Unlocking Scientific Intuition & Reasoning at Digital Speed

Pushmeet Kohli

Throughout history, scientific progress has been defined by the tools we build to extend our perception. Today, artificial intelligence stands as our newest instrument, not merely as a computational engine but as a revolutionary lens for deciphering the nature of reality. Drawing on my research journey at Google DeepMind, this essay explores how scientists are deploying AI to overcome the hardest "root node" bottlenecks in science. By applying machine intuition, we are uncovering breakthroughs across biology, materials science, and mathematics, from solving the grand challenge of protein folding with AlphaFold to charting new universes of stable materials with GNoME. The essay concludes by reflecting on the transition toward open-ended, agentic AI systems that actively generate novel hypotheses and their implications for the nature of scientific inquiry. We are evolving from the solvers of intricate puzzles into the architects of profound questions, embarking on the most exciting journey of discovery ever undertaken.

Humanity has come far. In a blink of planetary time, we have filled libraries, split atoms, built computing machines, and mapped genomes. On the scaffold of hard-won scientific insight we have constructed a thriving, global civilization. The mechanism that made all this progress possible – the human brain, with its astonishing abilities of perception, intuition, and reasoning – has enabled us to understand many of the processes and laws of our universe and use them to transform the world we live in.

Despite these advances, the challenges facing humanity today – climate change, pandemics, the need for sustainable energy – show that we need a deeper understanding of the world than we have ever had. The quest to understand the nature of reality is challenging the limits of our intuition and our reasoning abilities.

Throughout history, our tools have helped us transform civilization by extending what we could perceive, understand, and control. The telescope revealed the vastness of the cosmos; the microscope, the hidden architectures of life. The computer, the bedrock of modern society, made it possible to calculate, simulate, and accumulate a universe of scientific data whose vastness and subtlety can be overwhelming.

Artificial intelligence is the latest instrument in this lineage: a new lens on our perception of reality itself; a telescope that not only is powerful enough to scan the

https://doi.org/10.1162/DAED.a.974

vastness of the data-universe but can focus on what matters within that universe, delivering new levels of understanding and transforming research in areas like biology, materials science, and mathematics.

AI naturally dovetails with science because many of the hardest scientific problems are, at their core, search and reasoning problems. When we engage in scientific inquiry, we observe the world and generate and test hypotheses, refining or discarding them based on how well they explain what we see. But as the volume and complexity of scientific data have increased, this process has hit two roadblocks: 1) the scale of what we are trying to understand and 2) the speed at which we can test our ideas. AI overcomes these barriers because it not only can process an extraordinary volume of observations about the world but can propose and evaluate novel hypotheses at digital speeds.

Where traditional methods of inquiry relied on human ingenuity, often guided by sparse data, AI can now detect structure in overwhelming abundance. What emerges is something akin to *machine intuition*: a powerful, learned strategy for navigating immense complexity.

Because it is domain agnostic, AI also bridges disciplines. Unlike human experts, whose training is often deep but narrow, AI systems can ingest and integrate information across fields. Polymaths are rare; experts in both genomics and high-energy physics are virtually nonexistent. But because AI can draw on insights from both areas at once, it has a unique ability to make new connections. This, in turn, accelerates the interdisciplinary breakthroughs on which modern science increasingly depends.

A senior scientist recently confided to me, "I was skeptical of AI, but having experienced the speed and depth of insights AI tools bring, I find myself using them extensively in my work." That admission brought to mind how early rocket scientists must have felt when they first got access to digital computers. That leap in capability – from extensive manual trajectory calculations using ranks of human "computers" to digital computation – must have been as empowering as it was irresistible.

This is not to say that the current generation of AI systems is a panacea or the sudden solution to all the world's problems. They remain a tool in the hands of scientists; a powerful tool, to be certain, but a tool nonetheless.

For over two decades, I have immersed myself in machine learning and artificial intelligence. My research career began with a focus on perception and reasoning, endowing computers with the ability to "see" the world and building self-correcting systems that ensure software reliability. But while I found the mathematical and technical puzzles exciting, I harbored a lingering desire to understand the deeper nature of our world.

