Types of AI and their use in science

CENTRE FOR SCIENCE FUTURES

International Science Council

To cite this report: International Science Council (September 2025). *Types of AI and their use in Science.* DOI:10.24948/2025.09

Funding acknowledgement: This work was conducted with the aid of a grant from the International Development Research Centre (IDRC), Ottawa, Canada. The views expressed herein do not necessarily represent those of IDRC or its Board of Governors.

Author: Mariette Awad

Reviewers: Tapiwa Chiwewe**,** Stefan Harrer

Project coordination: Dureen Samandar Eweis, Vanessa McBride

Project chair: David Castle

Design: Mr Clinton

Cover photo: DKosig

About the International Science Council

The ISC is an international non-governmental organization with a unique global membership that brings together 250 international scientific unions and associations, national and regional scientific organizations, including science academies, research councils, regional scientific organizations, international federations and societies, and academies of young scientists and associations.

The ISC works at the global level to catalyse change by convening scientific expertise, advice and influence on issues of major importance to both science and society.

Contents

Acronyms

AI	artificial intelligence
AutoML	automated machine learning
CNN	convolutional neural network
DT	digital twin
GAN	generative adversarial network
GNN	graph neural network
GNS	graph network-based simulator
LLM	large language model
NLP	natural language processing
PINN	physics-informed neural network
RAG	retrieval-augmented generation
RL	reinforcement learning
RNN	recurrent neural network
SHAP	SHapley Additive exPlanations
SVM	support vector machine
VAE	variational autoencoder

Key takeaways

- Artificial intelligence (AI) is reshaping the scientific enterprise beyond computational assistance: AI is contributing to hypothesis generation and extending traditional methods of theory and experimentation.
- AI is not a monolithic technology, but encompasses a constellation of paradigms, each defined by distinct approaches to learning, inference or knowledge generation with different applications in the scientific context.
- A multitude of different AI techniques address descriptive, predictive, generative and optimization challenges within science.
- AI has found applications in research on medicine, climate science, genomics, social science and more.
- The reliance of AI on data means that computer science, mathematics and statistics intersect with social issues, including ethical questions and integrity of research.

About this paper

This paper provides an overview of artificial intelligence (AI) from the perspective of its use in science, although the techniques described have broad applications in many contexts. The paper examines how AI is contributing to the research process itself and provides some contextualized examples that illustrate applications of the techniques described. The paper aims to inform the rich debate within the science and science policy community by clarifying some of the concepts and techniques encompassed within the term 'AI'.

Techniques are presented in rough chronological order of development. Given the rapid rate of developments in AI, this paper provides a snapshot, rather than a comprehensive technological assessment.

The paper is part of a series of three primers that explore various technical dimensions of AI and its impact on science. The other primers are "Considerations on the environmental impact of AI in science" and "Data and AI for science".

Introduction

Artificial intelligence (AI) is reshaping the scientific enterprise, not only as a computational assistant, but as a generative force in knowledge production. It now contributes across disciplines to hypothesis generation, pattern recognition and predictive modelling, extending traditional methods of theory and experimentation (LeCun et al., 2015; Jordon and Mitchell, 2015).

This transformation signals not just technological progress, but a deeper shift in how science is done. As the World Economic Forum (2024) notes, "AI for scientific discovery" is a cross-cutting enabler of innovation which prompts institutions to rethink research methodologies and invest in computational infrastructure (National Academies of Sciences, Engineering, and Medicine et al., 2023). Alongside this, AI is contributing to scientific reasoning: classical approaches grounded in deduction and induction by humans are increasingly complemented and sometimes replaced by statistical and abductive inference (a logical process of forming the simplest and most likely explanation for a set of observations). In this way insights emerge from data patterns rather than theoretical premises.

Philosophers of science have framed this as a reorientation. Computational systems do not just enhance human enquiry; instead, they restructure it by generating knowledge through processes that may elude intuitive interpretation (Humpherys, 2004; Leonelli, 2019). In this light, AI systems act not only as tools but as epistemic agents, potentially influencing what counts as evidence, how explanations are formed and who participates in the construction of knowledge.

To understand this shift, this paper examines the methodological diversity within AI itself. Far from a monolithic technology, AI encompasses a constellation of paradigms, each defined by distinct approaches to learning, inference and knowledge representation. Foundational paradigms include **descriptive AI**, which extracts and structures patterns from data; **predictive AI**, which anticipates outcomes based on historical trends; **generative AI**, which produces novel content and scientific hypotheses. Decision-oriented paradigms build on these foundations: **optimization AI** identifies efficient solutions within complex constraints, while **prescriptive AI** goes further by recommending concrete courses of action, balancing trade-offs with policy and societal goals. Finally, cross-cutting principles such **privacy-aware AI**, enables responsible analysis of sensitive information; **causal AI**, seeks to uncover underlying cause-effect relationships; and **explainable AI** enhances the transparency of computational models and as such ensures responsible, transparent and trustworthy deployment.

Together, these paradigms support a broad range of scientific objectives – from classification and simulation to exploration and explanation. By performing tasks such as outcome prediction, scientific modelling, pattern discovery, content generation, model interpretability, experiment optimization and privacy preservation, AI tools collectively enhance the analytical and creative dimensions of scientific work.

In Section 1, this paper highlights these AI paradigms, and in Section 2, presents applications across diverse scientific domains to demonstrate how AI is being integrated into research workflows and accelerating knowledge discovery and validation.

Section 1: Types of AI in the context of science

Expanding on the role of AI in scientific enquiry, this section explores some types of AI within the scientific domain, spanning the discovery and generation of new scientific knowledge to the optimization and automation of experimental processes.

1.1 Predictive AI: Scientific modelling and forecasting

Predictive AI refers to a category of AI methods that use data to anticipate outcomes, identify trends or simulate behaviours across digital and physical systems. At the heart of predictive AI is machine learning – a family of approaches that enable systems to learn from examples. The process of machine learning typically unfolds in two stages. First, the system goes through a training phase, where it learns patterns from existing data. When the data is labelled, learning is called supervised learning. When the data is unlabelled, learning is referred to as unsupervised learning. Second, in the prediction phase, the trained system is applied to new, unseen information to make inferences about future outcomes or unknown variables. The core objective of predictive AI is to generalize what it has learned and to provide accurate and reliable predictions even in new or changing environments. This capacity is essential in many public-facing domains, such as forecasting public health trends, responding to climate shifts or optimizing infrastructure planning.

