

<https://doi.org/10.1038/s42005-026-02676-7>

Decoding complexity through machine learning is redefining scientific discovery



Ricardo Vinuesa^{1,2,3} , **Paola Cinnella**⁴, **Jean Rabault**⁵, **Hossein Azizpour**^{3,6}, **Stefan Bauer**^{7,8}, **Bingni W. Brunton**⁹, **Arne Elofsson**^{3,10}, **Elias Jarlebring**^{3,11}, **Hedvig Kjellström**^{3,6}, **Stefano Markidis**^{3,12}, **David Marlevi**^{13,14}, **Javier García-Martínez**¹⁵  & **Steven L. Brunton**¹⁶ 

As scientific instruments and the literature generate ever larger volumes of data, machine learning (ML) has become essential for organizing, analyzing and interpreting complex information. This Perspective examines how ML accelerates discovery across disciplines, with examples such as brain mapping and exoplanet detection. It also considers situations with different levels of prior knowledge about the underlying phenomenon, outlining strategies to address limitations and exploit ML effectively. Although growing reliance on ML raises challenges for research practice and validation, it is reshaping scientific methods and expanding what can be studied. We also highlight foundation models as a promising route to faster, broader scientific discovery.

Machines have played a critical role in scientific discovery (i.e. to obtain fundamental and formalized knowledge about Nature) by providing the tools to observe, measure, and analyze natural phenomena. Scientific instruments, such as telescopes and microscopes, have historically enabled groundbreaking discoveries by revealing details invisible to the naked eye, expanding our understanding of the universe and the microscopic world¹. With the advent of modern scientific instruments, including DNA sequencers, astronomical observatories, and high-resolution imaging devices, research facilities are producing terabytes or even petabytes of information. As data volumes grow, computers play a critical role in organizing, analyzing, and interpreting this information. Advanced computational methods help to reduce the complexity of the data, making it possible to extract meaningful insights^{2,3}. However, even with the most advanced computers, the sheer volume of data generated by large-scale projects such as the Large Hadron Collider (LHC)⁴ and the Square Kilometer Array (SKA)⁵, and the vast amount of information available in the scientific literature, make traditional analysis methods impractical. Complex problems such as weather forecasting, drug discovery, and genomic analysis often involve highly complex data sets and processes that cannot be efficiently managed without the assistance of

machine learning (ML), which can help sift through massive data streams, identify patterns, and extract valuable insights that would be impossible for humans or traditional computational methods alone. Complexity in these problems arises from non-linearity, high dimensionality, and multiscale dynamics, posing significant challenges for traditional mathematical tools and even recent simulation-based approaches. Despite advancements in simulations and big data, we still struggle to fully understand phenomena like turbulence or biological processes at a deeper level. Advanced ML tools are also improving decision-making, enabling faster and more accurate interpretations of complex phenomena, and addressing challenges in numerous scientific fields⁶. However, it also introduces some challenges and the need for ethical guidelines to ensure the appropriate use of ML for scientific research⁷. A key issue is algorithmic bias, which can distort outcomes and lead to incorrect conclusions, particularly in areas like health care, where biased predictions can impact patient care or drug development. Additionally, the black-box nature of ML models complicates transparency, making it hard for researchers to verify decisions. Lastly, data ownership and privacy concerns arise, especially with sensitive data like genetic or health records. These issues require the development of robust ethical

¹Department of Aerospace Engineering, University of Michigan, Ann Arbor, MI, USA. ²FLOW, Engineering Mechanics, KTH Royal Institute of Technology, Stockholm, Sweden. ³Swedish e-Science Research Centre, (SeRC), Stockholm, Sweden. ⁴Institut Jean le Rond D'Alembert, Sorbonne Université, Sorbonne, France. ⁵IT Department, Norwegian Meteorological Institute, Oslo, Norway. ⁶Robotics, Perception and Learning, KTH Royal Institute of Technology, Stockholm, Sweden. ⁷TUM School of Computation, Information and Technology, Technical University Munich, Munich, Germany. ⁸Helmholtz AI, Helmholtz Center Munich, Munich, Germany. ⁹Department of Biology, University of Washington, Seattle, WA, USA. ¹⁰Dept. of Biochemistry and Biophysics and Science for Life Laboratory, Stockholm University, Solna, Sweden. ¹¹Dept. Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden. ¹²Department of Computer Science, KTH Royal Institute of Technology, Stockholm, Sweden. ¹³Dept. Molecular Medicine and Surgery, Karolinska Institutet, Stockholm, Sweden. ¹⁴Inst. for Medical Engineering and Science, Massachusetts Institute of Technology, Cambridge, MA, USA. ¹⁵Departamento de Química Inorgánica, Universidad de Alicante, Alicante, Spain. ¹⁶Department of Mechanical Engineering, University of Washington, Seattle, WA, USA. ✉e-mail: rvinuesa@umich.edu; j.garcia@ua.es; sbrunton@uw.edu

guidelines to ensure that ML contributes to science in a fair, transparent, and socially responsible way⁸.

This Perspective advances three main messages: first, that machine learning is transforming scientific discovery by enabling researchers to address complexity across different levels of prior knowledge; second, that the form and limits of ML-driven discovery depend critically on the availability of mechanistic understanding; and third, that realizing the full potential of ML in science requires addressing challenges related to data, bias, interpretability and validation. In this work, we explore the potential of ML and artificial intelligence (AI) in three types of scientific problems: (i) those where all governing equations are known, (ii) those with partial knowledge and (iii) those where little is known. We illustrate this with examples from the physical and life sciences, including turbulent flows, dark matter, drug discovery, and brain research. Figure 1 summarizes how different uses of ML help address complexity in these areas. As discussed below, as system complexity increases, the distinction between full, partial, and no knowledge of the governing equations becomes increasingly blurred, but ML plays a crucial role in tackling these challenges. A good example of a complex system with vast amounts of data that traditional tools cannot process efficiently is brain research. ML enables the reconstruction of countless brain slices into highly accurate three-dimensional (3D) maps. In a recent study⁹, Google researchers used AI to process 300 million brain images from Harvard, creating the largest-ever interactive 3D brain tissue model, now available online (the database is available here: <https://h01-release.storage.googleapis.com/landing.html>). This innovation is crucial for understanding neurological disorders, as ML can detect patterns that traditional methods might miss, supporting early diagnosis and treatment planning¹⁰. One of the most widespread applications of ML is in drug discovery, addressing the time and cost challenges of traditional methods. The drug-development process can take over a decade and billions of dollars due to the complexity of identifying viable candidates. ML is transforming this by rapidly analyzing vast biological and chemical data, uncovering patterns that might remain hidden, and streamlining the identification of promising candidates. Additionally, ML enhances predictive modeling, allowing researchers to forecast drug efficacy and safety earlier in the process. By enabling virtual screening of drug-target interactions, ML reduces the need for costly lab experiments and helps design more efficient clinical trials, accelerating development and optimizing resources¹¹. Figure 2 provides a comprehensive visualization of the key themes discussed in this paper, focusing on the role of ML in scientific discovery and the challenges it

presents, particularly with respect to data. The top-left panel illustrates the significant increase in data generated by modern scientific instrumentation over time. Initially, humans organized and interpreted the data through manual analysis and generalization. Since then, computational methods have facilitated the management of large amounts of data, contributing significantly to fields such as genomics, chemistry and mathematics. More recently, with the advent of large-scale scientific facilities such as CERN's LHC⁴, ML techniques have become essential for processing vast amounts of data. This shift has reduced human involvement in direct observation and interpretation, raising concerns about our diminishing understanding of the discoveries made by these technologies. The top-right panel outlines four key challenges related to ML in science: data quality and availability, inherent biases, explainability, and the risk of overfitting. The bottom-left panel highlights a conceptual dilemma: while ML accelerates discovery, there is growing debate about what constitutes a true scientific breakthrough, and whether ML can only “rediscover” existing concepts rather than uncover new insights. Finally, the bottom-right panel emphasizes the increasing need for high-quality data to enable new discoveries through ML, while highlighting the limitations imposed by our current gaps in scientific knowledge. All of these issues are discussed in detail in the following sections of this paper.

Recent applications of AI/ML providing scientific breakthroughs

Artificial Intelligence is playing a transformative role in physics by improving data analysis, model development, and experimental interpretation. In astronomy, ML is improving the search for exoplanets by boosting the accuracy and efficiency of data analysis¹². AI-powered algorithms, particularly convolutional neural networks, can process massive data sets from telescopes to detect Earth-like exoplanets in noisy signals more precisely than traditional methods. The transit method, which detects exoplanets by observing mini-eclipses as they pass in front of their stars, can be complicated by planetary interactions that disrupt periodicity. To address this, researchers from the University of Geneva, University of Bern, and Disaitek applied ML and image recognition techniques to predict these interactions. By training a neural network on numerous examples, they built a model capable of detecting subtle exoplanet signals that might be missed by traditional methods. Their work led to the discovery of exoplanets Kepler-1705b and Kepler-1705c, advancing our understanding of planetary systems¹³. ML is also crucial in areas like the Standard Model of particle

Fig. 1 | Conceptual framework illustrating how machine learning (ML) and artificial intelligence (AI) contribute to scientific discovery as a function of the amount of prior theoretical knowledge available. The horizontal axis represents increasing understanding of the governing equations, from **a** limited knowledge, through **b** partial knowledge, **c** to significant knowledge. Each column shows **d** a representative scientific domain, **e** the corresponding state of knowledge and **f** the primary role of ML. When few or no governing equations are known, as in brain research, ML is primarily used for representation learning, manifold learning, and pattern discovery from high-dimensional data. In systems with partial knowledge, such as drug discovery, ML complements incomplete or effective theories through generative modeling, learning effective equations, and hypothesis generation. When the governing equations are well established, as in turbulent fluid flows, ML is mainly applied to surrogate modeling, explainable AI, and control and optimization.





