Mathematics and Artificial Intelligence

Discovery, Understanding and Trust

Executive Statement

International Council for Industrial and Applied Mathematics (ICIAM), September 2026

A moment of extraordinary opportunity, and responsibility

Artificial intelligence is rapidly changing mathematical research, scientific computing and engineering. ICIAM welcomes these developments and believes that applied and industrial mathematics must have a strong voice in shaping them. AI can search enormous spaces, generate proofs and counterexamples, accelerate simulation and support scientific discovery. But mathematical progress is more than the production of answers, and scientific usefulness requires more than impressive benchmark performance. Understanding, validation, reliability, attribution and human judgement remain essential. In industry, AI is increasingly being combined with mathematical models and simulation to enable faster design, real-time prediction, digital twins and decision support. For ICIAM, the central opportunity is therefore a partnership in which AI extends human capabilities while mathematics provides structure, insight and trust.

1. AI can accelerate discovery, but discovery is more than producing answers

Proofs, counterexamples and failed approaches are all part of mathematical progress. A counterexample does not end a question: it asks us why a conjecture fails, which assumptions are missing, and what stronger or more meaningful formulation should replace it. The same principle should apply to AI-assisted mathematics. Knowing where an AI system failed, which approaches it explored unsuccessfully and under which conditions its reasoning broke down can be as informative as knowing where it succeeded. Selective reporting of successes alone gives an incomplete picture. Openness about failures, limitations and unsuccessful avenues will strengthen understanding, reproducibility and responsible evaluation.

2. Mathematics is essential for making AI trustworthy

When AI and machine learning become part of scientific and engineering computation, reliability is not optional. Applied mathematics provides tools for stability, robustness, error control, uncertainty quantification, physical consistency, causality and validation. These properties are particularly important when AI-generated predictions are embedded in engineering workflows, where an apparently accurate model may nevertheless lead to unstable simulations, non-physical behaviour or unreliable decisions. Incorporating known mathematical structure, such as conservation laws, symmetries and positivity properties, can reduce the learning problem, improve data efficiency and provide stronger guarantees. Mathematical analysis can help establish which properties are essential and how they can be preserved. Mathematics can help turn powerful data-driven models into scientific instruments whose behaviour and limitations can be understood and trusted.

3. Mathematics provides a powerful toolbox for AI

Decades of computational mathematics have produced highly efficient methods for differential equations, optimisation, inverse problems, multiscale computation, uncertainty quantification, model reduction and large-scale simulation. These methods are often fast, robust, interpretable and supported by mathematical analysis. They provide a powerful foundation for developing AI approaches that are efficient, reliable and interpretable. Indeed, many AI surrogate approaches echo ideas long used in model order reduction: substantial offline computation is used to construct a low-dimensional or learned representation that can then be evaluated rapidly online. Combining these established mathematical ideas with the capabilities of AI creates opportunities for methods that neither approach could achieve alone.

4. The future is a partnership of human insight, mathematics, data and AI

The interaction between mathematics and AI creates major new opportunities for science, engineering and industry. Mathematical knowledge can guide learning, reduce complexity and data requirements, improve reliability, and help determine when predictions can be trusted. AI, in turn, brings powerful new capabilities for search, approximation, synthesis and the exploitation of data. For industry, combining mathematical models, simulation, data and AI can make reliable mathematical prediction an integrated component of design, control and decision-making. Human insight determines which questions matter and how results should be interpreted; mathematics provides structure, models and constraints; data connect models to observations; AI provides new computational capabilities. The age of AI creates not less need for applied mathematics, but more.

5. The mathematical community must help shape the rules and infrastructure

ICIAM and its member societies have a responsibility to engage actively. Applied and industrial mathematicians have a particularly important role in building these bridges, because they work at the interfaces between mathematical models, numerical methods, data, scientific computing and real-world applications. AI-assisted mathematics and science require reliable verification, careful attribution, reproducibility and transparent provenance. The community should help develop open and robust mathematical validation environments, encourage disclosure of computational methods and resources, and create opportunities to learn from both successes and failures. We should also consider equitable access to the AI and computational infrastructure increasingly required for frontier research. Finally, we must educate mathematicians not merely to use AI tools, but to understand, analyse and improve them, and to recognise when established mathematical methods, AI approaches or a combination of both provide the best solution.

Our message

The appropriate response to AI is neither resistance nor uncritical enthusiasm, but active engagement. Mathematicians should use these tools, study them, improve them and help shape the conditions under which they are used. We should be ambitious about what human and artificial intelligence can achieve together, while remaining rigorous about correctness, transparency, reproducibility and scientific meaning.

AI can accelerate discovery. Mathematics can provide understanding and trust. Together, human and artificial intelligence can open new directions for mathematics, science, engineering, industry and society.