Traditional machine learning predictive methods (see Appendix 1 for a selection of tools) have been successfully used in scientific and industrial settings for tasks such as classifying observations, estimating values or detecting irregularities in datasets (Huang et al., 2018; Nassif et al., 2021). While these tools can be effective for simpler or smaller-scale problems, they often require manual input by experts to select and shape the data that goes into the models. This process, known as **feature engineering**, can constrain flexibility by requiring manual selection of input variables. While some models, like decision trees, offer intuitive interpretability, others function more like "black boxes", producing results without transparently revealing how those results were derived. This opacity makes it more difficult to explain or justify their outputs. Moreover, as the size and complexity of datasets grow, these traditional methods struggle to maintain high performance. In high-dimensional or large-scale environments, their performance can plateau due to limitations in computational efficiency and the ability to capture complex relationships within the data (Zhou et al., 2017a). When datasets are too small, models fail to capture meaningful patterns and instead learn from noise, resulting in overfit models and unreliable outputs.

In response to these challenges, more advanced techniques known as **deep learning methods** have emerged. These AI techniques use layered networks, which are modelled loosely on how the human brain processes information. They automatically learn patterns from raw data and thus reduce the need for manual data preparation and feature engineering by experts. Deep learning has proven especially powerful in working with large, complex

datasets such as satellite imagery, genetic sequences or environmental simulations, where conventional tools often fall short. A foundational type of deep learning architecture, known as **convolutional neural networks (CNNs)** (Krizhevsky et al., 2012), is designed to recognize patterns in image data by scanning the images in layers to detect edges, shapes and textures. Beyond basic image analysis, CNNs play a central role in advanced scientific applications. These networks have been widely adopted for tasks such as identifying objects in images, sorting image categories and dividing images into meaningful parts (a process known as segmentation), which has made them central for analysing medical images such as MRIs and CTs (Gulsha et al., 2026).

Another contribution of deep learning to scientific modelling is **graph network-based simulators (GNS)**, a class of neural simulators designed to learn and predict the behaviour of complex physical systems. GNS are built upon **graph neural networks (GNNs)**, a type of deep learning architecture that operates on data structured as graphs. In this framework, physical systems are represented as graphs where nodes correspond to particles or entities (e.g. elements of a fluid or parts of a rigid body), and edges represent interactions or relationships between them (Figure 1). GNS leverage a computational mechanism known as message passing, in which nodes iteratively exchange information with their neighbours to update their own state (Sanchez-Gonzalez et al., 2020). Through this process, GNS can model and compute the dynamics of interacting entities, including rigid bodies, fluids and deformable materials, which are often computationally expensive or difficult to scale.

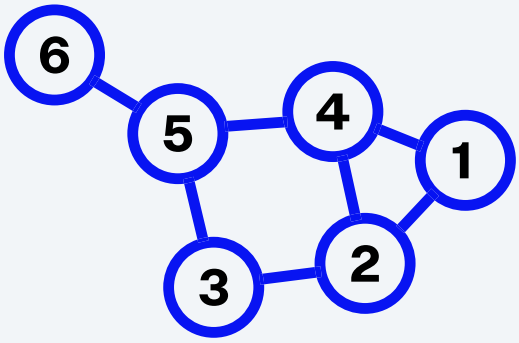

Figure 1. A simplistic depiction of a graph showing six nodes and seven edges (by Michel Bakni – Own work, CC BY-SA 4.0)

For temporal or sequential data such as climate records, sensor logs or gene sequences, **recurrent neural networks (RNNs)** are more appropriate. They model time-dependent phenomena by maintaining a hidden state that evolves across input sequences. In many scientific domains, RNNs are combined with attention mechanisms to focus on relevant data points in the context of long sequences. Also, some emerging biologically inspired models, such as liquid neural networks and spiking neural networks are being sparsely researched for temporal and real-time scientific data (Roy et al., 2019; Hasani et al., 2021).

Autoencoders are a type of deep learning used to uncover patterns in data without the need for labelled examples. They work by compressing the data into a simplified form and then reconstructing it, making them useful for tasks like reducing noise, detecting anomalies or simplifying complex datasets (Hinton et al., 2006).

Physics-Informed Neural Networks (PINNs) are a special type of deep learning that combines data-driven learning with established physical laws, such as those governing heat, motion or fluid flow (Raissi et al., 2019). By embedding these laws directly into the model, PINNs ensure that the results stay consistent with real-world scientific principles. This makes them especially valuable for solving complex problems in fields like climate modelling, engineering and energy systems – whether to simulate how a system behaves or to estimate unknown conditions based on limited observations.

While deep learning models are powerful, they typically require significant expertise to design, adjust and optimize. To make these tools more accessible, **automated machine learning (AutoML)** has emerged as a solution that automates many complex steps – such as preparing data, choosing the right model and fine-tuning settings (Feurer et al., 2015). AutoML uses advanced search strategies to find effective model configurations with minimal human input, helping researchers and practitioners apply AI more efficiently across fields without needing deep technical knowledge. AutoML is a type of meta-technology (Section 1.7).

1.2 Descriptive AI: Pattern discovery and textual interpretation in scientific research

Descriptive AI is a type of AI that helps researchers explore and understand large, complex, multimodal datasets. Instead of making predictions, it looks for patterns, connections and underlying themes – making it especially valuable in the early stages of scientific research, when the goal is to uncover what insights the data might hold.

A central area of descriptive AI is **natural language processing (NLP)**, which allows computers to read and analyse large volumes of text, such as academic papers, technical documents, policy reports and social media posts. NLP has been used in many contexts. One of them is to reveal thematic structures and track the evolution of discourse in areas like research and journalism (Jacobi et al., 2016). NLP can also identify key actors, institutions and scientific phenomena for structured mapping of knowledge from text. NLP methods can further be used to assess public opinion on issues such as climate change or health crises. However, studies (Caliskan et al., 2017) have shown that NLP tools may perpetuate and reflect unintentional societal biases, including race and gender stereotypes. This highlights the need for ethical awareness and examination when used in scientific and human contexts.

Descriptive AI also includes a rich set of unsupervised learning and visualization techniques, particularly useful when the data does not come with labels or predefined categories. Such techniques are used to, for example, identify subgroups in ecological surveys, sort patients with similar characteristics in clinical studies or group similar cells in single-cell RNA sequencing (Kiselev et al., 2019). They also simplify and visualize high-dimensional data to reveal patterns in domains from neuroimaging to materials science (López et al., 2009; Ehiro, 2023). In studies that involve data changing over time, like astronomy, geophysics, clinical medicine or climatology, AI models can be used to uncover repeating cycles, seasonal patterns or trends. Meanwhile, scientific knowledge graphs bring together both structured data (like databases) and unstructured information (like text) to map relationships between scientific

concepts, research papers, institutions and findings. AI can then offer visualizations of how research areas evolve over time and across disciplines (Ammar et al., 2018).