My research journey changed during a chance encounter with Demis Hassabis, the cofounder and CEO of DeepMind. When he shared his vision of AI not just as a computational tool but as the ultimate lens for deciphering the nature of reality, I joined the mission.

DeepMind had already demonstrated the power of deep reinforcement learning: first by mastering Atari 2600 games and then through AlphaGo. By defeating one of the world's greatest Go players with moves that defied centuries of human wisdom – such as the now-famous "move 37" – AlphaGo proved that AI could transcend human intuition.

Believing the technology was finally ready to tackle the hardest challenges in science, Demis encouraged me to establish and lead the AI for Science program at Google DeepMind. Drawing on my experience with AI for Science, this essay traces the trajectory of our recent breakthroughs and the principles that have guided them. It offers a field guide to how AI currently integrates with the scientific method – where it excels and where it must improve – and explores how the emergence of agentic systems will fundamentally reshape the future of discovery. The following sections share my reflections on this new era of scientific discovery through deep dives into three key fields – biology, materials science, and mathematics – before zooming out to consider how AI is reshaping the scientific process itself, where its current limits lie, and what the future holds.

Biology is a natural starting point to test the utility of AI. It is staggeringly complex and full of emergent patterns and hidden structures that challenge the limits of perception and intuition. From understanding proteins – the building blocks of life – to understanding the genome – the code of life – biology has challenged us with profound mysteries. Among these, one of the most challenging was the fifty-year "grand challenge" of predicting the structure of proteins.

Proteins are the molecular machines of life. Made of intricately folded chains of amino acids, they are responsible for nearly every biological function, from digesting food to fighting infection. To understand disease, design better drugs, and ultimately engineer the machinery of life itself, we must know how proteins fold.

While a protein's amino acid sequence is easy to determine from its genetic code, its folded 3D shape is another matter entirely: each chain can potentially twist and bend in an astronomical number of ways, requiring analysis beyond the capabilities of any classical computation system. Discovering the structure of even a single protein through lab work can take years of effort using specialized equipment, at a cost of hundreds of thousands of dollars. For decades, scientists had tried and failed to accurately predict a protein's 3D structure directly from its DNA sequence using computation alone.

As a problem for AI, protein folding has all the right ingredients: a vast search space of potential structures, a clear and measurable goal (atomic-level accuracy), and a rich supply of protein structure training data gathered through decades of painstaking experimental work.

Over several years, my team at DeepMind worked closely with structural biologists and domain experts, including collaborators at the Protein Structure Prediction Center, whose Critical Assessment of Protein Structure Prediction (CASP) experiment set the global benchmark for protein prediction. In 2018, our first AI model, AlphaFold, became the global leader in protein-structure prediction, but it was not sufficiently accurate to be scientifically useful.[1] That system was radically re-architected by our team, led by John Jumper, and at CASP14 in 2020, AlphaFold 2 essentially solved the grand challenge of predicting protein structure to atomic-level accuracy.[2] That this breakthrough happened during the COVID-19 pandemic was particularly significant: even in this context, humanity leaped forward and in the process prepared itself with more powerful tools to deal with future pandemics.

Since then, the AlphaFold Protein Structure Database has cataloged predictions for more than two hundred million proteins – nearly every cataloged protein known to science – and made them freely available to researchers everywhere.[3] What was once a specialized and often painstaking pursuit was suddenly a navigable landscape, brought into sharp focus by a new scientific lens. AlphaFold 2 revealed what was always present but hidden from view: the structural biology of life, now visible at atomic resolution.

More than three million scientists in 190 countries have already used the AlphaFold database to accelerate their work, saving many millions of dollars and substantial human effort. It is accelerating drug development and advancing scientific progress across a wide range of fields, from combating malaria to discovering enzymes that break down plastic pollution to supporting research to increase honeybees' chances of survival.[4] AlphaFold is a clear demonstration of how AI can expand access to the tools of discovery, democratizing and accelerating science on a global scale.

