

Mathematics in the age of artificial intelligence

A primer on key topics in the
mathematical sciences that underpin AI

June 2026

© Academy for the Mathematical Sciences 2026

This work is licensed under the Creative Commons BY-NC-ND 4.0 International licence.



A copy of the licence may be viewed [here](#).

When using or quoting this work, **credit must be attributed to the Academy for the Mathematical Sciences**. In academic publications, the work may be cited as:

Leslie *et al.* (2026). Mathematics in the age of artificial intelligence: A primer on key topics in the mathematical sciences that underpin AI. Academy for the Mathematical Sciences. DOI: [10.5281/zenodo.20538963](https://doi.org/10.5281/zenodo.20538963).

Any enquiries regarding this work should be sent to contact@AcadMathSci.org.uk.

The Academy for the Mathematical Sciences (AcadMathSci) was established in 2022 to be the single, authoritative voice for the UK mathematical sciences community, spanning education, academia, and business, industry, and government. It was registered as a charity in 2023; the first cohort of Fellows was appointed in late 2025 and announced in early 2026.

The Academy engages actively with national policy and strategic priorities—responding to government consultations and reviews and contributing to debates on issues from curriculum reform to AI. We work across all four nations to influence policy and champion the role of the mathematical sciences. Through the Learned Societies Forum (LSF), we bring together the five learned societies in the mathematical sciences to collaborate on shared priorities and amplify our collective voice.

Contents

Introduction	4
Algebraic geometry	5
Bayesian statistics	6
Complex systems theory	7
Complexity theory	8
Control theory	9
Dynamical systems theory	10
Experimental design	11
Fluid dynamics	12
Graph theory	14
Information theory	15
Inverse problems	16
Machine learning	18
Mathematical analysis	19
Mathematical modelling, validation, and verification	20
Numerical analysis	20
Optimisation theory	22
Physics-informed modelling	23
Statistical theory	24
Stochastic analysis	25
Topology and geometry	26
Uncertainty quantification	27

Introduction

Artificial intelligence (AI) is prompting every field of human endeavour to reconsider its future: Is human engagement and investment still valuable? What kinds of expertise will remain essential? Where should future effort be directed? This document argues that mathematics provides the foundations of modern AI, and further mathematical innovation will be essential to AI's future development.

As AI systems become more capable, more pervasive, and more consequential, progress will depend not simply on larger models, greater computational power, or wider data access, but on deeper theoretical understanding. Questions of reliability, interpretability, optimisation, uncertainty, safety, and robustness are all, in substantial part, mathematical questions. A strong mathematical sciences research base is therefore essential both to the UK's capacity to develop the next generation of AI and to the evaluation, deployment, and regulation of AI—responsibly and for the benefit of society as a whole.

This “primer” brings together short contributions from across the mathematical sciences to show why continued investment in mathematics, statistics, and data science is critical if we are to reap the benefits of AI in the future. Each section introduces a mathematical topic, notes some of its established applications, and outlines its relevance to AI. We could have included many more. The reader is referred to the Royal Statistical Society's Policy piece, [AI is Statistics](#), and the SIAM report [The Role of Applied Mathematics in a New Era of Artificial Intelligence](#) for further examples of the importance of mathematical science for current and future developments in AI.

Taken together, these pieces are intended to provide foundational evidence to support a wider advocacy effort within and beyond the mathematical sciences community: to help policymakers, funders, university leaders, and other decision-makers recognise that sustained investment in the mathematical sciences is vital to the future of AI—and to the public benefit that future AI may deliver.

Contributing authors

Prof. David Leslie , Prof. Miguel Anjos , Dr Olga Anosova , Prof. Rachel Bearon , Prof. Dame Muffy Calder DBE FREng FRSE , Prof. José Carrillo , Prof. Alan Champneys , Prof. Tom Coates , Prof. Sam Cohen , Dr Jessica Enright , Prof. Dame Alison Etheridge DBE FRS , Prof. Des Higham FRSE , Prof. Sidharth Jaggi , Prof. Vitaliy Kurlin , Dr Robert Leese, Dr Vinesh Maguire-Rajpaul , Dr Christie Marr , Prof. Chris Nemeth , Dr Dan Nicolau , Dr Christian Offen , Dr Lorenzo Pareschi , Prof. Marcelo Pereyra , Prof. Yasser Roudi , Prof. Carola-Bibiane Schönlieb , Zakhar Shumaylov , Prof. David Silvester , Prof. Peter Tino , Dr Sara Veneziale , Benjamin Walker , Prof. Dave Woods , Prof. Yi Yu .

Algebraic geometry

What is algebraic geometry?

Algebraic geometry studies the shapes defined by polynomial equations. A simple example is the (algebraic) equation $x^2 + y^2 = 1$ which defines a circle (a geometric object). These objects—called algebraic varieties—range from simple shapes to more complex systems involving many variables and higher-degree polynomials. Algebraic geometry provides the language to work on abstract spaces beyond our three-dimensional experience. Moreover, algebraic geometry establishes deep connections between different areas of pure mathematics (algebra, geometry, topology, and number theory) whilst having a rich history of applications across physics, statistics, cryptography, and AI.

Established applications

In physics, algebraic geometry gives researchers the language to describe the foundations of string theory and mirror symmetry, where different geometric spaces can describe the same physical phenomena [1]. In statistics and data science, algebraic geometry gives new ways of interpreting and studying the structure of statistical models: the space of all possible Gaussian mixtures, phylogenetic trees, or hidden Markov models can be understood as algebraic varieties, with their singularities (those bad areas that make their geometry degenerate) corresponding to models that are difficult to distinguish statistically [2, 3, 4]. Coding theory and cryptography rely heavily on algebraic curves and their properties, with elliptic curve cryptography providing security for modern digital communications.

Algebraic geometry meets AI

Today, ideas from algebraic geometry are increasingly central in establishing a theoretical basis to the development and, critically, understanding of modern AI systems. There are many areas of interactions between these fields; here we report only a few of them.

Singular learning theory, developed by Sumio Watanabe, applies algebraic geometry to analyse learning systems by studying the singularities in the loss function landscape, points where multiple parameters produce identical outputs. This geometric perspective explains why neural networks generalise well despite being overparameterised, resolving puzzles that traditional statistical theory could not address [5].

LoRA (Low-Rank Adaptation) is an empirical technique for efficiently and effectively fine-tuning large language models to specific tasks by optimising only a subset of the parameters (leading to huge reductions in the compute needed). The technique relies on low-rank matrix approximation and tensor decomposition, deeply geometric problems since the space of low-rank objects can be understood as an algebraic variety [6].

Recent work on tropical geometry, a piecewise-linear version of algebraic geometry, has revealed that neural networks naturally compute tropical rational functions, providing a new algebraic framework for understanding deep learning architectures. The expressivity and optimisation landscape of neural networks can be analysed through the lens of polynomial systems and the corresponding algebraic varieties [7].

Even fundamental questions about optimisation algorithms for machine learning, related to how to sample during stochastic gradient descent, require deep results from real algebraic geometry. The Recht-Ré conjecture, a long-standing conjecture in machine learning, was disproved using the noncommutative Positivstellensatz, a result from algebraic geometry [8].

Looking ahead

As AI systems grow more complex, sophisticated tools of algebraic geometry from this and the last century are proving indispensable for giving a theoretical understanding of how and why these systems work, and how to make them better.

The integration of algebraic geometry into AI demonstrates that fundamental mathematical research, often conducted decades before possible applications are discovered, remains essential for technological progress. Singular learning theory builds on algebraic geometry developed in the 1960s, while tropical geometry originated in the 1990s.

References

- [1] P. Candelas et al. “Vacuum Configurations for Superstrings”. In: *Nuclear Physics B* 258 (1985), pp. 46–74. DOI: [10.1016/0550-3213\(85\)90602-9](https://doi.org/10.1016/0550-3213(85)90602-9).
- [2] M. Drton, B. Sturmfels, and S. Sullivant. *Lectures on Algebraic Statistics*. Vol. 39. Oberwolfach Seminars. Birkhäuser Basel, 2009. ISBN: 978-3-7643-8904-8. DOI: [10.1007/978-3-7643-8905-5](https://doi.org/10.1007/978-3-7643-8905-5).
- [3] L. D. Garcia, M. Stillman, and B. Sturmfels. “Algebraic Geometry of Bayesian Networks”. In: *Journal of Symbolic Computation* 39.3–4 (2005), pp. 331–355. DOI: [10.1016/j.jsc.2004.11.007](https://doi.org/10.1016/j.jsc.2004.11.007).
- [4] L. Pachter and B. Sturmfels, eds. *Algebraic Statistics for Computational Biology*. Cambridge: Cambridge University Press, 2005. ISBN: 978-0-521-85700-0. DOI: [10.1017/CB09780511610684](https://doi.org/10.1017/CB09780511610684).
- [5] S. Watanabe. “Algebraic Information Geometry for Learning Machines with Singularities”. In: *Advances in Neural Information Processing Systems 13*. Ed. by T. K. Leen, T. G. Dietterich, and V. Tresp. MIT Press, 2000, pp. 329–335.

- [6] F. J. Király, L. Theran, and R. Tomioka. “The Algebraic Combinatorial Approach for Low-Rank Matrix Completion”. In: *Journal of Machine Learning Research* 16.41 (2015), pp. 1391–1436.
 - [7] L. Zhang, G. Naitzat, and L.-H. Lim. “Tropical Geometry of Deep Neural Networks”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by J. Dy and A. Krause. Vol. 80. Proceedings of Machine Learning Research. PMLR, 2018, pp. 5824–5832.
 - [8] Z. Lai and L.-H. Lim. “Recht-Ré Noncommutative Arithmetic-Geometric Mean Conjecture is False”. In: *Proceedings of the 37th International Conference on Machine Learning*. Ed. by H. Daumé III and A. Singh. Vol. 119. Proceedings of Machine Learning Research. PMLR, 2020, pp. 5608–5617.
-

Bayesian statistics

What is Bayesian statistics?

Bayesian statistics is a mathematical framework for reasoning under uncertainty. Instead of producing a single “best guess”, it represents uncertainty explicitly through probability distributions over unknown quantities, such as parameters, hidden states, or future outcomes. This allows us to say not only what seems likely, but also how confident we are and how that confidence should change as new evidence arrives.

The key mathematical tool is Bayes’ rule, which provides a principled way to combine prior knowledge with observed data. What is distinctive about the Bayesian approach is that all unknown quantities are treated probabilistically throughout, so that uncertainty is represented explicitly and updated systematically as evidence grows.

Bayes’ rule allows us to update beliefs in a coherent and transparent way as information accumulates. To make this framework usable in practice, Bayesian statistics relies heavily on computational methods: algorithms and theory that turn probabilistic models into workable methods for real data.

Established applications

Bayesian ideas have long been important in applications where decisions must be made under uncertainty. Many such problems are inverse problems: we observe incomplete or noisy data and must infer the underlying state of a system. Examples include tracking disease progression from partial measurements, estimating risk in complex engineering systems, calibrating scientific models, and making forecasts in changing environments.

In these settings, it is often not enough to issue a prediction alone. We also need to know how reliable that prediction is, how sensitive it is to assumptions, and what the consequences of error might be. A clinical risk model, for example, should not only predict an outcome but also indicate the uncertainty around that prediction, since this may affect treatment decisions. Bayesian methods provide a disciplined way to propagate uncertainty through the entire analysis and into downstream decisions.

At the computational core of practical Bayesian inference are Monte Carlo methods, which generate representative samples from complicated probability distributions. Markov chain Monte Carlo (MCMC), together with more recent developments, makes it possible to work with distributions that are easy to define mathematically but impossible to compute exactly.

Bayesian statistics meets AI

These ideas now sit at the centre of many modern AI and machine learning pipelines. Alongside Monte Carlo methods, optimisation-based approximations such as variational inference and particle-based methods such as sequential Monte Carlo are widely used, especially in large-scale, online, and time-dependent settings. A major strength of Bayesian and computational statistical methods is that they come with mathematical theory and diagnostic tools, helping us assess when these methods are reliable, when they are well calibrated, and when they may fail.

The connection to AI is especially clear in generative modelling. Generative AI can often be understood as learning a complex probability distribution from data and then sampling from it. Diffusion models, for example, generate new images and other forms of data by learning how to reverse a gradual noising process. More broadly, score-based methods and flow-matching approaches learn dynamics that describe how distributions evolve, and then generate samples by simulating those dynamics.

This means that many advances in generative AI rely on mathematical ideas that are already central to Bayesian computation and computational statistics: stochastic simulation, high-dimensional probability, numerical stability, calibration, convergence analysis, and principled ways to impose constraints or condition on available information.

Looking ahead

AI therefore increases, rather than reduces, the importance of Bayesian statistics and computational statistics. Reliable uncertainty quantification is essential for safety and robustness, for example when a system should defer a decision or request more information.

It also matters for security, where one wants to detect distribution shift or model mismatch, and for accountability, where outputs need to be explainable and auditable.

Just as importantly, the Bayesian workflow encourages model criticism: checking whether a model's predictions are consistent with reality, whether its uncertainties are well calibrated, and whether conclusions are sensitive to modelling assumptions. If AI systems are to be deployed responsibly in areas such as healthcare, finance, defence, science, and critical infrastructure, then we need mathematically grounded methods for inference, uncertainty quantification, calibration, and stochastic computation. These are what make learning from data dependable, rather than merely powerful.

