

A Policy Framework for Building the Future of Science with AI



Authors

Ellie Sakhaee, Alice Friend, Kent Walker, James Maniyka

Acknowledgements

We are grateful for the many Google colleagues who contributed to the ideas and insights put forward in this whitepaper and made this paper possible, including (in alphabetical order) Dan Altman, Michael Appel, Guy Ben-Ishai, Jennifer Beroshi, Jared Bomberg, Nicholas Bramble, Zoë Brammer, Erica Brand, Michael Brenner, Matt Carona, Roxanne Carter, Jenn Crider, Alex Diaz, Lizzie Dorfman, Omid Ghaffari-Tabrizi, Brigitte Hoyer Gosselink, Conor Griffen, Rachel Gruner, Marsden Hanna, Joseph Hassine, Julie Hootkin, Eunice Huang, Vilma Kaza, Andrew Kim, Ross Koningstein, Sebastian Krier, Sami Lachgar, Jennifer McCloskey, Kerry McHugh, Aalok Mehta, Sydney Mild, Sepi Hejazi Moghadam, Sasha Moss, Mellisa Nitti, Bakul Patel, Sam Pettit, Michael Pisa, John Platt, Paul Plumeri, Ashley Rivera, Tanya Rodchenko, Juliet Rothenberg, Jessica Sapick, Hanna Schieve, Ondrej Socuvka, Nick Swanson, Rocío Vives, Don Wallace, Kate Weber, and David Weller.

We appreciate the extensive body of scientific work conducted by various teams, organizations and individuals, which has informed and enriched this whitepaper, including the published work we cite.

01	Executive Summary
02	Introduction: A New Era of AI-enabled Scientific Discovery
03	Pillar One: Infrastructure — Increase Access to AI Infrastructure
	1.1 National and International Pools of Data, Models, Compute, and Software
	1.2 Data Accessibility and Software Interoperability
04	Pillar Two: Investment — Invest in the Science of AI
	2.1 AI Programs and Initiatives that Drive Socially Beneficial Scientific Innovations
	2.2 Interdisciplinary Research to Strengthen Collaboration
	2.3 Building a Strong Pool of Talent
05	Pillar Three: Innovation — Implement Pro-Science and Innovation Legal Frameworks
	3.1 Harmonized Regulations for Responsible and Reasonable Use of Data
	3.2 Regulations that Facilitate Responsible AI Progress
06	Conclusion
07	Appendix: Illustrations of AI's Transformational Impact on Science



Executive Summary

Artificial intelligence (AI) is ushering in a new era of scientific discovery, revolutionizing how we conduct research, and accelerating the pace of scientific breakthroughs. AI's ability to process vast amounts of data, identify patterns, and generate novel hypotheses can help scientists answer questions about problems that were once thought to be unsolvable in our lifetimes.

These are not promises projected into a hazy future. AI is already advancing science and changing lives today. Machine learning and prediction technologies are helping diagnose [genetic diseases](#), [identify cancer](#) earlier, [develop vaccines](#), [predict weather](#), [wildfires](#), and [floods](#), and accelerate [progress towards sustainable development goals \(SDGs\)](#). Behind these practical tools are AI-based approaches to fundamental scientific research in numerous fields, including [physics](#), [neuroscience](#), [earth sciences](#), [oceanography](#), [space exploration](#), [fusion energy](#), and [climate science](#), among others.

Following the successes of today, there is much more to be gained from AI in scientific breakthroughs tomorrow. Governments around the world can take critical steps to empower scientists to discover new drugs, find treatments for cancer, detect and mitigate natural disasters, drive toward sustainable development goals (SDGs), and much more.

This paper outlines three specific policy pillars – the “three I’s” – that governments can pursue to ensure AI positively impacts science and benefits everyone:

- **Infrastructure — Increase Access to AI Infrastructure:**
High-quality scientific data and computational resources can unlock AI-powered scientific solutions. To address the uneven global distribution of these resources for scientists and to foster international collaboration, policymakers should establish national—and an international—centers to provide AI tools and resources for advancing science. These centers would provide access to data, computational, and educational resources for leveraging AI for scientific work.
- **Investment — Invest in the Science of AI:**
Sustained government funding is crucial for ambitious, long-term scientific research, such as space exploration or developing carbon management technologies. To get there, governments should create a list of priority areas to direct their funding and create public challenges for solving the most pressing issues. Novel public-private partnerships and funding models can play an important role in fostering a thriving AI-for-science ecosystem and building a strong pool of talent.
- **Innovation — Implement Pro-Science and Innovation Legal Frameworks:**
Science thrives on collaboration and the open exchange of information. Global legal frameworks can protect the marketplace of scientific ideas. Governments should work together to create harmonized, interoperable regulations to encourage continued scientific progress, especially in areas like copyright, privacy, and cross-border data flows. Governments aiming to lead on the use of AI for science should also craft immigration policies to attract and retain AI-skilled talent.


The recommendations in this paper are based on lessons learned from many years of scientific research across Google and Google DeepMind, and build on a [recent essay](#)¹, which explores how AI is transforming every scientific discipline. By taking decisive action along these “three I” policy pillars, governments can harness AI to unlock scientific advancement, establish their countries as leaders and beneficiaries of this progress, and work together to tackle global challenges. If we build this enabling policy environment, we can realize AI’s remarkable potential to achieve scientific breakthroughs.

¹ As part of our effort to expand and deepen public dialogue around science and AI, Google DeepMind co-hosted the [AI for Science Forum](#) with the UK Royal Society in late November. As part of this event, Google also announced a \$20 million fund for AI in Science to drive AI-enabled scientific discovery and incubate more achievements that will improve the lives of millions of people.

02

Introduction:

A New Era of AI-enabled
Scientific Discovery



Science has been the cornerstone of human progress, from Galileo’s telescope revealing the cosmos to the Large Hadron Collider’s potential to revolutionize physics and our understanding of matter. It has empowered us to understand our world, whether at the subatomic level or its place in the vast expanse of the universe. Science is not just a pursuit of knowledge; it is fundamental to improving everyday life, driving innovation, and shaping the future for everyone.

Many of the confounding problems in science benefit from the use of advanced computational techniques enabled by AI, thus making AI a powerful tool to assist scientists in their endeavors. The current revolution in science driven by AI is not about replacing human scientists, but about augmenting their capabilities to an unprecedented degree. AI’s true power lies in its ability to dramatically accelerate the scientific process, sometimes condensing hundreds or even thousands of years of traditional experimentation and research into

a few months or days, in areas like [drug discovery](#), [materials science](#), and [disease diagnosis](#). Moreover, AI tools are making complex analyses of vast amounts of data possible and accessible to a wider community of researchers than ever before, fostering collaboration and opening up entirely new avenues of inquiry. In essence, AI is not just making science faster, it’s making it broader, deeper, and more accessible, ushering in an era of discovery with the potential to address some of humanity’s most pressing challenges.

The global public sees value in current applications of AI and is most excited by AI’s ability to benefit society by powering scientific discovery, advancing healthcare and medicine. **7 in 10 people globally expect AI to have a positive impact on science (72%) and medicine (71%)** leading a list of promising applications, such as agriculture (60%) and cybersecurity (57%).