Progress did not cease at understanding structures of individual proteins. While AlphaFold 2 essentially decoded the enormous "alphabet" of 3D protein structures, the subsequent breakthrough model, AlphaFold 3, revealed how that alphabet combines into words and sentences in the language of life. The updated 2024 model accurately showed how proteins interact with DNA, RNA, ligands, and each other to carry out coordinated biological processes.[5] We packaged both models into our free online tool, AlphaFold Server, for the benefit of the global scientific community, following extensive risk assessments and consultation with over fifty experts in biosecurity, research, and industry to ensure safe and responsible deployment.

Yet AlphaFold is not the final word in the understanding of proteins; many more research questions remain open. Scientists on my team at Google DeepMind and around the world are going beyond protein structures to learn to predict how proteins change shape (protein dynamics) and what roles they can play (protein function). AI plays a crucial role in these pursuits and I am confident that we will see many more breakthroughs on these problems.

If mathematics is the language of physics – precise, formal, and universal – then biology speaks in something far messier. It is a language of emergence, interaction, context, and complexity. The Human Genome Project, which allowed us to see our complete genetic code – all three billion base pairs – revealed the complexity of this task.

The genome is our cellular instruction manual. It's the complete set of DNA that guides nearly every part of a living organism, from appearance and function to growth and reproduction. Small variations in a genome's DNA sequence can alter an organism's response to its environment or its susceptibility to disease. But deciphering how the genome's instructions are read at the molecular level – and what happens when a small DNA variation occurs – is among our grandest scientific problems.

AI is the ideal translator for this task, not only because it can go beyond human abilities to learn this language from the massive amount of genetic data that we have collected about living organisms but also because its reasoning abilities can scale to handle the complexity of interactions that result from changes in the genetic code.

My team's own research on this problem led to the development of AlphaGenome and Enformer, new artificial intelligence tools that more comprehensively and accurately predict how single variants or mutations in human DNA sequences impact a wide range of biological processes regulating genes. This was enabled by, among other factors, technical advances allowing the model to process long DNA sequences and output high-resolution predictions. These models can integrate information from across vast stretches of DNA, identifying regulatory relationships previously invisible to researchers and offering a powerful new lens on the mechanisms underlying health and disease.

The genomics team at DeepMind has also leveraged AlphaFold, our protein structure model, to develop models like AlphaMissense, which focuses on the most common type of genetic mutation, whereby a single amino acid in a protein is replaced. The system predicts whether each change is likely to be benign or cause disease. It does this by combining evolutionary conservation – how vital that genetic site has been across millions of years – with structural context, assessing how the change might alter the protein's function. Validated against gold-standard datasets, AlphaMissense outperforms other computational methods, deciphering biology at a scale no human could and translating raw genomic code into functional insight.[6]

AI tools that help us interpret the genome in functionally meaningful ways are increasingly critical. As tools such as the gene-editing technology CRISPR push the boundaries of how we edit and write in the language of life, the ability to merely read DNA is no longer enough. Whether we are designing climate-resilient crops or engineering enzymes to break down plastic waste, the power to edit life depends on the power to understand its underlying code. AI is helping us connect the dots, showing how changes in DNA lead to changes in function, health, and disease.

Researchers are using this same capability to anticipate future health threats. In 2023, at Harvard Medical School, for example, computational biologist Debora Marks and her team developed EVEscape, a generative AI system trained on viral sequences from hundreds of thousands of nonhuman species.[7] By forecasting which potential mutations are most likely to help viruses evade immunity, EVEscape supports the design of vaccines capable of targeting both circulating strains and those likely to emerge. Just a few years ago, who could have imagined such forward-facing vaccines, designed not just for today's pathogens but for tomorrow's? This marks a fundamental transition in medicine: from a discipline of reaction to one of anticipation.

In 2024, our work on AlphaFold earned Demis and John Jumper, our DeepMind colleague and colead of the AlphaFold project, the Nobel Prize in Chemistry. They shared the award with David Baker, recognized for his contributions to computational protein design. The prize reflects not only the impact of AlphaFold but also a broader shift underway, a wave of AI-driven research transforming discovery across science and industry worldwide.