References

- [1] N. Metropolis et al. "Equation of State Calculations by Fast Computing Machines". In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092. doi: [10.1063/1.1699114](https://doi.org/10.1063/1.1699114).
- [2] W. K. Hastings. "Monte Carlo Sampling Methods Using Markov Chains and Their Applications". In: *Biometrika* 57.1 (1970), pp. 97–109. doi: [10.1093/biomet/57.1.97](https://doi.org/10.1093/biomet/57.1.97).
- [3] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. "Variational Inference: A Review for Statisticians". In: *Journal of the American Statistical Association* 112.518 (2017), pp. 859–877. doi: [10.1080/01621459.2017.1285773](https://doi.org/10.1080/01621459.2017.1285773).
- [4] Y. Song et al. "Score-Based Generative Modeling through Stochastic Differential Equations". In: *International Conference on Learning Representations*. 2021.
- [5] Y. Lipman et al. "Flow Matching for Generative Modeling". In: *International Conference on Learning Representations*. 2023.
- [6] P. Fearnhead et al. *Scalable Monte Carlo for Bayesian Learning*. Institute of Mathematical Statistics Monographs. Cambridge University Press, 2025. ISBN: 9781009288446. doi: [10.1017/9781009288460](https://doi.org/10.1017/9781009288460).

Complex systems theory

What is complex systems theory?

Complex systems theory examines systems composed of many interconnected components, where the interactions between these components give rise to behaviours and properties that are not easily deducible from the individual parts [1]. AI systems often exhibit intricate dynamics that are essential for understanding and interpreting their outputs. As AI continues to advance, recognising and modelling this complexity becomes increasingly important if we are to understand how these systems function and perform.

Established applications

The principles of complex systems theory have found applications in numerous fields in which interactions and connectivity between components are crucial. For example:

Biological systems. The study of ecosystems highlights how species interact and adapt, providing insights into biodiversity and ecological stability [e.g., 2, 3, 4].

Social networks. Analysis of social systems helps us understand how information spreads and how individual behaviours influence group dynamics.

Economic models. Complex systems theory informs the study of market behaviours, helping to clarify how economic agents interact and respond to various stimuli [e.g., 5, 6].

These applications underscore the theory's significance in analysing dynamic systems and uncovering emergent behaviours that cannot be understood through reductionist approaches.

Complex systems theory meets AI

While machine learning has shown great promise, much recent progress has come from scaling models to unprecedented sizes, resulting in very large numbers of tuneable parameters during training. This is facilitated by vast amounts of available data. However, because of the sheer complexity and opaqueness of these models, we still do not fully understand what they have learned or how that knowledge is transformed into outputs. Despite substantial efforts to interpret complex black-box machine learning models [e.g., 7, 8], our understanding remains only partial.

To reconcile the need for transparency with the advantages of machine learning, one possibility is to combine mechanistic modelling with machine learning [e.g., 9, 10]. Mechanistic modelling provides an interpretable framework for representing the underlying processes that generate the data and shape predictions. Machine learning, by contrast, can uncover complex dependencies in large data sets without requiring a complete prior mechanistic specification.

The core idea of combining mechanistic modelling with machine learning lies in using available data to learn within the space of mechanistic models. This approach allows mechanistic models to be refined using knowledge extracted from the data, while

maintaining overall system transparency [e.g., 11, 12]. However, the mathematics underpinning such a fusion poses significant challenges, especially when mechanistic models live in structured spaces with non-standard geometry, constraints, or notions of similarity.

Looking ahead

Going forward, developing a coherent mathematical framework to describe and analyse the fusion of mechanistic modelling and machine learning is crucial. Such a framework would improve understanding and help provide mathematical guarantees about the reliability and performance of these models in deployment. By addressing the complexities inherent in these mechanistic model spaces—such as their geometry and the non-standard metrics governing similarity and distance—we can strengthen efforts to make machine learning systems safer, more transparent, and more interpretable. Ultimately, this closer integration of methodologies will help move AI towards systems that do not merely fit data, but also reflect a deeper understanding of the structural dynamics of the problems they are used to solve.

References

- [1] J. Ladyman, J. Lambert, and K. Wiesner. “What Is a Complex System?” In: *European Journal for Philosophy of Science* 3.1 (2013), pp. 33–67. doi: [10.1007/s13194-012-0056-8](#).
- [2] J. Bascompte et al. “The Nested Assembly of Plant–Animal Mutualistic Networks”. In: *Proceedings of the National Academy of Sciences* 100.16 (2003), pp. 9383–9387. doi: [10.1073/pnas.1633576100](#).
- [3] S. Saavedra, F. Reed-Tsochas, and B. Uzzi. “A Simple Model of Bipartite Cooperation for Ecological and Organizational Networks”. In: *Nature* 457.7228 (2009), pp. 463–466. doi: [10.1038/nature07532](#).
- [4] J. J. Lever et al. “The Sudden Collapse of Pollinator Communities”. In: *Ecology Letters* 17.3 (2014), pp. 350–359. doi: [10.1111/ele.12236](#).
- [5] G. Orlando and G. Zimatore. “Recurrence Quantification Analysis of Business Cycles”. In: *Chaos, Solitons & Fractals* 110 (2018), pp. 82–94. doi: [10.1016/j.chaos.2018.02.032](#).
- [6] S. Battiston et al. “The Price of Complexity in Financial Networks”. In: *Proceedings of the National Academy of Sciences* 113.36 (2016), pp. 10031–10036. doi: [10.1073/pnas.1521573113](#).
- [7] V. Hassija et al. “Interpreting Black-Box Models: A Review on Explainable Artificial Intelligence”. In: *Cognitive Computation* 16.1 (2024), pp. 45–74. doi: [10.1007/s12559-023-10179-8](#).
- [8] Y. Zhang et al. “A Survey on Neural Network Interpretability”. In: *IEEE Transactions on Emerging Topics in Computational Intelligence* 5.5 (2021), pp. 726–742. doi: [10.1109/TETCI.2021.3100641](#).
- [9] W. Liu et al. “Structured Kolmogorov-Arnold Neural ODEs for Interpretable Learning and Symbolic Discovery of Nonlinear Dynamics”. In: (2025). doi: [10.48550/arXiv.2506.18339](#). arXiv: [2506.18339 \[cs.LG\]](#).
- [10] N. Schmid et al. “Assessment of Uncertainty Quantification in Universal Differential Equations”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 383.2293 (2025), pp. 1–14. doi: [10.1098/rsta.2024.0444](#).
- [11] Y. Shen et al. “Spatial-Temporal Modelling of fMRI Data Through Spatially Regularized Mixture of Hidden Process Models”. In: *NeuroImage* 84.1 (2014), pp. 657–671. doi: [10.1016/j.neuroimage.2013.09.003](#).
- [12] Y. Shen, P. Tiño, and K. Tsaneva-Atanasova. “Classification Framework for Partially Observed Dynamical Systems”. In: *Physical Review E* 95.4 (2017), p. 043303. doi: [10.1103/PhysRevE.95.043303](#).

Complexity theory

What is complexity theory?

Complexity theory studies the resources—especially time and space (memory)—needed to solve computational problems. Time is often the central concern; time complexity is measured by the maximum number of steps required by an algorithm as the size of the input grows. The scaling of that number of steps with input size is important, and two key cases are commonly distinguished. For input of size n , an algorithm has polynomial time complexity if the number of steps required grows as n^k , for a constant k ; such problems are generally regarded as solvable. By contrast, exponential-time algorithms, whose running time grows like k^n for some constant k , are usually considered unsolvable for large problem instances.

One of the most important complexity classes of decision problems is called NP-complete. This is the class of decision problems for which no polynomial-time solution is known, but for which a proposed solution can be checked efficiently.

Established applications

The existence of NP-complete and other intractable problems defines a fundamental constraint: no increase in hardware or use of heuristic pattern recognition removes the underlying complexity barrier. Examples arise in everyday settings, such as routing problems in logistics; in the *travelling salesman problem*, a proposed route is easy to check, but finding an optimal route requires computation

time that grows exponentially with the size of the input. This is not simply a matter of needing better software or faster hardware; it reflects a fundamental mathematical wall.

Complexity theory meets AI

In the context of modern AI, it is tempting to assume that enough computational power and enough data can eventually overcome any problem by “brute force”. But AI systems are bound by the same laws of computational complexity as any other computational method. Even the most advanced AI models often operate as sophisticated heuristic systems rather than exact solvers. They may excel at finding useful approximate (“good enough”) solutions or patterns, but that is not the same as overcoming the underlying hardness of NP-complete problems. As AI becomes increasingly integrated into critical applications—from logistics and drug discovery to cybersecurity—it becomes ever more important to understand when probabilistic “guesses” are adequate, and when stronger guarantees are required.

Looking ahead

Mathematics generally, and complexity theory specifically, helps identify the structure in problems that determines whether they are tractable or intractable. When we understand the underlying structure of a specific NP-complete instance, we can sometimes transform it into a tractable problem with polynomial time complexity. Continued advances in complexity theory will be essential for understanding where AI can provide safe and practically useful approximations, where exact solutions remain out of reach, and where problem structure can be exploited to make progress.

Control theory

What is control theory?

Control theory is the mathematical discipline which focuses on the analysis and design of systems that make decisions over time. Control theory emerged from engineering challenges, but its fundamental questions are precisely those faced in AI: how to act optimally in uncertain, dynamic environments; how to balance exploration and exploitation; how to guarantee stability and robustness; and how to reason over long time horizons.

Control theory provides the language and tools for sequential decision-making under uncertainty. Concepts such as feedback, stability, optimality, robustness, and observability are central both to engineered systems and to AI agents.

Established applications

Control theory has a distinguished history throughout engineering applications. It is deployed extensively in flight control systems to stabilise aircraft and provide autopilot capabilities, in automotive control such as anti-lock braking and engine management, and in electrical power systems to manage load on the grid. A classic and common application is industrial process control, in which control-theoretic methods are used to monitor and control chemical reactions, temperatures, pressures, and flow rates to ensure efficient operation. Less obvious applications arise in telecommunications, where bandwidth management and communication protocols use control theory to ensure quality of service, and in medical systems, where medical devices use control theory to maintain desired physiological parameters in patients. Control theory is even used in the monitoring of medical unit failure rates.

Control theory meets AI

A traditional way of organising the fields of AI and machine learning is to distinguish between three core problem classes. *Supervised learning* asks: how can we learn a relationship from examples? *Unsupervised learning* asks: how can we uncover the intrinsic structure of data? *Reinforcement learning* asks: how can we learn to make good decisions through interaction with an environment, when outcomes are uncertain and experimentation is required?

As models grow more capable and obtaining new data becomes more costly, reinforcement learning has become increasingly important for building reasoning, planning, and adaptive decision-making abilities. At its core, reinforcement learning is not an isolated invention of modern AI; it can be understood as a modern, data-driven, and computationally scalable extension of classical control theory, adapted to high-dimensional and unknown environments. Without the mathematical foundations, in particular the combination of control theory and modern probability theory, we would not know how to develop solutions to large-scale problems, especially where reliability is critical.

Several challenges in contemporary AI are, at their heart, control-theoretic problems:

Memory and long-term dependencies. Processing long texts, maintaining conversational coherence, or planning over extended horizons requires models that can incorporate and manage memory. In control theory, this relates to state representation, partial observability, and filtering, which are tools to describe internal representations of a system over time.

Stability and safety. As AI systems are deployed in real-world settings, ensuring stable and predictable behaviour becomes critical. Control theory has developed powerful tools to analyse stability, robustness to perturbations, and resilience to noise—precisely the issues that arise in AI safety and privacy guarantees.

Randomness and uncertainty. Randomness can be both a benefit and a risk in AI systems. Exploration strategies and randomised training methods drive learning, yet uncertainty must be carefully controlled to avoid catastrophic outcomes. Control theory offers a rigorous framework within which to reason about uncertainty, risk-sensitive objectives, and worst-case guarantees.

Scalability and effective use of computation. Modern AI systems rely on enormous computational resources. Control theory helps clarify how these resources can be used effectively: how to design algorithms that converge reliably, how to exploit structure in high-dimensional systems, and how to balance model complexity against stability and performance.

Looking ahead

As AI systems become more autonomous, adaptive, and embedded in society, the need for strong mathematical foundations only intensifies. Empirical progress alone is not sufficient: without theory, we lack guarantees, interpretability, and principled methods for improvement. Control theory serves as a bridge between abstract mathematics and practical AI engineering, translating theoretical insights into tools for designing reliable intelligent systems.

Continued progress in AI therefore depends not only on larger models and more data, but also on advances in the underlying mathematics of decision-making.

Dynamical systems theory

What is dynamical systems theory?

Dynamical systems theory studies how systems evolve over time. The ancients were interested in the motion of the sun, moon, and stars, in the phases of the tides, and in the changing of the seasons. A more systematic understanding emerged in the Age of Enlightenment through the work of Kepler, Newton, and many others. This understanding came from formulating laws of physics as dynamical equations that can, at least in favourable cases, be analysed or solved to make predictions [1]. These ideas can be generalised to a broad range of systems, with the common idea of a state—a snapshot of the system at a given moment—and a rule that determines how that state evolves.

Established applications

In the modern era, dynamical equations underpin computer simulations of increasingly complex systems, such as the motion of a Formula 1 car or of a passenger aircraft. Enormous computer simulations are also used to predict future global climate under different emissions and land-use scenarios.