Source: [Google-Ipsos study on global public perceptions of AI](#) (Jan 2025)

Today, AI used in **natural sciences** is accelerating discoveries at an unprecedented pace. For example, in structural biology, it used to be that x-ray crystallography was the fastest route to insights about the structure of proteins. A single experiment could take years of work and cost \$100,000, depending on the protein. Now, Google DeepMind’s [AlphaFold](#), which predicts the 3D structure of proteins, has led to a publicly-available [Protein Structure Database](#) that provides instant access to 200 million predicted protein structures for free. This database is available publicly and has received more than 2.5M users from over 190 countries today. [AlphaMissense](#), which builds on AlphaFold, predicts the pathogenicity of missense variants—single letter substitutions in DNA, potentially accelerating diagnosis and the development of life-saving treatments for conditions like cancer. Furthermore, AI is transforming materials science, as demonstrated by Google DeepMind’s [GNoME](#), which has discovered millions of new crystal structures, accelerating advancements in areas like battery and semiconductor technology. GNoME successfully discovered 2.2M new crystals – equivalent to nearly 800 years worth of knowledge from classical research techniques. Imagine AI powering more scientific discovery tools like AlphaFold, AlphaMissense, and GNoME in the years to come and its effect on accelerated progress.

In **climate science and sustainability**, AI is proving to be an invaluable tool for modeling, forecasting, and mitigation efforts. AI enhances the analysis of vast climate data from sources like satellites and weather stations, leading to more accurate climate models, such as [NeuralGCM](#), which significantly outperforms traditional physics-based models in speed and accuracy. NeuralGCM models are publicly-available to the scientific community to build upon. AI also aids in predicting natural disasters like [floods](#), enabling advanced warnings and preparedness. Google’s flood forecasting AI model predicts floods in over 100 countries, helping a population of more than 700 million people. Additionally, AI is being used

for [wildfire detection](#). Google’s wildfire detection AI models provide a map of real-time boundaries of large wildfires every 10–15 minutes, currently available in 27 countries and expanding to more. [FireSat](#), an ongoing project to build an AI-powered satellite constellation, aims to detect wildfires anywhere around the world within twenty minutes and to allow scientists and AI experts to study fire propagation. AI can accelerate climate research to move us toward a sustainable future and [accelerate progress towards the SDGs](#). In fact, AI has the potential to help [mitigate 5-10%](#) of global greenhouse gas (GHG) emissions by 2030 — the equivalent of the total annual emissions of the European Union.

AI is also opening new avenues to scientific applications in the **energy sector**, from harnessing renewable energy sources to the promise of nearly limitless, clean energy from nuclear fusion. The variability and unpredictability of renewable energy sources like wind and solar has traditionally been a challenge in optimizing the power grid. AI is addressing this challenge by enabling real-time data analysis and predictive modeling, forecasting energy production from renewable sources with high accuracy. Smart grids use AI to dynamically balance supply and demand, optimize energy storage, and manage distributed energy sources, thereby mitigating the impact of fluctuations in renewable energy generation. AI is also playing a crucial role in nuclear fusion research, from discovering how to [contain the plasma in a tokamak](#) to [optimizing plasma performance](#) in a nuclear fusion reactor.

Healthcare presents some of the most promising use cases for AI. Advances in AI promise a step-change improvement in both disease detection and treatment, especially in low-resource communities and regions of the world. AI-powered technologies have already helped scientists [develop vaccines](#), diagnose [genetic diseases](#), and [identify cancer earlier](#). AI can accurately interpret retinal scans to detect diabetic retinopathy, a preventable cause of blindness that affects roughly 100 million people.

Google’s [flood forecasting AI model](#) predicts extreme riverine events in [over 100 countries](#) and areas where 700 million people live. This model is integrated into [Google’s Flood Hub platform](#), Google Search, and Google Maps for everyone to use.

[MedGemini](#) is another example of democratizing access to high-quality, personalized healthcare with [generative AI](#).

Quantum is another area where we’re beginning to see constructive feedback loops between AI and scientific discovery. Recent achievements in [quantum processors](#) and quantum computing are opening the possibility of studying questions that were previously the realm of science fiction, like studying the characteristics of [traversable worm-holes](#) and opening up new possibilities for testing quantum gravity theories. AI is advancing our progress in quantum physics and quantum computing, even as quantum is helping advance research in AI.

Incredibly, these remarkable advancements, unprecedented in human history, represent just the beginning of AI’s potential contributions to science. There is much more that the global scientific community can do to unlock the full potential of this technology to benefit humanity.

But scientists can’t achieve these benefits alone. They rely on an ecosystem that governments, in

partnership with academia and industry, can cultivate—an environment that provides much-needed infrastructure, robust funding, a steady pipeline of skilled talent, and relevant education and training programs. Beyond infrastructure and other resource investments, governments can address societies’ most pressing issues by establishing pro-innovation policies that enable foundational scientific and technological breakthroughs.

Without policies aimed at fostering AI-driven scientific progress, countries risk being left behind as other nations drive innovation and economic growth. Because the underlying AI technology empowering scientific advancements evolves rapidly, regulatory frameworks need to be technology-agnostic and adaptable to keep pace and avoid becoming obsolete. We share policymakers’ commitment to managing AI’s risks with appropriate guardrails and we emphasize the [need for a balanced regulatory approach](#) that [prioritizes research, innovation, science and technology](#) while establishing regulations to address risks.

This policy paper is designed to serve as a blueprint for policymakers around the world to establish policies that unlock AI’s benefits for scientific progress. We outline policy recommendations across three pillars:

1 **Infrastructure** Increase Access to AI Infrastructure

National and international resource centers that provide AI tools and resources for advancing science


2 **Investment** Invest in the Science of AI

Sustained government funding and novel public-private partnerships for pursuing ambitious scientific goals

3 **Innovation** Implement Pro-Innovation Legal Frameworks

Harmonized, interoperable regulations to encourage continued scientific progress and attracting investment

Together, we can work to solve challenges that we once thought were insurmountable. But first, we need a policy environment that enables the transformative power of AI for scientific discovery.




*For additional examples of AI-driven scientific breakthroughs see [Appendix](#)

03

Pillar One:

Infrastructure — Increase Access to AI Infrastructure

- 1.1 National and International Pools of Data, Models, Compute, and Software
- 1.2 Data Accessibility and Software Interoperability



AI adoption in scientific contexts is particularly challenging due to several factors, such as access to high-quality, specialized data, constraints on access to computing power, and the need for interdisciplinary collaboration among AI and domain experts.

Lacking an orchestrated infrastructure for AI-powered scientific research and development (R&D), scientists must spend significant time and effort to coordinate data and model access, secure computing power, and become proficient with AI tools, all of which detracts from their core research activities. Although most scientists will not need to train their own large AI model, they will nevertheless need access to resources to fine-tune large models on scientific knowledge, run simulations to generate high-quality data, or [train \(often relatively small\)](#) AI models on their specialized data.