Mathematics and Artificial Intelligence

Discovery, Understanding and Trust

A statement by the International Council for Industrial and Applied Mathematics (ICIAM)

September 2026

A moment of extraordinary opportunity, and responsibility

Artificial intelligence is beginning to change the way mathematics is discovered, developed, communicated and used, and the way mathematical models and computations contribute to science, engineering and industry. Systems based on large language models, formal proof assistants, machine learning and increasingly autonomous collections of agents have moved from solving routine exercises to contributing to research level mathematics. They can search large spaces of possible arguments, suggest conjectures, find counterexamples, formalise proofs, check logical steps and combine information at a scale that no individual researcher can match. These developments are remarkable, and the international mathematical community should neither underestimate them nor turn away from them.

At the same time, the speed of these developments forces us to ask a more fundamental question. What do we actually mean by progress in mathematics? Is progress measured by the number of open problems that are closed, by the number of proofs that can be produced, or by something richer? Mathematics is not only a collection of correct statements. It is also a way of understanding structures, mechanisms and relations. It is a cumulative human enterprise in which ideas are tested, simplified, generalized, connected to other ideas, taught to students, applied to the world, and eventually absorbed into the common language of the discipline.

This concern was expressed forcefully in September 2026 in *A Severe Misalignment of AI in Mathematics*, a statement by 25 Fields medallists that was also published in Le Monde. Their warning should not be read as a rejection of artificial intelligence. Rather, it is a reminder that the objectives of commercial AI development and the objectives of the mathematical community need not automatically coincide. A system may be optimised to produce an answer, while a mathematician seeks understanding, explanation, generality and new questions. If we fail to distinguish these goals, we risk confusing mathematical output with mathematical progress.

The same concern lies at the heart of the *Leiden Declaration on Artificial Intelligence and Mathematics*, published in June 2026 and endorsed by the International Mathematical Union. The Declaration stresses values that have long been central to mathematics: correctness, transparency, attribution, independent verification, human responsibility, shared standards of depth and significance, and the development of understanding and judgement within the mathematical community. These values are not obstacles to technological progress. They are part of the foundation on which meaningful progress rests.

A similar perspective was expressed by the London Mathematical Society following the recent Navier–Stokes development. The LMS emphasized both the enduring importance of human creativity and understanding and the extraordinary potential of new technologies when mathematical insight and technological innovation work together.

For ICIAM, there is an additional dimension to this discussion. Mathematics is not only a field in which AI is beginning to make discoveries. It is also one of the foundations on which modern science, engineering and technology are built. Applied mathematics translates our understanding of the world into models, algorithms, simulations and decisions. Here the central question is therefore not only whether AI can solve a mathematical problem, but whether AI-enhanced methods can be understood, validated and trusted

when they are used to predict physical systems, design technologies or support decisions. This gives mathematics a central role in the age of AI: as a discipline whose own methods of discovery are being transformed, as a source of the structure, analysis and guarantees needed to make AI reliable in science and engineering, and as a repository of powerful computational methods that can contribute directly to making new AI-enhanced approaches more efficient, reliable and interpretable.

A proof is not the end of a mathematical story

The history of mathematics repeatedly shows that solving a problem is often the beginning of a new chapter rather than the end of one. A proof may introduce a new technique that turns out to be more important than the theorem for which it was invented. A counterexample may reveal that a conjecture was formulated at the wrong level of generality. A failed argument may expose a hidden assumption and thereby lead to a better theory. The same principle should apply to AI-assisted mathematics. Knowing where an AI system failed, which approaches it explored unsuccessfully, and under which conditions its reasoning broke down can be scientifically as informative as knowing where it succeeded. Selective reporting of successful outcomes alone gives an incomplete picture of both the capabilities and the limitations of these systems. Greater openness about unsuccessful approaches will therefore benefit mathematical understanding, reproducibility and the responsible evaluation of AI for research.

This is particularly clear when a conjecture is disproved. The natural mathematical response to a counterexample is not simply to write “false” next to the conjecture and move on. We ask why the conjecture fails. Which hypothesis is missing? Is the counterexample exceptional, or does it reveal a general mechanism? Can the conjecture be reformulated so that the intended phenomenon is captured correctly? Is there a lower dimensional case that remains true? Is there a structural condition under which the original conclusion can be recovered? The counterexample therefore generates new mathematics.

Recent AI-assisted results make this point vivid. A machine may search an enormous space and discover a configuration that humans had not found. That is valuable. But once the configuration is found, the mathematical work changes character. We need to understand the mechanism, identify the essential ingredients, simplify the construction, connect it to existing theory and decide what new questions it creates. The speed with which an answer is found does not eliminate this work. In many cases it makes the work more urgent.