Mathematical analysis has uncovered many important aspects of these systems. One important example is the sensitive dependence on initial conditions (the *butterfly effect*) associated with chaos. In 1963, Ed Lorenz [2] showed that even a ridiculously simplistic model of the atmosphere exhibits sensitive dependence on its initial conditions. To predict the outcome a year, say, into the future would require us to know the current state of the system to more digits of accuracy than there are atoms in the universe. Making a tiny error in the gazillionth digit behind the decimal point can make a difference between sunshine or snow in a year's time.

This property of chaos theory put an end to the idea of a clockwork universe that can be determined simply from knowing the current state and the laws of evolution. If that were the case, then we would all be able to make our fortunes predicting the stock market with perfect accuracy. But we can't. In many important systems, long-term prediction is fundamentally limited.

A further issue is that many dynamical systems of interest are safety-critical. Algorithms used in fly-by-wire aircraft, devices providing medical diagnoses or therapies, and other safety-critical systems rely on feedback control algorithms that can respond to input data in a certifiably correct way. Such certifiability depends on mathematical analysis of the qualitative behaviour of the dynamics and its uncertainty.

More generally, dynamical systems theory helps identify the qualitatively different behaviours a system can exhibit, and in combination with probability theory this helps us reason about the likelihood of different outcomes.

Dynamical systems theory meets AI

There are clear opportunities to apply AI tools to the study of dynamical systems; one recent example concerns smart grids for electricity delivery [3]. There is also substantial international activity aimed at learning dynamical laws directly from data using AI tools [4].

Dynamical systems theory also provides a powerful framework for understanding how some AI models behave, enabling us to derive new methods and to optimise them. For example, recurrent neural networks (RNNs) are AI tools that are used to make predictions of how key signals evolve over time. They work by having hidden states that evolve as discrete-time dynamical systems; aspects of

their behaviour such as their ability to learn can be understood using core ideas from dynamical systems theory such as attractors, sensitivity and transients.

Deep learning models that are used in predictive AI algorithms like ChatGPT can be understood in terms of the dynamics of large networks of switches. Each switch behaves predictably, but because of the huge number of layers of switches, small differences in inputs or prompts can in some settings lead to materially different trajectories through the model, making sensitivity and stability important analytical questions.

The training process of deep networks can also be viewed through a dynamical-systems lens. So-called gradient descent methods for finding the strengths of each connection between the switches can be thought of as the dynamics of rolling a ball downhill over a complicated (multi-dimensional) landscape. Mathematical methods help us to analyse the geometry of this landscape, and to optimally explore the different valleys and plateaus, while trying to reach the optimal set of strengths for each training input.

When AI systems act in the physical world—e.g., in robots, drones and autonomous vehicles—they must respect the dynamics of their environment. AI algorithms such as reinforcement learning are used to reward desired outcomes and penalise failure. In such settings, reinforcement learning must often be combined with ideas from optimisation and control theory in order to meet safety and stability constraints.

Looking ahead

Machine-learning methods are increasingly capable not only of making predictions, but also of helping to infer governing principles and to construct reduced-order dynamical models for complex phenomena directly from observations. This could have applications in fields such as climate science, high-performance engineering, and systems biology. Such models may in turn be used within reinforcement-learning or control frameworks to support more complex human-machine systems, including advanced medical treatments, active materials, and energy management.

There is also growing interest in dynamically inspired architectures for AI algorithms, using ideas such as periodic or chaotic attractors, continuous-time neurons, and so-called reservoir computers. These designs may offer new forms of robustness, efficiency, or reproducibility.

As AI systems become more autonomous, the need for formal safety and stability guarantees will grow. Dynamical systems theory—especially so-called Lyapunov methods and other tools from control theory—may provide just the right mathematical foundations for future trustworthy AI.

It is also a dream of many practitioners to be able to turn *black-box* AI into so-called *grey-box* systems. That is, even if we do not have a fully detailed account of how an AI model produces a given answer, we may still obtain a low-order understanding of what it is likely to do—and, crucially, what it won't do; one promising route is the idea of learning reduced-order dynamical-systems models of the black box.

References

- [1] F. Diacu and P. Holmes. *Celestial Encounters: The Origins of Chaos and Stability*. Princeton University Press, 1999. ISBN: 9780691005454.
- [2] E. N. Lorenz. “Deterministic Nonperiodic Flow”. In: *Journal of the Atmospheric Sciences* 20.2 (1963), pp. 130–141. doi: [10.1175/1520-0469\(1963\)020<0130:DNF>2.0.CO;2](https://doi.org/10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2).
- [3] M. Ahmadi, H. Aly, and J. Gu. “A Comprehensive Review of AI-Driven Approaches for Smart Grid Stability and Reliability”. In: *Renewable and Sustainable Energy Reviews* 226 (2026), p. 116424. doi: [10.1016/j.rser.2025.116424](https://doi.org/10.1016/j.rser.2025.116424).
- [4] S. L. Brunton and J. N. Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. 2nd ed. Cambridge University Press, 2022. ISBN: 9781009098489. doi: [10.1017/9781009098517](https://doi.org/10.1017/9781009098517).

Experimental design

What is experimental design?

Experimental design is the mathematical science of active data collection. It addresses a basic question: if we want to understand how a system behaves, what data should we collect, and under what interventions or conditions? Core principles—randomisation, control, and sequential or adaptive design—help distinguish genuine causal effects from noise, bias, and coincidence.

Established applications

The principles of experimental design are central in settings where experimental costs are high (e.g., manufacturing), where statistical rigour is a regulatory requirement (e.g., clinical trials), or where causal reasoning underpins decision-making (e.g., A/B testing in technology). When decisions involve safety, fairness, cost, or policy, understanding associations alone is insufficient; we need to understand the impact of interventions. Experimental design provides the most credible basis for causal inference.

Experimental design meets AI

Experimental design supplies the mathematical framework needed to test whether an AI system genuinely improves outcomes and whether apparent gains survive deployment. This distinction between prediction and causation is widely recognised as a central limitation of current AI methods [1, 2, 3].

Empirical evidence shows that many state-of-the-art models become overconfident or unreliable when conditions change [4, 5], and such failures often cannot be detected without targeted experimental evaluation. Benchmark accuracy alone is insufficient to establish reliability under distributional shift, rare events, or feedback effects induced by deployment. Carefully designed experiments determine which data reveal these weaknesses, and how to learn from them efficiently. This matters not only for *post hoc* testing, but for designing data-collection and evaluation regimes that expose problems before systems are deployed at scale. This is particularly important in high-stakes domains such as healthcare, finance, transport, and energy, where errors can propagate rapidly once systems are deployed at scale. These issues have become more salient as regulators increasingly demand explanations that go beyond correlations, particularly for automated decision systems that affect individuals and communities [6].

Bias and unfairness in AI systems likewise cannot be resolved simply by using larger datasets or more complex models. Poorly designed data collection and evaluation can entrench discrimination even when headline fairness metrics improve [7]. Experimental design provides the tools for rigorous auditing of AI systems and for evaluating interventions intended to improve fairness. Such methods are increasingly central to policy-relevant evaluation of algorithmic systems and are reflected in emerging AI governance frameworks [8].

Looking ahead

Many of the hardest problems in AI, such as generalisation, fairness, and accountability, are now constrained less by modelling capacity than by the absence of rigorous experimental design as standard practice. Experimental design already plays a central role in modern AI, but there remains substantial scope for wider and more systematic use of its classical principles. Work on active learning, Bayesian optimisation, and causal discovery formulates the choice of data, interventions, and experimental allocation as optimisation problems under budget and uncertainty constraints. These methods extend classical ideas from sequential and optimal design to the high-dimensional models and complex data-generating processes used in modern AI systems [9, 10, 11, 12].

References

- [1] J. Pearl. *Causality: Models, Reasoning, and Inference*. 2nd ed. Cambridge University Press, 2009.
- [2] L. Bottou et al. “Counterfactual Reasoning and Learning Systems: The Example of Computational Advertising”. In: *Journal of Machine Learning Research* 14.1 (2013), pp. 3207–3260.
- [3] P. Tigas et al. “Interventions, Where and How? Experimental Design for Causal Models at Scale”. In: *Advances in Neural Information Processing Systems* 35. Ed. by S. Koyejo et al. Curran Associates, Inc., 2022, pp. 24130–24143.
- [4] B. Lakshminarayanan, A. Pritzel, and C. Blundell. “Simple and Scalable Predictive Uncertainty Estimation Using Deep Ensembles”. In: *Advances in Neural Information Processing Systems* 30. 2017.
- [5] Y. Ovadia et al. “Can You Trust Your Model’s Uncertainty? Evaluating Predictive Uncertainty Under Dataset Shift”. In: *Advances in Neural Information Processing Systems* 32. 2019.
- [6] High-Level Expert Group on Artificial Intelligence. *Ethics Guidelines for Trustworthy AI*. Tech. rep. European Commission, 2019.
- [7] S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press, 2023.
- [8] National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. Tech. rep. National Institute of Standards and Technology, 2023. doi: [10.6028/NIST.AI.100-1](https://doi.org/10.6028/NIST.AI.100-1).
- [9] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. 2nd ed. MIT Press, 2018.
- [10] B. Settles. *Active Learning*. Springer, 2012. doi: [10.1007/978-3-031-01560-1](https://doi.org/10.1007/978-3-031-01560-1).
- [11] B. Shahriari et al. “Taking the Human Out of the Loop: A Review of Bayesian Optimization”. In: *Proceedings of the IEEE* 104.1 (2016), pp. 148–175. doi: [10.1109/JPROC.2015.2494218](https://doi.org/10.1109/JPROC.2015.2494218).
- [12] T. Rainforth et al. “Modern Bayesian Experimental Design”. In: *Statistical Science* 39.1 (2024), pp. 100–114. doi: [10.1214/23-STS915](https://doi.org/10.1214/23-STS915).

Fluid dynamics

What is fluid dynamics?

Classical mathematical fluid dynamics studies the motion of fluids (liquids and gases), treating the fluid as a continuum and analysing quantities such as velocity and pressure using vector calculus under a range of physical assumptions. The basic equations governing fluid motion have been studied by mathematicians ever since Claude-Louis Navier in Paris and George Stokes in Cambridge formulated them about 200 years ago. Motivated by their extraordinarily broad range of applications, the search for exact and approximate solutions of the Navier–Stokes equations became one of the dominant activities of applied mathematicians in the UK in the 20th century.

Established applications

Computational simulations of mathematical models of fluid flow are essential in applications ranging from weather prediction to assessment of nuclear reactor safety. Over the past 50 years, advances in simulation capability have transformed the design of fuel-efficient aircraft; indeed, computer-based aerodynamic modelling is now integral to the design of all modern vehicles. Sail configurations on racing yachts can now be optimised in real time, providing the marginal gains needed to win races in the America's Cup. Optimised aerodynamics helps modern cyclists ride faster, golf balls travel further, and elite swimmers reduce drag. Computational fluid dynamics enables the modelling of the flow of blood in the human heart, making the provision of patient-specific surgery possible.

One important success of conventional modelling approaches is the prediction of fluid-flow bifurcations. Laboratory experiments of fluids that are heated from below can easily be designed to show that a tiny change in the ambient temperature can result in a completely different flow profile. Vortices seemingly created by magic in air flowing over a flat terrain can be predicted by studying solutions to the Navier–Stokes equations.

Fluid dynamics meets AI

Driven by the AI revolution, the use of high-performance computers to simulate fluid flow has exploded over the past decade. A central ambition is that fluid flows may be predicted more efficiently by learning from data than by relying solely on conventional approaches based on generating approximate solutions to the Navier–Stokes equations.

Weather forecasting provides a good example of combining fluid-dynamical modelling with AI: rather than numerically integrating the governing equations on a fine grid, an AI model can learn an efficient approximation to the time-evolution operator governing atmospheric dynamics. The training data comes from decades of numerical weather-prediction outputs that have already been corrected through data assimilation, meaning they encode both the governing equations and the statistical structure of real observations. In effect, the AI model learns a reduced, data-driven representation of the underlying flow map, one that preserves the large-scale dynamical behaviour while bypassing the computational cost of solving the full system of governing equations.

A key issue still to be faced, however, is that pure AI approaches to fluid-flow simulation do not necessarily respect the underlying physics. Such inconsistencies are evident in computer games, where it is often obvious to the naked eye that the motion of the fluid is unrealistic. The risk is that naïve AI predictions of flows generated solely from data may violate physical constraints or exhibit impossible scenarios. This is especially concerning in the prediction of extreme or safety-critical events, such as tornadoes or tsunamis.

Physics-informed neural networks (PINNs) have been a prominent development at the interface of AI and scientific computing over the past decade. This illustrates a recurring pattern in scientific AI: the greatest gains often come not from replacing mathematical models, but from combining them with data-driven approximations in carefully structured ways. When applied to fluid flows, these methods incorporate the Navier–Stokes equations directly into the training objective by enforcing them at a suitably dense grid of points in space and time. In realistic fluid settings, however, these methods have often struggled to compete with conventional approaches: accuracy, convergence, and stability can all be sensitive to the choice of training strategy and to the selection of spatial and temporal collocation points.

Looking ahead

A promising strategy is to combine physics-based modelling with observational or imaging data. A variety of such combination strategies are being actively developed. Such hybrid strategies may be perfectly suited to biomedical and biological problems where fluid flows cannot be measured or imaged using conventional technology. This marriage of established mathematical modelling and new AI methods holds promise for the efficient computation of physically realistic fluid flows across a wide range of practical settings. The continued development of refined mathematical models of complex fluids is likely to be critically important if the promise of AI in this area is to be realised.