Lowering the entry barriers to developing, using, and deploying science-focused AI techniques is essential for maximizing access to resources for scientists around the world, including in less-well-resourced institutions or countries. As the Royal Society explained in its recent [AI in Science report](#), “Scientists from diverse disciplines require access to infrastructure to adopt more complex AI techniques, process higher volume[s] and types of data, and ensure quality in AI-based research”. There is therefore a growing need to expand access to compute capacity, data, AI models, software, and tools to a wider range of researchers in order to accelerate collaborative, AI-driven scientific research.

Worldwide, many governments have recognized this need and have begun creating national resources and data repositories for AI research. For instance, [INDIAai](#) serves as a centralized repository of AI datasets across critical sectors like health-care, agriculture, finance, and education, aiming to democratize access to high-quality data and spur innovation. In the U.S., [Data.gov](#) aggregates datasets from various government agencies. Efforts from the private sector, such as Google’s [Data Commons](#), has also contributed to democratizing data access by centralizing and streamlining publicly available data from diverse sources. Academic institutions also have collaborated and established joint data repositories, like the [National Data Platform](#), to boost researcher access to both data and funding. More recently, an industry-academia collaboration has resulted in a large-scale scientific data repository, called [The Well](#), which contains 15TB of physics simulations covering diverse scientific domains. Although these efforts have increased data availability, there is still much to be done to streamline access to various data sites, models, prior scientific art, and simulation tools for scientific research.

Scientific discovery typically relies heavily on simulating physical phenomena. High-performance computational models and simulations, combined with data from experiments and observations, are essential to advance our understanding of physical systems and processes. Building solutions to simulate complex systems has been one of the primary missions of the science and engineering community over the past 50 years. Scientists are [already leveraging AI](#) to generate reliable simulations across various fields, such as aerospace, automotive, electronics, and optics. They can also accelerate data selection, curation, and even scientific data generation through simulations. Although many tools exist for building simulations (e.g., Ansys [multiphysics simulation](#)), the power of AI is not yet fully harnessed to transform large-scale simulation of complex phenomena in physics, biology, chemistry, and other scientific domains.

Platforms that enable the sharing of AI models are also essential for fostering an open ecosystem that accelerates AI adoption and encourages collaboration across diverse fields. For instance, [VertexAI Studio](#) and [HuggingFace](#) facilitate access to AI models and rapid prototyping. However, to fully leverage AI for scientific progress, new innovations are needed to provide seamless, AI-ready platforms for scientific research and discovery. Below, we recommend policy options to establish and expand an infrastructure supportive of AI-powered scientific research.

1.1 National and International Pools of Data, Models, Compute, and Software

Create national AI for science resource centers to make data, AI models, compute capacity, software and tools accessible for scientific research. These national resource centers could also coordinate with an international center designed to facilitate the use of AI for global scientific collaboration and progress.

Countries must build the infrastructure necessary to empower their scientists with the tools they need to be dramatically more effective in solving pressing challenges like disease eradication, sufficient supplies of energy, and food security. Accomplishing this will require making AI-enabled research tools and resources *more accessible to more scientists in more places*. Public-private

partnerships to build infrastructure for science is an enabler for AI-powered scientific progress by reducing unnecessary obstacles to innovation and fostering broader collaborations.

The U.S. government, in partnership with the private sector, has started to pilot the [National AI Research Resource \(NAIRR\)](#), a research infrastructure designed to democratize access to AI R&D. By providing computational, data, software, and training resources, NAIRR aims to facilitate AI research. While currently in its pilot phase and requiring further funding to achieve its full potential, NAIRR serves as a valuable model for nations seeking to establish a robust infrastructure for AI-powered scientific research. Recently, the UK in its [AI Opportunities Action Plan](#) laid out its strategy to expand the capacity of its AI Research Resource (AIRR) by at least 20x by 2030. At the global level, the UN has [recommended](#) a Global Fund for AI and a Capacity Development Network, which aim to address the underlying capacity and collaboration gaps for countries with limited AI resources and provide access to AI expertise, computing power, and training data. Governments could build on these models to increase access to resources for scientific research both nationally and globally.

Infrastructure supportive of scientific research, however, requires resources beyond just compute to empower scientists from diverse backgrounds to effectively access data, conduct experiments, compare existing solutions, and build upon prior scientific work. National science resource centers—and an international analogue—could encompass high-performance computing (HPC) clusters, shared scientific datasets, AI models (e.g., those used for climate modeling), tools for simulations and data generation, and utility software that enable seamless integration of AI solutions in scientific research (e.g., [managing data dependencies](#)). Although the scientific community has made strides towards integrated data delivery and access and shared storage and computing resources for scientific data (for example, see the [National Science Data Fabric](#)), governments should not stop at these few efforts, but continue to dedicate funding and work toward expanding such efforts to meet scientific demand.

Many scientific breakthroughs have resulted from **international collaboration**. A global pool of shared resources for scientific research would foster such collaboration in different fields. This pool would enable scientists to access data, collaborate globally, efficiently execute AI-based experiments, and advance resource-heavy scientific discoveries towards solving the most pressing local and global issues. An international AI for science resource center would provide infrastructure for national centers to share best practices and coordinate resource needs in the pursuit of AI-powered scientific discovery.

AI for science resource centers could undertake a range of activities. For example, they could:

- Issue periodic requests for proposals in priority research areas that would allow scientists and organizations to apply for compute time.
- Solicit proposals focused on creating high-quality scientific datasets where gaps exist.
- Host public challenges that necessitate collaboration, creativity, and access to diverse data.
- Provide educational resources to empower scientists who are novices in AI to leverage AI tools.
- Pool data from diverse geographies, economies, and cultures to create representative, high quality datasets for use in key disciplines e.g. healthcare.

Leverage public-private partnerships to build training materials for scientists to use AI infrastructure.

Navigating advanced AI tools, national or global computational resources, or hyper-scale supercomputers is not trivial for those whose expertise lies outside of AI. Tutorials aimed at scientists could guide learners from basic AI skills to managing large-scale data and software. This allows science students and professionals to start on local devices or small cloud platforms and gradually progress to larger cloud or institutional resources, eventually working with national and global-scale infrastructure. [Systematic upskilling is crucial](#) to enable widespread access and proficient use of AI tools across the scientific community.

Public-private partnerships can play a critical role in building training programs for scientists to efficiently leverage AI in their workflows. Government science agencies should partner with companies and universities developing AI technology to build training programs for scientists on how to use AI. For example, Google DeepMind partnered with EMBL-EBI to provide an [accessible training course](#) for structural biologists on how to use AlphaFold2 in their work, including understanding its limitations. Similar to NAIRR, which provides [educational resources](#) for trustworthy AI research, **AI for science resource centers should also include educational resources, such as courses and training programs**, to help users easily run scientific experiments using the platform.

1.2 Data Accessibility and Software Interoperability

Increase the accessibility of government datasets and data from government-funded research.