Terence Tao made a closely related point in his public lecture, *Mathematics in the Age of AI*, at the 2026 International Congress of Mathematicians. He asks the mathematical community to look beyond the narrow goal of producing solutions. A result must be verified, explained, understood, digested by the community and eventually incorporated into the definitive theory of a field. Tao calls attention to the transition from an era of proof scarcity to a possible era of proof abundance. If proofs can be generated faster than mathematicians can read, understand and integrate them, then the scarce resource is no longer the proof itself. It is expert human attention, judgement and understanding.

The value of mathematics does not reside solely in reaching the destination. It also resides in the map that is created on the way, in the concepts that make the landscape intelligible, and in the ability of the next generation to travel further.

Use artificial intelligence, but do not stop using your own intelligence.

Navier–Stokes, a spectacular result and a source of new questions

The events surrounding the Navier–Stokes Millennium Prize Problem in September 2026 provide an unusually clear illustration. OpenAI reported that an internal multiagent system, involving on the order of

10,000 concurrent agents, had produced after approximately 88 hours a proposed analytical construction for finite-time singularity in the three-dimensional incompressible Navier–Stokes equations under the forced formulation of the Clay problem. A formalisation and verification in Lean took a further 17 hours. The computational scale was extraordinary, with millions of messages exchanged and roughly 130 billion output tokens used in the Navier–Stokes effort.

The achievement did not arise in an intellectual vacuum. Important human work had already opened the route. Diego Córdoba and Luis Martínez-Zoroa developed a cascade of vorticity approach, with structures acting across successively smaller scales. Tristan Buckmaster and Levent Alpöge were pursuing closely related questions with AI assistance and obtained results for forced Euler and other fluid models. There is continuing discussion about priority, attribution and the precise relationship between these strands of work. For an international mathematical organization, the right lesson is not to adjudicate that discussion prematurely, but to insist on careful attribution, transparent disclosure and independent mathematical scrutiny.

The result is striking precisely because it demonstrates what can happen when deep human mathematical ideas are combined with computational search at a scale previously unavailable. The AI system can explore enormous numbers of possibilities, coordinate many parallel lines of reasoning and formalise the outcome. Human mathematicians provide the conceptual landscape in which that search becomes meaningful, the previous theory on which it builds, and the judgement required to understand what the result actually says.

The physical interpretation is equally instructive. The Navier–Stokes equations are mathematical models of fluid motion. OpenAI's announced construction starts with a smooth fluid at rest and applies a smooth external force while keeping finite energy, yet produces unbounded velocity in finite time. If confirmed through the normal processes of mathematical scrutiny, this answers the corresponding mathematical formulation. It does not, however, end our scientific interest in the equations. It immediately raises further questions. How physically realizable is the mechanism? Which features of the construction are essential? What does the example teach us about the relationship between the continuum model and real fluids? Which classes of forcing or initial data exclude such behaviour? Are there physically motivated restrictions under which global regularity can still be established?

These questions do not weaken the mathematical achievement. They give it scientific meaning. Applied mathematicians are especially familiar with this distinction. A model is not nature itself. A model is a mathematically formulated representation of selected features of nature. The task of applied mathematics includes deriving and solving models, but also determining their domains of validity, identifying their limitations, comparing them with observations and refining them when necessary.

The Navier–Stokes episode therefore offers a positive lesson. AI may close a particular mathematical question at unprecedented speed, while simultaneously opening a much larger scientific conversation. That is exactly why mathematicians must remain actively involved.

From proof production to mathematical understanding

If AI systems become capable of generating large numbers of correct mathematical arguments, the community will have to reconsider where it places its effort and its recognition. Verification remains essential, but verification alone is not understanding. A formally checked proof may establish that every logical step is correct, yet still leaves a human reader without a clear view of the central idea. Conversely, an insightful explanation may reveal why a theorem is true, why a particular hypothesis matters, and how the result connects to a broader theory.

Tao's analysis is useful here. Mathematical problem solving can be viewed as a sequence. First a solution is found. Then it is verified. Then it is explained. Then it is digested and accepted by the community. Finally, the result is canonicalized, meaning that it is placed in its natural setting, connected to neighbouring results, taught in courses and textbooks, and turned into part of the shared mathematical toolkit. AI may accelerate the first stages dramatically. The later stages remain deeply social and intellectual.

Indeed, the success of present AI systems depends heavily on the centuries of canonicalization already carried out by mathematicians. Large language models learn from a world in which mathematical definitions, notation, standard arguments, examples and theories have been organized into papers, books, lecture notes, software libraries and formal repositories. The apparent fluency of an AI system in mathematics is therefore built upon an enormous inherited structure of human mathematical thought.

This observation also clarifies the role of education. We should not train young mathematicians merely to compete with machines at tasks machines can perform faster. We should train them to ask good questions, recognize structure, judge significance, connect different areas, formulate models, challenge assumptions, explain ideas and understand why an argument works. These are not secondary skills. They are central to the future of mathematics.

