

2nd IMA Congress: AI Unlocked: Innovation, Insight and Impact

Abstracts

17th – 18th September 2026



INSTITUTE OF
**MATHEMATICS &
ITS APPLICATIONS**

Contents

Gauss is Wrong Part 2: AI, Mathematics, Strategy and Culture.	3
Imaging in the Age of AI: Combining Physics, Mathematics and Data.....	4
Precision by Numbers: The Mathematics of Robotic Knee Surgery	5
What (and why) is graph-based learning?	6
Should you Listen to your Knees ? Using A.I. with Acoustic Signals to Monitor the Condition of Knee Joints.....	7
AI Ethics -- Mathematical Challenges and Progress	8
From Equations to Enterprise: Legal Foundations for Commercialising AI	9
Refinement and the Universal Refinement Principle: Theory, Structure Emergence, and Cross-Domain Performance	10
Mathematics Education in the Age of AI.....	11
AI for Science.....	12
Engineering Living Brain Circuits	13
Confidence Is Not Correctness, An Algebraic Foundation for Trustworthy AI in Regulated Industries	14
Geometric Properties of Sentence Embeddings for Invoice Categorisation	15
How Certain is the AI Neural Network Conformal Inference for Financial Document Classification	16
A Quarter Century of AI Research in Defence and Security	17
The Underpinning Trifecta: Mathematics, AI, and Future Hardware	18
AI Unlocked in Autonomous Sensing: Bridging Foundation Models, Transformers, and Bayesian Tracking	19
Understanding The Confidence with which a Machine Learning Model Classifies an Input Datum	20
Trajectory-Free Parameter Estimation in Interacting Particle Systems via Topological- Frequencies.....	21
Scientific Machine Learning	22

Gauss is Wrong Part 2: AI, Mathematics, Strategy and Culture.

Nira Chamberlain OBE – AtkinsRéalis

The phrase "culture eats strategy for breakfast" means that an organization's culture—its shared values, beliefs, and behaviours—is more influential in determining success than its strategic plan. Even the best-laid strategies will fail if the people within the company don't have a culture that supports their execution. This quote, famously attributed to management consultant [Peter Drucker](#), highlights that a company's underlying culture drives daily actions and ultimately determines whether a strategy succeeds or fails.

When a company culture is not founded on mathematics and science, what happens to the AI Strategy?

What lessons can we learn from the world of mathematics, such as the monte-carlo optimisation and ant colonisation optimization algorithms?

In this talk Professor Nira Chamberlain will explore this and the interesting intersection between AI, Mathematics, Strategy and Culture.

Imaging in the Age of AI: Combining Physics, Mathematics and Data

Carola Bibiane Schönlieb – University of Cambridge

Modern imaging is central to science, medicine and technology, from medical scans and microscopy to environmental monitoring and cultural heritage. Yet images are rarely direct observations: they are reconstructed from incomplete, noisy and indirect measurements, and their reliability depends on the mathematical and computational models used to produce them.

Artificial intelligence is transforming this landscape. Deep learning methods can deliver striking improvements in reconstruction, enhancement and interpretation, but they also raise important questions about trust, robustness, generalisation and the risk of producing plausible but unsupported features. In this talk, I will discuss how mathematical imaging is evolving in response to these opportunities and challenges. I will highlight approaches that combine data-driven models with physical measurement processes, variational principles and mathematical structure, and show how this combination can lead to more reliable and interpretable imaging methods. The central message is that the future of imaging lies not in replacing physics and mathematics by AI, but in integrating physics, mathematics and data.

Precision by Numbers: The Mathematics of Robotic Knee Surgery

Akash Sharma – Royal Orthopaedic Hospital Birmingham

Knee replacement has been described as the operation of the decade — among the most common and successful procedures in modern medicine. Yet around one in five patients remains dissatisfied, the result of a combination of non-surgical and surgical factors. Strikingly, both are increasingly governed by mathematics. Among the surgical contributors are millimetre- and degree-level errors in implant positioning and alignment, and it is here that robotic-assisted surgery has emerged to help — underpinned by a dense layer of mathematics that neither patient nor surgeon usually sees.

