introduced a new method inspired instead by electrodynamics, and more recently unified this method and diffusion into a new family of physics-inspired AI tools that are superior to all previous methods for many image denoising tasks, and exhibit improved robustness against modeling errors.

Additionally, physics-based tools are being used to better understand AI problems. The IAIFI's Tegmark and LHCb groups collaborated to show how physics tools such as effective theories and thermodynamic phase diagrams can help illuminate many aspects of AI learning dynamics. This project was a highlighted contribution at NeurIPS 2022, AI's leading conference. To truly trust AI models, we need to gain a deeper understanding of how they work, thus we view these efforts as an important contribution from physics to the field of AI.

Cultivating early-career talent, educating at the AI+physics intersection and engaging the public

AI technologies are advancing rapidly, making it both important and challenging to train junior researchers at the intersection of physics and AI. The IAIFI is recruiting and training a talented and diverse group of early-career researchers, especially at the postdoctoral level through our IAIFI Fellows Program [6], modelled after the MIT Physics Department's flagship Pappalardo Fellowships program. By offering our fellows their choice of research problems, and the chance to focus on exciting challenges in physics and AI, we are preparing many talented young scientists to become future leaders in both academia and industry.

The IAIFI Fellows are sparking interdisciplinary and multi-investigator collaborations, generating new ideas and approaches, translating physics challenges beyond their native domains, and helping develop a common language across disciplines. IAIFI Fellows Siddharth

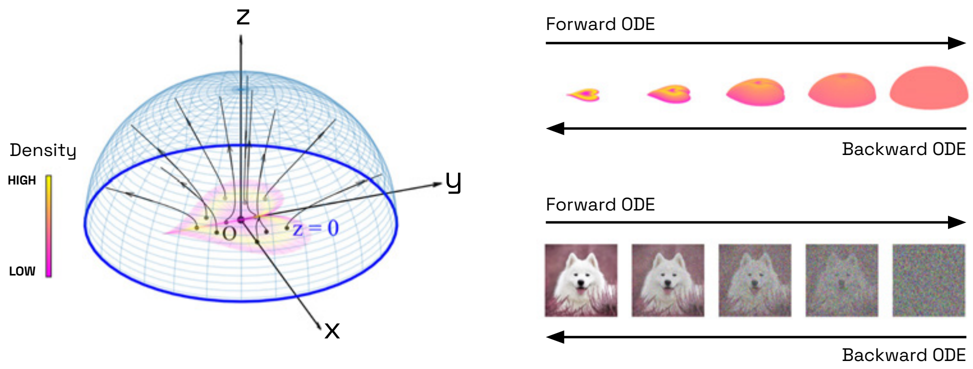

FIGURE 3:
(*From left*) Toy example showing how any 2-D density distribution, a heart-shaped one in this case, can be mapped to random noise on the surface of a 3-D hemisphere by some electric field (which is learned by the machine). The density can be viewed as composed of electrically charged particles. The right panel shows how one can use the learned electric field to generate noise on the hemisphere, which is then mapped to an image. In this way, the learned model can generate realistic images from scratch, or be used to denoise existing images.
Credit: Tegmark et al: *arxiv.org/abs/2209.11178*