References

- [1] D. Detomi, N. Parolini, and A. Quarteroni. “Mathematics in the Wind”. In: *Monografías de la Real Academia de Ciencias de Zaragoza* 31 (2009), pp. 35–56.
- [2] C. Wood. *Using AI, Mathematicians Find Hidden Glitches in Fluid Equations*. Quanta Magazine. Jan. 2026. URL: <https://www.quantamagazine.org/using-ai-mathematicians-find-hidden-glitches-in-fluid-equations-20260109/>.
- [3] K. A. Cliffe, A. Spence, and S. J. Tavener. “The Numerical Analysis of Bifurcation Problems with Application to Fluid Mechanics”. In: *Acta Numerica* 9 (2000), pp. 39–131. doi: [10.1017/S0962492900000018](https://doi.org/10.1017/S0962492900000018).
- [4] Z. Hao et al. “PINNacle: A Comprehensive Benchmark of Physics-Informed Neural Networks for Solving PDEs”. In: (2023). doi: [10.48550/arXiv.2306.08827](https://doi.org/10.48550/arXiv.2306.08827). arXiv: [2306.08827](https://arxiv.org/abs/2306.08827) [cs.LG].
- [5] M. Raissi, A. Yazdani, and G. E. Karniadakis. “Hidden Fluid Mechanics: Learning Velocity and Pressure Fields from Flow Visualizations”. In: *Science* 367.6481 (2020), pp. 1026–1030. doi: [10.1126/science.aaw4741](https://doi.org/10.1126/science.aaw4741).

Graph theory

What is graph theory?

Graph theory is the mathematics of connection. A graph is a mathematical abstraction of a network where dots, called vertices, are joined by lines called edges. Graphs can capture a wide variety of real-world systems; for example, they can represent a social network of people connected by friendships, a rail network of stations connected by rail lines, or a family tree. The edges may have a direction, for example in the case of a family tree where the direction indicates that one vertex is a parent and the other their child. Graph theorists study many different aspects of these graphs, including, for example, what causes different kinds of graphs to have important properties, or how to solve optimisation problems on them.

Established applications

Graph theory has contributed to a wide variety of real-world applications. Graphs have helped us model and understand biological systems in general [1, 2], and evolution in particular [3, 4]. Modelling the brain as a graph is an important tool in neuroscience [5] for understanding brains and developing interventions. Molecules and crystals can be modelled as graphs [6, 7], which helps us simulate and infer their structure and interactions. Modelling infectious disease continues to be a critical component of UK health security: accurate models that incorporate real contact patterns require graph theory [8]. Graph theory also contributes to information security via cryptography [9, 10]. The UK is a world leader in quantum computing, which presents a vast opportunity for the future of what it is possible to compute. Graph theory is foundational to both quantum algorithms and the technology of quantum computing itself [11, 12].

Graph theory meets AI

Graph theory is critical to modern AI: from efficiency to capability and explainability. For example, fundamentally, neural networks are graphs and understanding them and their behaviour requires graph-theoretic ideas. And considering data as graphs (e.g., as a graph database or via topological data analysis [13]) allows AI methods to get the best value from the data. Graph traversal methods and algorithms continue to be keystone building blocks for AI [14, 15]. Recent research shows that large language models—a form of generative AI that is seeing extremely broad application—can be improved in quality and efficiency by modelling parts of them as graphs [16].

Looking ahead

As AI methods make an ever-larger contribution to decision-making, trustworthiness, explainability, and understanding of bias become critical. Analysis using graph theory allows formal reasoning about the AI methods [17, 18], which can help us make guarantees about how they behave and why. In turn, AI has significant potential to accelerate graph theoretic discovery. AI tools are now assisting in mathematical investigation via supporting reasoning as well as simulation. Graph theory has particularly strong representation in formalised mathematical language libraries such as LEAN, which makes it especially suitable for application of AI methods to proving novel results.

References

- [1] C. Kadelka et al. “Modularity of Biological Systems: A Link between Structure and Function”. In: *Journal of the Royal Society Interface* 20.207 (2023), p. 20230505. doi: [10.1098/rsif.2023.0505](https://doi.org/10.1098/rsif.2023.0505).
- [2] A.-L. Barabási and Z. N. Oltvai. “Network Biology: Understanding the Cell’s Functional Organization”. In: *Nature Reviews Genetics* 5.2 (2004), pp. 101–113. doi: [10.1038/nrg1272](https://doi.org/10.1038/nrg1272).
- [3] E. Lieberman, C. Hauert, and M. A. Nowak. “Evolutionary Dynamics on Graphs”. In: *Nature* 433.7023 (2005), pp. 312–316. doi: [10.1038/nature03204](https://doi.org/10.1038/nature03204).
- [4] H. Ohtsuki et al. “A Simple Rule for the Evolution of Cooperation on Graphs and Social Networks”. In: *Nature* 441.7092 (2006), pp. 502–505. doi: [10.1038/nature04605](https://doi.org/10.1038/nature04605).
- [5] J. C. Pang et al. “Geometric Constraints on Human Brain Function”. In: *Nature* 618.7965 (2023), pp. 566–574. doi: [10.1038/s41586-023-06098-1](https://doi.org/10.1038/s41586-023-06098-1).
- [6] O. Delgado Friedrichs et al. “Systematic Enumeration of Crystalline Networks”. In: *Nature* 400.6745 (1999), pp. 644–647. doi: [10.1038/23210](https://doi.org/10.1038/23210).
- [7] A. T. Balaban. “Applications of Graph Theory in Chemistry”. In: *Journal of Chemical Information and Computer Sciences* 25.3 (1985), pp. 334–343. doi: [10.1021/ci00047a033](https://doi.org/10.1021/ci00047a033).
- [8] S. Eubank et al. “Modelling Disease Outbreaks in Realistic Urban Social Networks”. In: *Nature* 429.6988 (2004), pp. 180–184. doi: [10.1038/nature02541](https://doi.org/10.1038/nature02541).
- [9] P. Alikhani et al. “Experimental Relativistic Zero-Knowledge Proofs”. In: *Nature* 599.7883 (2021), pp. 47–50. doi: [10.1038/s41586-021-03998-y](https://doi.org/10.1038/s41586-021-03998-y).
- [10] M. Ben-Or et al. “Multi-Prover Interactive Proofs: How to Remove Intractability Assumptions”. In: *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*. Association for Computing Machinery, 1988, pp. 113–131. doi: [10.1145/62212.62223](https://doi.org/10.1145/62212.62223).

- [11] M. Howard et al. “Contextuality Supplies the ‘Magic’ for Quantum Computation”. In: *Nature* 510.7505 (2014), pp. 351–355. doi: [10.1038/nature13460](#).
- [12] X.-C. Yao et al. “Experimental Demonstration of Topological Error Correction”. In: *Nature* 482.7386 (2012), pp. 489–494. doi: [10.1038/nature10770](#).
- [13] N. Otter et al. “A Roadmap for the Computation of Persistent Homology”. In: *EPJ Data Science* 6.1 (2017), p. 17. doi: [10.1140/epjds/s13688-017-0109-5](#).
- [14] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. 4th ed. Pearson, 2021.
- [15] J. Schaeffer. “The Games Computers (and People) Play”. In: *Advances in Computers*. Vol. 52. Elsevier, 2000, pp. 189–266. doi: [10.1016/S0065-2458\(00\)80019-4](#).
- [16] M. Besta et al. “Graph of Thoughts: Solving Elaborate Problems with Large Language Models”. In: *Proceedings of the AAAI Conference on Artificial Intelligence* 38.16 (2024), pp. 17682–17690. doi: [10.1609/AAAI.V38I16.29720](#).
- [17] F. Lecue. “On the Role of Knowledge Graphs in Explainable AI”. In: *Semantic Web* 11.1 (2020), pp. 41–51. doi: [10.3233/SW-190374](#).
- [18] P. E. Pope et al. “Explainability Methods for Graph Convolutional Neural Networks”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019, pp. 10772–10781.

Information theory

What is information theory?

Information theory [1, 2] is the branch of mathematics that helps us understand the foundations of information processing tasks such as quantifying, encoding, and transmitting information. On the one hand, it provides principled mathematical methods to obtain fundamental limits of such problems (especially in the presence of noise and uncertainty). On the other hand, the theory provides insights that help us devise algorithms that approach these limits.

Established applications

Your mobile phone, your computer, internet communications, digital music and video, and even QR codes: all these rely on concepts from information theory to understand, on the one hand, the absolute limits beyond which performance for metrics such as storage/communication speed/data fidelity is impossible, and on the other hand, squeeze optimal performance out of these systems [3, 4, 5, 6]. This can be true even in fields with no obvious connection to information technology, such as evolutionary biology, linguistics, neuroscience, and economics [7, 8, 9, 10]. Often, understanding the flow of information through complex systems can provide an insightful lens into the behaviour of these systems.

Information theory meets AI

In a world in which AI plays an ever more important role, information theory presents an invaluable toolkit. Current AI systems, with their complex multi-layer and non-linear structure are notoriously hard to understand and optimise for performance. Often, they behave essentially like black boxes, resulting in a whole host of issues. For example, it is hard to gain insight into what architectures and operations might lead to better performance, or indeed what the fundamental limits of performance might be with a given amount of training data or computational power. Information theory can help us to understand the information-processing limits of systems, and hence whether improved performance is possible, and how to achieve those improvements.

Having explainability in AI is also important when these systems are used to make decisions impacting people’s lives—for instance when used within criminal justice. Also, data used during training can be inadvertently “memorised” by AIs, leading to intellectual property issues. Information theory provides tools to explore such issues [11].

Looking ahead

While information theory does not offer a silver bullet, it can nonetheless provide an intellectual framework to help us think about these issues and understand trade-offs. As systems grow more complex, information theory provides a way to measure what they “know” and how they “know” it; to understand why some models generalise well and others saturate/overfit; and to reason about noise tolerance and how to improve it [12, 13, 14].

References

- [1] C. E. Shannon. “A Mathematical Theory of Communication”. In: *Bell System Technical Journal* 27.3–4 (1948), pp. 379–423, 623–656.
- [2] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, 1991. ISBN: 9780471062592.
- [3] A. Goldsmith. *Wireless Communications*. Cambridge University Press, 2005. ISBN: 9780521837163. doi: [10.1017/CB09780511841224](#).

- [4] I. S. Reed and G. Solomon. "Polynomial Codes over Certain Finite Fields". In: *Journal of the Society for Industrial and Applied Mathematics* 8.2 (1960), pp. 300–304. doi: [10.1137/0108018](https://doi.org/10.1137/0108018).
- [5] S. B. Wicker and V. K. Bhargava, eds. *Reed-Solomon Codes and Their Applications*. Wiley-IEEE Press, 1999. ISBN: 9780780353916.
- [6] A. S. Downs, N. P. Sigmon, and R. E. Klima. "The Mathematics of QR Codes". In: *Proceedings of the Asian Technology Conference in Mathematics*. 2021, pp. 145–159.
- [7] C. Adami. "The Use of Information Theory in Evolutionary Biology". In: *Annals of the New York Academy of Sciences* 1256.1 (2012), pp. 49–65. doi: [10.1111/j.1749-6632.2011.06422.x](https://doi.org/10.1111/j.1749-6632.2011.06422.x).
- [8] F. Rieke et al. *Spikes: Exploring the Neural Code*. MIT Press, 1999. ISBN: 9780262681087.
- [9] C. E. Shannon. "Prediction and Entropy of Printed English". In: *Bell System Technical Journal* 30.1 (1951), pp. 50–64. doi: [10.1002/j.1538-7305.1951.tb01366.x](https://doi.org/10.1002/j.1538-7305.1951.tb01366.x).
- [10] C. A. Sims. "Implications of Rational Inattention". In: *Journal of Monetary Economics* 50.3 (2003), pp. 665–690. doi: [10.1016/S0304-3932\(03\)00029-1](https://doi.org/10.1016/S0304-3932(03)00029-1).
- [11] S. Bombari, M. H. Amani, and M. Mondelli. "Memorization and Optimization in Deep Neural Networks with Minimum Over-Parameterization". In: *Advances in Neural Information Processing Systems* 35. Curran Associates, Inc., 2022, pp. 7628–7640.
- [12] A. Xu and M. Raginsky. "Information-Theoretic Analysis of Generalization Capability of Learning Algorithms". In: *Advances in Neural Information Processing Systems* 30. Ed. by I. Guyon et al. Curran Associates, Inc., 2017, pp. 2524–2533.
- [13] R. Shwartz-Ziv and N. Tishby. "Opening the Black Box of Deep Neural Networks via Information". In: (2017). arXiv: [1703.00810](https://arxiv.org/abs/1703.00810) [[cs.LG](#)].
- [14] S. Bombari and M. Mondelli. "Spurious Correlations in High Dimensional Regression: The Roles of Regularization, Simplicity Bias and Over-Parameterization". In: *Proceedings of the 42nd International Conference on Machine Learning*. Ed. by A. Singh et al. Vol. 267. Proceedings of Machine Learning Research. PMLR, 2025, pp. 4839–4873.

Inverse problems

What are inverse problems?

Imagine you hear a crash in the kitchen. You didn't see it happen, but by observing the shattered glass on the floor, you deduce that the cat knocked over a vase. Congratulations, you have just solved an inverse problem: working backward from an observed effect to identify the hidden cause. While mathematics often considers "forward" problems, like predicting where the vase will fall when the cat hits it with a certain force, the real world cares almost exclusively about the "inverse". Doctors do not want to simulate how X-rays travel through a healthy body; they want to take the X-ray data and reconstruct the hidden image. Inverse problems underpin medical imaging, remote sensing from satellite data, materials discovery, and astrophysics [1]. In all these fields, they turn out to be the only means of observing or testing otherwise inaccessible phenomena, bringing together physics, data, optimisation and uncertainty. The discipline consists of a rigorous toolkit—regularisation theory, stability analysis, variational methods, and associated optimisation techniques [2]—developed to overcome the inherent ill-posedness of the inverse problems and select solutions that are not only mathematically admissible, but also physically meaningful.