This talk traces the mathematical pipeline from anatomy to controlled robotic action. Image-based systems begin with CT or MRI reconstruction, segmenting bone into a three-dimensional model; imageless systems instead build that model intraoperatively, digitising the bone surface and fitting statistical shape models to a sparse point cloud. Registration then aligns virtual and real anatomy using rigid-body transformations and iterative-closest-point algorithms, fusing optical and inertial tracking. Constrained optimisation sets implant size and position and weighs alternative alignment strategies, each a different mathematical trade-off, while computer vision preserves sub-millimetre accuracy and haptic boundary-control mathematics keeps the instrument within a safe envelope.

Crucially, the non-surgical determinants of outcome — metabolic, psychological and movement factors — are themselves a compound of mathematics, increasingly captured by predictive models and machine learning trained on large registry datasets. Healthcare sharpens this Congress's central question: surgical AI must be accurate, explainable, validated and safe. I will discuss how robotic surgery earns clinical trust — through transparent geometric reasoning, hard safety constraints and registry evidence — and where the field must still mature.

Mathematics, ultimately, is not a backdrop to AI in healthcare but the foundation on which precision, safety and patient trust are built.

What (and why) is graph-based learning?

Jeremy Budd, Assistant Professor – Birmingham University

In many data science applications, we are not only interested in the data in itself, but also (or even more so) in how the data relate to each other. A rich mathematical representation of interconnected data is that of a graph or network, which sees data as nodes (data points) linked by edges (describing their relationship with each other). Graph-based learning is all about using this representation to understand and make inferences from this sort of data. In recent years, graph neural networks (GNNs) have become the most well-known example of this idea, but in this talk I will highlight how there is much more to the area than just GNNs, and how in graph-based learning valuable practical tools meet deep mathematical theory.

Should you Listen to your Knees ? Using A.I. with Acoustic Signals to Monitor the Condition of Knee Joints.

Ivan Vatolik¹, Gordon Hunter², Mike Everington², Nicola Swann³, Chinedu Mbachu¹,

Andy Augousti¹

¹ School of Engineering, Faculty of Engineering, Computing and the Environment, Kingston University, U.K.

² School of Computer Science and Mathematics, Faculty of Engineering, Computing and the Environment, Kingston University.

³ School of Biosciences, University of Surrey, U.K.

Knee conditions such as osteoarthritis are major causes of disability and immobility worldwide, especially amongst older people and former participants in “high impact” sports such as volleyball, basketball, javelin throwing and fast bowling in cricket. Conventional methods for detecting and monitoring these conditions tend to require either invasive surgery (which can itself lead to problems) or require expensive equipment such as MRI scanners. However, in principle, monitoring Acoustic Emissions (AEs) – sounds produced by knee joints during exercise – could provide a low cost, non-invasive method of detecting and monitoring such conditions. In this paper, we describe a pilot study, combining AE and biomechanical data obtained for 30 volunteer participants (15 male, 15 female) of various ages whilst they underwent gentle exercises. The data was analysed using A.I. techniques such as Convolutional Neural Networks (CNNs) and introduced novel “collage plots” as a way to integrate both acoustic and biomechanical data into image representations for each participant. Using age as a proxy for knee condition, the CNNs proved successful classifiers for this based on each participant’s biomechanical – acoustic collage plot. Whilst the approach still needs further investigation- particularly using data from people with known clinical conditions – it does show promise as a possible method for early detection and monitoring of potentially debilitating knee conditions.

AI Ethics -- Mathematical Challenges and Progress

Richard Pinch – IMA Council Member and AI SIG

AI system implementation are subject to ethical, professional, regulatory and legal requirements, imposing demands in such areas as safety, security, privacy and fairness. Development had proceeded at a pace that leaves gaps in the theory, techniques, tools and tradecraft which need to be filled. This gives rise to mathematical challenges, many of which remain unexplored. In this talk I will illustrate what seem to be some important research gaps, and also describe progress which has been made, focussing on the area of explainable AI and techniques from game theory that can be used to support explainability and interpretability.