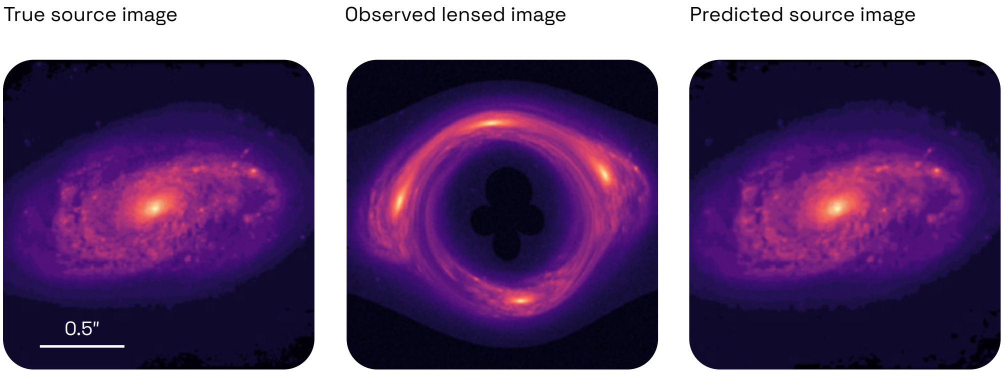

FIGURE 4:
(*Left*) Image of galaxy NGC2906. (*Center*) Simulated effect of gravitational lensing on this galaxy; this is what would be observed if this galaxy did undergo gravitational lensing when viewed from earth. (*Right*) AI prediction for what galaxy NGC2906 looks like given only the center lensed image. The agreement with the true image, which the AI did not have access to, is striking. Credit: Mishra-Sharma and Yang: *arxiv.org/abs/2206.14820*

Mishra-Sharma and Ge Yang, whose primary research focuses are astrophysics and robotics, respectively, collaborated on a project that pushed the boundaries of both AI and physics by making it possible to efficiently model high-resolution strong gravitational lensing observations at their full complexity [7].

Another related effort spearheaded by the authors and Alexander Rakhlin, an MIT professor of brain and cognitive science and IAIFI researcher, is the development of a new interdisciplinary PhD program in Physics, Statistics, and Data Science (PhysSDS). This is a collaborative effort between the Department of Physics and the Statistics and Data Science Center [8]. Statistics and data science are among the foundational pillars of AI. By providing physics PhD students formal training in these areas, we are fostering a new generation of leaders at the intersection of physics, statistics and AI. The first interdisciplinary PhD degree was awarded in Spring 2021, and thus far seven such degrees have been awarded, with more students joining the program every term. Roughly half of the students who have obtained this degree have gone into academia, with the other half now working in industry. To further bolster our industry connections, the IAIFI is hosting monthly lunches for our junior members with representatives from industry, and collating advice on how to get internships and jobs in industry from our alumni who have succeeded going that route.

In parallel with developing the PhysSDS PhD, MIT Physics faculty Phil Harris and Isaac Chuang, along with MITx Digital Learning Fellow Alex Shvonski, have created a new course, 8.16 *Data Science in Physics*, along with a concurrently run graduate version 8.316, which presents modern computational methods in the context of realistic contemporary examples of their applications to physics research. Students, for example, are taught how to confirm the recent Nobel Prize for the discovery of gravitational waves at LIGO,

FIGURE 5:
IAIFI postdoctoral associate Harold Erbin at the IAIFI Galaxy Zoo demonstration held at the Fall 2022 Cambridge Science Festival.

then improve on the published sensitivity using AI and fundamental physics principles. An MITx version of the course is partially online now [9], and planned to be fully available later in 2023. With projects like this, we hope to disseminate knowledge about, and enthusiasm for, physics, AI and their intersection.

To help ensure we are serving our junior members, we created an Early Career and Equity Committee (ECEC) within the IAIFI. The ECEC meets monthly to discuss issues related to the well-being and work environment of IAIFI members, and advises us on ways to create a more equitable, inclusive, welcoming and enjoyable place of work.

To reach beyond the MIT community, we launched a PhD Summer School in 2022 on the intersections of AI and physics. Its inaugural session was a huge success: more than twice as many students than we could accommodate applied, with truly exceptional ratings given by the students in the post-event survey. The weeklong school featured lectures by world-leading experts paired with hands-on tutorials run by our IAIFI Fellows, along with a hackathon, career panel and various networking opportunities. Our 2023 summer school is currently in preparation. At the Cambridge Science Festival in October 2022, Lina Necib hosted a galaxy zoo demonstration using labels created by citizen scientists to automate data analysis leveraging AI to better understand galaxy formation. IAIFI also ran a robotics demonstration, using a robot trained by IAIFI researchers using physics principles. In total, we interacted with about 1,000 people at this event, and another 600 using a similar demonstration at MIT's Winter Family Day.