In the age of AI, one might assume these classical tools are becoming obsolete. The reality is the opposite: skills developed in inverse problems are critical for AI in at least two fundamental ways.

Inverse problems and AI

Leaving aside the classical examples, training a neural network is, at its core, a massive, high-dimensional inverse problem: given examples of inputs and outputs (training data), we seek model parameters that could have generated them. A reason modern AI models can be trained at all is the pervasive use of regularisation—a core concept born and bred in inverse problem theory. Regularisation, both explicit and implicit, is what prevents models from simply memorising training data. Regularisation techniques such as weight decay, or early stopping in iterative optimisation algorithms for training neural networks, have direct analogues in classical inverse problem solvers. A striking example is Maximal Update Parametrisation (μ P), now widely used to transfer hyperparameters from small proxy networks to billion-parameter models [3]. This analysis of how parameters should scale is, in essence, a regularisation and stability argument. It enabled reliable training of massive models and underpinned the empirical "scaling laws" that drive today's AI [4]. Understanding why learning converges, what kinds of solutions AI systems are biased towards, and why they hallucinate requires exactly the inverse-problems toolkit: stability, convergence, and sensitivity analysis for highly ill-posed inference tasks. Without this mathematical foundation, scaling up AI is like building skyscrapers without understanding the physics of steel.

Beyond the training of models, we are witnessing a rush to apply Deep Learning to solve real-world inverse problems: creating faster MRI scans, lower-radiation CTs, higher-resolution microscopy and satellite imaging, and more efficient scientific experimentation [5]. These benefits are real and already in use: for instance, AI-accelerated MRI reconstruction has reduced scan times by factors of 2–10 in clinical pilots [6, 7], and AI-based deblurring and denoising are now standard in commercial microscopes and telescopes [8, 9, 10]. At the same time, using deep learning for inverse problems carries significant risks if it is not grounded in proper mathematical modelling. There are now many published examples where AI reconstructions of MRI and CT images introduce or remove structures that are not present in the raw data, yet the results look entirely plausible [11]. Such systems can also fail unpredictably on rare or

out-of-distribution cases and encode hidden biases that distort scientific or clinical conclusions. Crucially, these failures are often silent: the images appear convincing, and nothing in the output signals that anything has gone wrong.

Here, too, mathematics provides the guardrails. Inverse problem theory has already delivered concrete tools to make AI-based reconstructions safer: data-consistency layers and physics-informed network architectures enforce reconstructions that actually explain the measured data, ruling out many hallucinations by design [12]. It allows us to ask and answer questions such as: How sensitive is this reconstruction to noise or model mismatch? Under what conditions can we guarantee that an AI-generated image reflects reality rather than artefacts or statistical hallucinations? Without such guarantees, AI-based scanners and sensors risk giving us beautiful pictures that are scientifically or clinically misleading, leading to incorrect scientific discoveries, misdiagnoses, unnecessary or inappropriate medical interventions, and overreliance on unreliable automated decision systems. A language model hallucinating a citation is annoying; an AI system hallucinating a tumour is dangerous.

Looking ahead

AI is moving rapidly from impressive demonstrations to critical infrastructure for science, healthcare, and industry. Its true power does not lie in chatbots that write poetry, but in tools that accelerate scientific discovery and support high-stakes decisions: from AI systems attacking unsolved mathematical problems [13, 14, 15] to massive investments by industry leaders into biology, physics, and beyond [16, 17, 18, 19]. To make this transition safely, we need more than larger models and more data: we need the deep mathematics of inverse problems to understand when AI-based reconstructions are trustworthy, to design methods that respect physics and uncertainty, and to develop verifiable guarantees of robustness, fairness, and reliability [5].

References

- [1] C.-B. Schönlieb and Z. Shumaylov. “Data-Driven Approaches to Inverse Problems”. In: *arXiv preprint arXiv:2506.11732* (2025). doi: [10.48550/arXiv.2506.11732](https://doi.org/10.48550/arXiv.2506.11732).
- [2] H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*. Vol. 375. Mathematics and Its Applications. Dordrecht: Kluwer Academic Publishers, 1996. ISBN: 9780792341574. doi: [10.1007/978-94-009-1740-8](https://doi.org/10.1007/978-94-009-1740-8).
- [3] G. Yang et al. “Tensor Programs V: Tuning Large Neural Networks via Zero-Shot Hyperparameter Transfer”. In: *Advances in Neural Information Processing Systems* 34. 2021.
- [4] J. Kaplan et al. “Scaling Laws for Neural Language Models”. In: *arXiv preprint arXiv:2001.08361* (2020). doi: [10.48550/arXiv.2001.08361](https://doi.org/10.48550/arXiv.2001.08361).
- [5] S. Arridge et al. “Solving Inverse Problems Using Data-Driven Models”. In: *Acta Numerica* 28 (2019), pp. 1–174. doi: [10.1017/S0962492919000059](https://doi.org/10.1017/S0962492919000059).
- [6] J. Zbontar et al. “fastMRI: An Open Dataset and Benchmarks for Accelerated MRI”. In: *arXiv preprint arXiv:1811.08839* (2018). doi: [10.48550/arXiv.1811.08839](https://doi.org/10.48550/arXiv.1811.08839).
- [7] H. Wang et al. “Application of Deep Learning Image Reconstruction in Low-Dose Chest CT Scan”. In: *The British Journal of Radiology* 95.1133 (2022), p. 20210380. doi: [10.1259/bjr.20210380](https://doi.org/10.1259/bjr.20210380).
- [8] Nikon Instruments Inc. *NIS-Elements: Microscope Imaging Software Suite*. Accessed 2026-04-21. n.d. URL: <https://www.microscope.healthcare.nikon.com/products/software/nis-elements>.
- [9] Carl Zeiss Microscopy GmbH. *AI Microscopy Solutions*. Accessed 2026-04-21. n.d. URL: <https://www.zeiss.com/microscopy/en/products/ai-microscopy-solutions.html>.
- [10] J. Sanchez-Bermudez et al. “Revealing Io’s Surface Using JWST-NIRISS Aperture Masking Interferometry and Neural Network Deconvolution”. In: *Monthly Notices of the Royal Astronomical Society* 543.1 (2025), pp. 608–624. doi: [10.1093/mnras/staf1414](https://doi.org/10.1093/mnras/staf1414).
- [11] A. Breger et al. “A Study of Why We Need to Reassess Full Reference Image Quality Assessment with Medical Images”. In: *Journal of Imaging Informatics in Medicine* 38.6 (2025), pp. 3444–3469. doi: [10.1007/s10278-025-01462-1](https://doi.org/10.1007/s10278-025-01462-1).
- [12] J. Schwab, S. Antholzer, and M. Haltmeier. “Deep Null Space Learning for Inverse Problems: Convergence Analysis and Rates”. In: *Inverse Problems* 35.2 (2019), p. 025008. doi: [10.1088/1361-6420/aaf14a](https://doi.org/10.1088/1361-6420/aaf14a).
- [13] S. Bubeck et al. “Early Science Acceleration Experiments with GPT-5”. In: *arXiv preprint arXiv:2511.16072* (2025). doi: [10.48550/arXiv.2511.16072](https://doi.org/10.48550/arXiv.2511.16072).
- [14] T. Feng et al. “Semi-Autonomous Mathematics Discovery with Gemini: A Case Study on the Erdős Problems”. In: *arXiv preprint arXiv:2601.22401* (2026). doi: [10.48550/arXiv.2601.22401](https://doi.org/10.48550/arXiv.2601.22401).
- [15] A. Novikov et al. “AlphaEvolve: A Coding Agent for Scientific and Algorithmic Discovery”. In: *arXiv preprint arXiv:2506.13131* (2025). doi: [10.48550/arXiv.2506.13131](https://doi.org/10.48550/arXiv.2506.13131).
- [16] J. Abramson et al. “Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3”. In: *Nature* 630.8016 (2024), pp. 493–500. doi: [10.1038/s41586-024-07487-w](https://doi.org/10.1038/s41586-024-07487-w).
- [17] Ž. Avsec et al. “Advancing Regulatory Variant Effect Prediction with AlphaGenome”. In: *Nature* 649.8099 (2026), pp. 1206–1218. doi: [10.1038/s41586-025-10014-0](https://doi.org/10.1038/s41586-025-10014-0).
- [18] OpenAI. *OpenAI and Los Alamos National Laboratory Announce Bioscience Research Partnership*. Accessed 2026-04-21. July 2024. URL: <https://openai.com/index/openai-and-los-alamos-national-laboratory-work-together/>.
- [19] Google DeepMind. *AI to Accelerate National Renewal and Growth as Google DeepMind Backs UK Tech and Science Sectors*. UK Government news release. Dec. 2025. URL: <https://www.gov.uk/government/news/ai-to-accelerate-national-renewal-and-growth-as-google-deepmind-backs-uk-tech-and-science-sectors>.

Machine learning

What is machine learning?

Machine learning is the development and study of computer algorithms for collecting and processing data to draw conclusions and make decisions [1, 2]. There is an overlap with statistics, but with a particular focus on actively seeking data, making predictions and taking decisions [3]. Research in this area typically starts from a problem specification stating the type of data available, and the task to be performed. Outputs are algorithms to use the data to perform the task, and theory describing how well we might expect to be able to carry out the task.

Established applications (machine learning meets AI)

The engine of modern AI is the suite of methods delivered by machine learning. A prominent example is in the recommender algorithms used to power the personalised internet [4]: machine learning methods learn individuals' preferences and choose appropriate content, products, and adverts to display. Several different machine learning methods were brought together in the training cycle of DeepMind's breakthroughs, including reinforcement learning to update behaviours based on experience [5], Bayesian optimisation for parameter tuning [6], and deep neural network training for function approximation [7]. When combined with computer science contributions to gather and process vast quantities of data, these approaches have resulted in the large language models that are the most prominent current AI methods.

Looking ahead

The only constant in AI research is that nothing is constant. New machine learning techniques will be required to develop the next generation of AI methods. One stream of research is to ensure that learning systems are as efficient as possible: the arrival of DeepSeek [8] rocked our assumptions about how much compute is required to train a high-performance system [9], and efficient machine learning approaches can contribute significantly to further developments in this area. An obvious gap in the current machine learning toolkit is the ability to deal with data that has complex structure; future machine learning methods will be devised to work with structured data, such as architectural plans or computer systems designs [10, 11], in contrast to flat images and texts. And in current approaches there is usually an implicit assumption that each data point is independent of the others; a similar independence assumption underpinned the banking collapse in 2008 [12], and it is important that we develop methods to robustly handle non-independent, and perhaps even machine-generated training data, otherwise future AIs will be unpredictably fragile [13].

A further critical contribution from machine learning, alongside statistical theory, uncertainty quantification, and information theory, will be an understanding of the limitations of what can be reasonably achieved by a learned system [14]. This will allow us to deploy our AI tools in the knowledge of their limitations, instead of blindly assuming that they are completely capable.

References

- [1] T. M. Mitchell. *Machine Learning*. McGraw-Hill, 1997. ISBN: 0070428077.
- [2] K. P. Murphy. *Probabilistic Machine Learning: An Introduction*. MIT Press, 2022. ISBN: 9780262046824.
- [3] L. Breiman. "Statistical Modeling: The Two Cultures". In: *Statistical Science* 16.3 (2001), pp. 199–231. doi: [10.1214/ss/1009213726](#).
- [4] F. Ricci, L. Rokach, and B. Shapira, eds. *Recommender Systems Handbook*. 2nd ed. Springer, 2015. ISBN: 9781489976376. doi: [10.1007/978-1-4899-7637-6](#).
- [5] D. Silver et al. "A General Reinforcement Learning Algorithm That Masters Chess, Shogi, and Go through Self-Play". In: *Science* 362.6419 (2018), pp. 1140–1144. doi: [10.1126/science.aar6404](#).
- [6] J. Snoek, H. Larochelle, and R. P. Adams. "Practical Bayesian Optimization of Machine Learning Algorithms". In: *Advances in Neural Information Processing Systems* 25. 2012.
- [7] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. ISBN: 9780262035613.
- [8] DeepSeek-AI et al. "DeepSeek-V3 Technical Report". In: (2024). doi: [10.48550/arXiv.2412.19437](#). arXiv: [2412.19437 \[cs.CL\]](#).
- [9] J. Kaplan et al. "Scaling Laws for Neural Language Models". In: (2020). doi: [10.48550/arXiv.2001.08361](#). arXiv: [2001.08361 \[cs.LG\]](#).
- [10] W. L. Hamilton, R. Ying, and J. Leskovec. "Representation Learning on Graphs: Methods and Applications". In: *IEEE Data Engineering Bulletin* 40.3 (2017), pp. 52–74.
- [11] M. M. Bronstein et al. "Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges". In: (2021). doi: [10.48550/arXiv.2104.13478](#). arXiv: [2104.13478 \[cs.LG\]](#).
- [12] D. Brigo, A. Pallavicini, and R. Torresetti. *Credit Models and the Crisis: A Journey into CDOs, Copulas, Correlations and Dynamic Models*. Wiley, 2010. ISBN: 9780470665660. doi: [10.1002/9781118374733](#).
- [13] I. Shumailov et al. "The Curse of Recursion: Training on Generated Data Makes Models Forget". In: (2023). doi: [10.48550/arXiv.2305.17493](#). arXiv: [2305.17493 \[cs.LG\]](#).
- [14] D. H. Wolpert. "The Lack of A Priori Distinctions Between Learning Algorithms". In: *Neural Computation* 8.7 (1996), pp. 1341–1390. doi: [10.1162/neco.1996.8.7.1341](#).