1.3 Generative AI: Synthesis and analysis of scientific data

Building on descriptive approaches, **generative AI** goes a step further by not only analysing data but also creating new, realistic content based on information learned by the algorithm. This could include generating text, images, computer code or even data. In scientific research, generative AI can simulate experiments, model complex systems where data is limited or uncertain, and support discovery by helping researchers reason through information, explore innovative ideas more quickly, and generate new hypotheses and research directions.

One significant application of generative AI in science is the generation of synthetic datasets, particularly when real data are rare, expensive or constrained by privacy regulations. **Generative adversarial networks (GANs)** address this challenge by training two neural networks in a competitive setting: a network (generator) that tries to create fake data that looks real, and another (discriminator) that tries to tell the difference between real and fake data. As they compete, both get better and the result is highly realistic data. In genomics, for example, GANs have been used to create artificial DNA sequences that resemble real DNA while protecting individual privacy. They thus offer a valuable tool for population genetics and precision medicine (Yelen et al., 2021).

Another type of generative AI is **variational autoencoders (VAEs)**. VAEs, being more advanced than traditional autoencoders, learn to compress and reconstruct data and model the underlying patterns in the data by learning a "probabilistic map" of it – called a latent space. This capacity enables VAEs to generate new synthetic data instances by sampling from the learned space. In the sciences, VAEs are used for tasks such as generating medical images to expand small datasets, filling in missing data from scans or tests, and helping researchers explore hidden factors in complex biological datasets, becoming valuable in biomedical imaging and genomics (Kingma and Welling, 2013).

Diffusion models go further; they gradually refine random noise into highly detailed, reliable outputs. For science, this means, for example, more accurate climate simulations, better synthetic medical data to protect patient privacy and richer satellite imagery for environmental monitoring – offering policymakers a stronger evidence base for decision-making.

While diffusion models outperform both VAEs and GANs in scientific text-to-image and image-to-image generation (by preserving image structure and quality and semantic alignment), they are also more computationally and time-intensive. All three of these generative techniques suffer from a lack of tuning to scientific language and science-specific imaging details when employed in a scientific context (Sordo et al., 2025).

The major breakthrough in generative AI is **large language models (LLMs)** that are built on transformer architecture (Vaswani et al., 2017). Transformers are deep learning models that understand and generate sequences – like sentences or DNA strings – by learning

how each part of the sequence relates to every other part (attention mechanism). Unlike earlier AI models that processed data step by step, transformer models analyse entire input sequences at once. This parallel processing makes them significantly faster and more effective, especially for large-scale data tasks. In scientific tasks, transformers are used from protein structure prediction (e.g. AlphaFold 3), to storage and reasoning of scientific information, such as Galactica (Taylor et al., 2022) or ChemCrow (Bran et al., 2024) for various chemical tasks.

AI and trusted scientific results

While the **Galactica** model was trained on millions of scientific papers, it was withdrawn days after launch when it was found to produce fabricated citations and misleading claims (Heaven, 2022). The controversy underscored a core limitation of domain-specific large language models: while they can accelerate research, without strong validation, they risk spreading authoritative-sounding misinformation that can distort evidence and weaken trust in science.

Through the aid of transformers, LLMs are demonstrating success in language tasks (e.g. question answering, summarization, translation). In addition, LLMs are assisting scientists with code-related tasks. For example, Code Llama (Rozière et al., 2024) can generate and complete code, supporting tasks such as building simulation models, preprocessing data or running analysis. These tools can reduce time spent on routine coding, facilitating the development of computational experiments, especially for interdisciplinary researchers or those with limited programming proficiency.

Types of large language models

Major LLMs fall into three broad categories: proprietary, open weight and open source.

Proprietary LLMs, such as ChatGPT (OpenAI), Gemini (Google DeepMind) and Claude (Anthropic), are developed and maintained by private companies. These models are accessible only via APIs or user interfaces, with no public release of their underlying model weights, source code or training data.

In contrast, open-weight models release their pretrained model weights but under restrictive licenses – typically limiting commercial use, redistribution or fine-tuning. Notable examples include Llama (Meta), Gemma (Google) and Yi (01.AI). These models enable research and experimentation but are not fully open source in the traditional sense.

Open-source LLMs, on the other hand, provide both the model weights and source code under permissive licenses that support reuse, redistribution and commercial deployment. Key examples include DeepSeek, Falcon (Technology Innovation Institute), Mistral (Apache 2.0 release), Command R+ (Cohere) and Pythia (EleutherAI). BLOOM (BigScience) is also widely recognized, though it uses a Responsible AI License (RAIL) license that permits many forms of use while imposing ethical restrictions.

More advanced multimodal LLMs are helping scientists work with different types of data – such as text, images and diagrams – all at once. These models are designed to understand and connect information across multiple formats, making it easier to analyse complex scientific materials. In chemistry, ChemVLM (Li et al., 2025) combines visual and textual inputs to perform tasks like chemical optical character recognition, multimodal molecule understanding and diagram-based chemical reasoning. By combining what they "see" and "read", these models are improving how machines interpret scientific information that does not exist in plain text, opening new possibilities for automated scientific analysis.

A technique that extends the effectiveness of LLMs is **retrieval-augmented generation (RAG)** (Lewis et al., 2020). RAG combines two components: one that retrieves relevant information from large external sources (like scientific databases), and another that generates human-like text using this information. For example, when a researcher asks a question, the system first searches for related documents and then uses a pretrained large language model to generate an answer based on both the question and the retrieved data. This method helps make AI-generated responses more accurate and grounded in real facts, reducing the risk of producing incorrect or made-up information – also known as "hallucinations".

In the scientific domain, RAG is increasingly used to enhance scientific literature search and scholarly assistants that require citation fidelity. BIORAG was introduced as a RAG and LLM-based framework for biological question answering and reasoning (Wang et al., 2024). In protein science, multi-agent RAG LLM is employed for hypothesis generation (Bazgir et al., 2025) and illustrates the growing utility of these models in facilitating scientific discovery through data-informed reasoning.

1.4 Optimization AI: Optimization and automation of scientific experimentation

Beyond generating content or assisting with scientific reasoning, AI is also being used to make decisions and take actions in dynamic environments – thanks to a branch of AI known as **reinforcement learning (RL)**. RL enables an AI model to learn through interaction with its environment, receiving feedback in the form of rewards to gradually discover the best strategy for achieving a goal. In scientific settings, RL has been used to automate experimental design and control protocols. For example, deep RL was applied to chemical synthesis to learn which experimental conditions lead to the best outcomes and then suggest improved experiments based on what it has learned (Zhou et al., 2017b). RL has also powered autonomous laboratory systems such as AlphaFlow – a robotic fluidic system that uses RL to guide the creation of nanoparticles, reducing the need for human intervention and enhancing the speed and scale of chemical experimentation (Volk et al., 2023).