Fig. 2 | Visual summary of the main concepts discussed in this paper. **a** Growth in the volume of data produced by scientific instrumentation over time, accompanied by a transition from human-led data organization to computer-assisted methods and, more recently, to machine-learning-driven approaches. This shift is associated with a progressive reduction in direct human observation, intervention, and interpretability in the discovery process. **b** The four key challenges identified in applying machine learning (ML) to scientific discovery: data quality and availability, bias, explainability, and overfitting. **c** Conceptual distinction between genuine scientific discovery and rediscovery, highlighting the question of whether ML can generate fundamentally new ideas, concepts, or laws, or primarily recovers existing knowledge. **d** The inverse relationship between prior domain knowledge and data requirements, emphasizing the need for larger and higher-quality datasets when existing theoretical understanding is limited.



physics¹⁴. Automated algorithms were central to the discovery of the Higgs boson¹⁵, and future experiments will need to operate at higher energies and intensities, generating data volumes too large for traditional methods to handle. For instance, it is expected that the Large-Hadron Collider (LHC) will increase proton collision rates by an order of magnitude in the next decade, requiring data analysis tools such as ML, to identify trends, uncover hidden relationships, and design more effective experiments.

Another scientific area in which ML is increasingly assisting research is mathematics, in this case by improving the process of theorem proving, mathematical method development and discovery. ML systems have already demonstrated their ability to automate aspects of theorem proving, where they naturally complement existing formal theorem provers such as LEAN¹⁶: an AI can be used to formulate a proof, which can be verified by a formal proving tool, giving absolute trust in the validity of the result. For example, Meta AI's neural theorem prover¹⁷ successfully solved 10 International Math Olympiad (IMO) problems, far exceeding the performance of previous ML systems¹⁸. More recently, DeepMind's AlphaProof¹⁸ combined large language models with symbolic reasoning to generate and formally verify mathematical proofs, achieving performance at the level of top human competitors in advanced benchmark problems. Beyond automated proving, ML is also contributing to the discovery of new mathematical insights. Systems such as DeepMind's AlphaEvolve¹⁹ illustrate how ML can go beyond proof verification by iteratively refining mathematical ideas and conjectures in interaction with symbolic reasoning tools. DeepMind's collaborations with mathematicians have led to ML-assisted advances in knot theory and representation theory. In particular, ML has been used in a constructive way to suggest proof strategies and uncover previously unknown connections, such as in the study of Kazhdan-Lusztig polynomials. Similarly, formal results can be obtained in quantum physics using Logical AI as a way to derive formally provable results²⁰. More generally, the development of these more advanced and autonomous systems mirrors the rise of Agentic AI, which has recently become prominent in advanced question answering, coding tasks, and is now being applied to all categories of problem solving^{21,22}. Agentic AI is, in particular, diffusing into scientific applications, which gives AI systems the autonomy and "agency" to close the loop between hypothesis formulation, the design of strategies and experiments to test these, the implementation of such hypothesis testing and the analysis of the obtained results. This would lead to the confirmation or

falsification of the hypothesis to test, possibly resulting in successive loops of hypothesis formulation and testing²³. These results illustrate how AI can complement human intuition, accelerate mathematical research, and open new frontiers²⁴.

Embracing complexity

Machine learning comprises a growing set of algorithms^{25,26}, enabled by increasingly vast amounts of data and computing power²⁷, that show incredible promise for handling complexity spanning virtually all aspects of human endeavors, as we discuss in the following. Neural networks in particular, despite being governed by simple rules, can perform complex tasks for which no traditional algorithms exist. Although they are fully observable and deterministic, we often cannot explain their decisions. However, they have led to groundbreaking discoveries, such as a new class of antibiotics^{28,29}. This creates challenges, such as the need for explainable AI (XAI)³⁰⁻³². Symbolic approaches, such as gene-expression programming, sparse regression, and sparse Bayesian learning³³⁻³⁵, have been successful, but their complexity grows exponentially with the search-space size. Recent efforts to combine symbolic and deep-learning approaches have enabled advances, such as discovering new materials³⁶. This raises fundamental questions about the limits of ML in scientific discovery; for example, can a complex system understand its own complexity? And how much can AI discover beyond its training data³⁷? These issues highlight the opportunities and challenges that data-driven methods bring to science. Some believe in the "unreasonable effectiveness of data"³⁸, particularly in deep learning³⁹, but the practical implications for future discoveries remain uncertain. A key question is whether ML can provide not only computational solutions but also fundamental scientific understanding. Note that ML's applications in science are not limited to discovery. For example, AI is revolutionizing optimization, laboratory automation, and solving governing equations, such as partial differential equations (PDEs). Recent studies⁴⁰ have focused on the potential of self-supervised learning in experiments and simulations (including representations of scientific data), while others⁴¹ are exploring AI's role in formulating scientific questions. Note that the unique contribution of this work is the study of how ML can tackle complexity to reach scientific discoveries, acknowledging that different levels of knowledge of the governing equations (see Fig. 1) will require completely different ML approaches. This may constitute a revolution in how we organize disciplines

when using ML for scientific discovery, where apparently different communities may share many similarities in terms of how much is known regarding the governing equations and therefore in terms of the ML methods to be developed.

The emergence of scientific foundation models (SFM) and large language models (LLMs) is further pushing the boundaries of ML methods. Foundation models are large machine-learning or deep-learning generative models trained on vast amounts of data so they can be applied on a wide range of cases. They predict masked data regions to learn associations, excelling in multiple tasks without large, labeled datasets. Despite risks like bias and transparency issues⁴², foundation models have revolutionized AI, especially through chatbots like ChatGPT⁴³. LLMs are now able to assist with tasks from writing⁴⁴ and coding⁴⁵, to generating creative ideas and guiding scientific experiments^{46–51}. Foundation models also show promise in non-text-focused areas, such as protein-structure prediction⁵², protein design⁵³ and climate simulations⁵⁴. An interesting feature of foundation models, and particularly LLMs, is that they demonstrate so-called “emergent abilities”^{55,56}. The term refers to unexpected, not explicitly programmed, capabilities that arise as model scale increases, and do not consist in mere extrapolations of smaller models’ performance. Such abilities were first observed in complex natural systems, e.g. phase changes in materials, complex behavior in flocks of birds or fishes, as well as in computational systems such as cellular automata and agent-based models, and the term was originally popularized by the Nobel-prize winner P.W. Anderson⁵⁷. Examples of LLM emergent abilities include in-context learning, complex reasoning, and multi-step problem-solving, which are highly valuable for scientific research. Thanks to this, LLMs are found to generalize to new tasks without prior examples (zero-shot learning) or with minimal data (few-shot learning), making them particularly useful in domains where data is scarce or highly specialized. More generally, they have the potential to aid hypothesis generation, automate literature reviews, predict protein structures, and accelerate scientific discovery, especially when combined with reinforcement learning (RL) and a suitable system of rewards/penalties to optimize specific scientific tasks, such as designing experiments or iteratively refining hypotheses based on feedback. Synergy with scientists may allow to exploit the reasoning capabilities of the models in dynamic, real-world scenarios, enabling more efficient exploration of complex scientific problems. On the other hand, these abilities are often unpredictable and non-linear, raising challenges in reliability, interpretability, and ethical use. Careful oversight is then needed to ensure that these models align with the scientific goals and produce trustworthy results. Risks include the potential for generating misleading or incorrect outputs, amplifying biases present in training data, and over-reliance on automated systems without sufficient human validation. Additionally, the opacity of how these models arrive at their conclusions can hinder their adoption in critical scientific applications, raising important questions about their benefits and risks in science⁵⁸. Fajardo-Fontiveros et al.⁵⁹ discuss when it is possible to learn models from data and the acceptable noise levels for accurate learning.

Discovery versus re-discovery by machine learning

While ML has proven invaluable in refining existing knowledge, its real potential lies in detecting patterns and correlations that may not be immediately apparent due to the vast amounts of data involved. This raises the question of whether ML is truly discovering new insights or merely reinterpreting existing knowledge. Historically, scientific discovery has been rooted in inquiry, observation, and experimentation, with creativity playing a crucial role in generating new insights into phenomena or theories^{60,61}. While ML excels at analyzing large datasets, identifying patterns, and generating hypotheses (which are key components of discovery), its ability to make truly autonomous, original discoveries is still debated. Ground-breaking discoveries usually require creativity, broader contextual understanding and sometimes leaps beyond available data. ML is revolutionizing research by uncovering complex trends, and its potential to independently produce original scientific discoveries is a topic of active research in several studies⁶², and workshops, such as the one organized by the National

Academies on How ML Is Shaping Scientific Discovery⁶³. Furthermore, a recent AAAI overview⁶⁴ has synthesized progress and gaps for discovery with generative AI, calling for science-focused agents and better benchmarks/evaluations that test long-horizon reasoning and novelty. It is important to observe, from the examples listed throughout this letter, that ML has already contributed to significant scientific breakthroughs: this highlights AI’s potential to navigate uncharted scientific territory, and the value there is in leveraging AI to enable scientific discovery⁶⁵.