Government institutions often have access to (or can generate) scientific data that can accelerate innovation across disciplines. Publicly funded data initiatives have driven many advancements in different fields. For example, [NASA’s EMIT](#), [Germany’s EnMAP](#), and [Japan’s HISUI](#) offer invaluable data for developing AI models for climate science. Government data in medicine, agriculture, energy use, and space exploration can similarly accelerate research. Governments should adopt open data policies, making publicly funded research data and government datasets available through consolidated repositories to improve access and enable combining disparate (or siloed) data to create AI-ready, large and diverse datasets. The [National Cancer Institute](#) (NCI)’s [Cancer Research Data Commons](#) initiative, which has compiled more than 10 petabytes of data from hundreds of NCI-funded programs and studies as well as data from select external cancer research programs for cancer research, serves as another example of government facilitating access to scientifically useful data.

To encourage accessibility of government data, the UK government has adopted [an open data policy](#), known as the “**Open by Default**” policy, that encourages the release of public sector data to increase transparency and support economic growth. Additionally, the [UK AI Opportunities Action Plan](#) lays out the country’s strategic plan to create the National Data Library to make high-impact public datasets available to AI researchers and innovators and build public sector data collection infrastructure. In the U.S., [Data.gov](#) offers a large catalog of U.S. government datasets to drive innovation. We recommend that government-funded scientific research should by default result in publicly available data and methods (unless the risks of data sharing outweigh the benefits). Additionally, governments should allocate **grants to early-stage research that focuses on scientific data creation and curation**, rather than primarily funding research that has already shown promising results.

Governments should develop a national data strategy to dramatically expand the availability of data for training AI models. They could also work with the private sector to [establish scientific data-sharing consortia](#) wherein participants contribute their data in exchange for access to the collective datasets held by the consortium.

Although data sharing fuels scientific advancements, a cautious approach is crucial. We must acknowledge that not all data is suitable for widespread distribution. Sensitive data, such as healthcare records, and data that could compromise the privacy of individuals, confidentiality, or proprietary information or trade secrets should not be mandated for sharing. Additionally, significant caution is necessary when considering the sharing of data that could be misused in high-risk applications, particularly those with implications for national security or potential misuse in developing dangerous capabilities, such as biological or chemical weapons development. However, best practices in data management and information security can help governments responsibly manage sensitive datasets that can be leveraged for beneficial scientific research and applications.

Create a list of priority areas for which the scientific community lacks data.

Governments should establish clear roadmaps outlining their most pressing challenges and develop a list of high-priority scientific areas and problems to solve. Such a targeted approach not only focuses resources on the most critical areas but also encourages collaboration between public institutions, AI companies, and research organizations to address identified challenges. Governments can then direct their R&D funding or data-related efforts to create (or improve existing) curated, publicly available data for advancing scientific progress within identified priority areas. Priority areas could differ for each nation depending on the pressing issues that governments can be facing locally.

Facilitate the creation of anonymized, aggregated, and varied datasets in sensitive contexts.

Personal information must be safeguarded when building scientific datasets. In sensitive contexts, such as medical science and healthcare, governments should facilitate the creation of fully anonymized, broadly representative datasets that can be used to train AI models and address gaps in public goods and services. The sensitive nature of healthcare data and privacy regulations pose significant challenges for model developers seeking to collect representative datasets or assessing AI models across demographic groups. National institutions, such as the National Institutes of Health, with their capacity for data stewardship, are uniquely positioned to anonymize, aggregate, and disseminate high-quality, representative datasets in healthcare and medicine, along with guidelines and best practices in creating clinical datasets that address all patient populations, empowering scientists to develop effective health AI systems without violating privacy rights.

Encourage data portability and software interoperability.

Data portability and software interoperability are crucial for transferring and using data, models, and software on different computing platforms. Interoperability means data can be processed by different services on different systems due to common specifications, and portability means that data can be moved and applications can be ported and run on different systems. Governments should encourage cloud providers to offer portability and interoperability between providers (See Google Cloud's commitment to portability and interoperability [here](#)). This flexibility has several benefits, including reducing reliance on a single provider, and improving resilience against single provider outages. Additionally, portability across platforms allows easy access to new and emerging cloud technologies, a significant factor for those using cloud computing for scientific research.

Policymakers should facilitate portability of data and interoperability of software to **provide choice and avoid locking researchers into one specific service provider.**

04

Pillar Two:

Investment — Invest in the Science of AI

- 2.1 AI Programs and Initiatives that Drive Socially Beneficial Scientific Innovations
- 2.2 Interdisciplinary Research to Strengthen Collaboration
- 2.3 Building a Strong Pool of Talent




Government investments are crucial for sustaining an ambitious and long-term scientific research agenda, particularly when such efforts do not yield immediate commercial returns. In addition to increasing budgets across high-priority scientific domains, governments should invest in areas where private capital may not see sufficient near-term return on investment.

Large-scale projects with ambitious, decades-long goals often require more public investment. Constructing [extremely large telescopes](#), building high-speed [particle accelerators](#), [restoring coral reefs](#), developing and deploying [carbon management technologies](#) are examples of scientific areas that require public support due to the long-term nature of the efforts needed.

Public-private partnerships could also be instrumental, focusing on applying AI to critical areas like drug discovery, climate modeling, genomics research, materials science, and development of personalized medicine. These partnerships leverage the strengths of both sectors, combining public funding and academic expertise with the private sector's resources and ability to commercialize research findings.

Government investments can take several forms, such as **(a) creating programs and initiatives, (b) supporting interdisciplinary research, and (c) building a strong pool of talent.** Investing in these areas not only provides the necessary resources for scientific advancements but also fosters collaboration among interdisciplinary groups. Creating public grand challenges, for example, elevates important problems that need to be solved and encourages multidisciplinary research collaborations to invent novel solutions.

A critical area for government investment is talent development. It is vital to train and prepare existing and next-generation talent for innovation and entrepreneurship in the AI era. Governments should invest in multidisciplinary research and education to equip the next generation of scientists with the skills needed to drive progress in AI and other scientific fields. A holistic investment approach to talent development will ensure a robust pipeline of AI-skilled talent to tackle pressing societal challenges.



2.1 AI Programs and Initiatives that Drive Socially Beneficial Scientific Innovations

Extend existing and establish new financial support mechanisms for small businesses advancing AI-enabled scientific applications.

Financial support mechanisms, such as grants, public procurement contracts, [“jump start” programs](#), and matching funds, can stimulate private investment and encourage startups to engage in scientific research. Programs like the [Small Business Technology Transfer \(STTR\)](#) help small businesses and incentivize commercialization by providing funding opportunities to participate in R&D with nonprofit research institutions. Government funding is particularly vital for organizations tackling high-impact but commercially challenging issues like climate change. Funding and facilitating research collaborations between startups, academia, government labs, and the private sector, along with providing access to public datasets and computational resources, can greatly reduce entry barriers for businesses with limited resources, enabling them to innovate with AI.

Increase funding for high social-impact projects to fill gaps in private investment

For successes in research to reach the people who need them, it is imperative to speed the transition of innovative ideas into practical, accessible solutions. For example, the COVID-19 pandemic highlighted the potential for greater public sector involvement in late-stage drug development and production. Public funds to [late-stage biomedical projects](#) with high social impact could bridge the gap in funding where private investment is unlikely to drive commercialization. The [Investor Catalyst Hub](#), as part of the [ARPA-H Health Innovation Network](#), is an example of how governments can scale research and venture investment ecosystem capabilities to help innovators bring ideas to market.