Another essential method in experiment optimization is **Bayesian optimization**, especially when experiments are expensive, time-consuming or uncertain. Instead of testing every possible option, this method builds a smart model (surrogate model) that predicts which experiment is likely to be most informative. It then chooses the next experiment based on

this prediction – balancing between trying new possibilities and focusing on what already looks promising. Scientists use this approach in cases like climate modelling (Lian et al., 2022) or materials discovery (Diwale et al., 2022). It provides a principled way to balance the exploration of new possibilities with the exploitation of known promising configurations.

Active learning is an AI approach that makes experimentation and data use more efficient. Instead of requiring massive, labelled datasets, the system learns by identifying which data points are most important to label and analyse next. This reduces costs and minimizes the human effort needed for large-scale data preparation, while still improving accuracy and performance. This is useful in domains like biology, where experiments like gene editing can be labour-intensive. Instead of testing everything, active learning helps focus only on the most uncertain or potentially impactful cases or hypotheses. It is used in closed-loop systems, where the AI keeps updating its strategy as new results are collected – making research more targeted, efficient and adaptive (Yang et al., 2025).

AI tools like evolutionary algorithms and swarm intelligence are also integrated to optimize scientific experiments. These techniques are commonly used in optimization problems, where the goal is to find the best possible settings or configurations for a system. For example, they have been used to optimize parameters in synthetic gene circuits in biology or adjust robotic controls for performing delicate tasks in experimental setups. Although they often require a lot of computing power, these methods are highly valued because they can search more thoroughly and avoid getting stuck in suboptimal solutions – which is essential when exploring unknown scientific territories (Brambilla et al., 203; Rodrigo et al., 2012).

Another important application of AI in scientific research is **surrogate modelling**. Surrogates are simplified, AI-enabled versions of complex experiments or simulations – built with tools such as neural networks, decision trees or Gaussian processes – that can replicate system behaviour at a fraction of the cost and time. They allow researchers to test ideas, explore scenarios and generate predictions without running prohibitively expensive or time-consuming simulations (Radaideh et al., 2020). In fields like climate science and engineering, surrogate models have become indispensable for accelerating discovery and enabling experimentation at scale. Closely related, but distinct, is **digital twin (DT)** technology; a dynamic, data-linked virtual counterpart of a real system that evolves alongside it, providing continuous monitoring and decision support. Surrogates often operate inside DTs, but their roles differ: surrogates accelerate scientific research, while DTs transform real-world operations. Beyond manufacturing and energy, DTs are now opening new frontiers in medicine and the life sciences – for instance, simulating patient populations to design and test clinical trials more efficiently and ethically (Attaran and Celik, 2023).

1.5 Causal and interpretable AI: Explainability in scientific research

As AI becomes more involved in scientific research, it is important not just to get accurate outputs, but also to understand how and why those outputs are made. Two approaches that help scientists achieve this are **causal inference** and **explainable AI**. These methodologies make AI outputs more interpretable and reproducible and support hypothesis testing, intervention planning and knowledge extraction.

Causal inference helps scientists move beyond simply identifying patterns or correlations. Instead, it allows them to explore cause-and-effect relationships, which is essential in domains such as epidemiology, economics and climate science. Tools such as structural causal models (Pearl, 2009), potential outcomes frameworks, causal Bayesian networks and causal discovery algorithms make it possible to study interventions and even hypothetical scenarios (called counterfactuals) from observational or experimental data.

In parallel, **explainable AI** provides insight into complex "black box" models. This is important in critical domains like biomedicine and climate forecasting. Techniques like SHapley Additive exPlanations (SHAP) and local interpretable model-agnostic explanations (LIME) help interpret model predictions by showing which input features (such as symptoms, temperature or data points) influenced the result the most. However, these explanations can vary depending on the model and data used, so scientists must validate them carefully to avoid misleading interpretations (Salih et al., 2025).

Explainability for LLMs require different tools, that fall broadly into two categories. Mechanistic methods look inside the model to trace how it processes information (sometimes described as "opening the box"), using techniques like attention visualization or attribution to identify which data influenced a decision. Behavioural methods treat the model as a black box, testing only its inputs and outputs to reveal patterns, biases or blind spots. Combining both perspectives is essential: mechanistic tools help researchers diagnose why errors occur, while behavioural tools provide practical evidence of reliability, fairness and trustworthiness in real-world use.

Interpretable and explainable AI – is there a difference?

In the context of AI for science, it is important to distinguish between interpretable AI and explainable AI:

- **Interpretable AI** refers to **models that are transparent by design**, such as decision trees or linear regressions, where the logic behind each decision is directly understandable.
- **Explainable AI**, by contrast, is the broader concept that also **includes techniques developed to make complex "black box" systems** like deep learning models or LLMs **more understandable after the fact**.

For policymakers, both approaches matter: interpretable AI offers inherent simplicity and accountability, while explainable AI ensures that even advanced systems used in critical domains like healthcare, climate or security can be made transparent and trustworthy.

Beyond these methods, interpretable-by-design models such as **generalized additive models**, prototype-based networks and decision rule sets offer transparent alternatives. These models are built in a way that makes their decision-making process easy to follow, making them especially useful in scientific research where clarity, reproducibility and peer

review are essential (Rudi, 2019). Scientists can trace how each input leads to a result, which supports accountability and easier error checking.

Neuro-symbolic AI represents a growing class of methods which combine two strengths: the ability of neural networks to learn from large, complex datasets and the clarity of symbolic reasoning, such as logic and rules. This combination is especially useful in fields like the social sciences, where researchers need both data-driven insights and clear explanations of why something is happening. In symbolic regression, AI tools automatically generate easy-to-understand mathematical formulas from real-world data. These tools help scientists discover new relationships and build interpretable models which turn raw data into new scientific hypotheses or causal explanations. This integration of learning and reasoning advances explainability, enhances scientific accountability and positions neuro-symbolic AI as a promising frontier for formal knowledge production in interdisciplinary research.

1.6 Privacy-aware AI

As AI becomes more involved in scientific research, it is important not just to get accurate predictions but also to understand how and why those predictions are made, and to ensure that sensitive data can be protected while providing useful input.