For example, Baker's team at the Institute for Protein Design has shown that proteins can be designed from scratch as nanoscale machines with programmable shapes and functions: where nature evolves these structures over millions of years, AI-powered tools such as RFdiffusion can now generate and test thousands of candidates *in silico*.[8] The long-term ambition of Baker's lab is striking: to create bespoke molecular machines that neutralize viruses, repair tissues, catalyze new forms of chemistry, and much more.

That same acceleration is reaching therapeutics. In 2019, Regina Barzilay and her team at MIT used deep learning to discover that a molecule previously researched (and discontinued) as a treatment for diabetes was also a potent new antibiotic active against drug-resistant bacteria.[9] Identified from a vast digital library of molecules in just hours, the drug, which they renamed halicin, represented the kind of AI-accelerated discovery that is inaccessible or prohibitively difficult using traditional methods.

In neuroscience, AI is visualizing the brain in unprecedented detail. In a landmark collaboration in 2024, Jeff Lichtman at Harvard University and Viren Jain at Google Research used AI to reconstruct a detailed human brain circuit of a small tissue sample. Even though the sample was just half the size of a grain of rice, it required more than a petabyte of electron microscopy data and revealed tens of thousands of cells and over 150 million synapses, including previously unknown structures.[10] Their work heralds a new era in brain science defined by AI-enabled seeing.

So what other breakthroughs should the scientific community pursue? I have never thought of AI as a solution in search of a problem. At Google DeepMind, we begin with a bold question: Which foundational scientific bottlenecks, if solved, would unlock cascades of transformative, downstream discovery? We call these "root node" problems.

Having defined the target, I ask which of these challenges appear genuinely out of reach – the bottlenecks that scientists broadly agree won't yield to traditional methods within the next decade. I only want our team to pursue the hardest questions that require the broadest multidisciplinary approaches, the best AI research, and the best compute and data, all coming together to advance the state of the art farther and faster.

Consider materials science and mathematical discovery. Each represents a domain in which progress had slowed, not for lack of effort, but because the sheer complexity of their problems was outpacing available methods. Our strategy is to intervene at these points: assembling multidisciplinary teams, developing bespoke tools, and staying the course until we and the researchers our work supports achieve the necessary breakthroughs. This is not AI as a hammer looking for nails, but as keys carefully designed to unlock the gates of scientific progress.

Materials have shaped our civilizations. From the Stone Age to the Silicon Age, they have powered every major leap in human progress: bronze gave rise to tools and weapons, steel to cities and railways, silicon to the digital world. The next great leap may come from the discovery and deployment of materials with properties we have never seen before. For example, human advancement now depends on us finding new sources of energy and new approaches to efficiently store and transform energy. Many of the most promising approaches – from next-generation solar panels and better rechargeable batteries to room-temperature superconductors – hinge on the discovery of new materials.

Yet materials discovery remains painfully slow. For most of history, it has been driven by trial and error, and the low-hanging fruit was picked long ago. What remains is a largely uncharted search space, too immense for even the most determined researchers to navigate unaided. This bottleneck holds back multiple scientific fields and technologies. But the dam is starting to crack. AI is now changing how materials research is done, shifting away from trial and error and intuition toward prediction and design. AI models can now estimate key properties such as conductivity, stability, and reactivity directly from a material's atomic structure, having learned from decades of experimental and computational data.

Still, the tens of thousands of scientifically explored stable materials represent only a tiny fraction of the billions of possible atomic arrangements allowed by chemistry. To explore this universe, we need new tools of unprecedented speed and precision. That is why we built an AI system called Graph Networks for Materials Exploration (GNoME).[11] Trained on decades of open materials data – including the twenty thousand stable materials already discovered through experimentation and another twenty-eight thousand via conventional computation – GNoME explores the vast space of potential material structures and accurately predicts the most stable among them.

GNoME was trained to model the laws of quantum chemistry, the fundamental rules that govern how atoms bond and interact. These laws are well understood but very costly to simulate directly. GNoME's breakthrough was in learning to approximate them accurately and at scale, allowing it to quickly scan millions of potential compounds that would have taken centuries to evaluate by traditional means.