Create public Grand Challenges that focus on solving scientific problems.

“AI Grand Challenge” programs, such as those promoted in the [AI Grand Challenges Act \(S. 4236\)](#) in the U.S., spur innovation in science and technology by **(a) identifying and prioritizing urgent problems** and **(b) creating a competitive environment for groups of interdisciplinary researchers** to collaborate. Funded by governments or the private sector, these challenges should focus on technical innovation in applications of AI that would fundamentally transform the process of science, engineering, or medicine. [Research shows](#) that challenge prizes attract unconventional entrants, lead to novel R&D approaches, and induce innovations beyond what would have occurred without the prize.

Public and private grand challenges have had an important role in driving innovation. The [DARPA Grand Challenge](#) for autonomous vehicles in 2004 and 2005 dramatically accelerated the development of self-driving car technology. It spurred significant investment and innovation in the field, leading to the advanced driver-assistance and autonomous driving systems we have today. [The Ansari XPRIZE](#) successfully incentivized the development of private spaceflight. [Grand Challenges Canada](#) combines scientific, technological, social, and business innovation to address global health and humanitarian issues. It supports innovators in low- and middle-income countries and Canada. The UK Research and Innovation ([UKRI](#)) [Challenge Fund](#) addresses major societal challenges through 23 different challenges, aligned with the UK government’s industrial strategy, including clean growth, aging society, future of mobility AI and data economy.

Some areas that governments can launch public grand challenges for science and society [include](#) mitigating the interconnected issues of population growth and overconsumption, ecosystem destruction, climate change impacts (e.g., wildfires, extreme weather, sea-level rise, and ocean acidification), and pollution (air pollution and contaminated drinking water). [Additional areas](#) that challenge prizes could accelerate innovation include AI for healthcare, education, and national security. Furthermore, governments can incentivize the development of solutions for combating pandemics and disease outbreaks through grand challenges, or design challenges at the intersection of AI and Quantum Computing, similar to the [XPRIZE challenge on Quantum Applications](#), or AI and the UN’s Sustainable Development Goals to spur innovation in solving humanity’s long-standing problems.

Strengthen public sector procurement to drive AI innovations.

The benefits of integrating science-driven solutions into public services are already evident worldwide. Several cities worldwide have adopted Project [GreenLight](#), utilizing AI-powered traffic light control to curb traffic emissions and optimize urban mobility. Furthermore, innovative public-private partnerships, like the collaboration between [Chile’s national system operator \(CEN\)](#) and Alphabet’s [Tapestry](#) project aims to accelerate the country’s transition to renewables. This partnership empowers CEN to decarbonize Chile’s energy system and achieve its goal of becoming carbon neutral by 2050.

Project Cost

\$3b

estimated

Overall Impact

\$965b

economic output

The Human Genome Project was an international research effort, launched in 1990 and took 13 years to complete, enabling advances in medicine and disease diagnosis.

Ensure robust and sustained funding to pursue ambitious goals over an extended period.

A sustained research budget is essential as it would provide organizations and researchers the stability and investments necessary to pursue ambitious, long-term agendas and achieve meaningful breakthroughs. Policymakers should ensure that funding provided for AI research extends over enough time to allow for research agendas to evolve, mature, and yield tangible outcomes. For example, the [Human Genome Project](#) was an international research effort with the goal to generate the first sequence of the human genome. It launched in 1990 and took 13 years to complete. The knowledge gained from this project has accelerated the study of human biology, enabling advances in medicine and disease diagnosis. The cost of this project, while in the billions of dollars (approximately \$3 billion), has been greatly offset by the positive economic benefits that genomics has yielded in the ensuing decades. In 2013, the economic impact of the project was [estimated at \\$965 billion](#) (direct and indirect) along with 4.3 million job-years of employment. Similarly, the Internet, as we know it today, began as a government-funded research project called [ARPANET](#) in the 1960s. The U.S. government's sustained investment in basic research and infrastructure development was crucial for the internet's initial creation and subsequent growth. Long-term public funding commitment is not just a recommendation, but a prerequisite for transformative change. Policymakers should recognize that sustained funding in R&D is not merely an expense, but a strategic investment that promises exponential returns for the future.

Amplify the impact of nonprofit organizations that use scientific advances to benefit communities.

Another way for governments to encourage socially beneficial combinations of AI and scientific research is in the nonprofit sector. Governments should support and help scale up nonprofits that use scientific advances to benefit communities. Support can be in the form of grants and financial aid, or more importantly, collaborative efforts for scaling up solutions. For example, the Indian government is scaling a [nonprofit-developed precision agriculture tool](#), which has boosted farmer profits by 20% and reduced pesticide use by 25%. Additionally, equipping the nonprofit sector with AI tools can boost their efficiency and amplify their impact. Initiatives such as the [AI for Changemakers](#), [AI Collaboratives](#), and [Google for Startups](#) programs highlight the value of partnerships between tech companies and nonprofits, offering tailored AI tools and expertise to scale organizations' efforts.

2.2 Interdisciplinary Research to Strengthen Collaboration Across Disciplines

Establish and fund interdisciplinary research centers and programs in AI for science to tackle complex scientific challenges.

Interdisciplinary collaboration between AI and scientific domain experts is crucial for developing more effective and accurate AI solutions in science. However, this is sometimes [hindered by siloed research environments](#) and a lack of incentives for interdisciplinary work in career advancement. [MIT-Google Program for Computing Innovation](#) is an example of how industry-academia collaborations bring interdisciplinary teams together to accelerate research on global societal and sustainability challenges. By establishing interdisciplinary research centers and programs, governments can incentivize collaboration between AI experts and non-AI scientists, integrate knowledge from various fields, and strengthen cross-sector innovations. The U.S.'s National Science Foundation (NSF) has a successful model for establishing [interdisciplinary institutes](#) to motivate interdisciplinary research to solve complex research problems, although sustained funding is required to support research centers and pursue long-term research agendas.

Fund research at the intersection of AI and other emerging technologies.

Converging advancements in AI with fields such as quantum computing (QC), biotechnology, and nanotechnology present unprecedented opportunities for synergistic breakthroughs. QC is a multidisciplinary field at the intersection of computer science, physics, and mathematics that seeks to use the information processing power of quantum mechanics to solve otherwise difficult computational problems. Although the private sector has made significant strides in QC and [quantum chip development](#), government funding can catalyze further private funding and facilitate research on new applications that can benefit from the power of QC to address complex, global challenges in climate, sustainability, health or novel AI-driven solutions. What makes the intersection of AI and QC unique is that advances in AI can power progress towards a stable quantum computer, and breakthroughs powered by QC would advance AI itself by uncovering novel algorithms or solving computationally infeasible problems. Similarly, the integration of AI with nanotechnology could enable the development of intelligent nanoscale devices for [targeted drug delivery](#) and [precision agriculture](#). Investments at the intersection of AI and other emerging technologies would enable researchers to tackle big questions in areas such as neuroscience, precision medicine, cosmological physics, [quantum networking](#), and more.