Mathematical analysis

What is mathematical analysis?

Mathematical analysis explores the profound properties and behaviours of mathematical functions and structures. It provides the techniques to reason about whether objects (for example videos or audio files) are close together or far apart. By considering limits, continuity, and infinite processes, analysis uncovers principles that govern everything from the simplest shapes to the most complex systems.

Established applications

Analysis underpins the whole of the mathematical sciences, providing the core machinery upon which the discipline is built. Beyond this, some of the current most immediate applications of analysis include financial modelling, where analytical tools are used to evaluate financial instruments, and computational fluid dynamics, where analysis allows researchers to determine whether their simulations are being carried out at a sufficiently high fidelity to be meaningful.

Analysis meets AI

AI may appear to be powered purely by data and computing power, but beneath the surface lies a deep mathematical structure. Far from being abstract or detached from reality, analysis provides the language and tools that make modern AI possible. There are two main reasons why mathematical analysis is essential in today's AI-driven world.

First, AI systems work in extraordinarily high-dimensional spaces. A high-resolution image, for example, contains millions of pixels; a short video contains vastly more data; even a sentence is represented internally as a large collection of numbers. To generate new images, write text, or simulate realistic data, AI systems rely on probability. They learn a complex probability distribution that describes what “realistic” data looks like, and then they sample from that distribution to produce new content. But high-dimensional probability is notoriously subtle. How do we represent such complex distributions? How can we efficiently generate new samples from them? How can we ensure that the generated output truly reflects the patterns in the data? Answering these questions requires powerful tools from probability theory and mathematical analysis.

Second, AI uses optimisation in a fundamental way. When a model learns to recognise handwritten digits or translate between languages, it does so by adjusting itself to reduce error. In essence, it measures how far its predictions are from reality and then tries to minimise that discrepancy. Mathematics provides the framework for defining what “error” or “distance” means and for systematically reducing it. Fields such as optimisation theory, calculus of variations, and partial differential equations (PDEs) give us the machinery to analyse and solve these problems. They allow us to understand not just how to train a model, but why certain training procedures work, when they converge, and how stable they are.

Looking ahead

One particularly difficult challenge in AI is sampling from complex probability distributions in very high dimensions. Many modern algorithms tackle this problem by simulating large collections of interacting “particles” whose collective behaviour approximates the desired distribution. While this may sound abstract, the idea is intuitive: instead of trying to describe a complicated distribution directly, we simulate many simple agents whose interactions produce the right global pattern. The mathematics behind these methods studies how these particle systems behave over time, whether they converge, how fast they do so, and under what conditions they remain stable. By minimising measures of discrepancy between probability distributions, such as the Kullback–Leibler divergence from information theory, researchers unify ideas from optimisation, sampling, and data filtering into a single conceptual framework.

Even the architectures behind large language models, such as Transformers, can be viewed through this analytical lens. These systems consist of many layers of interacting components that transform data step by step. Recent research suggests that their internal dynamics can be interpreted as large systems of interacting elements governed by equations familiar from mathematical physics. When examined at a larger scale, these systems often give rise to partial differential equations, revealing hidden structures and regularities. Why does this matter? Because understanding these mathematical structures is the key to moving beyond trial and error. Rigorous analysis can help us design more reliable algorithms, predict their behaviour, improve their stability, and ultimately provide guarantees about their performance.

References

- [1] J. A. Carrillo et al. “The Mean-Field Ensemble Kalman Filter: Near-Gaussian Setting”. In: *SIAM Journal on Numerical Analysis* 62.5 (2024), pp. 2549–2587. doi: [10.1137/24M1628207](https://doi.org/10.1137/24M1628207).
- [2] S. Chewi, J. Niles-Weed, and P. Rigollet. *Statistical Optimal Transport*. Vol. 2364. Lecture Notes in Mathematics. Cham: Springer, 2025. doi: [10.1007/978-3-031-85221-3](https://doi.org/10.1007/978-3-031-85221-3).
- [3] K. Craig, N. García Trillos, and D. Slepčev. “Clustering Dynamics on Graphs: From Spectral Clustering to Mean Shift through Fokker–Planck Interpolation”. In: *Geometric and Numerical Foundations of Movements*. Cham: Birkhäuser, 2022, pp. 105–151. doi: [10.1007/978-3-030-96179-4_4](https://doi.org/10.1007/978-3-030-96179-4_4).

- [4] N. García Trillos et al. “Error Estimates for Spectral Convergence of the Graph Laplacian on Random Geometric Graphs toward the Laplace–Beltrami Operator”. In: *Foundations of Computational Mathematics* 20.4 (2020), pp. 827–887. doi: [10.1007/s10208-019-09436-w](https://doi.org/10.1007/s10208-019-09436-w).
 - [5] B. Geshkovski et al. “A Mathematical Perspective on Transformers”. In: *Bulletin of the American Mathematical Society* 62.3 (2025), pp. 427–479. doi: [10.1090/bull/1863](https://doi.org/10.1090/bull/1863).
 - [6] J. Lu, Y. Lu, and J. Nolen. “Scaling Limit of the Stein Variational Gradient Descent: The Mean Field Regime”. In: *SIAM Journal on Mathematical Analysis* 51.2 (2019), pp. 648–671. doi: [10.1137/18M1187611](https://doi.org/10.1137/18M1187611).
 - [7] G. Peyré and M. Cuturi. “Computational Optimal Transport”. In: *Foundations and Trends in Machine Learning* 11.5–6 (2019), pp. 355–607. doi: [10.1561/22000000073](https://doi.org/10.1561/22000000073).
-

Mathematical modelling, validation, and verification

What is mathematical modelling?

Mathematical modelling involves building a mathematical representation of a real-world situation or system and using it to inform decisions or assess performance. The full breadth of mathematical sciences is employed in modelling, from discrete and continuous mathematics to logics and statistical methods.

Established applications

Examples are very broad, from the life sciences to finance, transportation, computer networks, and public health, to name just a few. With any model, it is important to have confidence that it is fit for purpose as designed (validation) and that it has been implemented as intended (verified). Validation and verification improve trustworthiness of a model, and thus decisions taken based on them. Models are very often approximate, in that they do not include every piece of detailed knowledge about the system, and are always provisional, in the sense that they may need to be updated or modified as new knowledge becomes available. Validating a model involves looking carefully at its assumptions and checking that they are justified in the context in which the model is intended to be used. Verification involves carefully testing the implementation of the model (almost always in software) and correcting any bugs or inaccuracies that could compromise the confidence in decisions that are guided by the model’s output. It is often very difficult to be sure that all ‘edge cases’ have been explored.

Mathematical modelling meets AI

AI provides powerful new tools for mathematical modelling. It offers the prospect of creating more complex models, creating them more quickly, and being able to iterate or update them more quickly. However, the capability to create models is going to outstrip the capability to validate and verify them. The traditional tasks of validation and verification make a distinction between the model (which is to be validated) and the implementation (which is to be verified). In an AI world, there is a very real sense in which the model IS the implementation. The distinction between validation and verification then becomes blurred. What does it even mean to validate or verify an AI model? These remain important open questions.

Looking ahead

In applications where decision-making processes require rigorous governance and audit trails, the successful widespread take-up of AI-based modelling will be held back unless these questions about validation and verification are answered. Even seemingly simpler questions currently have no accepted answer. For example, how should we assess whether two AI-generated models are different from each other? Questions such as these are stepping stones on the way to addressing the key challenges of validation and verification. The relevant areas of mathematics will be drawn from multiple branches, but especially statistics, optimisation and graph theory.

Numerical analysis

What is numerical analysis?

Numerical analysis is a pivotal field of mathematics that has evolved alongside digital computing, focusing on the design, testing, and rigorous investigation of algorithms that solve mathematical problems. This discipline assesses the efficiency, reliability, and limitations of algorithms, providing crucial insights into their capability to handle complex calculations. The advances in numerical analysis over the past half-century have underpinned significant developments in various applications, including weather forecasting, aerospace, and civil engineering. However, unlike the algorithms used in these traditional areas, those in AI often lack the same level of understanding and quality guarantees. This results in vulnerabilities where AI systems can be misled into providing incorrect outputs, a situation analogous to hallucinations or optical illusions. Examples include:

- Automated driving systems misinterpreting traffic signs due to minor alterations.
- AI-based facial recognition systems being deceived by altered glasses.
- Medical imaging AI being tricked into recognising spurious patterns.
- Interpretation methods for AI decisions being attacked.
- Large language models (LLMs) producing inappropriate content through crafted prompts.

By applying the tools of numerical analysis, researchers can better understand these vulnerabilities, assess their inevitability, and explore potential remedies, which is particularly relevant for enhancing the security and reliability of AI systems in practice.

Established applications

Numerical analysis has established itself as foundational in many critical fields, driving innovations such as:

Weather forecasting: algorithms that simulate weather patterns for accurate predictions.

Aerospace engineering: techniques that ensure the safety and efficiency of flight systems.

Generative AI: methods that support creation of text, images, and scientific designs, which continue to grow in complexity and size.

Demand for efficient algorithms is increasing due to the resource-intensive nature of generative AI; recent developments like Neural Ordinary Differential Equations (Neural ODEs) showcase the capacity to streamline memory and energy use during model training.

Numerical analysis meets AI

As AI systems increasingly permeate multiple sectors, the principles of numerical analysis have become integral to their development and optimisation. The challenges of growing model sizes and computational demands necessitate smarter algorithms, leading to significant improvements in memory and energy efficiency, thus enabling large-scale applications without overwhelming hardware requirements. Specific examples include:

- *MALI*: A memory-efficient integrator that reduces computational costs during training [1].
- *ClimODE*: A framework leveraging Neural ODEs for climate forecasting with fewer parameters [2].
- *EARTH* (Epidemiology-Aware Neural ODE): A model that outperforms traditional methods in predicting disease spread [3].

Looking ahead

Innovations not only enhance efficiency but also foster compliance with emerging regulations, supporting verifiability and explainability in AI systems. Investing in numerical analysis methods is essential to avoid inefficiencies and ensure a competitive edge in responsible AI development. The potential in the context of AI continues to expand, with several key areas for future exploration:

Development of advanced algorithms: focusing on creating more efficient numerical methods that can address the computational demands of complex AI systems.

Regulatory compliance: ensuring that AI applications meet emerging regulations, particularly in high-stakes domains, in which understanding and transparency are critical [4, 5].

Democratisation of AI tools: techniques like model reduction and approximation theory can lead to smaller, more agile AI systems, making large-scale tools accessible to a broader range of developers and applications.

In summary, numerical analysis is vital for the continued advancement of AI, providing the necessary frameworks and guarantees that enhance the adaptability, reliability, and transparency of AI systems while addressing emerging challenges in the field.

References

- [1] J. Zhuang et al. “MALI: A Memory Efficient and Reverse Accurate Integrator for Neural ODEs”. In: *International Conference on Learning Representations*. 2021.
- [2] Y. Verma, M. Heinonen, and V. Garg. “ClimODE: Climate and Weather Forecasting with Physics-Informed Neural ODEs”. In: *International Conference on Learning Representations*. 2024.
- [3] G. Wan et al. “Epidemiology-Aware Neural ODE with Continuous Disease Transmission Graph (EARTH)”. In: *International Conference on Machine Learning*. 2025.
- [4] Department for Science, Innovation and Technology. *Implementing the UK’s AI Regulatory Principles: Initial Guidance for Regulators*. Tech. rep. GOV.UK, 2024.
- [5] Information Commissioner’s Office. *Guidance on AI and Data Protection*. 2023. URL: <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>.
- [6] A. Bastounis, I. Y. Tyukin, and D. J. Higham. “A Survey of Inevitability Results in AI Instability”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* (2026). Accepted/In press.

- [7] A. Bastounis et al. “On the Consistent Reasoning Paradox of Intelligence and Optimal Trust in AI: The Power of ‘I Don’t Know’”. In: (2024). doi: [10.48550/arXiv.2408.02357](https://doi.org/10.48550/arXiv.2408.02357). arXiv: [2408.02357](https://arxiv.org/abs/2408.02357) [[cs.AI](#)].
- [8] A. Grivas, A. Vergari, and A. Lopez. “Taming the Sigmoid Bottleneck: Provably Argmaxable Sparse Multi-Label Classification”. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 38. 11. 2024, pp. 12014–12022. doi: [10.1609/aaai.v38i11.29110](https://doi.org/10.1609/aaai.v38i11.29110).
- [9] N. Savage. “Always Improving Performance”. In: *Communications of the ACM* 65.6 (2022), pp. 16–18. doi: [10.1145/3531221](https://doi.org/10.1145/3531221).
- [10] S. Maslovskaya et al. “Adaptive Higher Order Reversible Integrators for Memory Efficient Deep Learning”. In: (2024). doi: [10.48550/arXiv.2410.09537](https://doi.org/10.48550/arXiv.2410.09537). arXiv: [2410.09537](https://arxiv.org/abs/2410.09537) [[math.NA](#)].
-

Optimisation theory

What is optimisation theory?