From Equations to Enterprise: Legal Foundations for Commercialising AI

Sarah Coe, Partner, Collyer Bristow LLP.

Most delegates at the AI Congress are creating extraordinary things: models, datasets and methods. However, the uncomfortable reality is this: innovation alone does not create value; control does.

In a commercial context, that control is largely a function of legal structure.

Every successful AI business stands on three core legal pillars: ownership, control and trust. What do you actually own? What rights have you granted to others? And can third parties safely and confidently deploy what you have built?

This presentation will explore the legal foundations underpinning successful AI commercialisation, offering a practical and commercially grounded framework for researchers, founders and technical leaders. It focuses on these three pillars as the determinants of whether AI innovation can translate into sustainable and scalable value.

First, the session examines the nature of AI assets, including code, models, training data and outputs, highlighting the frequent disconnect between perceived and actual sources of value. Particular attention is given to data provenance and usage rights, which are often the most critical, yet least understood, elements.

Secondly, it considers how value is realised through commercial arrangements, with a focus on licensing structures, Software as a Service models, and the allocation of rights in data and model improvements. It demonstrates how contractual frameworks, rather than intellectual property rights alone, drive scalability and revenue generation.

Finally, the presentation addresses the increasing importance of governance and regulatory alignment in enabling real-world deployment, particularly in regulated sectors where trust, auditability and data integrity are essential.

Subject to time, the session will also adopt an investor lens, highlighting key legal risks arising in due diligence.

I shall conclude with audience discussion on practical legal challenges encountered in building AI businesses.

Refinement and the Universal Refinement Principle: Theory, Structure Emergence, and Cross-Domain Performance

Lance Charles Carr

Independent Researcher, Nottingham, UK

Email: lance.charles.carr@outlook.com

Statistical learning systems often fail to generalise when distributions, tasks, or representational structures shift across domains. Existing approaches—such as transfer learning, domain adaptation, and large foundation models—typically rely on domain-specific heuristics that do not provide a principled account of how knowledge should transform between heterogeneous settings. This work introduces the concept of refinement, a formal, domain-agnostic process for transforming representations while preserving structural invariants, and the Universal Refinement Principle (URP), a unifying schema whose application shows how statistical models can be decomposed, re-expressed, and recomposed without loss of functional coherence.

Our objective is to establish a theoretically grounded mechanism for cross-domain generalisation that operates by structural alignment rather than surface-level adaptation. We develop a mathematical formulation of refinement as a sequence of representation-preserving transformations and formalise URP as a recursive optimisation framework that maintains informational sufficiency while eliminating domain-specific artefacts. To illustrate this, we train a neural network on a source domain and show that, despite the domains being mathematical analogues, its accuracy collapses when interrogating related domains B, C, and D.

Applying URP, we extract the shared underlying structure across all domains and recompile the model, demonstrating substantial improvements in cross-domain accuracy. We also visualise the refinement process, showing graphically how the unified structure emerges as domain-specific distortions are removed. Finally, we quantify the operational benefits of refinement by demonstrating that a locally computed, structurally aligned model can replace repeated cloud-based retraining, yielding significant real-world cost savings.

Our results show that refinement improves robustness to distributional shift, enhances interpretability by exposing structural correspondences, and increases transfer performance across tasks. We conclude that refinement provides a universal, theoretically principled alternative to current domain-specific strategies, positioning it as a foundational tool for the next generation of statistical learning systems.

Mathematics Education in the Age of AI

Stephen Lynch – Loughborough University

There is growing value in helping students and pupils understand the mathematical ideas that underpin modern AI systems. In this talk, I will share practical experiences of introducing neural networks and machine learning concepts to Foundation-level students and undergraduates from non-Computer Science disciplines. I will demonstrate how concepts such as neurons, weights, prediction, and learning can be taught without advanced mathematics, with particular emphasis on an intuitive explanation of the backpropagation algorithm and how neural networks learn through error correction. The talk will end with the example presented to AS-Level Mathematics pupils as part of this congress. School pupils can explore stock market prediction using deep learning and TensorFlow by changing just five parameters in a Python program. I will argue that AI offers a powerful and engaging vehicle for developing mathematical thinking and AI literacy, enabling students to engage with sophisticated concepts much earlier than is often assumed.