2.3. Building a Strong Pool of Talent

Invest in AI education across all scientific disciplines to develop the next generation of AI-skilled scientists.

Governments should increase investments in scholarships, fellowships, and grants for graduate students and early-career professionals working at the intersection of AI and scientific fields through both public funding and public-private partnerships. The U.S. [NSF AI Education Act of 2024](#) highlights the private sector's role in supporting AI education by creating fellowships and scholarships for students and professionals, including those at community colleges. The Google [PhD Fellowship Program](#), for example, has supported 800+ aspiring researchers from 210+ universities in 53 countries,

since 2009. Additionally, governments should lead in training educators in AI concepts through updated professional development programs. Universities and colleges also have a pivotal role to play by integrating AI into their curriculums. By offering dedicated AI programs and interdisciplinary courses that blend AI with other scientific disciplines, academic institutions can equip the next generation of scientists with the skills necessary to navigate the AI-driven landscape of the future.

Compete for talent across STEM fields, wherever it is.

Students, researchers, and scientists across the Science, Technology, Engineering, and Mathematics (STEM) fields come from everywhere, and governments should give themselves every opportunity to attract candidates from the highly competitive global pools of both science and AI talent. To compete for talent, governments could increase access to training programs and research infrastructure, proactively recruit international researchers by making visa programs easy to navigate, and educate the next generation of AI and scientific experts at home. [Google's AI Opportunity Fund](#), for example, aims to train one million Americans of all backgrounds to provide them with AI skills at no cost and to increase equitable access to AI training and resources.

Support nonprofit organizations that provide AI skilling programs.


Increasing AI skills across national economies will boost societies' capacity to unlock and adopt scientific applications. Nonprofits that focus on AI upskilling and reskilling play an important role in increasing the penetration of AI skills within communities, as well as preparing the next generation of AI innovators, developers, and adopters to build with and use AI. [Apolitical](#), as an example, is a nonprofit that has launched the [Government AI Campus](#) to bring AI skills to policymakers, whereas [aiEDU](#) and [Mind the Gap](#) train students and entrepreneurs on AI skills. Governments should support nonprofits that offer AI educational and skilling programs. This support can be in the form of grants and contracts, providing training/educational materials, or simplifying grant application and evaluation processes, and data-driven grant allocation to empower education-focused nonprofits.

05

Pillar Three:

Innovation — Implement Pro-Science and Innovation Legal Frameworks

- 3.1 Harmonized Regulations for Responsible and Reasonable Use of Data
- 3.2 Regulations that Facilitate Responsible AI Progress



Policymakers around the world have attempted to keep pace with digital innovation by writing new laws for AI even as they examine the suitability of existing sectoral regulations. The task is challenging not only because AI is a general-purpose technology that will affect a wide range of economic and social activities, but also because the implications of AI vary depending on the context for its use. And as countries consider new rules about AI technologies, they risk creating a patchwork of conflicting regulations at home and around the world.

Despite efforts among some policymakers to ensure research and development is not over-burdened by new requirements, scientists are not exempt from regulatory ambiguity. In fact, uncertainty is especially burdensome for scientific researchers, who are already grappling with the challenges of simple access to AI tools. Meanwhile, regulatory fragmentation can slow scientific collaboration across borders.

Clarifying the applicability of existing laws (e.g. around copyright, privacy, and the use of data) and harmonizing national and global legal frameworks would provide the certainty innovators and investors need to ensure compliance and move from research to marketable applications confidently and at scale. **Risk-based, sector-specific, and technology-agnostic legal frameworks** would facilitate global collaboration, data sharing, and compliance, enabling institutions and businesses worldwide to innovate swiftly—and safely.

3.1. Harmonized Regulations for Responsible and Reasonable Use of Data

Create (or maintain) copyright frameworks that enable the safe use of publicly available information for training and testing AI-powered systems.

To accelerate the advancement of AI in science, particularly in applications that [sift through scientific literature for insights](#) or [hypothesis generation](#), balanced copyright laws that ensure access to publicly available scientific papers are essential. Copyright laws, such as the fair use doctrine and text and data mining (TDM) exceptions, allow for uses of copyrighted, publicly-available material for AI training without lengthy negotiations with data holders during model development or scientific experimentation. Where the local laws are not prescriptive, as in the United States, the judicial branches may be called upon to provide certainty for how intellectual property laws interact with AI.

Several countries, including [Singapore](#), [the US](#), and [Japan](#) have recognized the importance of limitations and exceptions to copyright laws for advancing AI technology development. Further, the European Union’s Copyright Directive includes provisions for TDM for both research and commercial purposes, recognizing the substantial potential benefits of AI, while also giving content owners the ability to opt-out of having their content used for AI training. Fair use and TDM exceptions around the world support innovation by carefully balancing protections for creators with the need for innovation and cumulative creativity, while ensuring that developers are able to assemble the building blocks needed for the development of AI.

It’s important to clarify that AI models are not simply compressed copies of their training data. Training models is a transformative, non-expressive use of training data that should be recognized by copyright law as fair in jurisdictions that have fair use, or as eligible for a TDM exception in jurisdictions that rely on enumerated exceptions.

Establish harmonized data privacy laws that focus on responsible and reasonable data collection and use.

Pro-innovation privacy laws should balance safeguarding personal data and enabling technological progress. Policymakers should apply privacy protections proportionally, ensuring user privacy is protected while taking the benefits of AI into account. It is especially crucial for policymakers to weigh data-protection rights against the societal benefits of scientific advances. Privacy regulations should aim to be adaptive, risk-based, and technologically neutral, and focus on mitigating the potential harms of outputs, rather than regulating the inputs used in development. They should stay future-oriented, engaging with scientists and industry as new privacy issues arise with evolving technologies and their scientific applications. Additionally, policymakers should harmonize efforts internationally and across regulatory domains to avoid fragmentation.

To further bolster privacy and innovation, governments should actively encourage the development and adoption of privacy-enhancing technologies (PETs). For example, [Singapore’s National AI strategy](#) emphasizes the need to invest in PETs to address barriers around data protection and sharing, and to encourage more data flows. Investing in PETs can unlock the value of data for scientific developments, for example, in medical science and healthcare, while respecting individuals’ privacy and data confidentiality.

We refer readers to Google’s Policy Recommendations on [Generative AI and Privacy](#) for additional recommendations and details.

Strengthen trade policies and support cross-border data flows.

Trade is critical both to the development of AI research and to the diffusion of scientific applications of AI around the world. Nearly a quarter of [scientific research papers](#) published today involve authors from multiple countries – and that rate of cross-border collaboration goes up to 40% in countries like the US that are leading on the development of AI. Cross-border collaboration is further enhanced by AI software libraries that are publicly available, such as Google’s [TensorStore](#) and [Flood-Filling Networks](#) which are widely used in neuroscience, or [DeepVariant](#) and [DeepConsensus](#)

frequently used in genomics. These libraries enable researchers to rapidly prototype ideas through working code rather than just research papers.