The field of optimisation theory develops tools to find the best solution to a problem from a set of choices. Finding the maximum of a function, such as a quadratic equation, is an optimisation task. Finding the routing tables that make the internet as robust and stable as possible is also an optimisation task. Of course, one of these tasks is much simpler than the other, but both can be addressed using optimisation methods. The field uses ideas from calculus, analysis, graph theory and numerical analysis, and provides algorithms that find the best possible inputs given a careful description of the problem.

Established applications

Optimisation has been used in a vast range of important areas, including:

- deriving the least cost diet that meets nutritional needs for the US Army in World War II;
- investment diversification in portfolio optimisation, being a solution of the optimisation problem of maximising the expected return for a given level of risk;
- warehouse and distribution optimisation for retail giants such as Tesco and Amazon;
- medical treatment optimisation, finding the regime that delivers minimal side effects while maximising the clinical benefit; and
- finding the statistical model that best fits a given set of data.

Optimisation meets AI

Large language models (LLMs), of which ChatGPT is the most well-known, are one of the AI tools most used by individuals and companies. An LLM is an example of a neural network, in other words, a computer construct that simulates the operation of the human brain. The success of a neural network in providing good answers depends on finding effective settings for its internal connections, and this is achieved during training using optimisation algorithms. This leads to the question: How can we find the best-possible set of connections for a neural network? This is an optimisation problem because the network can be completely described using mathematical formulae, and an optimisation algorithm can then be used to find the optimal (i.e., best-possible) choice of connections. As AI grows in importance, so do the optimisation algorithms that are vital for its success, and progress on the algorithms directly improves the performance of AI tools. Recommended references are the book [1] and the article [2] on the most popular optimisation method, called Adam, cited more than 230,000 times.

Once we have an AI tool that works well, we may be interested in understanding why the answer it gives may not be the answer we expected. There is the known problem of “hallucinations” which means that the answer is nonsense. However many applications of AI are based on giving yes/no answers to questions, for example deciding whether the bank should approve a loan application or whether a job application should be shortlisted for interview. Suppose an application is rejected; we may want to know (or be legally required to justify) why the AI made this decision. One way to explain the decision is by asking for a counterfactual: What part or parts of the application need to change for the outcome to change? For the example of a loan application, the counterfactual may involve having a higher income or owing less money on other loans. It is often of interest to ask more specifically: What is the smallest change to the application that would change the decision? This is again an optimisation problem. Recommended references are the surveys [3] and [4].

A similar but different question is: Is there a way to change the connections in the AI so that it makes decisions that are “wrong”? For example, can we mislead the AI into approving a loan that should be rejected? Moreover, is it possible to do this in such a way that the AI user, the loan provider in our example, is unaware of the change? The answer is yes, and one way to do this is via poisoning attacks. These attacks take place by changing the data informing the optimisation that sets up the connections, in such a way that the outcome of the optimisation and the resulting AI network is different. Perhaps surprisingly, researchers have shown that optimisation can be used to design such “attacks”, or alternatively, to test the ability of the AI to resist such attacks. Recommended references are the survey [5] and the very recent article [6].

Looking ahead

Since optimisation lies at the core of AI, future developments in AI will rely on, and inspire, new developments in optimisation theory. In fact, very many of the current questions around AI safety, reliability and security can, and will be, formulated as optimisation problems.

References

- [1] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press, 2016. ISBN: 9780262035613.
- [2] D. P. Kingma and J. Ba. “Adam: A Method for Stochastic Optimization”. In: (2014). DOI: [10.48550/arXiv.1412.6980](https://doi.org/10.48550/arXiv.1412.6980). arXiv: [1412.6980](https://arxiv.org/abs/1412.6980) [cs.LG].
- [3] R. Guidotti. “Counterfactual Explanations and How to Find Them: Literature Review and Benchmarking”. In: *Data Mining and Knowledge Discovery* 38.6 (2024), pp. 2770–2824. DOI: [10.1007/s10618-022-00831-6](https://doi.org/10.1007/s10618-022-00831-6).
- [4] M. A. Prado-Romero et al. “A Survey on Graph Counterfactual Explanations: Definitions, Methods, Evaluation, and Research Challenges”. In: *ACM Computing Surveys* 56.7 (2024), 171:1–171:37. DOI: [10.1145/3618105](https://doi.org/10.1145/3618105).
- [5] Z. Tian et al. “A Comprehensive Survey on Poisoning Attacks and Countermeasures in Machine Learning”. In: *ACM Computing Surveys* 55.8 (2022), 168:1–168:35. DOI: [10.1145/3551636](https://doi.org/10.1145/3551636).
- [6] Z. Şuvak et al. “Design of Poisoning Attacks on Linear Regression Using Bilevel Optimization”. In: *Data Science and Optimization*. Ed. by S. Dang et al. Vol. 91. Fields Institute Communications. Cham: Springer Nature Switzerland, 2026, pp. 41–65. DOI: [10.1007/978-3-032-03844-9_2](https://doi.org/10.1007/978-3-032-03844-9_2).

Physics-informed modelling

What is physics-informed modelling?

There is a strong history of mathematical models being developed to describe and understand physical phenomena: from Newtonian mechanics to quantum physics. Recently, the terminology “physics-informed modelling” has been used to describe a powerful framework for combining modern data-driven machine learning with physical laws that describe how the real world works. Rather than relying on data alone, these models are guided by physical principles that are formulated mathematically, often as differential equations, and that have been tested and refined over centuries.

Established applications

The approach of physics-informed modelling is especially important in real physical systems, such as robots or vehicles, or to understand and predict complex systems, including climate and Earth systems, biological organs, or the behaviour of plasma in nuclear fusion reactors.

Many critical real-world problems are constrained by limited, costly, or imperfect data. Measurements may be expensive (for example, medical imaging or environmental monitoring), difficult or impossible to obtain directly (such as inside the Earth, a living cell, a star, or a nuclear reactor), and are often noisy or incomplete. In some cases, multiple explanations may fit the available data equally well. By embedding physical laws and statistical inference principles directly into the modelling process, physics-informed approaches rule out physically implausible solutions, reducing the risk of fitting data while violating physical consistency. This leads to predictions that are more reliable when data are scarce, and more resilient under changing or uncertain conditions. As a result, physics-informed modelling is essential for building robust, efficient, and credible decision-support tools.

Physics-informed modelling meets AI

Whereas statistical inference focuses on identifying the correct model that connects input to output, machine learning approaches focus on making output predictions based on inputs. When models respect physical reality, they are easier for scientists and engineers to trust, interpret, explain, improve, and adopt in practice, making them particularly valuable for scientific and engineering applications of machine learning.

Physics-informed modelling also plays a central role in emerging uses of generative AI for physical systems, including so-called digital twins: virtual replicas of real-world systems. Digital twins are transforming engineering and manufacturing by enabling the design, optimisation, and evaluation of new materials, products, and services in simulated environments that respect physical constraints.

Looking ahead

Physics-informed modelling brings together scientific knowledge, statistical inference, and modern machine-learning capabilities to support decisions based not only on data, but also on the fundamental laws that govern the world. By combining machine learning with the mathematical sciences that underpin reliable modelling, these approaches help produce predictions that are more robust,

credible, and trustworthy in real-world settings. For example, future robotic systems and autonomous vehicles will increasingly be trained and tested using data generated in realistic, physically consistent virtual environments.

References

- [1] T. De Ryck and S. Mishra. “Numerical analysis of physics-informed neural networks and related models in physics-informed machine learning”. In: *Acta Numerica* 33 (2024), pp. 633–713.
- [2] G. E. Karniadakis et al. “Physics-informed machine learning”. In: *Nature Reviews Physics* 3.6 (2021), pp. 422–440.
- [3] N. Baker et al. *Workshop report on basic research needs for scientific machine learning: Core technologies for artificial intelligence*. Tech. rep. USDOE Office of Science (SC), Washington, DC (United States), 2019.
- [4] H. Wang et al. “Scientific discovery in the age of artificial intelligence”. In: *Nature* 620.7972 (2023), pp. 47–60.
- [5] J. Tachella and M. Davies. “Self-supervised learning from noisy and incomplete data”. In: *Foundations and Trends® in Signal Processing* 20.2 (2026), pp. 85–184.

Statistical theory

What is statistical theory?

Statistical theory is the branch of mathematics that explains how we can learn reliably from data when certainty is impossible and the cost of error is real. Its value becomes clearest not in calm conditions, but under pressure. As documented in [1], during the Second World War, statisticians were asked questions that could not be answered by experimentation alone: how to armour aircraft so pilots were more likely to return alive; how many weapons must be tested before concluding a new design was safer or more dangerous; how to set proximity fuses so artillery shells exploded at exactly the right moment. Their answers were not academic. Fighter planes entered combat using their calculations; rockets were approved or rejected on their advice. As one participant later recalled, when recommendations were made “things happened—and maybe the pilots came back or maybe they didn’t.” Statistical theory transformed fragmentary data, uncertainty and urgency into decisions that could be justified and lived with, rather than guessed at.

Established applications (statistical theory meets AI)

Today’s AI systems operate under the same conditions: imperfect data, uncertainty and irreversible consequences, but now decisions are executed automatically, at scale, and often without meaningful human intervention. This is why statistical theory is central to AI safety: it is the discipline that asks not only whether a system works on average, but how often it fails, how badly it fails, and who bears the cost when it does.

The most visible risks of AI today—loss of privacy, unfair or discriminatory outcomes, and overconfident systems deployed beyond their limits—are not peripheral concerns but safety failures in precisely this sense. An AI system that leaks sensitive personal data, systematically disadvantages particular groups, or produces confident but wrong outputs in unfamiliar settings, is not merely imperfect, it is unsafe. Statistical theory provides the tools to confront these risks before they cause harm. It makes privacy quantifiable rather than aspirational, reveals unavoidable trade-offs between different notions of fairness, and equips practitioners to measure uncertainty and detect when models are operating outside the conditions for which they were designed. Without these foundations, calls for trustworthy AI amount to intentions without enforcement.

Privacy provides a clear illustration of why statistical theory is indispensable rather than optional. Modern AI models can unintentionally memorise and reveal sensitive information, making informal assurances of privacy meaningless. Statistical theory makes privacy precise by defining how much information about any individual can be inferred from a system’s outputs, even by an adversary with substantial background knowledge. Techniques such as differential privacy arise directly from this work, offering provable bounds on information leakage while preserving utility [2]. Without continued advances in statistical theory, privacy protections cannot keep pace with increasingly powerful models, and assurances of safety collapse into trust-based promises.

Fairness poses an even more uncomfortable challenge. In high-stakes applications, AI systems influence access to jobs, credit, healthcare, and public services. Statistical theory shows that fairness is not a single objective but a set of competing criteria that cannot all be satisfied simultaneously. This is not a moral failure of AI, but a mathematical reality. The value of theory is that it exposes these trade-offs clearly, enabling transparent governance rather than hidden bias. Statistical methods also reveal how unfairness can arise even when sensitive attributes are excluded, through correlations embedded in data and feedback loops created by deployment [3].

At the same time, statistical theory is not only defensive: it is also a driver of innovation in AI itself. Diffusion models, which underpin some of the most striking recent advances in generative AI, are built on explicitly statistical foundations. Their success relies on understanding how noise can be injected and removed in a controlled way, how probability distributions can be learned through score matching, and how stochastic processes evolve over time [4]. Statistical theory explains why these models are stable, when they generalise, and how they can be extended beyond images to scientific simulation and inverse problems. Without this theoretical understanding, diffusion models would remain fragile heuristics; with it, they become reliable, extensible tools.

Looking ahead

Taken together, these examples point to a single conclusion. As AI systems grow more capable, they also grow more consequential, and the margin for error shrinks. Statistical theory ensures that AI systems can be trusted not just when they succeed, but when they fail; not just on average, but in the rare cases that matter most. Investment in statistical theory is therefore not an investment in abstract mathematics for its own sake. It is an investment in AI systems that are safe to deploy, governable at scale, and worthy of the confidence society increasingly places in them.

References

- [1] W. A. Wallis. “The Statistical Research Group, 1942–1945”. In: *Journal of the American Statistical Association* 75.370 (1980), pp. 320–330. doi: [10.1080/01621459.1980.10477469](https://doi.org/10.1080/01621459.1980.10477469).
- [2] R. Cummings et al. “Advancing Differential Privacy: Where We Are Now and Future Directions for Real-World Deployment”. In: *Harvard Data Science Review* 6.1 (2024). doi: [10.1162/99608f92.d3197524](https://doi.org/10.1162/99608f92.d3197524).
- [3] S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning: Limitations and Opportunities*. MIT Press, 2023. ISBN: 9780262048613.
- [4] Y. Song et al. “Score-Based Generative Modeling through Stochastic Differential Equations”. In: *International Conference on Learning Representations*. 2021.

Stochastic analysis

What is stochastic analysis?

Stochastic analysis studies how random or highly irregular processes evolve over time and influence the systems in which they appear. Its tools can be used to describe pollen grains moving in water, noisy signals from a sensor, or fluctuating financial prices. Many central examples are too rough for ordinary calculus. Brownian motion is the classical case, with paths that are, almost surely, continuous but nowhere differentiable. Itô’s breakthrough was to give rigorous meaning to integration with respect to Brownian motion, leading to the theory of stochastic differential equations. Central objects in the field include Markov processes, martingales, stochastic partial differential equations, filtering equations, interacting particle systems, and rough paths [1, 2, 3, 4, 5]. A recurring theme is the search for a state that captures what the past tells us about the future. The point is not to eliminate randomness, but to represent the information available so far in a form that tells us what can be predicted, inferred, or controlled. Sometimes the observed present is enough. Often it is not. Hidden variables, noisy measurements, and path dependence mean that relevant information may be distributed across the past rather than visible in the latest observation. Stochastic analysis encodes that information through filtrations, conditional laws, filters, Markovian representations, path signatures, and other enriched states.