One of the most significant advancements in AI-driven scientific discovery is AI-Hilbert. Developed by IBM researchers, AI-Hilbert acts as an “AI scientist” that transforms existing theories and data into new, consistent mathematical models. Its goal is to accelerate scientific discovery by automating hypothesis generation and testing. AI-Hilbert helps scientists uncover new knowledge by analyzing large scientific datasets and revealing patterns overlooked by traditional methods. It also refines theories by managing conflicting data⁶⁶. Furthermore, Cornelio et al.⁶⁷ have developed a new ML tool which combines axiomatic knowledge with experimental data to derive scientific models. By integrating logical reasoning and symbolic regression, it has rediscovered laws like Kepler’s third law and Einstein’s time-dilation law. This tool can distinguish between competing formulas, even with limited data. While efficient at replicating human discoveries, these tools often confirm known theories rather than offering new insights. ML can be biased towards existing patterns, limiting its ability to generate new hypotheses. In these cases, AI’s role is more supportive, validating established knowledge rather than challenging it⁶⁸. ML holds immense potential for making groundbreaking discoveries, particularly in fields like drug development and astrophysics. However, for ML to drive new knowledge, it must evolve beyond confirming human findings. This will require developing ML models with an intrinsic understanding of the mechanisms it analyzes, where causality is key⁶⁹. ML discoveries can complement human expertise, relying on human interpretation and creativity to fully appreciate and exploit the insights gained. Explainability tools will be key for this^{32,70,71}.

Machine-learning-driven scientific discovery when complete information is available

There are several applications within Physical Sciences where the underlying governing equations are known perfectly but the high-level global dynamics are still not well understood. For instance, while we know the quantum equations behind biomolecular dynamics, complexity makes biology a partial-knowledge problem since only limited biological systems can be measured or simulated, and simulating the full human brain is impossible. Similarly, turbulent flows, described by the Navier–Stokes equations, are only partially understood due to their chaotic nature as the Reynolds number increases⁷². While the large scales in the flow can be simulated, the smaller scales are much more difficult to simulate, making turbulence a major challenge⁷³. However, some ML techniques, such as symbolic regression and reduced-order modeling, can help uncover unknown flow features and physical properties from large direct-numerical-simulation (DNS) datasets^{71,74}, as indicated in Fig. 3, as well as reveal hidden biases and structures present in experiments⁷⁵. Advances in turbulence modeling using supervised ML have also improved closure models for Reynolds-averaged Navier–Stokes (RANS) and large-eddy-simulation (LES) turbulence models^{76,77}. In astrophysics, the supervised classifier SPOCK predicts long-term stability in multi-planet systems (which requires integration of the laws of gravitation over billions of orbital periods) using short-term simulations and effectively generalizing to larger systems⁷⁸.

Optimal control of complex physical and biological systems is another challenging area where ML, particularly deep reinforcement learning (DRL), shows promise. DRL has led to breakthroughs in quantum physics, astronomy, turbulence control, and tokamak-instability control, offering insights into complex systems and discovering new strategies^{79–84}. DRL has even uncovered previously unknown thermodynamic cycles⁸⁵. ML can dramatically accelerate simulations of complex systems with known equations. Autoencoders can be used to discover latent-space representations

Fig. 3 | Schematic representation of ML directions to enable scientific discoveries when complete information about the governing equations is available. In such a case, both supervised, unsupervised, and reinforcement-learning methodologies can be used. Supervised and unsupervised methodologies are made possible by generating large datasets of synthetic data simulated from the governing equations. This allows the deployment of a variety of ML techniques that can discover complex hidden relations, nonlinear coordinate systems, hidden dynamics or solve otherwise intractable problems. Reinforcement learning can also be used by coupling it to the physics simulator, which has already proven successful at discovering previously unknown control strategies and regimes of complex systems or generating high-quality heuristic guesses that can be tested in the case of problems where solution verification is easy, but the suggestion of good candidate solutions is hard.



that enable faster time integrators and simulations, leading to better optimization and systematic studies^{86,87}. For instance, high-energy physics relies on comparing observed particle detector data with simulations, which are computationally intensive. Fast generative models like generative adversarial networks (GANs) and variational autoencoders (VAEs) offer a faster alternative, although on-going research aims at ensuring that the required accuracy is achieved^{88,89}. In quantum computing, AI is envisioned as a possible tool to tame some of the complexity that makes designing and operating a real-world quantum computer challenging⁹⁰. In climate science, foundation models for weather forecasting are revolutionizing climate studies, offering potential breakthroughs in understanding climate change and paleoclimates⁹¹. In applied mathematics, DRL and large language models (LLMs) are generating new algorithms and optimizations. Notable examples include matrix operations⁹² and combinatorial problems, where AI aids in discovering new strategies that can be validated with classical algorithms⁹³. In these cases, AI provides valuable heuristic methods, although it must be highlighted that finding adequate solutions still remains challenging.

When governing equations are known, ML primarily acts as an enabling and revealing tool: it accelerates simulations, explores large parameter spaces, uncovers latent structures, and identifies effective control strategies. In this regime, ML does not replace physical laws, but enhances our ability to exploit them, providing computational efficiency, improved modeling, and new heuristic insights into complex dynamics.

Machine-learning-driven scientific discovery when only partial information is available

In contrast to systems with well-understood governing equations, in some scientific problems we only have access to partial knowledge of the underlying mechanisms, as seen in complex materials (e.g., composites or textured materials) and certain fluids (e.g., granular or multiphase flows). These systems can exhibit simple microscopic behaviors yet result in complex macroscopic phenomena, such as the spread of infectious diseases or “active turbulence” in biological matter^{94,95}. In these cases, inductive biases, like frame invariance or symmetry constraints, can be incorporated into ML models to improve the discovery process^{70,96}. Physical constraints (e.g., thermodynamics) help create more generalizable models. For example, Moya et al.⁹⁷ proposed a neural network constrained by known thermodynamic properties, enabling broader applications across physical systems. It is also important to highlight digital twins, where only some data are available, and they combine data-driven aspects with simulations while preserving generalization properties⁹⁸. Another example of taking advantage of known physical properties of the system is embedding symmetries in


autoencoders, intending to develop reduced-order models (ROMs) in physical systems invariant to input transformations^{99,100}. These data-driven and physics-informed techniques are crucial for scientific discovery; for instance, ML is helping to derive constitutive laws for materials with complex rheologies. De Lorenzis and collaborators¹⁰¹ introduced a hybrid framework (EUCLID) to learn constitutive equations for hyperelastic solids, and the SpaRTA framework for data-driven turbulence modeling¹⁰² has been adapted to model elastic solids¹⁰³. Such approaches have also been used to infer constitutive equations from crystal structures and rheology of complex fluids¹⁰⁴. (see Fig. 4).

In the context of Life Sciences, structural biology is a field where ML has made significant advances despite partial knowledge of the phenomena. AlphaFold¹⁰⁵, for instance, folds proteins into their three-dimensional (3D) native form from a one-dimensional (1D) amino acid sequence using deep learning with embedded biases like multiple sequence alignment (MSA) and 3D equivariance. AlphaFold has also spurred innovations in ML, including single-sequence methods like ESMfold and generative models for protein design^{106,107}. Diffusion models are also being used to generate protein-backbone structures and molecular ensembles, bypassing the need for simulations^{108,109}. Furthermore, generative models have also been used to discover physical models, integrating prior knowledge with data to generalize to new scenarios¹¹⁰. Transformers, for instance, are now used in Chemistry to complete chemical reactions and predict reaction yields¹¹¹. In quantum technologies, ML methods, particularly reinforcement learning, are increasingly used for closed-loop quantum control and automated experiment design, enabling data-driven protocols for quantum-state preparation and stabilization^{20,90,112–114}. ML is also being explored for device calibration, noise characterization and error mitigation to reduce the impact of hardware noise in near-term processors^{90,115,116}. In addition, quantum computing has been proposed as a platform that could accelerate specific ML primitives (quantum ML), but whether this yields practical advantage for real-world workloads remains an open question^{90,117,118}. In climate modeling, ML has developed new LES models to ensure stable long-term forecasts¹¹⁹. ML has also been shown to enhance traditional weather-prediction systems¹²⁰.

Machine-learning-driven scientific discovery when little information is available

There are numerous phenomena across scientific disciplines whose origins and underlying principles remain elusive. In these cases, the absence of well-established governing equations or foundational physical models makes it difficult to fully capture and understand their critical dynamics. For example, neuroscience has no first-principle equations, as there are no

Fig. 4 | Example of machine learning applied to a case where partial knowledge is available. Here we illustrate a model (for instance a flow with complex rheology or a flow through a porous medium, top right of the picture) which depends on a set of known inputs x (e.g. geometry, boundary conditions, etc.) as well as on a set of hidden (unobservable) variables α describing, e.g., the fluid constitutive behavior. The latter may involve small-scale phenomena that can be difficult or impossible to describe. In such conditions, experimental or numerical data for observable quantities (e.g. velocity fields or stresses y) can be used to infer the unknown field by training a machine learning model (here represented as a neural network, although other ML approaches are possible), subjected to physical constraints (e.g. positivity, symmetries or invariances). The whole process allows, on the one hand, to train a data-driven closure model for the hidden variables α and, on the other hand, to gain a-posteriori physical knowledge of the fluid constitutive properties.



known conservation laws or symmetries to derive generalizable differential equations. Even with equations describing molecules or cells, the complexity of the brain makes full-scale simulation unfeasible. Despite this, advances in neural data acquisition, such as large-scale neural recordings and connectomics, produce alternative datasets, promising a new era of ML-driven discovery in neuroscience and behavior^{121,122}. Although full-brain simulations remain beyond reach, data-driven models can replicate key input-output relationships, generating predictions vital for discovery. In neuroscience, perturbation experiments (such as activating neuron populations during visual tasks) provide insights into visual perception. Data-driven models can generate testable hypotheses and refine their predictions iteratively through experiments. ML also synthesizes diverse experimental data, such as the MICrONS dataset, which links functional imaging with structural reconstructions of cortical neurons¹²³.