Preserving cross-border data flows is particularly crucial for advancing AI-powered scientific discoveries, as they enhance the volume, variety, and diversity of data. AI depends on the ability to rapidly access large amounts of data (much of which may be stored overseas) and rapidly exchange information between different research groups, often located around the world. But barriers to data flows, such as forced data localization – which requires organizations to manage, store, or process data locally – can significantly increase costs for researchers while slowing down collaboration.

Using and aggregating data from different regions also allows researchers to access a broader range of samples, variables, and conditions, leading to robust solutions that generalize well to real-world scenarios. This is particularly important in fields such as medical science and drug discovery, where data diversity and representativeness are key to reducing potential biases and ensuring that scientific breakthroughs benefit all populations.

Collaboration between partners with shared values facilitates scientific breakthroughs by strengthening trade, investment, and trusted cross-border data flows. For example, Asia-Pacific economies have been incorporating AI provisions into digital trade agreements. Non-binding commitments to collaborate on safe and unbiased use of AI have been included in agreements such as the UK-Singapore Digital Economy Agreement, the Korea-Singapore Digital Partnership Agreement, and the UK-New Zealand Free Trade Agreement. Additionally, bilateral AI partnerships are emerging, such as the UK's efforts to collaborate with Korea and Singapore on identifying “trustworthy” AI uses, that could benefit all participating nations by accelerating AI research and development.

At the same time, governments should explore next-generation trade control policies for specific applications of AI-powered software that are deemed security risks, and on specific entities that provide support to AI-related research and development in ways that could threaten global security.

3.2. Regulations that Facilitate Responsible AI Progress

Adopt a proportionate and risk-based approach to regulation that is focused on the context of use and actual risks, rather than the underlying technology.

National legislation should enable countries to benefit from the economic, scientific, and societal opportunities unlocked by AI, while mitigating potential risks. By focusing regulation on potentially harmful use cases where known AI risks are likely or could have significant impacts, countries can implement a proportionate approach that protects people and makes room for scientific advancements. The [EU AI Act](#), for example, acknowledges that most AI use cases are benign and recognizes high-risk areas where systems should be scrutinized with greater rigor. Although the EU AI Act exempts research and development from the requirements of the act, its risk-based approach is especially suitable for bringing scientific research to market, where the context of use of a scientific solution may implicate different types and levels of risk depending on the final application.

Take steps to mitigate risks of dangerous outcomes across the entire scientific ecosystem.

A concern shared by policymakers in many countries is that advanced AI models may present dangerous capabilities that could lower the barrier to entry for non-experts to design harmful materials or compounds. It is important to address this concern through targeted and consequential interventions across the scientific enterprise, recognizing that the use of an AI system could be just one of many activities in a complex effort to build illicit weapons. Scrutinizing access to certain materials and high-level labs is part of a comprehensive approach to research security while preserving the possibility for AI to support beneficial use cases in chemistry and biology.

Exempt pre-market AI R&D for scientific purposes from AI regulations.

Early-stage scientific research using AI should not be burdened by the same regulations intended for market-ready products. Regulations should exempt activities related to R&D for scientific purposes prior to placing products on the market. The European Union AI Act, for example, includes exemptions for “any research, testing or development activity regarding AI systems or AI models prior to their being placed on the market.” This approach ensures that regulations address the real-world risks associated with applications deployment, not the exploratory nature of scientific inquiry. Regulatory efforts should prioritize use cases that pose significant risks of harm or misuse due to deployment, rather than imposing blanket restrictions during testing and development, which could stifle innovation and delay beneficial technologies from reaching the market.

Foster international cooperation frameworks and embrace scientific community guidelines for AI-powered research.

International cooperation frameworks can provide a forum for participating nations to collaborate on responsible and safe development of AI-driven technologies. Through collaboration, countries can explore mutually beneficial approaches that prioritize safety, security, and preventing significant societal harm.

Community-driven guidelines, such as the [Hague Ethical Guidelines in chemistry](#) and professional codes of conduct, such as the [international medical code of ethics](#), have traditionally helped develop responsible and ethical technologies without the need for prescriptive laws. In medical science, for example, biomedical ethicists have a long history of assessing benefits and risks to inform decision-making for new technologies. Life sustaining technologies, such as ventilators and feeding tubes, and reproductive technologies, such as IVF, IUI, PGD, MRT, CRISPR-Cas9, are two classes of medical scientific research where the pace of technological innovation was far ahead of regulation, and the bioethics field helped to inform responsible and equitable use. Governments should embrace

these community-established guidelines as guiding principles that scientists follow to ensure their research is safe and trustworthy, and develop regulations as technologies mature and in consultation with the scientific community, industry, academia, and nonprofit organizations.


Establish vendor-agnostic regulatory sandboxes for fast experimentation.

Regulatory sandboxes offer a key way to accelerate the transfer of research into tangible products and services. Participating firms obtain a waiver from specific legal provisions or compliance processes to be able to rapidly innovate. Sandboxes are particularly beneficial for resource-constrained small businesses and startups, helping them navigate complex regulatory landscapes. To that effect, many economies such as [the UK](#), [Spain](#), [Singapore](#), and [Norway](#) have already piloted AI regulatory sandboxes to enable businesses to experiment with new products or services and regulators to gain practical insights from the early stage technology when considering different regulatory frameworks.

Similarly, to encourage climate solutions, [Spain](#), [Brazil](#), and [Australia](#) have launched regulatory sandboxes for their electricity sectors that “aim to help energy innovators and start-ups navigate complex regulatory frameworks and enable the trial of new products and services that will deliver greater choice and cheaper energy options for consumers.” The Australian Energy Innovation Toolkit framework enables regulators to issue trial waivers, which “temporarily exempt an innovator from having to comply with specific rules that are acting as regulatory barriers to allowing an innovative trial to proceed.”

Sandboxes not only enable fast development and trial of new products, but also enable regulators to explore different regulatory frameworks. For example, The US Food and Drug Administration (FDA) piloted a [software pre-certification program](#) regulatory sandbox that helped inform the development of an adaptive regulatory approach. It is crucial, however, that these sandboxes remain vendor-neutral to avoid locking users into a specific vendor’s platform and limiting their choices and flexibility.

06 Conclusion



AI is a breakthrough in how we make breakthroughs. But the scientific benefits of AI are not guaranteed. Public policies will play a significant role in shaping whether and how societies achieve AI's full scientific potential, and whether we can make new discoveries and apply them toward practical solutions to health, energy, and other key challenges.

AI has the potential to democratize access to science and empower more researchers to participate in scientific research. Bringing more scientists onboard widens the range of problems that can be tackled at digital speed and at a scale which offers solutions to humanity's longstanding challenges. To foster a thriving ecosystem for AI-powered scientific progress, policymakers can increase access to AI infrastructure, invest in the future of AI for science, and implement pro-science and innovation legal frameworks. This approach should begin with the assumption that existing law is a strong foundation for the safe use of AI, and build toward global alignment on national-level regulation based on common standards for AI safety across borders. Open data policies, coupled with consistent privacy laws and permissive copyright frameworks, are essential for enabling broad researcher access to the data that fuels AI innovation. To ensure the

greatest possible chance for scientific breakthroughs, nations should establish AI for science resource centers, along with a multinational infrastructure, providing safe access to scientific data, computational resources, and educational materials. And governments should invest in AI education across all scientific disciplines, direct funding towards interdisciplinary research projects, and create public grand challenges focused on solving critical scientific problems to encourage collaborative innovation and develop a strong pool of talent.