It returned with extraordinary results: 2.2 million new materials, 380,000 of which were predicted to be stable enough to potentially power future technologies. Among these were 52,000 new layered compounds akin to the wonder-material graphene and more than five hundred lithium-ion conductors – twenty-five times more than previously reported – that could help improve the performance of rechargeable batteries. We released the stable candidates to the global scientific community, estimating that GNoME delivered the equivalent of eight hundred years' worth of scientific progress that day.

Like a telescope scanning the sky, GNoME mapped a constellation of stable materials, guiding experimentalists toward the brightest regions where breakthroughs are most likely to emerge. In this way, AI has become an invention for invention, a tool for building better tools and making more discoveries, with each spin of the flywheel amplifying the next.

Of course, akin to our early work on protein folding, this is just the beginning. While GNoME can predict potentially stable materials, the most promising of these candidates must also be synthesized and tested in the laboratory and in the real world to verify their properties and potential. While autonomous "self-driving" labs are accelerating this important step, from the A-Lab at Lawrence Berkeley National Laboratory to the Argonne National Laboratory in Lemont, Illinois, we still have further to go before we can fully create bespoke substances *in silico*, tailored to the needs of future technologies.[12]

In materials science, what we need may not yet exist, but the elements to create it are all around us. With AI as our guide, we will begin to conjure new matter from the vast possibility space of nature itself. Yet materials are just one part of building a sustainable future. We also need to pursue every viable path for generating clean, abundant energy. Among the most promising paths – and most difficult – is nuclear fusion, the reaction that powers the stars.

Controlling the superheated plasma inside a fusion reactor is one of science's toughest challenges, requiring precise, rapid adjustments to shifting magnetic fields. In collaboration with EPFL (École Polytechnique Fédérale de Lausanne, home of the Swiss Plasma Center), we developed an AI system that learns to control and shape plasma in real time, opening new paths for fusion research and bringing this clean, abundant energy source closer to reality.[13]

Elsewhere in high-energy physics, detectors at the Large Hadron Collider generate petabytes of data from trillions of collisions, a scale that overwhelms traditional analysis. Physicist Kyle Cranmer at the University of Wisconsin–Madison is doing pioneering work using machine learning to extract rare signals from noisy backgrounds and accelerate the search for new physics.[14] His research exemplifies a growing trend across the physical sciences: combining deep domain knowledge with AI systems that are not just data driven but increasingly aware of physical principles.

Mathematics is the invisible foundation of modern science and technology, from computing and communications to physics, engineering, and medicine. It also is the perfect test bed to study and evaluate the intuition and reasoning abilities of AI systems.

Early large language model (LLM)–based AI systems struggled with even basic mathematical reasoning. Their proclivity to hallucinate made it challenging for them to reason over the long sequences of reasoning steps typically required for nontrivial mathematical problems. Disappointed by these earlier investigations, many experts felt that the mathematical intuition and reasoning required to solve the kinds of complex problems featured at math competitions were many years away.

To overcome these challenges, we built AlphaProof, a specialized agent for solving math problems. Unlike normal LLM-based conversational agents, AlphaProof trains itself to prove mathematical statements in the formal open-source programming language Lean. It couples a pretrained language model with our AlphaZero reinforcement learning algorithm, which previously taught itself how to master the games of chess, shogi, and Go.

Formal languages offer the critical advantage that proofs involving mathematical reasoning can be formally verified for correctness. Their use in machine learning has, however, previously been constrained by the very limited amount of human-written data available. In contrast, natural language–based approaches can hallucinate plausible but incorrect intermediate reasoning steps and solutions, despite having access to orders of magnitude more data. We established a bridge between these two complementary spheres by fine-tuning a Gemini model to automatically translate natural language problem statements into formal statements, creating a large library of formal problems of varying difficulty.

When presented with a problem, AlphaProof generates solution candidates and then proves or disproves them by searching through possible proof steps in Lean. Each proof found and verified is used to reinforce AlphaProof's language model, enhancing its ability to solve subsequent and more-challenging problems. We trained AlphaProof to compete in the International Mathematical Olympiad, a prestigious competition for young mathematicians and a useful benchmark, by proving or disproving millions of problems, covering a wide range of difficulties and mathematical topic areas over a period of weeks leading up to the competition. AlphaProof also applied this training loop during the contest, reinforcing proofs of self-generated variations of the contest problems until it could find a full solution.