ML can learn dynamics where no a-priori knowledge exists, as seen in the Hodgkin–Huxley equations modeling neural dynamics¹²⁴. Approaches like SINDy (sparse identification of nonlinear dynamics)³⁵, genetic algorithms¹²⁵ and reinforcement learning¹²⁶ can help uncover system dynamics. Neural ordinary differential equations (NODEs)¹²⁷ are another method for modeling continuous dynamics, although they often lack scientific insight. Hybrid models relying on transformers¹²⁸ or SINDy-inspired architectures¹²⁹ aim to bridge this gap. Interpretable-ML models can also discover biomarkers for diseases or predict treatment outcomes from patient data¹³⁰. In systems lacking governing equations, interventional data like CRISPR–Ko experiments in single-cell biology¹³¹ are enabling ML to identify some of the underlying mechanisms. Automated setups are advancing¹³², leading to research in designing experiments for system identification, especially for non-linear models like NODEs¹³³. These capabilities are increasingly integrated into closed-loop, agentic systems that propose experiments, execute them and update hypotheses^{64,134}. Representation-learning techniques, such as variational autoencoders¹³⁵, are essential for simplifying the characterization of complex systems by identifying latent spaces that reduce dimensionality and reveal causal structures¹³⁶. Causality and its application to dynamical systems have gained prominence¹³⁷, especially in Earth sciences¹³⁸ and molecular biology¹³⁹, with some studies using invariance from heterogeneous experiments as a signal to identify causal ODEs¹⁴⁰. Learning structured latent spaces is of crucial importance since it provides an effective coordinate system in which the dynamics have a simple representation, which is a key requirement for generalization and interpretability¹⁴¹. In Fig. 5 we provide a schematic representation of the identification of an underlying causal structure from observations where the variables of interest are not directly observed.

Recent research has also explored improving diffusion models and solvers¹⁴², which are useful for generating state-of-the-art results in areas like image generation¹⁴³, protein modeling¹⁴⁴ and materials science¹⁴⁵. Despite their success, large pre-trained models provide limited scientific insights¹⁴⁶, and therefore constitute an opportunity for future research. Machine learning has also enhanced data collection and processing in systems with indirect or incomplete measurements. ML imputation and generative modeling can complete time-series data, improving downstream applications¹⁴⁷. Furthermore, computer-vision techniques have also automated previously manual tasks, such as segmentation in microscopy, enabling large-scale analysis of cell populations¹⁴⁸. In ethology, ML has transformed video data into animal kinematics and poses, linking behavior to underlying neural computations¹⁴⁹. Notably, large language models (LLMs) and scientific foundation models (SFMs)¹⁵⁰ are opening new avenues for extracting scientific insights directly from data, in particular when combined with agentic AI^{32,151}. Genome-scale language models (GenSLMs), for instance, are helping to learn the evolutionary landscape of SARS-CoV-2 genomes¹⁵², while LLMs are enhancing neuroscience research by integrating different datasets and summarizing insights across isolated subfields¹⁵³.

Some agentic systems are being used to coordinate the full discovery loop: hypothesis → code/experiment → analysis → write up. Early demonstrations include “The AI Scientist”¹³⁴, which auto-generates ideas, runs experiments and drafts papers, and an AI co-scientist that proposes and iteratively improves hypotheses with multi-agent debate and different validations in biomedicine⁴⁸.


In the absence of governing equations, ML enables discovery by **identifying structure in data itself**, through representation learning, causal inference, and hypothesis generation. Rather than uncovering explicit laws, ML provides testable models and organizing principles that guide experimentation and theory development, making it possible to explore systems that are otherwise beyond current mechanistic understanding.

In sum, it is hard to overstate the ongoing impact of machine learning as a critical tool that, when used in conjunction with other approaches (e.g. experiments, causality analysis, development of an adequate coordinate system, etc.), catalyzes advances in scientific fields where no information on the phenomenon under study is available.

The drawbacks, limitations and challenges of machine learning for scientific discovery

Despite the significant advances that ML has brought to scientific discovery, there are key areas that need to be addressed to accelerate its contributions

Fig. 5 | Schematic of a model where the behavior as observed in data depends on unknown structure. The observed behavior or dynamics might occur on several different spatial and temporal scales, and the observed data might reflect more or fewer aspects of the underlying system. In such conditions, *representation-learning* methods can be employed to distill out an explanation of the observed data in the form of a system of ODEs or as a causal-graph representation.



and realize its full potential while minimizing some clear risks, as discussed below:

- **Data-related challenges:** the success of ML in scientific discovery relies on large high-quality, structured datasets, but scientific data is often incomplete, noisy, or imbalanced, leading to biased models. Unstructured data, especially in fields like biology, chemistry, and geology, complicate the use of ML applications, since these algorithms are not inherently designed to handle such data. However, it was recently shown that foundation models could help for prediction tasks in small general datasets¹⁵⁴. Note that the absence of labeled data complicates the use of supervised-ML techniques¹⁵⁵.
- **Bias and ethical issues:** ML models are prone to data bias, distorting results and hindering scientific discovery. Bias in data collection or model training can reinforce existing assumptions instead of revealing true insights¹⁵⁶. This is especially concerning in fields such as drug discovery, where existing models might overlook demographic groups, making the discoveries less generalizable¹⁵⁷. Ethical concerns, especially in medicine, highlight risks in applying ML to decisions affecting human life¹⁵⁸. It is important to note that the policy implications¹⁵⁹ extend beyond research integrity: economic models of AI-assisted discovery predict faster, more profitable search only with parallel investment in hypothesis-testing infrastructure (laboratories, compute for simulation and data governance). Regulators and funders should therefore couple support for AI agents with capacity-building and auditing of closed-loop pipelines.
- **Explainability and interpretability:** ML models, particularly those based on deep learning, often function as “black boxes”, making accurate predictions without revealing their decision-making processes and introducing significant challenge for the extraction of scientific understanding for example in drug discovery, where understanding why a model predicts a certain interaction is crucial¹⁶⁰, medical applications in general, where decisions need to be communicated with the patient¹⁶¹, and more generally for any attempt at gaining fundamental understanding about a given system¹⁶². Advances in explainable AI (XAI) have improved interpretability in areas like materials science¹⁶³, chemistry¹⁶⁴, and medicine¹⁶⁵.
- **Overfitting and generalization:** ML models often overfit their training data, performing poorly on unseen data, a fact that limits their ability to generalize. In scientific contexts, this can produce misleading results and fail to capture complex, nonlinear relationships, as seen in

chemistry, biology, and astronomy^{166–168}. Overfitting restricts the potential of ML to develop universal theories essential for broad scientific understanding¹⁶⁹.

- **Evaluation and governance of agentic discovery systems:** it is essential to develop task-grounded benchmarks⁶⁴, as well as policy that matches AI-driven hypothesis generation¹⁵⁹ with sufficient testing capacity; otherwise any potential gains will not materialize. This complements economic models of prioritized search showing that AI raises success probability only if experimentation capacity scales alongside it.

Improving data quality and diversity is critical for ML models to generalize across disciplines. Initiatives like the Open Reaction Database¹⁷⁰ and the Crystallography Open Database^{171,172} aim to enhance data management under the FAIR (findability, accessibility, interoperability and reusability) principles¹⁷³. Incorporating bias detection and fairness-aware algorithms can reduce biased data impacts¹⁷⁴. XAI methods are being developed to improve the transparency of ML models, particularly in healthcare applications like brain-tumor segmentation¹⁷⁵. Hybrid approaches that combine ML with physics-based models yield promising results by ensuring that the predictions adhere to known scientific principles^{176,177}. It is also important to note that ethical frameworks are also needed to ensure fairness, transparency, and accountability, especially in medicine^{178,179}.

While AI has shown great potential in scientific discovery, significant challenges remain. Issues related to data quality, bias, interpretability, and overfitting must be addressed to harness the full potential of AI in advancing science. By improving data access, enhancing model transparency, and developing hybrid models, the scientific community can overcome these obstacles and drive meaningful, trustworthy discoveries using AI which may play an instrumental role in areas as far-reaching as gravitational waves¹⁸⁰, space exploration¹⁸¹ or even the discovery of extraterrestrial life¹⁸².