One of the most impactful approaches is **federated learning,** which enables scientists to train AI models collaboratively without sharing sensitive data. Each institution keeps its data locally and shares only model updates, preserving privacy and sovereignty. This approach is vital in areas like health and climate science, where data cannot be centralized, allowing secure, compliant and cooperative scientific research (Sheller et al., 2019; Li et al., 2022).

Complementing federated learning is **differential privacy**, a formal mathematical framework that ensures individual-level data cannot be reverse-engineered or identified from AI model outputs. Differential privacy makes it possible to publish research findings or train AI models on sensitive datasets without risking privacy breaches. This has been critical in public health studies and government census work (Dwork and Roth, 2014).

Additionally, AI can be combined with privacy-enhancing technologies like **homomorphic encryption** and **secure multiparty computation**, which enable computations to be performed directly on encrypted data or jointly across multiple institutions without revealing raw data. Although these methods are not types of AI themselves, they are integrated in AI-driven scientific systems, particularly in genomics and collaborative epidemiological modelling (Brisimi et al., 2018).

Another emerging technique used for data synthesis is **differentially private generative adversarial networks**, which create synthetic datasets that resemble real-world information but do not expose any actual individual's records (Cheng et al., 2021). These synthetic datasets are useful when real data cannot be openly shared, enabling broader collaboration while respecting data protection laws.

1.7 Meta-scientific AI: Automation of scientific discovery

Meta-scientific AI refers to a class of AI systems designed to support or enhance the scientific process itself – by generating hypotheses, designing experiments, analysing data or connecting insights across different scientific fields. When these systems operate with a degree of autonomy – able to plan, reason and take actions towards scientific goals – they are described as agentic. These AI "agents" can make decisions, adapt their behaviour based on outcomes, and even prioritize tasks similar to a junior researcher. The convergence of meta-scientific and agentic capabilities marks a profound shift in how science is conducted. Rather than being just advanced tools, AI systems are beginning to act as active collaborators in the discovery process.

Recent advances have enabled partial automation of the scientific discovery pipeline, especially through the use of LLMs integrated with symbolic reasoning, knowledge graphs and simulation engines such as LilaSciences, Google Co-Scientist, Future House's ether0, CSIRO's Sciansa (Lohr, 2025; Gottweis and Natarajan, 2025; Savitsky, 2025; Harrer et al, 2024). These systems are capable of identifying unexpected relationships, refining theoretical models and accelerating materials discovery by uncovering latent design principles. Beyond LLMs, meta-scientific AI also includes symbolic regression tools. These developments point towards a growing ecosystem where different forms of AI – neural, symbolic and hybrid – contribute to diverse stages of the scientific method, from data structuring to theory generation. Examples include IBM's Generative Toolkit for Scientific Discovery (GT4SD) and SciAgents (Ghafarollahi and Butler, 2024), which combines LLMs with large ontological knowledge and other agents for bio-inspired materials development.

Section 2: Selected AI applications in scientific research and discovery

To illustrate the impact of AI across scientific disciplines, this section presents selected applications in biomedicine, climate science and social science to demonstrate how AI tools are being integrated into domain-specific workflows to support scientific discovery, advance knowledge production, enhance forecasting and inform evidence-based decision-making.

2.1 Predictive and generative AI in biomedicine

One of the most impactful applications of AI in science is in the field of drug discovery, where the complexity, cost and pace of research have long posed major barriers to innovation. In drug discovery, high failure rates are often attributed to an incomplete understanding of disease mechanisms at the molecular level (Qui et al., 2024). A critical early step in the discovery process involves identifying a biological target – typically a protein – that plays a central role in a disease pathway. Designing a therapeutic molecule that interacts with this target requires knowledge of its three-dimensional structure, which has traditionally depended on resource-intensive techniques such as X-ray crystallography or nuclear magnetic resonance imaging.

Advances in AI are accelerating this process by predicting protein structures directly from amino acid sequences. DeepMind's AlphaFold2 (Jumper et al., 2021), a deep learning-based framework, has predicted over 200 million protein structures, covering nearly all catalogued proteins. More recently, AlphaFold3 (Abramson et al., 2024) has extended this capability to model interactions not only between proteins, but also with DNA, RNA, ligands and small molecules. AI-enabled drug discovery has already progressed to clinical trials, with companies like Insilico Medicine (Zitnik, 2025) advancing multiple AI-conceived candidates – demonstrating how rapidly AI is moving from structural prediction to real-world medical innovation.

2.2 Predictive AI for noncoding variants in genomics

Beyond proteins, AI is also shedding light on the vast noncoding regions of the human genome, which play a critical but often overlooked role in regulating health and disease. While noncoding regions of the genome do not encode proteins, they play a crucial role in regulating gene expression – and their disruption can contribute to complex diseases. Identifying the functional impact of these noncoding variants has long posed a major challenge in human genetics due to their scale and regulatory complexity. AI has significantly advanced this field by enabling models that can predict the regulatory effects of noncoding DNA from sequence data. A key example is DeepSEA (Zhou and Troyanskaya, 2015), a deep learning model that uses CNNs to analyse raw genomic sequences. Trained on large-scale chromatin-profiling datasets, DeepSEA can predict how single-nucleotide

changes affect transcription factor binding, chromatin accessibility and other regulatory features – to prioritize noncoding variants that may influence disease risk. Tools like DeepSEA demonstrate how AI, particularly deep learning, can support the interpretation of the vast, uncharted regulatory genome. By linking genetic variation to gene regulation, AI is contributing to advances in functional genomics and laying the groundwork for more personalized approaches in medicine.

2.3 Predictive AI for climate modelling and preparedness

As global temperatures rise, the scientific community faces growing pressure to deliver accurate, actionable climate predictions. Traditional physics-based general circulation models are essential tools for simulating atmospheric dynamics. However, their reliance on simplified representations of small-scale processes – such as cloud formation or turbulence – limits their accuracy, particularly in regional forecasts.

Advances in machine learning are now reshaping climate modelling. Google's NeuralGCM (Kochkov et al., 2024) highlights a new class of predictive AI that combines physical simulators with learned components. It integrates a differentiable numerical solver – grounded in atmospheric physics – with neural networks trained on large-scale weather and climate datasets. This hybrid architecture enables more accurate and efficient short-to-medium-term weather forecasting, as well as longer-term climate scenario modelling. By incorporating complex climate subsystems such as the carbon cycle and ocean dynamics, NeuralGCM provides scientists with a powerful tool to test hypotheses and refine regional projections of climate impacts. As such, predictive AI is not only accelerating discovery in climate science but also contributing to global preparedness and decision-making for climate adaptation.