Real intelligence is needed to make artificial intelligence work

There is a second sense in which mathematics is indispensable to artificial intelligence itself. Modern AI, including large language models, is not a mysterious technology that sits outside mathematics. Its foundations are profoundly mathematical.

At the most elementary level, a large language model represents tokens by vectors in high-dimensional spaces. Linear algebra governs the transformations of these vectors through layers of the network. The transformer architecture, which underlies modern large language models, uses attention mechanisms built from matrix products, inner products, normalization and the softmax function. Probability enters because the model represents conditional probability distributions over possible next tokens. Statistics enters through learning patterns from very large data sets. Optimisation is central to training, where enormous numbers of parameters are adjusted to reduce loss functions. Calculus and automatic differentiation provide gradients. Numerical analysis and high-performance computing determine whether these operations can be carried out accurately and efficiently at scale.

There is also deeper mathematics. Geometry helps us think about representations in high-dimensional spaces. Approximation theory asks what classes of functions networks can represent and how efficiently. Probability theory and statistical learning theory address generalisation and uncertainty. Dynamical systems viewpoints help interpret deep networks and continuous time models. Information theory is closely connected to representation, compression and prediction. Spectral analysis appears in the study of training behaviour and frequency bias. Optimisation theory is needed to understand loss landscapes and the behaviour of gradient based methods. Even AI tools that appear to communicate in ordinary language are therefore supported by a large mathematical infrastructure.

The statement “Real intelligence is needed to make artificial intelligence work” has two meanings. Human intelligence is needed to formulate worthwhile questions, design systems, select data, judge outputs, detect failure, interpret results and take responsibility for decisions. Mathematical intelligence is also needed inside the technology itself. Mathematics provides the language in which the models are built and, increasingly, the tools by which their limitations can be understood and overcome.

This becomes even more important when AI is used not for conversation or text generation, but as part of scientific and engineering computation.

Real intelligence is needed to make artificial intelligence work. [1]

Scientific Machine Learning: where mathematics and AI must meet

Scientific Machine Learning, usually abbreviated as SciML, is one of the areas in which the relationship between mathematics and AI is becoming most consequential. SciML combines data driven learning with mathematical models, physical laws, numerical simulation and scientific computing. Its promise is enormous. Surrogate models can replace expensive high-fidelity simulations in repeated evaluations. Operator learning can approximate mappings between entire function spaces. Physics-informed neural networks can incorporate differential equations into learning. Reduced-order models can compress systems with millions of degrees of freedom into models that run in real time. Hybrid approaches can retain trusted physics while learning only those components that are unknown or too expensive to model directly.

Industry has a strong need for these capabilities. Classical simulation has achieved extraordinary levels of fidelity, but high-fidelity computations may still require hours or days, substantial computing resources and specialist expertise. Modern engineering increasingly requires decisions in seconds or milliseconds, continuous updates from sensor data, exploration of huge design spaces, digital twins, inverse design, uncertainty assessment and real-time control. The opportunity is therefore not simply to make the old simulation workflow somewhat faster. It is to move from simulation toward decision intelligence. This creates a paradox for modern computational science: we can simulate an extraordinary range of complex systems with great fidelity, yet these simulations are often too slow or too specialised to be used at the moment when an engineering decision must actually be made. The challenge is therefore not only to accelerate simulation, but to make mathematical modelling and simulation an integrated, increasingly automated component of real-time prediction, design and decision making.

This transition is already clearly visible in industry. A 2025 McKinsey/NAFEMS report, aptly entitled *On the brink of a revolution? Engineering simulation in the age of AI*, describes engineering simulation as entering a new phase driven by AI and machine learning. Earlier McKinsey/NAFEMS survey results had already shown that around two thirds of respondents were experimenting with AI/ML-based simulation, although only a small fraction had yet deployed these methods at scale. The more recent study confirms rapid, though uneven, progress and points to faster time-to-market as an increasingly important driver for simulation. At the same time, the report shows that wider industrial adoption still faces important challenges: engineers need confidence in the reliability of AI-based results, organisations need the expertise to develop and assess these methods, and AI tools need to be integrated effectively into existing engineering workflows. The message is significant for applied mathematics: AI/ML is moving from experimental demonstrations towards industrial engineering practice, and the question is increasingly not whether these methods will be used, but how to make their use reliable, scalable and trustworthy.