In 2024, AlphaProof (along with AlphaGeometry, our specialized system for geometry problems) made history by being the first AI system to solve four out of six problems from that year's International Mathematical Olympiad, achieving the same level as a silver medalist in the competition. In 2025, an advanced version of Google's Gemini Deep Think model (which incorporates some findings from AlphaProof) solved five of the problems perfectly, achieving a gold-medal level of performance.[15] This result was repeated by a specialized version of OpenAI's ChatGPT system.

Impressive as it is, solving Olympiad-level problems doesn't alone advance mathematics. To achieve that, AI must produce results that we have not yet discovered. A good example is my team's collaboration with mathematicians to use AI in making significant progress on solving partial differential equations – a key step on the path to solving the Navier Stokes equations – a long-standing fundamental question in fluid dynamics and one of the Millennium Prize Problems.[16] While we are yet to see one of these grand challenge problems in mathematics fall to AI, it is only a matter of time.

A hint at this future for mathematics research comes in the form of our AlphaEvolve system, which can generate novel algorithms and solutions across both computing and mathematics. AlphaEvolve combines the creativity of large language models with automated evaluators to evolve algorithms for math and practical applications in computing. It operates as an evolutionary process at digital speed.

Already, AlphaEvolve has delivered practical gains – improving code and resource efficiency across Google's global computing infrastructure – while also demonstrating capabilities in fundamental mathematics. Tested on more than fifty open problems in areas including geometry, number theory, and combinatorics, AlphaEvolve rediscovered the best-known solutions in 75 percent of cases; in 20 percent of cases, it went further, improving on existing results and advancing the frontier of mathematical knowledge.[17]

Crucially, AlphaEvolve does not operate as a black box. It expresses solutions as human-readable programs, providing a transparent chain of reasoning that allows scientists to inspect and build upon its discoveries. In a subsequent study, our team worked together with mathematicians Terence Tao and Javier Gómez-Serrano to evaluate AlphaEvolve on a number of other math problems. They found that AlphaEvolve can help discover new results across a range of problems, including discovering a new construction for the finite-field Kakeya conjecture.

While both AlphaEvolve and AlphaProof have limitations in terms of their efficiency and ability to handle even more complex problems, they signal some of the ways computer scientists and mathematicians will leverage AI to accelerate research.

While AI is already impacting mathematics and science through breakthrough models like AlphaEvolve and AlphaFold, there is a broader transformation now underway: the rise of agentic AI. This new generation of AI systems moves beyond narrow task execution to take a more autonomous role in discovery, not just assisting human inquiry but actively advancing it. Unlike specialist models built for specific domains like protein prediction or weather forecasting, agentic systems are more general and capable of open-ended exploration.

Google's AI co-scientist, built on Gemini 2.0 and launched in 2025, is a case in point. It is a multi-agent system, guided by human-defined goals, that draws on existing literature to generate novel research hypotheses and experimental protocols.[18] I think of it as a symphony of synthetic scientists working to increase the cadence of discovery.

With collaborators including the Fleming Initiative, Imperial College London, and Stanford University, we evaluated AI co-scientist's ability to generate novel hypotheses in biomedicine – including drug repurposing, treatment targets, and antimicrobial resistance – and validated these hypotheses through laboratory experiments. The results were promising, suggesting that AI will help shape the research agenda itself by surfacing salient questions before anyone thinks to ask them. Future systems will be increasingly general. They will propose their own hypotheses, design and refine experiments, and seek out the data they need, perhaps even navigating the physical world to gather it (with human oversight, of course).

We are moving toward AI that doesn't just facilitate science but begins to *do* science. For now, however, our human collaborators face a specification problem. They are finding that as these hypotheses-generating AI agents become more powerful, it is increasingly important for human scientists to be extremely precise in specifying the problem they want these agents to solve. Otherwise, the AI agents tend to find a solution that violates assumptions the human scientists were making but had not included in the original problem description. The burden of reasoning is shifting; we are moving from being the solvers of puzzles to the architects of questions.