AI for Science

Marika Taylor – University College London

In this talk we will explore how AI is being used to accelerate scientific research. AI companies and universities all around the world are developing AI scientists, to automate and accelerate the research process. We will explain the various approaches being used and discuss what progress has been made. A key element in AI for science is understanding the scientific discovery process, for example, how laws of Nature are inferred. Leading companies in AI such as Anthropic have been founded and are led by mathematical scientists. We will discuss how Anthropic's research on AI for science and frontier AI draws on mathematical sciences research, especially the founders' experience of working on the laws of Nature.

Engineering Living Brain Circuits

Paul Roach – Loughborough University

Using principles from biomaterials, microfabrication, and cell surface interactions, neuronal growth can be precisely guided through tailored chemistry and topographical cues. This talk will explore how control across multiple length scales, from molecular ligands to mesoscale architectures, enables reproducible assembly of functional neural networks in vitro. I will present examples showing how subtle changes in biomaterial properties can significantly influence the engineering of neural circuitry.

These design strategies are implemented within microfluidic platforms to create brain-on-a-chip systems that emulate key aspects of the neural microenvironment. Such systems offer enhanced capability to observe cellular-level function in living brain tissue and can be developed using human stem cells to model the human brain, providing a superior alternative to animal models. Recent advances in fabricating bespoke cellular architectures will be highlighted, including patterned axon guidance, compartmentalised connectivity, and mixed neuronal–glial co-cultures that support long-term activity and plasticity. Coupling these living networks to high-density microelectrode arrays enables closed-loop stimulation and real-time recording.

Beyond modelling, these brain-on-a-chip platforms are now being adapted for biological neuromorphic computing. Adaptive living networks provide hardware-efficient solutions for learning, inference, and control. Hybrid systems that link biological circuits to spiking silicon offer new strategies for energy-efficient computation and for understanding brain-like processing. By merging materials science, neurobiology, microfabrication, and machine intelligence, these innovations promise transformative impacts on neurotechnology, accelerate therapeutic development, and deepen our understanding of complex neural processes. Ultimately, harnessing the inherent computational power of biological systems could redefine intelligent technologies, creating platforms that think and adapt with living matter. Such systems point toward ethical, low-resource, and scalable innovation worldwide.

Confidence Is Not Correctness, An Algebraic Foundation for Trustworthy AI in Regulated Industries

Dr Anjulika Salhan FIMA, CSci Cmath

AI can be confident and wrong at the same time. In regulated industries, that gap costs real money. HMRC attributes £21.5 billion of the UK's annual tax gap to error and carelessness at the bookkeeping step, where transactions are categorised against the chart of accounts.

This talk introduces an algebraic validation framework that constrains AI outputs to satisfy structural conditions drawn from five centuries of mathematical work from Pacioli's double-entry system in 1494 through the modern theory of Abelian groups. It shows why algebra, not statistical confidence, is the right foundation for trustworthy AI in regulated domains.

The framework is not a research proposal. It runs in production on a live UK platform serving over 1,600 SMEs since October 2025, and is, to the best of public knowledge, the first AI system of its kind for the SME segment with structurally validated automation.

The talk presents the mathematics, the architecture, and the case for algebraic validation as part of UK national infrastructure for trustworthy AI.

Geometric Properties of Sentence Embeddings for Invoice Categorisation

Emma Ceccherini

AI is increasingly being adopted across industries to automate routine tasks, including invoice categorisation. This process is at the base of both financial reporting and tax compliance, and improving it would benefit businesses and regulators. However, it is still unclear how pretrained SLMs represent financial documents in their internal representations -embeddings- and how that organisation interacts with input formatting and pooling strategies.