These developments also help to clarify what is needed if AI/ML-based simulation is to fulfil its promise. Scientific Machine Learning creates opportunities to bring established simulation methods and data-driven learning together. Its strength lies in combining two complementary worlds. Scientific computing contributes differential equations, mathematical modelling, discretization, stability, convergence, conservation laws, error control and decades of accumulated scientific and engineering knowledge. Machine learning contributes flexible approximation, pattern recognition, large-scale data processing and new forms of adaptation. The challenge is to combine these strengths in methods that are not only fast, but also reliable, interpretable and trustworthy. Mathematical knowledge can guide learning by providing structure, reducing complexity and data requirements, and identifying properties that reliable scientific

[1] Phrase coined by Wil Schilders and used in his presentations for many years.

models should preserve. Learning, in turn, can exploit data and approximate components that are unknown, difficult to model explicitly or computationally expensive.

Why present machine learning is not yet enough

Achieving this combination is far from straightforward. Rapid industrial interest and impressive empirical performance should not be confused with scientific maturity. A network may perform extremely well on benchmark data while behaving unpredictably outside the training regime. Small perturbations may produce unexpectedly large changes. Long-time predictions may drift or become unstable. A model may violate conservation laws, symmetries, monotonicity, positivity, causality or other structural properties that are known in advance. Uncertainty may be poorly quantified. The choice of network architecture and training procedure may rely heavily on trial and error, while training itself can be computationally expensive, ill conditioned and highly sensitive to choices for which a solid mathematical understanding is still lacking.

One important example is spectral bias, sometimes called the frequency principle. Neural networks trained by gradient-based methods often learn low frequency, smooth components before high frequency, oscillatory components. In many scientific problems, however, the high frequency behaviour is precisely where important phenomena occur, for example in multiscale systems, wave propagation, boundary layers and highly oscillatory partial differential equations. This is not merely an implementation inconvenience. It is a mathematical property of the approximation and training process, and it calls for mathematical remedies.

Another issue is implicit bias. Learning algorithms do not explore all possible solutions neutrally. The architecture, parametrization, initialization and optimisation method favour some representations over others. Sometimes this bias is beneficial, but in scientific applications we need to understand it. If a model systematically favours smooth solutions, particular frequencies, low complexity representations or correlations present in the training data, then its predictions can be misleading when the underlying physics lies outside those preferences.

Physics-informed neural networks illustrate both the promise and the difficulty. Adding the residual of a differential equation to a loss function can strongly guide learning, but placing physics in a loss function is not the same as guaranteeing that the final solution satisfies the physical structure. Competing loss terms can be difficult to balance, training can become ill conditioned, causality may not be respected automatically, and long-time or high-dimensional problems remain challenging. The lesson is not that PINNs are a failure. The lesson is that more mathematics is needed to turn an excellent research concept into a robust scientific technology.

The same applies to operator learning. DeepONets, Fourier neural operators and related approaches can learn mappings from input functions to solution fields and can make repeated evaluations extraordinarily fast. Yet questions of generalisation, stability, physical consistency, data efficiency and behaviour outside the training distribution remain. Functional analysis, operator theory, spectral methods and approximation theory are not peripheral to these questions. They are central.

Mathematics can provide the missing structure

Mathematics has spent centuries identifying structures that distinguish meaningful scientific solutions from arbitrary functions. Conservation of mass, energy or momentum, symmetries, invariants, positivity, monotonicity, causality, Hamiltonian and port-Hamiltonian structure, geometric constraints and differential algebraic relations are examples. In classical numerical analysis we learned that preserving such structure can be essential for stability, long-time accuracy and physical fidelity. This knowledge provides a powerful foundation for constructing reliable AI-based scientific models.

Known laws and mathematical constraints provide valuable prior information for learning. Incorporating such knowledge into an AI model can reduce the space of admissible solutions and focus learning on what is genuinely unknown. Structure-preserving neural networks, Hamiltonian neural networks, symplectic networks, equivariant networks and physics-motivated activation functions are examples of this philosophy.

There is also an important consequence for data efficiency. Purely data-driven models may require very large amounts of training data to learn behaviour that is already known from mathematics or physics. If conservation laws, symmetries, governing equations, known constitutive relations or other structural information are incorporated into the model, the learning problem can become much smaller. Data are then needed primarily to identify the unknown or inadequately modelled components, rather than to rediscover the entire physical system from examples. This can reduce data requirements substantially, which is especially important in scientific and industrial applications where high-fidelity simulations, experiments or measurements may be expensive or difficult to obtain.

Mathematical analysis often provides rigorous properties that a model must satisfy in order to be physically meaningful and numerically reliable. These may include conservation laws, monotonicity, positivity, symmetry, stability, causality, or other structural conditions. Such knowledge does not cease to be relevant when a conventional model is replaced, or partly replaced, by an artificial neural network or another AI/ML model. On the contrary, these learned models should respect the same proven mathematical properties. If they do not, they may produce apparently accurate results on training or benchmark data while leading to unphysical behaviour, instability, or convergence problems when embedded in a numerical simulation. An important task for mathematics is therefore to identify the essential properties of the underlying problem and to ensure, preferably by construction, that AI/ML models preserve them.