2.4 Generative AI for social science and human behaviour simulation

Historically, sociologists have used tools like surveys and social network analysis to understand how people relate and how society works. AI is transforming those methods by giving researchers the ability to study social behaviour on a deeper and larger scale. This shift began when researchers realized that digital data – like social media, online searches and digital records – could help us observe society in real time (Lazer et al., 2009), and with AI advancements social research methods are further extended.

LLMs are simulating human behaviour, allowing researchers to simulate human samples that would otherwise require real people. These simulations help explore social dynamics and test hypotheses at scale (Grossman et al., 2023).

A recent study used AI systems comprising multiple LLM agents to simulate how people cooperate in social situations that require working together for the common good (prosocial behaviour). The researchers considered experiments of the "public goods game", a common experiment that studies how individuals decide to contribute resources to a shared pool that

benefits the whole group. The AI agents were able to mimic real human behaviour under conditions such as transparency, priming and varying endowments. The simulations even demonstrated behaviours seen outside the lab, like working together and cheating (Sreedhar et al., 2025).

This kind of AI simulation can assist policymakers by allowing them to test how different policy decisions might encourage people to act more cooperatively in society – without needing to run costly or challenging experiments with real people. Such approaches may provide a flexible way to develop policies that promote positive social behaviour.

Closing thoughts

Together, AI approaches are not just accelerating science – they are transforming how it is conducted, interpreted and shared. From descriptive models that help organize complex datasets, to predictive systems that enhance planning and preparedness, to generative tools that open new paths to discovery, AI is influencing every stage of the scientific method. Causal and explainable AI strengthen trust and transparency, while collaborative AI fosters new kinds of partnerships – across institutions, disciplines and even between humans and machines.

But AI's influence goes deeper than speed or scale. It is changing how scientific knowledge is created and validated. Advanced systems are now supporting hypothesis generation, causal reasoning and even aspects of experimental design – sometimes with minimal human input. These developments raise important questions about scientific norms, including how findings are interpreted, how researchers interact with AI collaborators, the future of scientific work and how responsibility is shared. Understanding and guiding these shifts will be essential to ensuring AI enhances, rather than disrupts, the integrity and purpose of scientific research (Harrer, 2025).

As Thomas Kuhn (1962) argued, paradigm shifts redefine the underlying assumptions and methods of enquiry. The integration of AI into science may mark such a shift – not just in instrumentation, but in the epistemic foundations of research. As AI systems become more autonomous, the need for transparency, reproducibility and accountability becomes urgent. Who verifies machine-derived knowledge? How do we detect subtle biases embedded in AI-generated hypotheses? And how do we safeguard the tacit, intuitive and ethical dimensions of science that AI cannot replicate?

To ensure that this transformation strengthens rather than undermines human understanding, we must pair technical innovation with ethical foresight. This includes designing systems that allow for human interpretability, supporting inclusive datasets that reflect diverse worldviews, and building governance frameworks for auditing, validating and contesting AI-generated knowledge. AI should not only accelerate discovery – it should deepen science trustworthiness, widen its accessibility and extend its capacity for epistemic inclusiveness. Only then can it serve as a true generative partner in the future of knowledge.

References

Abramson, J. et al. 2024. Accurate structure prediction of biomolecular interactions with AlphaFold 3 *Nature*, 630(8016), pp. 493–500. https://doi.org/10.1038/s41586-024-07487-w.

Ammar, W. et al. 2018. Construction of the Literature Graph in Semantic Scholar. *arXiv.* https://doi.org/10.48550/ARXIV.1805.02262.

Attaran, M. and Celik, B.G. 2023. Digital Twin: Benefits, use cases, challenges, and opportunities, *Decision Analytics Journal*, 6, p. 100165. https://doi.org/10.1016/j.dajour.2023.100165.

Bazgir, A. et al. 2025. ProteinHypothesis: A Physics-Aware Chain of Multi-Agent RAG LLM for Hypothesis Generation in Protein Science, in *Towards Agentic AI for Science: Hypothesis Generation, Comprehension, Quantification, and Validation*. *ICLR 2025 Workshop AgenticAI*. https://openreview.net/forum?id=5XQIbNIhAW.

Board on Mathematical Sciences and Analytics et al. 2023. *Opportunities and Challenges for Digital Twins in Biomedical Research: Proceedings of a Workshop-in Brief*. Edited by L. Casola. Washington, D.C.: National Academies Press, p. 26922. https://doi.org/10.17226/26922.

Brambilla, M. et al. 2013. Swarm robotics: a review from the swarm engineering perspective, *Swarm Intelligence*, 7(1), pp. 1–41. https://doi.org/10.1007/s11721-012-0075-2.

Brisimi, T.S. et al. 2018. Federated learning of predictive models from federated Electronic Health Records, *International Journal of Medical Informatics*, 112, pp. 59–67. https://doi.org/10.1016/j.ijmedinf.2018.01.007.

Caliskan, A. et al. 2017. Semantics derived automatically from language corpora contain human-like biases, *Science*, 356(6334), pp. 183–186. https://doi.org/10.1126/science.aal4230.

Cheng, V. et al. 2021. Can You Fake It Until You Make It?: Impacts of Differentially Private Synthetic Data on Downstream Classification Fairness, in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. *FAccT 21: 2021 ACM Conference on Fairness, Accountability, and Transparency*, Virtual Event Canada: ACM, pp. 149–160. https://doi.org/10.1145/3442188.3445879.

Cranmer, M. 2023. Interpretable Machine Learning for Science with PySR and Symbolic Regression.jl. *arXiv.* https://doi.org/10.48550/arXiv.2305.01582.

Dey, R. and Salem, F.M. 2017. Gate-Variants of Gated Recurrent Unit GRU. Neural Networks. *arXiv.* https://doi.org/10.48550/ARXIV.1701.05923.

Diwale, S. et al. 2022. Bayesian optimization for material discovery processes with noise, *Molecular Systems Design & Engineering*, 7(6), pp. 622–636. https://doi.org/10.1039/D1ME00154J.

Dwork, C. and Roth, A. 2013. The Algorithmic Foundations of Differential Privacy, *Foundations and Trends® in Theoretical Computer Science*, 9(3–4), pp. 211–407. https://doi.org/10.1561/0400000042.

Ehiro, T. 2023. Feature importance-based interpretation of UMAP-visualized polymer space, *Molecular Informatics*, 42(8–9), p. 2300061. https://doi.org/10.1002/minf.202300061.

Feurer, M. et al. 2015. Efficient and Robust Automated Machine Learning, in C. Cortes et al. eds. *Advances in Neural Information Processing Systems*. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2015/file/11d0e6287202fced83f79975ec59a3a6-Paper.pdf.