This last hurdle foreshadows a future question: What happens when our tools begin to uncover truths we struggle to grasp? Science has always aimed not only to predict but to understand – to build coherent theories that reveal how the world works. Yet as AI grows more capable, it may pose questions or offer insights that resist easy explanation or even human comprehension. So our task as scientists will evolve from interpreting what AI finds to understanding why it matters and acting on it responsibly.

Until now, science has been a strictly human endeavor. The discovery of the scientific method and its focus on dispassionate hypotheses set humanity on the course to modernity. We now stand at an inflection point; the tools we built using the scientific method are promising to propel us into a vast new unknown for the first time in centuries.

As humanity moves ahead and draws the first maps of this new world, we must tread carefully. Long term, if we become dependent on systems we do not fully understand, we risk accepting their answers without sufficient scrutiny. More than a century ago, E. M. Forster's *The Machine Stops* imagined a future in which humanity lives passively under a vast machine's care, having long forgotten how the world – and the Machine – works.[19] As we build future generations of AI scientific assistants, we must remain clear-eyed about what we are gaining – and what we could lose along the way.

Given the pace of progress, making predictions more than even a few months ahead feels like a fool's errand; but I can say with confidence that the breadth and scale of advances AI will soon unlock are difficult to overstate. Within five years, AI will have enabled dozens, if not more, Nobel Prize–level breakthroughs. This is what happens when you give scientists superpowers.

AI will tackle problems that have resisted human insight for decades, from the deepest questions in mathematics (such as the Millennium Prize Problems) to the design of high-temperature superconductors that could radically improve how we generate and transfer energy. In biology, future AI systems may simulate entire cells, design synthetic organisms with myriad applications, and eliminate entire categories of disease.

Of course, the pace of progress is really the pace of understanding. For most of human history, new understandings took hold across centuries, from Copernicus to Newton, Darwin to Watson and Crick. Over time, that cadence has quickened, from decades to years to – now, with AI – months or even days. As our new scientific tools collapse the distance between question and answer, we find ourselves at an inflection point not just in capability but in tempo. And for those of us building these tools, we may yet have farther to go, but it feels like we are on the most exciting scientific journey ever undertaken.

At the precipice of change, some have wondered, "When does a machine win the Nobel Prize? And after that, will humans ever win it again?" These questions capture the common anxiety that AI might one day eclipse human contributions altogether. This is understandable but misses the point of why we do science. We do not just seek answers; we seek meaning. In that sense, AI will not replace scientists. Instead, it will expand what they can imagine, understand, and achieve. After all, the telescope didn't make astronomers obsolete. It gave them the stars.

ABOUT THE AUTHOR

Pushmeet Kohli is the Chief Scientist of Google Cloud and Vice President of Science at Google DeepMind. He previously served as Director of Research at Microsoft Research. He has recently published in *Nature*, *Science*, *Nature Medicine*, and *Advances in Neural Information Processing Systems*.

ENDNOTES

[1] Andrew W. Senior, Richard Evans, John Jumper, et al., "Improved Protein Structure Prediction Using Potentials from Deep Learning," *Nature* 577 (7792) (2020): 706–710.

[2] John Jumper, Richard Evans, Alexander Pritzel, et al., "Highly Accurate Protein Structure Prediction with AlphaFold," *Nature* 596 (7873) (2021): 583–589.

[3] Ewen Callaway, "'The Entire Protein Universe': AI Predicts Shape of Nearly Every Known Protein," *Nature*, July 29, 2022, https://www.nature.com/articles/d41586-022-02083-2.

[4] "Stopping Malaria in Its Tracks," Google DeepMind, October 13, 2022, https://deepmind.google/discover/blog/stopping-malaria-in-its-tracks; "Creating Plastic-Eating Enzymes That Could Save Us from Pollution," Google DeepMind, July 28, 2022, https://deepmind.google/discover/blog/creating-plastic-eating-enzymes-that-could-save-us-from-pollution; and "How the Honeybee Could Help Protect Species around the World," Google DeepMind, July 28, 2022, https://deepmind.google/discover/blog/how-the-honeybee-could-help-protect-species-around-the-world.