By working together to promote responsible innovation, international collaboration, and broad-based access to AI resources, governments can unlock the transformative potential of AI to address humanity's most pressing scientific challenges and create a brighter future for everyone.

07

Appendix:

Illustrations of AI's Transformational
Impact on Science

Natural Sciences

Connectomics

[Science 2024](#)

- Mapping a 3D, nanoscale-resolution map of a piece of the human brain to a level of detail never previously seen.
- This work may change our understanding of how the brain works which could help researchers better understand neurological diseases such as Alzheimer's and also answer fundamental questions (eg. how memories form).

AlphaFold

[Nature 2024](#), [Science 2023](#), [Nature 2021](#)

- Predicts the structure of all of life's molecules including proteins, DNA, RNA, ligands and how they interact.
- Researchers have used AlphaFold for scientific discovery - ranging from developing new malaria vaccines, to tackling antibiotic resistance to developing new gene therapy methods for treatment diseases like Cancer, tackling pathogens that blight our crops, and developing plastic-eating enzymes.

AlphaMissense

[Science 2023](#)

- Predicts the pathogenicity of missense variants by integrating knowledge gained from both protein structure and evolutionary characteristics.
- Helps researchers determine whether a specific genetic variant is likely to cause a disease, e.g., to help unpick the genetic drivers of epilepsy.

Human Pangenome

[Nature 2023](#)

- The first draft human pangenome that combines assembled genomes from 47 people from diverse ancestries around the world.
- This draft pangenome is a new resource that better represents human genetic diversity, allowing scientists and doctors to more accurately [diagnose and treat diseases](#) and develop new therapeutics.

GNoME

[Nature 2023](#)

- Helps generate novel candidate crystals and predict their stability.
- Discovered 2.2M new crystals - equivalent to nearly 800 years worth of knowledge from classical research techniques.

AlphaProteo

- Generates new protein binders for diverse target proteins.
- Can lead to the discovery of new drugs, the development of biosensors and improve our understanding of biological processes.

Climate Science & Sustainability

Flood Forecasting

[Nature 2024](#)

- AI model that achieves reliability in predicting extreme riverine events at up to a five-day lead time.
- Integrated into [Google's Flood Hub platform](#) Google Search, Google Maps, and available in over 100 countries, covering 700M people.

NeuralGCM

[Nature 2024](#)

- produces ensemble weather forecasts.
- It can simulate over 70,000 days of the atmosphere in the time it would take a physics-based model to simulate only 19 days; it is 1K to 1M times more computationally efficient than SOTA physics models.
- [Openly available](#), which, combined with its ease of use and efficiency, could make climate modeling more accessible to researchers.

GraphCast

[Nature 2023](#), [Science 2023](#)

- AI model that predicts weather conditions up to 10 days in advance.
- Openly available model, which predicts the tracks of cyclones with great accuracy further into the future, identifies atmospheric rivers associated with flood risk, and predicts the onset of extreme temperatures.

Contrails

[arXiv 2023](#), [arXiv 2023](#)

- AI model that identifies areas where airplane contrails are likely to form, allowing for flight rerouting to reduce the climate impact of air travel.
- Reducing the frequency of contrail formation could have a significant impact on emissions from air travel as they account for ~35% of the global warming impacts of the aviation industry.

Wildfire Detection

[ArXiv 2022](#)

- AI model that analyzes satellite imagery to map real-time boundaries of large wildfires.

FireSat

- An AI-powered global satellite constellation designed to detect and track wildfires the size of a classroom (5x5 meters) within 20 minutes.

Energy

Magnetic Plasma Control

[Nature 2022](#)

- The first deep reinforcement learning system that autonomously discovers how to stabilize and shape the plasma within an operational tokamak. Stabilizing plasma is a critical step on the path toward stable fusion.

TORAX

- An open source plasma core simulator, which enables new directions for plasma scenario design and accelerates the research in the fusion space.

Health Sciences

Breast Cancer Prediction

[Nature 2020](#)

- An AI-powered system, which integrates into breast cancer screening workflows to help radiologists identify breast cancer earlier and more consistently.

Lung Cancer Detection

[Nature Medicine 2019](#)

- Lung cancer leads to over 1.8 million deaths per year world wide, accounting for almost one in five cancer deaths, and is the largest cause of cancer mortality.
- This research shows AI can help physicians more accurately screen for lung cancer and identify the disease.

Preventing blindness

[JAMA 2016](#)

- Automated Retinal Disease Assessment (ARDA) uses AI to detect diabetic retinopathy.
- Currently being used to detect diabetic retinopathy in India and the European Union.
- Almost 3k new screenings are supported by ARDA weekly.

Multimodal Medical AI

[MedGemini](#)

- A Gemini-based multimodal medical model, which has demonstrated important advances in clinical reasoning, multimodal, and long-context capabilities across various modalities such as images, surgical videos, genomics, ultra-long health records, ECGs, and more.

MedLM

- A family of foundation models fine-tuned for healthcare.
- Encompasses a range of applications, including answering medical queries, summarizing complex medical information, and extracting insights from unstructured data.

Mathematics

AlphaGeometry

- Solved 83% of all historical International Mathematical Olympiad (IMO) geometry problems.
- Prior version demonstrated AI performance on geometry problems approaching the level of a human Olympiad gold-medalist.

AlphaProof

- A reinforcement-learning-based system that trains itself to prove mathematical statements.
- A significant advancement for formal math reasoning.

Quantum Computing

Willow

[state-of-the-art quantum chip](#)

- Can reduce errors exponentially as we scale up using more qubits.
- Cracks a key challenge in quantum error correction that the field has pursued for almost 30 years.

Continuous Quantum Error Correction

[Nature 2023](#), [arXiv 2022](#)

- Researchers developed an ML algorithm for continuous quantum error correction that uses a recurrent neural network to identify bit-flip errors.
- This breakthrough, and others like it, will accelerate progress towards a large-scale error-corrected quantum computer.

Quantum Gravity

[Nature 2022](#)

- Researchers were able to explore quantum gravity by replicating the dynamics of a traversable holographic wormhole on a 9-qubit quantum computer.
- Represents a step towards being able to study quantum gravity in a laboratory setting.

Quantum Chemistry Simulations

[Nature 2022](#)

- The largest chemistry simulations to date on a quantum computer using Fermionic quantum Monte Carlo (QMC) methods.
- These simulations will offer accurate predictions of chemical reactivity and kinetics.

Education

LearnLM

- A family of models fine-tuned for learning, based on learning science principles.
- Helps simplify and improve the process of lesson planning to help teachers discover unique activities, find engaging materials, and differentiate their lessons and content to meet each of their students where they are.

A Policy Framework for Building the Future of Science with AI