Outlook

Karl Popper famously stated in *The Open Universe: An Argument for Indeterminism from the Postscript to The Logic of Scientific Discovery* that “science can be described as the art of systematic oversimplification”¹⁸³ and this remains true due to the historical limitations in processing and analyzing vast amounts of data. As a result, scientists have long been forced to simplify their objects of study by focusing on isolated phenomena or simple models. While simplification has clear benefits, for example in the teaching and the systematization of science, oversimplification carries significant risks, as essential information and key aspects of problems may be lost in the

Table 1 | Summarizing overview of the opportunities for machine learning in scientific discovery

LEVEL OF KNOWLEDGE	MODE OF DISCOVERY	METHODS	APPLICATIONS	
			Domain (examples)	Activity (examples)
SIGNIFICANT <i>Systems with pre-existing deterministic governing equations</i>	Enabling	Supervised ML - Symbolic regression - Autoencoders - Reduced order models - Foundation models - Surrogate modelling Unsupervised ML	Fluid mechanics	Surrogate modelling, computational speed-up, and data generation
			Control and quantum physics	Control and optimization
			Mathematics	Computational speed
			Chemistry	Parametric compound sweeps
			Climate science	Forecasting and surrogate modelling
	Revealing	Explainable ML Deep Reinforcement Learning	Material science	Stability prediction, compound discovery, and self-correcting process identification
			Fluid mechanics	Turbulence dynamics, and flow control
			Quantum physics	Ground state predictions. Chemical reactivity
			Mathematics	Algorithmic optimization
			Life science	Molecular dynamics
PARTIAL <i>Systems with known input behaviour, but where model complexity prohibits complete knowledge</i>	Enabling	Generative modelling Deep Reinforcement Learning	Biology	Drug discovery, including protein and structural design
			Chemistry	Reaction prediction and synthesis planning
			Physics	Shifting known equations into unknown settings (multiscale; multidimensional)
	Revealing	Explainable ML Physics-constrained learning	Biology	Unveiling analytical equations of complex systems (multiscale; multidimensional)
			Material science	Constitutive material discovery
	Inducing	Reduced order models Physics- constrained learning	Biology	Drug development through digital twinning of known systems
			Quantum physics and quantum chemistry	Material and compound discovery through digital twinning
LIMITED <i>Systems with complex multiscale nature going beyond current deterministic knowledge</i>	Causal inference and induction	Interpretable ML	Multidisciplinary	Multiorgan, multiscale, or multidomain systems
		Explainable ML Deep Reinforcement Learning	Systems biology	Biomarker and pattern discovery

Based on the notion of decoding complexity, the level of prior deterministic knowledge (*left*) can be used to differentiate methods (*second right*) and applications (*right*) through which machine learning is fueling scientific discovery. Notably, this includes various modes of discovery (*second left*), ranging from cases where machine learning enables discovery by allowing for efficient computational usage, parametric sweeps, etc., to cases where machine learning is used to causally infer underlying mechanistic behaviors in complex multidisciplinary systems. We also note a cross-cutting agentic mode (end-to-end discovery agents) spanning all knowledge regimes by orchestrating hypothesis generation, experiment design/execution and reporting^{48,134}.

process. For example, in ecological research, studies often focus on individual species interactions without considering the broader dynamics of ecosystems or other effects, such as climate change¹⁸⁴.

ML methods have already enabled several scientific and technological advancements, from solving image classification to strategic decision-making, to winning at Go. This article has highlighted how modern data-driven methods are enabling breakthroughs in scientific discovery, focusing on state-of-the-art techniques that push beyond previous limitations. We have categorized these approaches by how much knowledge about the

underlying mechanisms is available, ranging from well-understood systems to those with unknown governing principles. As summarized in Table 1, the potential of ML spans a wide spectrum of applications, including discovering physical laws, evaluating complex systems, inferring unknown behaviors, and uncovering multiscale mechanisms. These advances apply across fields like Physics, Mathematics, Chemistry, and Life Sciences, promising faster scientific progress than ever before. At the same time, using ML for scientific discovery raises important challenges. A key advantage of ML is its ability to model complex systems, but this often requires extensive

training data. In areas like astrophysics, rare diseases, or new pharmaceuticals, data scarcity is a frequent obstacle. Fortunately, complementary ML methods can either generate needed data or bypass large datasets through techniques like self-supervised learning. Even when ML does not directly result in discoveries, it can facilitate breakthroughs where data are limited. Validation is another challenge, especially in cases where the governing principles are unknown. But ML-driven discoveries can be confirmed using traditional scientific methods, such as hypothesis testing, observational confirmation and benchmarking. Additionally, ML models often function as “black boxes”, making it hard to derive formal knowledge from their results. This is a significant issue for science, where understanding is key. However, explainable- and interpretable-ML methods offer solutions, helping to achieve discoveries in the context of established scientific principles⁷¹. Despite these challenges, ML is becoming an essential tool across disciplines, and its continued evolution promises even more opportunities for scientific discovery. With further refinement, ML techniques are likely to address the current limitations, enabling even greater advancements.

The recent Nobel Prizes in Physics¹⁸⁵ and Chemistry¹⁸⁶ highlight the transformative impact of machine learning, demonstrating its ability to revolutionize scientific discovery. ML is accelerating breakthroughs across disciplines and enabling researchers to tackle levels of complexity that were previously inaccessible through powerful tools for data analysis and pattern discovery. Embracing this complexity is essential: only by doing so can science capture the complicated interdependencies of natural systems and enable additional discoveries. The real challenge (and opportunity) of AI-driven scientific discovery lies in moving beyond solving narrow, well-defined problems to tackling complex, open-ended questions. In this broader context, debates surrounding artificial general intelligence (AGI) and artificial superintelligence (ASI)^{187,188}, although not the focus of this article, provide a useful conceptual backdrop for current developments in AI-assisted science. One prominent trajectory towards greater generality relies on increasingly large foundation models trained on massive, heterogeneous datasets¹⁸⁹. Scaling up the number of parameters (and the amount of training data) is then supposed to cause the appearance of so-called “emergent abilities”⁵³, i.e. complex, often unexpected, behaviors or phenomena that arise from the interaction of components within the model, making it able to perform tasks well beyond those for which it has been trained. While some of the existing large foundation models exhibit impressive versatility, they also expose fundamental limitations related to interpretability, causal reasoning, and robustness under distributional shifts, which are properties that are central to scientific reasoning and discovery. In addition, their training requires substantial amounts of data that are yet far from being available in a number of scientific disciplines. A complementary trajectory emphasizes the development of world models^{190,191}, in which learning systems aim to build structured, predictive representations of their environment through a combination of unsupervised and reinforcement deep learning. This approach is more naturally aligned with the goals of scientific modeling, as it seeks to capture dynamics, constraints and latent structure rather than purely input-output correlations. However, world models also introduce new risks, including implicit modeling assumptions, accumulated biases and challenges in validating the learned representations against physical reality¹⁹².

From the standpoint of scientific discovery, neither scaling alone nor increasing model autonomy is sufficient. Progress will depend on striking a careful balance between expressiveness and structure, data-driven learning and physical priors, predictive accuracy and falsifiability. Framed in this way, AGI and ASI should be seen less as imminent objectives than as conceptual limits that help clarify the challenges involved in integrating AI meaningfully into the scientific process. The real opportunity of AI-driven scientific discovery lies not in replacing human inquiry, but in augmenting it by enabling connections across scales, disciplines, and representations that would otherwise remain inaccessible, while preserving the epistemic foundations on which science ultimately depends.

Received: 2 September 2025; Accepted: 27 April 2026;

Published online: 14 May 2026

References

1. Tucci, P. History of scientific instrumentation and history of science. In Pisano, R. (ed.) *A History of Physics: Phenomena, Ideas and Mechanisms*, vol. 42 of *History of Mechanism and Machine Science* (Springer, 2025).
2. Szalay, A. & Gray, J. Science in an exponential world. *Nature* **440**, 413–414 (2006).
3. Eisenstein, M. Big data: the power of petabytes. *Nature* **527**, S2–S4 (2015).
4. Evans, L. The large hadron collider. *Annu. Rev. Nucl. Part. Sci.* **61**, 435–466 (2011).
5. Dewdney, P. E., Hall, P. J., Schilizzi, R. T. & Lazio, T. J. L. The square kilometre array. *Proc. IEEE* **97**, 1482–1496 (2009).
6. Huang, L. & Peissl, W. Artificial intelligence—a new knowledge and decision-making paradigm? In Hennen, L. et al. (eds.) *Technology Assessment in a Globalized World* (Springer, 2023).
7. Resnik, D. & Hosseini, M. The ethics of using artificial intelligence in scientific research: new guidance needed for a new tool. (AI Ethics, 2024).
8. UNESCO. Recommendation on the ethics of artificial intelligence. *SHS/BIO/PI/2021/1* 1–43 (UNESCO, 2022).
9. Shapson-Coe, A. et al. A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science* **384**, eadk4858 (2024).
10. Badrulhisham, F., Pogatzki-Zahn, E., Segelcke, D., Spisak, T. & Vollert, J. Machine learning and artificial intelligence in neuroscience: a primer for researchers. *Brain Behav. Immun.* **115**, 470–479 (2024).
11. Hutson, M. How ai is being used to accelerate clinical trials. *Nature* **627**<https://www.nature.com/articles/d41586-024-00753-x> (2024).
12. Malik, A., Moster, B. P. & Obermeier, C. Exoplanet detection using machine learning. *Monthly Not. R. Astronomical Soc.* **513**, 5505–5516 (2022).
13. Leleu, A. et al. Alleviating the transit timing variation bias in transit surveys - i. rivers: method and detection of a pair of resonant super-earth around kepler-1705. *Astron. Astrophys.* **655**, A66 (2021).
14. Radovic, A. et al. Machine learning at the energy and intensity frontiers of particle physics. *Nature* **560**, 41–48 (2018).
15. Radovic, A., Williams, M., Rousseau, D. et al. Machine learning at the energy and intensity frontiers of particle physics. *Nature* **560**, 41–48 (2018).
16. De Moura, L., Kong, S., Avigad, J., Van Doorn, F. & von Raumer, J. The lean theorem prover (system description). In *International Conference on Automated Deduction*, 378–388 (Springer, 2015).
17. Lample, G. et al. Hypertree proof search for neural theorem proving. *Adv. Neural Inf. Process. Syst.* **35**, 26337–26349 (2022).
18. Hubert, T. et al. Olympiad-level formal mathematical reasoning with reinforcement learning. *Nature* **651**, 607–613 (2025).
19. Novikov, A. et al. Alphaevolve: A coding agent for scientific and algorithmic discovery. *arXiv preprint arXiv:2506.13131* (2025).
20. Cervera-Lierta, A., Krenn, M. & Aspuru-Guzik, A. Design of quantum optical experiments with logic artificial intelligence. *Quantum* **6**, 836 (2022).
21. Acharya, D. B., Kuppan, K. & Divya, B. Agentic AI: autonomous intelligence for complex goals—a comprehensive survey. *IEEE Access* **13**, 18912–18936 (2025).
22. Sapkota, R., Roumeliotis, K. I. & Karkee, M. Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenges. *arXiv preprint arXiv:2505.10468* (2025).
23. Nägele, M. & Marquardt, F. Agentic exploration of physics models. *arXiv preprint arXiv:2509.24978* (2025).
24. Davies, A., Veličković, P., Buesing, L. et al. Advancing mathematics by guiding human intuition with AI. *Nature* **600**, 70–74 (2021).