[5] Josh Abramson, Jonas Adler, Jack Dunger, et al., "Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3," *Nature* 630 (8016): 493–500.

[6] Jun Cheng, Guido Novati, Joshua Pan, et al., "Accurate Proteome-Wide Missense Variant Effect Prediction with AlphaMissense," *Science* 381 (6664) (2023), https://doi.org/10.1126/science.adg7492.

[7] Nicole N. Thadani, Sarah Gurev, Pascal Notin, et al., "Learning from Prepandemic Data to Forecast Viral Escape," *Nature* 622 (7984): 818–825.

[8] Joseph L. Watson, David Juergens, Nathaniel R. Bennett, et al., "De Novo Design of Protein Structure and Function with RFdiffusion," *Nature* 620 (7976) (2023): 1089–1100.

[9] Jonathan M. Stokes, Kevin Yang, Kyle Swanson, et al., "A Deep Learning Approach to Antibiotic Discovery," *Cell* 180 (4) (2020): 688–702.

[10] Alexander Shapson-Coe, Michał Januszewski, Daniel R. Berger, et al., "A Petavoxel Fragment of Human Cerebral Cortex Reconstructed at Nanoscale Resolution," *Science* 384 (6696) (2024), https://www.science.org/doi/10.1126/science.adk4858.

[11] Amil Merchant, Simon Batzner, Samuel S. Schoenholz, et al., "Scaling Deep Learning for Materials Discovery," *Nature* 624 (7990) (2023): 80–85.

[12] Michael Eisenstein, "Self-Driving Laboratories, Advanced Immunotherapies and Five More Technologies to Watch in 2025," *Nature* 637 (8047) (2025): 1008–1011; Lauren Biron, "Meet the Autonomous Lab of the Future," News from Berkeley Lab, April 17, 2023, https://newscenter.lbl.gov/2023/04/17/meet-the-autonomous-lab-of-the-future; and "Autonomous Discovery: Reimagining the Future of Scientific Discovery Through

AI, Robotics and Labwide Collaboration," Argonne National Laboratory, https://www.anl.gov/autonomous-discovery (accessed December 18, 2025).

[13] Jonas Degrave, Federico Felici, Jonas Buchli, et al., "Magnetic Control of Tokamak Plasmas through Deep Reinforcement Learning," *Nature* 602 (7897): 414–419.

[14] The ATLAS Collaboration, "An Implementation of Neural Simulation-Based Inference for Parameter Estimation in ATLAS," *Reports on Progress in Physics* 88 (6) (2025): 067801.

[15] Thang Luong and Edward Lockhart, "Advanced Version of Gemini with Deep Think Officially Achieves Gold-Medal Standard at the International Mathematical Olympiad," Google DeepMind, July 21, 2025, https://deepmind.google/discover/blog/advanced-version-of-gemini-with-deep-think-officially-achieves-gold-medal-standard-at-the-international-mathematical-olympiad.

[16] Yongji Wang and Sam Blackwell, "Discovering New Solutions to Century-Old Problems in Fluid Dynamics," Google DeepMind, September 18, 2025, https://deepmind.google/blog/discovering-new-solutions-to-century-old-problems-in-fluid-dynamics; and Minglang Yin, Nicolas Charon, Ryan Brody, et al., "A Scalable Framework for Learning the Geometry-Dependent Solution Operators of Partial Differential Equations," *Nature Computational Science* 4 (12) (2024): 928–940.

[17] Alexander Novikov, Ngân Vũ, Marvin Eisenberger, et al., "AlphaEvolve: A Coding Agent for Scientific and Algorithmic Discovery," arXiv (2025), https://doi.org/10.48550/arXiv.2506.13131.

[18] Juraj Gottweis and Vivek Natarajan, "Accelerating Scientific Breakthroughs with an AI Co-Scientist," Google Research, February 19, 2025, https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist.

[19] E. M. Forster, "The Machine Stops," *Oxford and Cambridge Review* 8 (1909).