Ghafarollahi, A. and Buehler, M.J. 2024. SciAgents: Automating scientific discovery through multi-agent intelligent graph reasoning. *arXiv.* https://doi.org/10.48550/arXiv.2409.05556.

Gibney, E. 2025. DeepMind unveils "spectacular" general-purpose science AI, *Nature*, 641(8064), pp. 827–828. https://doi.org/10.1038/d41586-025-01523-z.

Gottweis, J. and Natarajan, V. 2025. Google Research, *Accelerating scientific breakthroughs with an AI co-scientist*, 19 February. https://research.google/blog/accelerating-scientific-breakthroughs-with-an-ai-co-scientist/.

Grossmann, I. et al. 2023. AI and the transformation of social science research, *Science*, 380(6650), pp. 1108–1109. https://doi.org/10.1126/science.adi1778.

Gulshan, V. et al. 2016. Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs, *JAMA*, 316(22), p. 2402. https://doi.org/10.1001/jama.2016.17216.

Harrer, S. 2025. *AI agents: The scientist's new superpower* (TEDx Talks). https://www.youtube.com/watch?v=z7shPnkFjt0.

Harrer, S., et al. 2024. Generative AI agents are transforming biology research: high resolution functional genome annotation for multiscale understanding of life, *eBioMedicine*, 109, p. 105446. https://doi.org/10.1016/j.ebiom.2024.105446.

Hasani, R. et al. 2021. Liquid Time-constant Networks, *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(9), pp. 7657–7666. https://doi.org/10.1609/aaai.v35i9.16936.

Heaven, W.D. 2022. *Why Meta's latest large language model survived only three days online*, *MIT Technology Review*. https://www.technologyreview.com/2022/11/18/1063487/meta-large-language-model-ai-only-survived-three-days-gpt-3-science/.

Hinton, G.E. and Salakhutdinov, R.R. 2006. Reducing the Dimensionality of Data with Neural Networks, *Science*, 313(5786), pp. 504–507. https://doi.org/10.1126/science.1127647.

Hochreiter, S. and Schmidhuber, J. 1997. Long Short-Term Memory, *Neural Computation*, 9(8), pp. 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735.

Huang S. et al. 2018. Applications of Support Vector Machine SVM. Learning in Cancer Genomics, *Cancer Genomics & Proteomics*, 15(1). https://doi.org/10.21873/cgp.20063.

Humphreys, P. 2004. *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. 1st edn. Oxford University Press New York. https://doi.org/10.1093/0195158709.001.0001.

Jacobi, C. et al. 2016. Quantitative analysis of large amounts of journalistic texts using topic modelling, *Digital Journalism*, 4(1), pp. 89–106. https://doi.org/10.1080/21670811.2015.1093271.

Jordan, M.I. and Mitchell, T.M. 2015. Machine learning: Trends, perspectives, and prospects, *Science*, 349(6245), pp. 255–260. https://doi.org/10.1126/science.aaa8415.

Jumper, J. et al. 2021. Highly accurate protein structure prediction with AlphaFold, *Nature*, 596(7873), pp. 583–589. https://doi.org/10.1038/s41586-021-03819-2.

King, R.D. et al. 2009. The Automation of Science, *Science*, 324(5923), pp. 85–89. https://doi.org/10.1126/science.1165620.

Kingma, D.P. and Welling, M. 2022. Auto-Encoding Variational Bayes. *arXiv.* https://doi.org/10.48550/arXiv.1312.6114.

Kiselev, V.Y. et al. 2019. Challenges in unsupervised clustering of single-cell RNA-seq data, *Nature Reviews Genetics*, 20(5), pp. 273–282. https://doi.org/10.1038/s41576-018-0088-9.

Kochkov, D. et al. 2024. Neural General Circulation Models for Weather and Climate, *Nature*, 632(8027), pp. 1060–1066. https://doi.org/10.1038/s41586-024-07744-y.

Krizhevsky, A., et al. 2017. ImageNet classification with deep convolutional neural networks, *Communications of the ACM*, 60(6), pp. 84–90. https://doi.org/10.1145/3065386.

Kuhn, T.S. 1962. *The structure of scientific revolutions* (3rd edition). University of Chicago Press.

Lazer, D. et al. 2009. Computational Social Science, *Science*, 323(5915), pp. 721–723. https://doi.org/10.1126/science.1167742.

LeCun, Y., Bengio, Y. and Hinton, G. 2015. Deep learning, *Nature*, 521(7553), pp. 436–444. https://doi.org/10.1038/nature14539.

Leonelli, S. 2016. *Data-centric biology: a philosophical study*. Chicago; London: The University of Chicago Press.

Lewis, P. et al. 2021. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *arXiv.* https://doi.org/10.48550/arXiv.2005.11401.

Li, J. et al. 2025. ChemVLM: Exploring the Power of Multimodal Large Language Models in Chemistry Area, *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(1), pp. 415–423. https://doi.org/10.1609/aaai.v39i1.32020.

Li, Q. et al. 2021. Federated Learning on Non-IID Data Silos: An Experimental Study. *arXiv.* https://doi.org/10.48550/arXiv.2102.02079.

Lian, Y. et al. 2021. Climate-driven Model Based on Long Short-Term Memory and Bayesian Optimization for Multi-day-ahead Daily Streamflow Forecasting. In Review. https://doi.org/10.21203/rs.3.rs-216524/v1.

Lohr, S. 2025. The Quest for A.I. "Scientific Superintelligence", *The New York Times*, 10 March. https://www.nytimes.com/2025/03/10/technology/ai-science-lab-lila.html?smid=url-share.

Lu, Chris et al. 2024. The AI Scientist: Towards Fully Automated Open-Ended Scientific Discovery. *arXiv.* https://doi.org/10.48550/arXiv.2408.06292.

M. Bran, A. et al. 2024. Augmenting large language models with chemistry tools, *Nature Machine Intelligence*, 6(5), pp. 525–535. https://doi.org/10.1038/s42256-024-00832-8.

Manhaeve, R. et al. 2018. DeepProbLog: Neural Probabilistic Logic Programming, in S. Bengio et al. eds. *Advances in Neural Information Processing Systems*. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2018/file/dc5d637ed5e62c36ecb73b654b05ba2a-Paper.pdf.

Nassif, A.B. et al. 2021. Machine Learning for Anomaly Detection: A Systematic Review, *IEEE Access*, 9, pp. 78658–78700. https://doi.org/10.1109/ACCESS.2021.3083060.