This is a broader lesson for scientific AI. Structure should be designed, not merely hoped for. Guarantees should be established where possible, not inferred solely from successful benchmark tests. Robustness should be a design objective. Uncertainty should be quantified. Mathematical knowledge can transform a black box that interpolates data into a scientific model whose behaviour can be analysed and trusted.

Mathematics can also change how AI learns

The contribution of mathematics is not limited to imposing physical constraints. It can change the learning algorithms themselves. Spectral bias suggests the use of localization, multilevel methods and basis functions with compact support so that globally high-frequency phenomena become easier local problems. Domain decomposition, familiar from numerical analysis, can turn one difficult global approximation problem into many smaller problems that can be solved in parallel.

Training can also be reconsidered. Gradient descent is extraordinarily successful, but it is not a law of nature. For some classes of scientific problems, random feature methods, least squares, classical linear algebra and ordinary differential equation solvers can replace large nonlinear optimisation procedures. Recent work on Frozen PINNs is one example, separating space and time, freezing spatial features and evolving coefficients causally in time. The resulting methods can be faster, more stable and more faithful to the temporal structure of differential equations.

Network architecture itself need not always be determined by trial and error. For dynamical systems, ideas from state space theory, eigenvalue analysis, model order reduction, Koopman operators and system identification can guide or even determine the topology of a neural representation. Dynamic neurons can incorporate differential equations directly. Reduced-order modelling can compress both physical models and neural models. In this way, decades of knowledge from numerical linear algebra, dynamical systems and control can become part of the design of new AI methods.

This is what we mean by Mathematics Guided Artificial Intelligence. The aim is not to constrain AI until it can do only what we already know. The aim is to use mathematical structure to remove unnecessary guesswork, improve efficiency, guarantee essential properties and focus learning on the genuinely unknown part of the problem.

The mathematical toolbox provides powerful opportunities for AI

Computational mathematics has developed an extraordinary body of knowledge since the advent of digital computers. Over many decades, applied mathematicians have created highly efficient and increasingly sophisticated methods for solving differential equations, optimisation problems, inverse problems and large-scale systems, together with techniques for adaptivity, decomposition, preconditioning, multilevel computation, uncertainty quantification and model reduction. Many of these methods are fast, robust, interpretable and supported by mathematical analysis. They provide a powerful foundation for developing AI-enhanced approaches that combine computational efficiency and mathematical structure with the flexibility and data-processing capabilities of learning. One of the great achievements of computational mathematics has been to replace brute-force computation by algorithms that exploit mathematical structure; the interaction with AI creates new opportunities to take this principle further.

Indeed, some of the computational principles now associated with AI-enhanced simulation have close parallels in established applied mathematics. Model order reduction is a striking example. Its classical offline/online paradigm invests substantial computational effort offline, using high-fidelity models, simulations or other structural information to identify a low-dimensional representation and construct a reduced model. Once this has been done, the reduced model can be evaluated online extremely rapidly, often in real time. Many AI-based surrogate approaches follow a conceptually similar pattern: expensive high-fidelity data is generated offline, a lower-dimensional or learned representation is constructed, and a computationally inexpensive model is then used online for prediction, optimisation or control.

These parallels point to substantial opportunities for combining established computational mathematics with new AI capabilities. Classical algorithms can provide efficiency, structure, guarantees and interpretability; learning can add flexibility, exploit data and approximate components that are difficult to model explicitly. In some applications a classical reduced-order model may remain the best solution; in others a learned surrogate may be superior; and increasingly the strongest approach may be a hybrid of the two. The criterion should not be whether a method is labelled “AI”, but whether it is accurate, efficient, robust, understandable and appropriate for the scientific or industrial task.

From simulation to discovery

Building on this combination of established computational mathematics and new AI capabilities, the role of simulation itself is beginning to change. Traditionally, simulation has been used primarily in a forward direction. We specify a model and parameters, solve the equations and predict what the system will do. Increasingly, the goal is inverse and exploratory. We want to discover materials with desired properties, infer unknown parameters, explore enormous design spaces, control complex systems and connect phenomena across scales. The latter is a particularly important challenge: many scientific and industrial systems span scales from molecular or microscopic processes to components, complete devices and systems-of-systems. Mathematics provides essential tools for bridging these scales, including multiscale analysis, homogenisation, model reduction and structure-preserving approximation, while AI and data can help identify and represent interactions that are difficult to derive explicitly.

Hybrid approaches are especially powerful because mathematical models, computational methods, experimental information and learning can contribute complementary strengths. A high-fidelity model may provide accurate data in a limited number of expensive simulations. Reduced-order models and surrogates can then make millions of evaluations possible. Experimental and sensor data can correct model deficiencies. Uncertainty quantification can tell us where predictions are reliable and where new information is needed. AI can propose candidates, while mathematical models, analysis and data help assess whether they are meaningful and reliable.