25. Mohri, M., Rostamizadeh, A. & Talwalkar, A. Foundations of machine learning (MIT press, 2018).
26. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *nature* **521**, 436–444 (2015).
27. Sevilla, J. et al. Compute trends across three eras of machine learning. In *2022 international joint conference on neural networks (IJCNN)*, 1–8 (IEEE, 2022).
28. Wong, F. et al. Discovery of a structural class of antibiotics with explainable deep learning. *Nature* **626**, 177–185 (2024).
29. Stokes, J. M., Yang, K., Swanson, K. et al. A deep learning approach to antibiotic discovery. *Cell* **180**, 688–702 (2020).
30. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **1**, 206–215 (2019).
31. Vinuesa, R. & Sirmacek, B. Interpretable deep-learning models to help achieve the sustainable development goals. *Nat. Mach. Intell.* **3**, 926 (2021).
32. Vinuesa, R., Brunton, S. L. & Mengaldo, G. Explainable AI: Learning from the learners. *Preprint arXiv:2601.05525* (2026).
33. Ferreira, C. Gene expression programming: mathematical modeling by an artificial intelligence, vol. 21 (Springer, 2006).
34. Schmidt, M. & Lipson, H. Distilling free-form natural laws from experimental data. *Science* **324**, 81–85 (2009).
35. Brunton, S. L., Proctor, J. L. & Kutz, J. N. Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *Proc. Natl. Acad. Sci. USA* **113**, 3932–3937 (2016).
36. Merchant, A. et al. Scaling deep learning for materials discovery. *Nature* **624**, 80–85 (2023).
37. Leslie, D. Does the sun rise for ChatGPT? Scientific discovery in the age of generative AI. (AI and Ethics, 2023).
38. Halevy, A., Norvig, P. & Pereira, F. The unreasonable effectiveness of data. *IEEE Intell. Syst.* **24**, 8–12 (2009).
39. Sejnowski, T. J. The unreasonable effectiveness of deep learning in artificial intelligence. *Proc. Natl. Acad. Sci.* **117**, 30033–30038 (2020).
40. Wang, H. et al. Scientific discovery in the age of artificial intelligence. *Nature* **620**, 47–60 (2023).
41. Zenil, H. et al. The future of fundamental science led by generative closed-loop artificial intelligence. *Front. Artif. Intell.* **10** (2026).
42. Van Dis, E. A., Bollen, J., Zuidema, W., Van Rooij, R. & Bockting, C. L. ChatGPT: five priorities for research. *Nature* **614**, 224–226 (2023).
43. OpenAI. Chatgpt: Optimizing language models for dialogue. <https://openai.com/blog/chatgpt/> (2022).
44. Granjeiro, J. M. et al. The future of scientific writing: Ai tools, benefits, and ethical implications. *Braz. Dent. J.* **36**, e25–6471 (2025).
45. Daniotti, S., Wachs, J., Feng, X. & Neffke, F. Who is using AI to code? global diffusion and impact of generative AI. *Science* ead9311 (2026).
46. Bellemare-Pepin, A. et al. Divergent creativity in humans and large language models. *Scientific Reports* **16**, 1279 (2026).
47. M Bran, A. et al. Augmenting large language models with chemistry tools. *Nat. Mach. Intell.* **6**, 525–535 (2024).
48. Gottweis, J. et al. Towards an AI co-scientist. *arXiv preprint arXiv:2502.18864* (2025).
49. Su, H. et al. Many heads are better than one: Improved scientific idea generation by a llm-based multi-agent system. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 28201–28240 (ACL, 2025).
50. Boiko, D., MacKnight, R., Kline, B. & Gomes, G. Autonomous chemical research with large language models. *Nature* **624**, 570–578 (2023).
51. Wang, Q., Downey, D., Ji, H. & Hope, T. Scimon: Scientific inspiration machines optimized for novelty. *arXiv preprint arXiv:2305.14259* (2023).
52. Lin, Z. et al. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science* **379**, 1123–1130 (2023).
53. Hsu, C. et al. Learning inverse folding from millions of predicted structures. *ICML* <https://www.biorxiv.org/content/early/2022/04/10/2022.04.10.487779> (2022).
54. Chen, S., Long, G., Jiang, J., Liu, D. & Zhang, C. Foundation models for weather and climate data understanding: a comprehensive survey. *arXiv preprint arXiv:2312.03014* (2023).
55. Wei, J. et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682* (2022).
56. Schaeffer, R., Miranda, B. & Koyejo, S. Are emergent abilities of large language models a mirage? Advances in Neural Information Processing Systems 36 (NIPS, 2024).
57. Anderson, P. W. More is different: Broken symmetry and the nature of the hierarchical structure of science. *Science* **177**, 393–396 (1972).
58. Birhane, A., Kasirzadeh, A., Leslie, D. & Wachter, S. Science in the age of large language models. *Nat. Rev. Phys.* **5**, 277–280 (2023).
59. Fajardo-Fontiveros, O. et al. Fundamental limits to learning closed-form mathematical models from data. *Nat. Commun.* **14**, 1043 (2023).
60. Bechtel, W. & Richardson, R. Discovering Complexity (Princeton University Press, 1993).
61. Ben-Menahem, Y. Causation in Science (Princeton University Press, 2018).
62. Krenn, M., Pollice, R., Guo, S. Y. et al. On scientific understanding with artificial intelligence. *Nat. Rev. Phys.* **4**, 761–769 (2022).
63. of Sciences, N. A. How AI is shaping scientific discovery <https://www.nationalacademies.org/news/2023/11/how-ai-is-shaping-scientific-discovery> (2023).
64. Reddy, C. K. & Shojaei, P. Towards scientific discovery with generative AI: progress, opportunities, and challenges (ACM, 2025).
65. Nature Editorial AI will transform science – now researchers must tame it. *Nature* **621**, 658 (2023).
66. Cory-Wright, R., Cornelio, C., Dash, S. et al. Evolving scientific discovery by unifying data and background knowledge with AI Hilbert. *Nat. Commun.* **15**, 5922 (2024).
67. Cornelio, C., Dash, S., Austel, V. et al. Combining data and theory for derivable scientific discovery with AI-deduction. *Nat. Commun.* **14**, 1777 (2023).
68. Bianchini, S., Müller, M. & Pelletier, P. Artificial intelligence in science: an emerging general method of invention. *Res. Policy* **51**, 104604 (2022).
69. Martínez-Sánchez, Á, Arranz, G. & Lozano-Durán, A. Decomposing causality into its synergistic, unique, and redundant components. *Nat. Commun.* **15**, 9296 (2024).
70. Cranmer, M. et al. Discovering symbolic models from deep learning with inductive biases. 34th Conference on Neural Information Processing Systems (NeurIPS, 2020).
71. Cremades, A. et al. Identifying regions of importance in wall-bounded turbulence through explainable deep learning. *Nat. Commun.* **15**, 3864 (2024).
72. Chapman, S. & Cowling, T. G. The mathematical theory of non-uniform gases: an account of the kinetic theory of viscosity, thermal conduction and diffusion in gases (Cambridge University Press, 1990).
73. Fefferman, C. L. Existence and smoothness of the Navier-Stokes equation. *Millennium Prize Probl.* **57**, 67 (2000).
74. Lozano-Durán, A. & Arranz, G. Information-theoretic formulation of dynamical systems: causality, modeling, and control. *Phys. Rev. Res.* **4**, 023195 (2022).
75. Angriman, S. et al. Active grid turbulence anomalies through the lens of physics informed neural networks. *Results Eng.* **24**, 103265 (2024).