Pearl, J. 2000. *Causality: models, reasoning, and inference*. Cambridge, U.K.; New York: Cambridge University Press.

Qiu, X. et al. 2024. Advances in AI for Protein Structure Prediction: Implications for Cancer Drug Discovery and Development, *Biomolecules*, 14(3), p. 339. https://doi.org/10.3390/biom14030339.

Radaideh, M.I. and Kozlowski, T. 2020. Surrogate modeling of advanced computer simulations using deep Gaussian processes, *Reliability Engineering & System Safety*, 195, p. 106731. https://doi.org/10.1016/j.ress.2019.106731.

Raissi, M., et al. 2019. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *Journal of Computational Physics*, 378, pp. 686–707. https://doi.org/10.1016/j.jcp.2018.10.045.

Riegel, R. et al. 2020. Logical Neural Networks. *arXiv.* https://doi.org/10.48550/arXiv.2006.13155.

Rodrigo, G., et al. 2012. De novo automated design of small RNA circuits for engineering synthetic riboregulation in living cells, *Proceedings of the National Academy of Sciences*, 109(38), pp. 15271–15276. https://doi.org/10.1073/pnas.1203831109.

Roy, K. et al. 2019. Towards spike-based machine intelligence with neuromorphic computing, *Nature*, 575(7784), pp. 607–617. https://doi.org/10.1038/s41586-019-1677-2.

Rozière, B. et al. 2024. Code Llama: Open Foundation Models for Code. *arXiv.* https://doi.org/10.48550/arXiv.2308.12950.

Rudin, C. 2019. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *arXiv.* https://doi.org/10.48550/arXiv.1811.10154.

Salih, A. et al. 2025. A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME, *Advanced Intelligent Systems*, 7(1), p. 2400304. https://doi.org/10.1002/aisy.202400304.

Sanchez-Gonzalez, A. et al. 2020. Learning to Simulate Complex Physics with Graph Networks. *arXiv.* https://doi.org/10.48550/arXiv.2002.09405.

Savitsky, Z. 2025. Exclusive: Start-up FutureHouse debuts powerful AI "reasoning model" for science, *Nature*, 642(8068), pp. 552–553. https://doi.org/10.1038/d41586-025-01753-1.

Sheller, M.J. et al. 2019. Multi-institutional Deep Learning Modeling Without Sharing Patient Data: A Feasibility Study on Brain Tumor Segmentation, in A. Crimi et al. eds. *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*. Cham: Springer International Publishing Lecture Notes in Computer Science), pp. 92–104. https://doi.org/10.1007/978-3-030-11723-8_9.

Shiba, K. and Inoue, K. 2024. Harnessing causal forests for epidemiologic research: key considerations, *American Journal of Epidemiology*, 193(6), pp. 813–818. https://doi.org/10.1093/aje/kwae003.

Solomon, M. 2013. The structure of scientific revolutions Thomas S. Kuhn, 1970, 2nd ed. Chicago, London: University of Chicago Press Ltd. 210 pages), *Philosophical Papers and Review*, 4(4), pp. 41–48. https://doi.org/10.5897/PPR2013.0102.

Sordo, Z. et al. 2025. Synthetic Scientific Image Generation with VAE, GAN, and Diffusion Model Architectures, *Journal of Imaging*, 11(8), p. 252. https://doi.org/10.3390/jimaging11080252.

Sreedhar, K. et al. 2025. Simulating Cooperative Prosocial Behavior with Multi-Agent LLMs: Evidence and Mechanisms for AI Agents to Inform Policy Decisions. *arXiv.* https://doi.org/10.48550/arXiv.2502.12504.

Taylor, R. et al. 2022. Galactica: A Large Language Model for Science. *arXiv.* https://doi.org/10.48550/arXiv.2211.09085.

Udrescu, S.-M. and Tegmark, M. 2020. AI Feynman: a Physics-Inspired Method for Symbolic Regression. *arXiv.* https://doi.org/10.48550/arXiv.1905.11481.

Vaswani, A. et al. 2023. Attention Is All You Need. *arXiv.* https://doi.org/10.48550/arXiv.1706.03762.

Volk, A.A. et al. 2023. AlphaFlow: autonomous discovery and optimization of multi-step chemistry using a self-driven fluidic lab guided by reinforcement learning, *Nature Communications*, 14(1), p. 1403. https://doi.org/10.1038/s41467-023-37139-y.

Wang, C. et al. 2024. BioRAG: A RAG-LLM Framework for Biological Question Reasoning. *arXiv*. https://doi.org/10.48550/arXiv.2408.01107.

World Economic Forum 2024. *Top 10 Emerging Technologies of 2024*. https://www.weforum.org/publications/top-10-emerging-technologies-2024/.

Yang, J. et al. 2025. Active learning-assisted directed evolution, *Nature Communications*, 16(1), p. 714. https://doi.org/10.1038/s41467-025-55987-8.

Yelmen, B. et al. 2021. Creating artificial human genomes using generative neural networks, *PLOS Genetics*. Edited by S. Mathieson, 17(2), p. e1009303. https://doi.org/10.1371/journal.pgen.1009303.

Zhou, J. and Troyanskaya, O.G. 2015. Predicting effects of noncoding variants with deep learning–based sequence model, *Nature Methods*, 12(10), pp. 931–934. https://doi.org/10.1038/nmeth.3547.

Zhou, L. et al. 2017a. Machine learning on big data: Opportunities and challenges, *Neurocomputing*, 237, pp. 350–361. https://doi.org/10.1016/j.neucom.2017.01.026.

Zhou, Z. et al. 2017b. Optimizing Chemical Reactions with Deep Reinforcement Learning, *ACS Central Science*, 3(12), pp. 1337–1344. https://doi.org/10.1021/acscentsci.7b00492.

Zitnik, M. 2025. AI-enabled drug discovery reaches clinical milestone, *Nature Medicine*, 31(8), pp. 2490–2491. https://doi.org/10.1038/s41591-025-03832-2.

Work with the ISC to advance science as a global public good.

About the International Science Council

The International Science Council (ISC) works at the global level to catalyse change by convening scientific expertise, advice and influence on issues of major importance to both science and society.

The ISC is a non-governmental organization with a unique global membership that brings together more than 250 international scientific unions and associations, national and regional scientific organizations including academies and research councils, international federations and societies, and young academies and associations.

Connect with us at:

council.science

secretariat@council.science

International Science Council

5 rue Auguste Vacquerie

75116 Paris, France

bsky.app/profile/sciencecouncil.bsky.social

threads.net/@council.science

facebook.com/InternationalScience

https://www.linkedin.com/company/international-science-council

instagram.com/council.science