The emerging paradigm is a loop in which mathematical models, data, experiments, analysis and learning continually inform one another. In materials discovery, digital twins, energy systems, mobility, health technology and advanced manufacturing, this loop may allow us to move from expensive isolated computations toward continuous scientific discovery and decision support.

For ICIAM, this development is particularly important because applied and industrial mathematicians are accustomed to connecting mathematical methods with scientific, engineering and industrial challenges. They understand modelling, numerical methods, optimisation, uncertainty, algorithms and applications, and can translate scientific questions into mathematical formulations. The age of AI therefore creates not less need for applied mathematics, but more.

A constructive agenda for the mathematical community

The appropriate response to AI is neither resistance nor uncritical enthusiasm. It is active engagement. Mathematicians should use these tools, study them, improve them and help shape the conditions under which they are used. We should be ambitious about what human and artificial intelligence can achieve together, while remaining rigorous about correctness, attribution, reproducibility, transparency and scientific meaning.

The Leiden Declaration provides an important framework. Human authors must remain responsible for the correctness of work produced with automated assistance. Tools used and computational resources should be disclosed. Prior human contributions should be traced and credited carefully. Peer review and independent scrutiny remain essential. Mathematical organizations should build expertise, develop policies and participate actively in public discussion rather than allowing the norms of mathematical research to be set elsewhere.

Tao's message adds another dimension. The community should decide what it values before external metrics decide for us. If the easiest quantity to measure is the number of problems solved, we should resist the temptation to identify that number with mathematical progress. Explanation, theory building, mentoring, synthesis, verification and canonicalization are equally important parts of the mathematical enterprise, even if they are harder to count.

ICIAM brings a specifically applied and industrial perspective. Mathematics is not only a domain in which AI can demonstrate its power. Mathematics has a central role in making AI itself more reliable, efficient, understandable and useful. This is particularly urgent in scientific machine learning, where decisions may affect engineered systems and eventually society. Robustness, stability, error control, physical consistency and quantified uncertainty are not optional refinements. They are prerequisites for trust.

There is also a broader challenge for the scientific method itself. As AI systems become active participants in research, generating hypotheses, writing code, constructing models, analysing data and interacting with other computational tools, reproducibility will require more than recording the final result. We will increasingly need transparent provenance of the computational and scientific process: which models and tools were used, which assumptions and data entered the analysis, how results were verified, and where human judgement intervened. Mathematical and scientific organizations should help develop

appropriate standards for documenting such AI-assisted research, so that results remain open to independent scrutiny, attribution and reproduction.

For ICIAM and its member societies, this creates a concrete responsibility. We should encourage research and education at the interface of mathematics, AI and scientific computing, while ensuring that established mathematical and computational knowledge remains an integral part of this development. We should bring mathematicians into interdisciplinary and industrial AI initiatives from their inception, and promote standards for validation, reproducibility, uncertainty and mathematical consistency. We should create opportunities for researchers, industry and AI developers to exchange experience, including examples where AI works spectacularly well and examples where it fails. At the same time, we should encourage the intelligent combination of new AI methods with the powerful modelling, numerical and computational techniques developed over many decades. We should also be attentive to equitable access to the computational and AI infrastructure increasingly required for frontier mathematical and scientific research, so that these developments broaden rather than narrow participation in science worldwide. And we should ensure that the next generation of applied mathematicians is equipped not only to use AI tools, but also to understand, analyse and improve them, and to recognize when established mathematical methods, AI approaches or a combination of both provide the best solution.

Human insight, mathematics and AI, together

The central message is therefore positive. Artificial intelligence can accelerate mathematical and scientific discovery. It can search spaces that humans cannot search, coordinate large numbers of attempts, reveal patterns, suggest conjectures, generate counterexamples, formalise arguments and accelerate simulation. These capabilities will become increasingly important.

But speed is not the same as understanding. A solution creates value when we understand why it works, how it relates to what was known before, what its limitations are, and which new questions it opens. A scientific prediction creates value when we know when it can be trusted. A learned model becomes a scientific instrument when its structure, uncertainty and failure modes are understood sufficiently well for the task at hand.

Humans bring curiosity, judgement, responsibility, creativity, scientific context and the ability to decide which questions matter. Mathematics brings abstraction, structure, rigor, models, algorithms, quality guarantees and a language for understanding complexity. AI brings new powers of search, approximation, synthesis and computation. The emerging paradigm is therefore a partnership of human insight, mathematics, data and AI.

The challenge is not simply to make artificial intelligence more powerful. It is to make it a trustworthy partner in mathematical and scientific discovery. This requires us to keep thinking, to keep questioning, and to bring the full strength of mathematics to the development and use of AI.