76. Duraisamy, K., Iaccarino, G. & Xiao, H. Turbulence modeling in the age of data. *Annu. Rev. Fluid Mech.* **51**, 357–377 (2019).
77. Brunton, S. L., Noack, B. R. & Koumoutsakos, P. Machine learning for fluid mechanics. *Annu. Rev. Fluid Mech.* **52**, 477–508 (2020).
78. Tamayo, D. et al. Predicting the long-term stability of compact multiplanet systems. *Proc. Natl. Acad. Sci. USA* **117**, 18194–18205 (2020).
79. Borah, S., Sarma, B., Kewming, M., Milburn, G. J. & Twamley, J. Measurement-based feedback quantum control with deep reinforcement learning for a double-well nonlinear potential. *Phys. Rev. Lett.* **127**, 190403 (2021).
80. Yi, K., Moon, Y.-J. & Jeong, H.-J. Application of deep reinforcement learning to major solar flare forecasting. *Astrophys. J. Suppl. Ser.* **265**, 34 (2023).
81. Guastoni, L., Rabault, J., Schlatter, P., Azizpour, H. & Vinuesa, R. Deep reinforcement learning for turbulent drag reduction in channel flows. *Eur. Phys. J. E* **46**, 27 (2023).
82. Seo, J. et al. Avoiding fusion plasma tearing instability with deep reinforcement learning. *Nature* **626**, 746–751 (2024).
83. Silver et al. D. Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
84. Degraeve, J. et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature* **602**, 414–419 (2022).
85. Beeler, C. et al. Optimizing thermodynamic trajectories using evolutionary and gradient-based reinforcement learning. *Phys. Rev. E* **104**, 064128 (2021).
86. Solera-Rico, A. et al. β -Variational autoencoders and transformers for reduced-order modelling of fluid flows. *Nat. Commun.* **15**, 1361 (2014).
87. Park, S. et al. Optimization of physical quantities in the autoencoder latent space. *Sci. Rep.* **12**, 9003 (2022).
88. Goodfellow, J. et al. Generative adversarial networks. *Preprint arXiv:1406.2661* (2014).
89. Albertsson, K. et al. Machine learning in high energy physics community white paper. In *Journal of Physics: Conference Series*, vol. 1085, 022008 (IOP Publishing, 2018).
90. Alexeev, Y. et al. Artificial intelligence for quantum computing. *Nat. Commun.* **16**, 10829 (2025).
91. Wong, C. How AI is improving climate forecasts. *Nature* **628**, 710–712 (2024).
92. Fawzi et al. A. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature* **610**, 47–53 (2022).
93. Romera-Paredes, B. et al. Mathematical discoveries from program search with large language models. *Nature* **625**, 468–475 (2024).
94. Souza, L. F., Rocha Filho, T. M. & Moret, M. A. Relating SARS-CoV-2 variants using cellular automata imaging. *Sci. Rep.* **12**, 10297 (2022).
95. Alert, R., Casademunt, J. & Joanny, J.-F. Active turbulence. *Annu. Rev. Condens. Matter Phys.* **13**, 143–170 (2022).
96. Liu, Z., Chen, Y., Du, Y. & Tegmark, M. Physics-augmented learning: A new paradigm beyond physics-informed learning. *arXiv preprint arXiv:2109.13901* (2021).
97. Moya, B., Badiás, A., González, D., Chinesta, F. & Cueto, E. A thermodynamics-informed active learning approach to perception and reasoning about fluids. *Comput. Mech.* **72**, 577–591 (2023).
98. Kapteyn, M. G., Pretorius, J. V. R. & Willcox, K. E. A probabilistic graphical model foundation for enabling predictive digital twins at scale. *Nat. Comput. Sci.* **1**, 337–347 (2021).
99. Kneer, S., Sayadi, T., Sipp, D., Schmid, P. & Rigas, G. Symmetry-aware autoencoders: s-PCA and s-NLPCA. *Preprint arXiv:2111.02893v3* (2022).
100. Otto, S. E., Zolman, N., Kutz, J. N. & Brunton, S. L. A unified framework to enforce, discover, and promote symmetry in machine learning. *J. Mach. Learn. Res.* **26**, 1–83 (2025).
101. Flaschel, M., Kumar, S. & De Lorenzis, L. Automated discovery of generalized standard material models with euclid. *Computer Methods Appl. Mech. Eng.* **405**, 115867 (2023).
102. Schmelzer, M., Dwight, R. P. & Cinnella, P. Discovery of algebraic Reynolds-stress models using sparse symbolic regression. *Flow. Turbulence Combust.* **104**, 579–603 (2020).
103. Wang, M., Chen, C. & Liu, W. Establish algebraic data-driven constitutive models for elastic solids with a tensorial sparse symbolic regression method and a hybrid feature selection technique. *J. Mech. Phys. Solids* **159**, 104742 (2022).
104. Mahmoudabadbozchelou, M., Kamani, K. M., Rogers, S. A. & Jamali, S. Digital rheometer twins: learning the hidden rheology of complex fluids through rheology-informed graph neural networks. *Proc. Natl. Acad. Sci. USA* **119**, e2202234119 (2022).
105. Jumper, J. et al. Highly accurate protein structure prediction with alphafold. *Nature* **596**, 583–589 (2021).
106. Madani, A. et al. Large language models generate functional protein sequences across diverse families. *Nat. Biotechnol.* **41**, 1099–1106 (2023).
107. Dauparas, J. et al. Robust deep learning-based protein sequence design using ProteinMPNN. *Science* **378**, 49–56 (2022).
108. Watson, J. et al. De novo design of protein structure and function with RFdiffusion. *Nature* **620**, 1089–1100 (2023).
109. Ingraham, J. et al. Illuminating protein space with a programmable generative model. *Nature* **623**, 1070–1078 (2023).
110. Cornelio, C. et al. Combining data and theory for derivable scientific discovery with ai-descartes. *Nat. Commun.* **14**, 1777 (2023).
111. Irwin, R., Dimitriadis, S., He, J. & Bjerrum, E. J. Chemformer: a pre-trained transformer for computational chemistry. *Mach. Learn.: Sci. Technol.* **3**, 015022 (2022).
112. Ma, H. et al. Machine learning for estimation and control of quantum systems. *Natl. Sci. Rev.* **12**, nwaf269 (2025).
113. Bukov, M. et al. Reinforcement learning in different phases of quantum control. *Phys. Rev. X* **8**, 031086 (2018).
114. Zhang, X.M., Wei, Z., Asad, R., Yang, X.-C. & Wang, X. When does reinforcement learning stand out in quantum control? a comparative study on state preparation. *npj Quantum Inf.* **5**, 85 (2019).
115. Cai, Z. et al. Quantum error mitigation. *Rev. Mod. Phys.* **95**, 045005 (2023).
116. Canonici, E., Martina, S., Mengoni, R., Ottaviani, D. & Caruso, F. Machine learning based noise characterization and correction on neutral atoms nisy devices. *Adv. Quantum Technol.* **7**, 2300192 (2024).
117. Biamonte, J. et al. Quantum machine learning. *Nature* **549**, 195–202 (2017).
118. Ciliberto, C. et al. Quantum machine learning: a classical perspective. *Proc. R. Soc. A* **474**, 20170551 (2018).
119. Frezat, H., Sommer, J., Fablet, R., Balarac, G. & Lguensat, R. A posteriori learning for quasi-geostrophic turbulence parametrization. *J. Adv. Modeling Earth Syst.* **14**, e2022MS003124 (2022).
120. Molina, M. J. et al. A review of recent and emerging machine learning applications for climate variability and weather phenomena. *Artif. Intell. Earth Syst.* **2**, 220086 (2023).
121. Steinmetz, N. A., Zarka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
122. Yao, S. et al. A whole-brain monosynaptic input connectome to neuron classes in mouse visual cortex. *Nat. Neurosci.* **26**, 350–364 (2023).
123. Consortium, M. et al. Functional connectomics spanning multiple areas of mouse visual cortex. *BioRxiv*2021-07 (2021).
124. Nelson, M. & Rinzel, J. The Hodgkin–Huxley model. *The Book of Genesis* (Springer, 1995).