Established applications

Stochastic analysis is used wherever uncertainty interacts with dynamics. In physics and chemistry, it describes Brownian motion, molecular fluctuations, Langevin dynamics, diffusion and jump processes, and random forcing. In engineering, it underpins filtering, signal processing, sensor fusion, and control from noisy measurements. In finance and insurance, it is used to model volatility, risk, rare losses, interest rates, and path-dependent contracts. In biology and medicine, it helps describe gene regulation, epidemics, neuronal activity, and biochemical networks. In climate and Earth systems, stochastic models help represent unresolved scales, propagate uncertainty, and reason about extreme events.

Stochastic analysis meets AI

Modern AI systems are stochastic at many levels. Neural networks are randomly initialised, training data are sampled into mini-batches, gradients are noisy estimates of ideal gradients, and regularisation may deliberately inject noise. Stochastic gradient descent can therefore be studied not only as an optimisation algorithm, but as a stochastic dynamical system moving through a high-dimensional loss landscape [6, 7, 8, 9]. Generative AI gives an even more direct connection. Diffusion models and score-based generative models add noise to data and train a model to reverse that noising process. In denoising diffusion probabilistic models, this is formulated as a stochastic Markov chain whose reverse transitions are learned from data [10]. Score-based generative models place the same idea in continuous time using stochastic differential equations [11]. Neural SDEs extend the stochastic-dynamical viewpoint to path-valued data, using Brownian motion as input noise and learning continuous-time generative models for time series [12]. Related approaches, including Schrödinger bridges and stochastic interpolants, view generation as the controlled movement of probability distributions [13, 14]. Understanding when these methods work, miss rare modes, or become unstable is a problem for stochastic analysis. Stochastic analysis is also important for AI systems that process information over time. A language model does not predict the next token from the current token alone. It builds an internal representation of the preceding context, and this representation is used to predict the probability distribution of what comes next [15]. Continuous-time machine learning models, such as neural controlled differential equations and structured linear CDEs, play an analogous role for irregular streams of data, such as medical records, sensor streams,

financial data, and climate observations [16, 17]. Rough path theory provides the mathematical foundations for analysing these continuous-time models [18].

Looking ahead

As AI systems become more autonomous and more closely connected to science, medicine, finance, and physical infrastructure, their internal randomness and external uncertainty will matter more. We need to know when stochastic training is stable, when sampling algorithms fail, how uncertainty accumulates through long chains of predictions and decisions, and what information a learned state has retained or discarded. Stochastic analysis connects optimisation with random dynamics, generative modelling with stochastic flows, sequence modelling with state construction, and autonomous decision-making with stochastic control. It supplies mathematical foundations for AI systems whose uncertainty, memory, stability, and failure modes can be understood rather than merely observed.

References

- [1] G. Grimmett and D. Stirzaker. *Probability and Random Processes*. 3rd ed. Oxford University Press, 2001.
- [2] J.-F. Le Gall. *Brownian Motion, Martingales, and Stochastic Calculus*. Springer, 2016.
- [3] L. C. G. Rogers and D. Williams. *Diffusions, Markov Processes and Martingales*. 2nd ed. Cambridge University Press, 2000.
- [4] A. Bain and D. Crisan. *Fundamentals of Stochastic Filtering*. Springer, 2009. doi: [10.1007/978-0-387-76896-0](https://doi.org/10.1007/978-0-387-76896-0).
- [5] P. K. Friz and M. Hairer. *A Course on Rough Paths: With an Introduction to Regularity Structures*. Springer, 2014. doi: [10.1007/978-3-319-08332-2](https://doi.org/10.1007/978-3-319-08332-2).
- [6] S. Mandt, M. D. Hoffman, and D. M. Blei. “Stochastic Gradient Descent as Approximate Bayesian Inference”. In: *Journal of Machine Learning Research* 18 (2017).
- [7] U. Şimşekli, L. Sagun, and M. Gürbüzbalaban. “A Tail-Index Analysis of Stochastic Gradient Noise in Deep Neural Networks”. In: *Proceedings of the 36th International Conference on Machine Learning*. 2019.
- [8] J. Sirignano and K. Spiliopoulos. “Mean Field Analysis of Neural Networks: A Law of Large Numbers”. In: *SIAM Journal on Applied Mathematics* 80.2 (2020), pp. 725–752.
- [9] Z. Xie, I. Sato, and M. Sugiyama. “A Diffusion Theory for Deep Learning Dynamics: Stochastic Gradient Descent Exponentially Favors Flat Minima”. In: *International Conference on Learning Representations*. 2021.
- [10] J. Ho, A. Jain, and P. Abbeel. “Denoising Diffusion Probabilistic Models”. In: *Proceedings of the 34th Conference on Neural Information Processing Systems*. 2020.
- [11] Y. Song et al. “Score-Based Generative Modeling through Stochastic Differential Equations”. In: *International Conference on Learning Representations*. 2021.
- [12] P. Kidger et al. “Neural SDEs as Infinite-Dimensional GANs”. In: *Proceedings of the 38th International Conference on Machine Learning*. 2021.
- [13] V. De Bortoli et al. “Diffusion Schrödinger Bridge with Applications to Score-Based Generative Modeling”. In: *Proceedings of the 35th Conference on Neural Information Processing Systems*. 2021.
- [14] M. S. Albergo, N. M. Boffi, and E. Vanden-Eijnden. “Stochastic Interpolants: A Unifying Framework for Flows and Diffusions”. In: *Journal of Machine Learning Research* 26.209 (2025), pp. 1–80.
- [15] A. Vaswani et al. “Attention Is All You Need”. In: *Proceedings of the 31st Conference on Neural Information Processing Systems*. 2017.
- [16] P. Kidger et al. “Neural Controlled Differential Equations for Irregular Time Series”. In: *Proceedings of the 34th Conference on Neural Information Processing Systems*. 2020.
- [17] B. Walker et al. “Structured Linear CDEs: Maximally Expressive and Parallel-in-Time Sequence Models”. In: *Proceedings of the 39th Conference on Neural Information Processing Systems*. 2025.
- [18] T. Lyons. “Differential Equations Driven by Rough Signals”. In: *Revista Matemática Iberoamericana* 14.2 (1998), pp. 215–310.

Topology and geometry

What are topology and geometry?

Both topology and geometry study “shapes” from different points of view. Topology considers the connectivity of a shape and ignores the distances between points. In this view a coffee mug and a doughnut are topologically equivalent because they both have exactly one hole. In contrast, geometry focuses on the precise relative positions of the different parts of the shape.

Established applications

In recent years, data analysis techniques have developed using both topology and geometry as the foundation. *Topological Data Analysis* studies topological features, such as connected components, loops and voids, across all scales in a cloud of unordered points. *Geometric data science* allows us to see and measure the difference between real objects, such as a coffee mug and a doughnut, and

gives rigorous answers [1] to the three fundamental questions about any real data objects: (1) Same or different? (2) If different, by how much? (3) Where do all real objects live?

Geometric data science: topology and geometry meet AI

Although the space of all objects of a particular type (e.g., all images of a given size) is extremely large, the *manifold hypothesis* posits that any realistic data actually lies in a much smaller, lower-dimensional subset of the space. This is what allows AI models to learn statistical patterns and to generalise from the observed data. Geometric data science allows us to map ambiguous representations of objects onto a smaller subspace of realisable and invariant data to distinguish what the observations actually represent.

Many objects, ranging from molecules to galaxies, are observed through photograph-style representations, and a newly observed object can be claimed “new” if this object cannot be related to existing objects. Here geometric methods can be used to measure distances between objects and find already known near-duplicates [2].

A recent correction in *Nature* [3] crossed out almost all words “novel” and “discovery” due to a mathematical review using geometric methods [4]. The problem even impinges on Google DeepMind’s protein folding prediction engine, AlphaFold, which can make physically unrealistic predictions [5] due to a black-box training on the Protein Data Bank, which geometric tools have shown to include many atomic clashes and duplicate entries [6].

Looking ahead

Data spaces are often high-dimensional, and understanding the “shape” of data is critical to understand important issues around safety and trustworthy AI. The huge size of data spaces means that a greedy optimisation can always find a way around manually designed “guardrails” [7]. Finding off-manifold data is an established mechanism for finding adversarial inputs to AI systems (e.g. [8]): inputs that look like one thing to humans but are processed as another by the system (a simple example being an image of a cat that gets classified as a dog). Further understanding and utilising the topology and geometry of data will enhance our understanding of attacks, defences, and guardrails in AI systems.

Geometric data science can be used to help AI systems to maintain semantic consistency, ensuring that AI responses stay within a logically sound “manifold”, and hence preventing hallucinations.

Furthermore, since states in forthcoming quantum computers are inherently geometric, future quantum machine learning will need geometry as its primary “language” of data.

References

- [1] O. D. Anosova and V. A. Kurlin. “Geometric Data Science”. In: (2025). DOI: [10.48550/arXiv.2512.05040](https://doi.org/10.48550/arXiv.2512.05040). arXiv: [2512.05040](https://arxiv.org/abs/2512.05040) [math.MG].
- [2] D. S. Chawla. *Duplicate Structures Haunt Crystallography Databases*. Chemical & Engineering News. Dec. 2025. URL: <https://cen.acs.org/research-integrity/Duplicate-structures-haunt-crystallography-databases/103/web/2025/12>.
- [3] N. J. Szymanski et al. “Author Correction: An Autonomous Laboratory for the Accelerated Synthesis of Inorganic Materials”. In: *Nature* 650.8100 (2026), E1. DOI: [10.1038/s41586-025-09992-y](https://doi.org/10.1038/s41586-025-09992-y).
- [4] D. Widdowson and V. Kurlin. “Geographic-Style Maps with a Local Novelty Distance Help Navigate in the Materials Space”. In: *Scientific Reports* 15 (2025), p. 27588. DOI: [10.1038/s41598-025-10672-0](https://doi.org/10.1038/s41598-025-10672-0).
- [5] D. Chakravarty and L. L. Porter. “AlphaFold2 Fails to Predict Protein Fold Switching”. In: *Protein Science* 31.6 (2022), e4353. DOI: [10.1002/pro.4353](https://doi.org/10.1002/pro.4353).
- [6] A. Wlodawer et al. “Duplicate Entries in the Protein Data Bank: How to Detect and Manage Them”. In: *Acta Crystallographica Section D: Structural Biology* 81.4 (2025), pp. 170–180. DOI: [10.1107/S2059798325001883](https://doi.org/10.1107/S2059798325001883).
- [7] Y. Bengio et al. “International AI Safety Report 2026”. In: (2026). DOI: [10.48550/arXiv.2602.21012](https://doi.org/10.48550/arXiv.2602.21012). arXiv: [2602.21012](https://arxiv.org/abs/2602.21012) [cs.AI].
- [8] D. Stutz, M. Hein, and B. Schiele. “Disentangling Adversarial Robustness and Generalization”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2019.

Uncertainty quantification

What is uncertainty quantification?

Whether deciding to take an umbrella out on a cloudy day or to buy shares in a company, most of our decisions are taken under uncertainty. After all, it may not rain, or the share price might drop tomorrow. Whether our sense of what is likely or not likely to happen stems from gut feelings or deep insight into past data, it is a remarkable ability we possess. It not only plays an integral part

in our daily lives, but many professional endeavours—from doctors to lawyers and financial advisers advising their clients—involve assessing and processing uncertainty.

However, there is often disagreement about how to assign probabilities to future events, and how to make decisions when faced with this uncertainty. Uncertainty quantification is the field of mathematics that focuses on the identification and characterisation of uncertainty in computational and experimental models. It provides methods and language with which to reason about uncertainty in our models of the world.

Established applications

Uncertainty quantification uses probability and sensitivity analyses, and simulation, to explore the uncertainty in a system. It is integral to enhancing the robustness and reliability of models and predictions across various scientific and engineering applications. It is used extensively in engineering design and optimisation, helping engineers to explore and express the uncertainty in the behaviour of their systems. It is also a key tool in climate modelling and other predictive models in environmental science; it provides the techniques that are used to quantify and express the uncertainties presented in future outcomes for our climate.

Uncertainty quantification meets AI

Ideally, like their remarkable ability to explain and reason with difficult logical concepts, AI systems should also possess the ability to quantify their uncertainty, reason with it, and advise accordingly. Currently, they do not do this well, instead preferring to confidently provide an answer rather than to exhibit uncertainty. While tools of probability theory and statistics are integral to the training and running of AI models, they have not yet been well incorporated into the presentation of responses.

Looking ahead

There is enormous untapped potential for uncertainty quantification to be deployed in AI tools. This would allow greater trust to be placed in the outputs of AI systems; if we have trust in the accuracy estimates a system provides in its answer, then when it tells us it is 99% certain of its answer we can act on that answer with high confidence. In contrast, we should be very sceptical of any answer given without trustworthy uncertainty quantification, since we simply do not know how much confidence to place in what is stated.

Proper uncertainty quantification is also critical to allow AI systems to become better team members; expressing confidence levels allows other (human or AI) members of the team to respond appropriately to the outputs that are provided.



Academy for the
Mathematical Sciences