If we succeed, the outcome will not be mathematics without mathematicians. It will be a richer form of mathematical practice in which human insight and artificial intelligence reinforce one another, producing not only faster answers, but deeper understanding, more reliable science and new possibilities for society.

AI can accelerate discovery. Mathematics can provide understanding and trust. Together, human and artificial intelligence can open new directions for mathematics, science, engineering, industry and society.

References

1. **Collective of 25 Fields Medallists**, *A Severe Misalignment of AI in Mathematics*, 11 September 2026. Full statement and list of initial signatories: Math and AI, A Severe Misalignment of AI in Mathematics. Also published in English by *Le Monde* as *25 Fields medalists warn: "The goals of the AI companies and the goals of the mathematical community are severely misaligned"*, 11 September 2026, available at Le Monde, English version.
2. **Leiden Declaration on Artificial Intelligence and Mathematics**, 2 June 2026, endorsed by the International Mathematical Union. DOI: 10.5281/zenodo.20302944. The declaration, signatories and endorsements are available at Leiden Declaration on Artificial Intelligence and Mathematics; the archived publication is available at Zenodo record and PDF.
3. **London Mathematical Society**, *Navier–Stokes Equations Breakthrough*, 9 September 2026. Statement by the London Mathematical Society, including comments by LMS President Mark Chaplain on human mathematical understanding and the role of AI in discovery. Available at London Mathematical Society, Navier–Stokes Equations Breakthrough.
4. **Terence Tao**, *Mathematics in the age of AI*, 2026, arXiv:2608.16753. Essay based on Tao's public lecture at the International Congress of Mathematicians 2026. Available at arXiv:2608.16753.
5. **OpenAI**, *On the Navier–Stokes Millennium Prize Problem*, 8 September 2026. The announcement describes the proposed analytical construction, multi-agent research process and Lean formalisation and verification. Available at OpenAI, On the Navier–Stokes Millennium Prize Problem.
6. **McKinsey & Company and NAFEMS**, *On the brink of a revolution? Engineering simulation in the age of AI*, 12 February 2025. The report discusses the adoption of AI/ML in engineering simulation, including the finding that two-thirds of survey respondents had used AI/ML simulation while only 5% were using it at scale. Available at McKinsey, Engineering simulation in the age of AI.
7. **McKinsey & Company**, *The next horizon: Engineering simulation for an AI-powered world*, 27 March 2026. The report examines how AI is changing engineering simulation and moving its impact from individual tasks towards engineering activities at the systems level. Available at McKinsey, The next horizon: Engineering simulation for an AI-powered world.

OpenAI recently started a math advisory group, see https://openai.com/index/advisory-group-on-mathematics-and-ai/

Statements, responses and discussions by ICIAM member societies

Several ICIAM member societies have responded to the recent developments in AI and mathematics, particularly the September 2026 Navier–Stokes announcement.

American Mathematical Society (AMS), 8 September 2026.
The AMS welcomed the reported Navier–Stokes progress while emphasizing the long history of human mathematical work underlying it.
x.com/amermathsoc/status/2097380478349463939

London Mathematical Society (LMS), 9 September 2026.
The LMS emphasized human creativity and mathematical understanding alongside the potential of AI-assisted discovery.
www.lms.ac.uk/news/navier-stokes-equations-breakthrough

European Mathematical Society (EMS), 10 September 2026.
The EMS highlighted questions of human–AI collaboration, attribution, authorship, open science and access to AI research tools.
euromathsoc.org/news/ems-statement-on-recent-navier-stokes-announcement-225

Unione Matematica Italiana (UMI), September 2026.
The UMI created a collection of reflections on AI and mathematics, including contributions responding to the Navier–Stokes development.
umi.dm.unibo.it/intelligenza-artificiale-e-comunita-matematica/

Société Mathématique de France (SMF), September 2026.
The SMF Board addressed the development and questions concerning attribution, research practices and the impact of AI on mathematics.
smf.emath.fr/index.php/actualites-smf/position-bureau-smf-navier-stokes-et-openai

Real Sociedad Matemática Española (RSME), 15 September 2026.
The RSME stressed transparency, attribution, independent verification and the continuing importance of human mathematical understanding.
rsme.es/declaracion-de-la-real-sociedad-matematica-espanola-rsme-ante-los-recientes-avances-en-inteligencia-artificial-y-la-resolucion-de-navier-stokes/

Institute of Mathematics and its Applications (IMA), September 2026.
The IMA announced a Virtual Maths Tea, *Navier-Stokes and OpenAI: A sign of what's to come?*, for 30 September, discussing AI, mathematical research and attribution of AI-assisted discoveries.
LinkedIn: Institute of Mathematics and its Applications