125. Chen, Y., Luo, Y., Liu, Q., Xu, H. & Zhang, D. Symbolic genetic algorithm for discovering open-form partial differential equations (sga-pde). *Phys. Rev. Res.* **4**, 023174 (2022).
126. Du, M., Chen, Y. & Zhang, D. Discover: deep identification of symbolic open-form PDEs via enhanced reinforcement-learning. *Phys. Rev. Res.* **6**, 013182 (2024).
127. Chen, R. T., Rubanova, Y., Bettencourt, J. & Duvenaud, D. K. Neural ordinary differential equations. *Advances in neural information processing systems* 31 (2018).
128. Becker, S., Klein, M., Neitz, A., Parascandolo, G. & Kilbertus, N. Predicting ordinary differential equations with transformers. In *International Conference on Machine Learning*, 1978–2002 (PMLR, 2023).
129. Sahoo, S., Lampert, C. & Martius, G. Learning equations for extrapolation and control. In *International Conference on Machine Learning*, 4442–4450 (PMLR, 2018).
130. Qiu, S. et al. Development and validation of an interpretable deep learning framework for Alzheimer's disease classification. *Brain* **143**, 1920–1933 (2020).
131. Ji, Y., Lotfollahi, M., Wolf, F. A. & Theis, F. J. Machine learning for perturbational single-cell omics. *Cell Syst.* **12**, 522–537 (2021).
132. MacLeod, B. P. et al. A self-driving laboratory advances the pareto front for material properties. *Nat. Commun.* **13**, 995 (2022).
133. Du, J., Futoma, J. & Doshi-Velez, F. Model-based reinforcement learning for semi-Markov decision processes with neural ODEs. *Adv. Neural Inf. Process. Syst.* **33**, 19805–19816 (2020).
134. Lu, C. et al. Towards end-to-end automation of AI research. *Nature* **651**, 14–919 (2026).
135. Bengio, Y., Courville, A. & Vincent, P. Representation learning: a review and new perspectives. *IEEE Trans. pattern Anal. Mach. Intell.* **35**, 1798–1828 (2013).
136. Schölkopf, B. et al. Toward causal representation learning. *Proc. IEEE* **109**, 612–634 (2021).
137. Camps-Valls, G. et al. Discovering causal relations and equations from data. *Phys. Rep.* **1044**, 1–68 (2023).
138. Runge, J. et al. Inferring causation from time series in Earth system sciences. *Nat. Commun.* **10**, 2553 (2019).
139. Lobentanzer, S., Rodriguez-Mier, P., Bauer, S. & Saez-Rodriguez, J. Molecular causality in the advent of foundation models. *Preprint arXiv:2401.09558* (2024).
140. Pfister, N., Bauer, S. & Peters, J. Learning stable and predictive structures in kinetic systems. *Proc. Natl. Acad. Sci. USA* **116**, 25405–25411 (2019).
141. Champion, K., Lusch, B., Kutz, J. N. & Brunton, S. L. Data-driven discovery of coordinates and governing equations. *Proc. Natl. Acad. Sci. USA* **116**, 22445–22451 (2019).
142. Lu, C. et al. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Adv. Neural Inf. Process. Syst.* **35**, 5775–5787 (2022).
143. Rombach, R., Blattmann, A., Lorenz, D., Esser, P. & Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10684–10695 (IEEE, 2022).
144. Watson, J. L. et al. De novo design of protein structure and function with rfdiffusion. *Nature* **620**, 1089–1100 (2023).
145. Zeni, C. et al. A generative model for inorganic materials design. *Nature* **639**, 624–632 (2025).
146. Nichani, E., Damian, A. & Lee, J. D. How transformers learn causal structure with gradient descent. *Preprint arXiv:2402.14735* (2024).
147. Vetter, J., Macke, J. H. & Gao, R. Generating realistic neurophysiological time series with denoising diffusion probabilistic models. *bioRxiv2023-08* (2023).
148. Kirillov, A. et al. Segment anything *arXiv: 2304.02643* (2023).
149. Mathis, A. et al. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).
150. Bommasani, R. et al. On the opportunities and risks of foundation models. *Preprint arXiv:2108.07258* (2022).
151. Carreon, A., Sharma, V. & Raman, V. Automated design optimization via strategic search with large language models. *Preprint arXiv:2511.22651* (2025).
152. Zvyagin, M. et al. Genslms: Genome-scale language models reveal sars-cov-2 evolutionary dynamics. *Int. J. High. Perform. Comput. Appl.* **37**, 683–705 (2023).
153. Bzdok, D. et al. Data science opportunities of large language models for neuroscience and biomedicine. *Neuron* **112**, 698–717 (2024).
154. Hollmann, N. et al. Accurate predictions on small data with a tabular foundation model. *Nature* **637**, 319–326 (2025).
155. Singh, S. & Hooda, S. A study of challenges and limitations to applying machine learning to highly unstructured data. In *2023 7th International Conference On Computing, Communication, Control And Automation (ICCUBEA)*, 1–6 (ICCUBEA, 2023).
156. Jones, C., Castro, D. C., De Sousa Ribeiro, F. et al. A causal perspective on dataset bias in machine learning for medical imaging. *Nat. Mach. Intell.* **6**, 138–146 (2024).
157. Vaidya, A., Chen, R. J., Williamson, D. F. K. et al. Demographic bias in misdiagnosis by computational pathology models. *Nat. Med.* **30**, 1174–1190 (2024).
158. van Giffen, B., Herhausen, D. & Fahse, T. Overcoming the pitfalls and perils of algorithms: a classification of machine learning biases and mitigation methods. *J. Bus. Res.* **144**, 93–106 (2022).
159. Agrawal, A., McHale, J. & Oettl, A. Artificial intelligence and scientific discovery: a model of prioritized search. *Res. Policy* **53**, 104989 (2024).
160. Savage, N. Breaking into the black box of artificial intelligence. *Nature* **66**, 101797 (2022).
161. Quinn, T. P., Jacobs, S., Senadeera, M., Le, V. & Coghlan, S. The three ghosts of medical AI: can the black-box present deliver? *Artif. Intell. Med.* **124**, 102158 (2022).
162. Mengaldo, G. Explain the black box for the sake of science: the scientific method in the era of generative artificial intelligence. *arXiv preprint arXiv:2406.10557* (2024).
163. Zhong, X., Gallagher, B., Liu, S. et al. Explainable machine learning in materials science. *npj Comput. Mater.* **8**, 204 (2022).
164. Gallegos, M., Vassilev-Galindo, V., Poltavsky, I. et al. Explainable chemical artificial intelligence from accurate machine learning of real-space chemical descriptors. *Nat. Commun.* **15**, 4345 (2024).
165. Moncada-Torres, A., van Maaren, M. C., Hendriks, M. P. et al. Explainable machine learning can outperform Cox regression predictions and provide insights in breast cancer survival. *Sci. Rep.* **11**, 6968 (2021).
166. Lever, J., Krzywinski, M. & Altman, N. Model selection and overfitting. *Nat. Methods* **13**, 703–704 (2016).
167. Welcome to the AI future? *Nat. Astron.* **7**, 1 (2023).
168. Sapoval, N., Aghazadeh, A., Nute, M. G. et al. Current progress and open challenges for applying deep learning across the biosciences. *Nat. Commun.* **13**, 1728 (2022).
169. Ridhawi, I. A., Otoum, S., Aloqaily, M. & Boukerche, A. Generalizing AI: Challenges and opportunities for plug and play ai solutions. *IEEE Netw.* **35**, 372–379 (2021).
170. The open reaction database. *J. Am. Chem. Soc.* **143**, 18820–18826 (2021).
171. Jacobsson, T. J. et al. An open-access database and analysis tool for perovskite solar cells based on the FAIR data principles. *Nat. Energy* **7**, 107–115 (2022).
172. Gražulis, S. et al. Crystallography open database—an open-access collection of crystal structures. *J. Appl. Crystallogr.* **42**, 726–729 (2009).

173. Barba, L. *Reproducibility and replicability in science* (National Academies Press, 2019).
174. Singh, P. Systematic review of data-centric approaches in artificial intelligence and machine learning. *Data Sci. Manag.* **6**, 144–157 (2023).
175. Yang, J., Soltan, A. A. S., Eyre, D. W. et al. Algorithmic fairness and bias mitigation for clinical machine learning with deep reinforcement learning. *Nat. Mach. Intell.* **5**, 884–894 (2023).
176. Abrámoff, M. D., Tarver, M. E., Loyo-Berrios, N. et al. Considerations for addressing bias in artificial intelligence for health equity. *npj Digit. Med.* **6**, 170 (2023).
177. Zheng, P., Zubatyuk, R., Wu, W. et al. Artificial intelligence-enhanced quantum chemical method with broad applicability. *Nat. Commun.* **12**, 7022 (2021).
178. Davies, J. Program good ethics into artificial intelligence. *Nature* **10**, 538291 (2016).
179. Mhasawade, V., Zhao, Y. & Chunara, R. Machine learning and algorithmic fairness in public and population health. *Nat. Mach. Intell.* **3**, 659–666 (2021).
180. Huerta, E. A., Khan, A., Huang, X. et al. Accelerated, scalable and reproducible AI-driven gravitational wave detection. *Nat. Astron.* **5**, 1062–1068 (2021).
181. Space missions out of this world with AI. *Nat. Mach. Intell.* **5**, 183 (2023).
182. Ma, P. X., Ng, C., Rizk, L. et al. A deep-learning search for technosignatures from 820 nearby stars. *Nat. Astron.* **7**, 492–502 (2023).
183. Popper, K. *The Open Universe: An Argument for Indeterminism From the Postscript to The Logic of Scientific Discovery* (Routledge, 1992).
184. Akesson, A., Curtsdotter, A., Eklöf, A. et al. The importance of species interactions in eco-evolutionary community dynamics under climate change. *Nat. Commun.* **12**, 4759 (2021).
185. Physics Nobel scooped by machine-learning pioneers. *Nature* <https://www.nature.com/articles/d41586-024-03213-8> (2024).
186. Chemistry Nobel goes to developers of AlphaFold AI that predicts protein structures. *Nature* <https://www.nature.com/articles/d41586-024-03214-7> (2024).
187. Mitchell, M. Debates on the nature of artificial general intelligence. *Science* **383**, 7869 (2024).
188. Nature Editorial More-powerful AI is coming. academia and industry must oversee it — together. *Nature* **636**, 273 (2024).
189. Menon, S. S. et al. On scientific foundation models: rigorous definitions, key applications, and a comprehensive survey. *Neural Netw.* **198**, 108567 (2026).
190. Ha, D. & Schmidhuber, J. World models. *CoRR* arXiv: 1803.10122 (2018).
191. LeCun, Y. A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27. *Open Rev.* **62**, 1–62 (2022).
192. Mitrokhov, K. Between world models and model worlds: on generality, agency, and worlding in machine learning. *AI Soc.* **40**, 5087–5099 (2025).

Acknowledgements

The following researchers are acknowledged for helpful discussions during the preparation of this article: Frida Bender, Annica Ekman, Inga Koszalka, Romit Maulik, Henrik Nielsen, Gunilla Svensson, Björn Wallner. Funding RV

and HA acknowledge SeRC and Digital Futures for funding the workshop that initiated this work. RV acknowledges financial support from ERC grant no. ‘2021-CoG-101043998, DEEPCONTROL’. PC acknowledges support by the French National Research Agency (ANR) under the France 2030 program, within the framework of the IA Cluster initiative (grant ANR-23-IACL-0007). DM acknowledges financial support from ERC grant no. 2022-StG-101075494, MultiPRESS. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible. AE was funded by the Vetenskapsrådet Grant No. 2021-03979 and the Knut and Alice Wallenberg Foundation and by SeRC. SLB acknowledges funding support from the US National Science Foundation AI Institute in Dynamic Systems (grant number 2112085) and from The Boeing Company.

Author contributions

R.V. and H.A. initiated the idea for this article following a workshop celebrated in November 2022 at KTH. R.V., P.C., J.R., H.A., S.B., B.W.B., A.E., E.J., H.K., S.M., D.M., J.G.-M. and S.L.B. contributed equally to the conceptualization, analysis, writing, editing, and revision of the manuscript. All authors reviewed and approved the final version.

Funding

Open access funding provided by Royal Institute of Technology.

Competing interests

All authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Ricardo Vinuesa, Javier García-Martínez or Steven L. Brunton.

Peer review information : *Communications Physics* thanks Anurag Daram and the other, anonymous, reviewer for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026