Building new connections at MIT and beyond

Leveraging MIT's culture of collaboration, the IAIFI is generating new connections and strengthening existing ones across MIT and beyond. Of the 26 current IAIFI senior investigators, 15 are at MIT and members of the Laboratory for Nuclear Science, Research Laboratory for Electronics, Kavli Institute for Astrophysics and Space Research, Computer Science and AI Laboratory, and Institute for Data, Systems, and Society. In addition, IAIFI investigators are members of related NSF-supported efforts at MIT, such as the Center for Brains, Minds, and Machines within the McGovern Institute for Brain Research and the MIT-Harvard Center for Ultracold Atoms.

More broadly, the IAIFI is making Cambridge, MA, and the greater Boston area a hub for collaborative efforts to advance both physics and AI. This past year, the IAIFI hosted Brian Nord, an astrophysicist at the Fermi National Accelerator Laboratory and a senior member of the Kavli Institute for Cosmological Physics at the University of Chicago, as a year-long visitor through MIT's Martin Luther King, Jr. Scholars program. Nord's work involves training machines to explore the cosmos and fighting for equity in research.

As we teach in 8.01 and 8.02 at MIT, part of what makes physics so powerful is that it provides a universal language that can be applied to a wide range of scientific problems. Through the IAIFI, we are creating a common language that transcends the intellectual borders between physics and AI to facilitate ground-breaking discoveries. In doing so, we are tackling two of the greatest mysteries of science: how our universe works and how intelligence works. By linking them, using physics to improve AI and AI to improve physics, we are advancing physics knowledge and galvanizing AI research innovation. More broadly, a revolution is brewing in AI+science, and our efforts are aimed at positioning MIT to be a leader in this electrifying field.

PROFESSOR MIKE WILLIAMS is the founder and leader of the LHCb group at MIT and the inaugural Deputy Director of the IAIFI. He works on advancing our knowledge of fundamental particles by both proposing and performing novel experimental measurements at cutting-edge facilities. Williams is primarily focused on searching for as-yet-unknown particles and forces, possibly components of the dark sector of matter, and on studying largely unexplored emergent properties of the strong nuclear force. The LHCb group at MIT is a leader in the LHCb real-time data-processing system. To enable his scientific pursuits, Mike also works on advancing the usage of machine learning algorithms and other state-of-the-art data science tools within the domain of particle physics research, and on advancing our understanding of AI itself.

PROFESSOR JESSE THALER is a theoretical particle physicist who fuses techniques from quantum field theory and machine learning to address outstanding questions in fundamental physics. His current research is focused on maximizing the discovery potential of the Large Hadron Collider through new theoretical frameworks and novel data analysis techniques. Thaler is an expert in jets, which are collimated sprays of particles that are copiously produced at the LHC, and he studies the substructure of jets to enhance the search for new phenomena and illuminate the dynamics of gauge theories. He joined the MIT Physics Department in 2010, and is based in the Center for Theoretical Physics. In 2020, Thaler became the inaugural Director of the NSF Institute for Artificial Intelligence and Fundamental Interactions.

REFERENCES

[1] https://iaifi.org

[2] See https://aiinstitutes.org/institutes/ for a complete list of the NSF AI institutes.

[3] https://journals.aps.org/prd/abstract/10.1103/PhysRevD.107.043015

[4] https://arxiv.org/abs/2211.07541

[5] See https://fastmachinelearning.org and https://a3d3.ai, both co-led by Prof. Phil Harris.

[6] https://iaifi.org/current-fellows.html

[7] https://arxiv.org/abs/2206.14820

[8] https://physics.mit.edu/academic-programs/graduate-students/psds-phd/

[9] https://mitxonline.mit.edu/courses/course-v1:MITxT+8.S50.1x/