We study the embedding geometry of two models, Sentence-BERT (all-mpnet-base-v2) and DeBERTa-v3, on a real-world dataset of invoices labelled with accounting categories. We compare three input formulations: vendor-only metadata (A), raw concatenated text including line items (B), and a human-readable structured input with field labels (C). We report three main findings: (1) DeBERTa's embeddings collapse under longer inputs (B/C), whereas S-BERT recovers semantically rich representations, resulting in S-BERT significantly outperforming DeBERTa in classification. (2) Unsupervised clustering of embeddings recovers vendor identity more strongly than accounting category, across all models and inputs. (3), we find that, counterintuitively, the human-readable input C displays poorer geometric properties and achieves lower classification performance than the raw input B for both models.

Overall, we find that the smaller S-BERT with raw textual input is the most geometrically interpretable and accurate combination. Our results reinforce the value of sentence-level over token-level embeddings, and offer guidance on automatic invoice classifications and describe how input formatting, training and pooling shape representation geometry in SLM.

How Certain is the AI Neural Network Conformal Inference for Financial Document Classification

Student Name: Neva Fradd, Da

Department of Mathematics

University of Bristol

University Supervisor: Prof Daniel Lawson

Industry Supervisor: Dr Anjulika Salhan

Neva Fradd, Emma Ceccherini, Daniel Lawson (University of Bristol) and Anjulika Salhan (Systems Holdings)

Large Language Models are a powerful tool for classification. Whilst performance is often good, classification accuracy will vary between documents of different types, from almost perfect to "hallucination" that there is any information at all. The most promising solution is "Conformal inference", a framework to obtain calibrated uncertainty from such a neural network, i.e. if we say that a datapoint is classified with 95% certainty, then we will be correct 95% of the time.

Conformal inference is easy to set up to be correct when averaged over all data, but becomes more challenging when we want to also be correct for a specific hard case. The high-dimensionality of neural network embeddings mean that it will sometimes be the case that we don't have a good representative set of training data. It is the "geometry" of the space that tells us this - are we locally somewhere "flat" and can we average over similar data? Or is this datapoint unusual, such as at a "curved" place in the space which doesn't look like an average of similar data? This work explores whether we can correctly predict which data are hard to classify, and whether understanding the internal geometry of neural network embeddings helps address this problem.

A Quarter Century of AI Research in Defence and Security

Richard Lane – QinetiQ

Artificial Intelligence has long played a role in defence, with early breakthroughs depending on deep domain expertise. Today, the rise of open-source tools and commercial AI success has shifted the landscape toward more general-purpose technologies that can be adapted for defence needs. This session explores the evolution of AI in defence over the past 25 years, examines cutting-edge research on managing uncertainty in AI systems, and offers a forward-looking perspective on emerging opportunities and challenges.

The Underpinning Trifecta: Mathematics, AI, and Future Hardware

Dr. David Snee

Defence Science and Technology Laboratory (Dstl)

The UK Ministry of Defence has highlighted a set of emerging, critical technologies that are redefining the way that UK Defence operates. These include, but are not limited to, AI & Data Science, Robotics & Autonomy, Cyber, Quantum, and Space-based capabilities. These technologies work together, offering significant leaps in operational capability, and are critically reliant on access to the compute hardware on which they run. Of the emerging technologies, AI is arguably the most immediate and transformative, capturing public attention daily. However, the most successful and widely adopted AI models pose significant economic and societal challenges due to their immense demand for computing power and energy. To fully harness an underpinning technology such as AI, we need to be smarter in our use and diversity of the underpinning compute.

In this talk, we will discuss how mathematics and rigorous theory can provide the fundamental understanding to efficiently design and deploy AI architectures & algorithms on the most relevant and novel compute hardware. Defence adoption of Future Compute technologies requires new mathematical constructions and representations to reliably inform their use and provide the necessary design and performance assurances.

© Crown copyright (2026), Dstl. This information is licensed under the Open Government Licence v3.0.

AI Unlocked in Autonomous Sensing: Bridging Foundation Models, Transformers, and Bayesian Tracking

James Hopgood – The University of Edinburgh

Artificial Intelligence is transforming autonomous multi-domain and multi-platform sensing across critical national infrastructure protection, search and rescue, smart environments and defence, by processing complex, high-volume sensor data with unprecedented speed and precision. While four decades of model-based training-free statistical signal processing techniques have produced excellent results, modern real-time demands require new paradigms.

This talk explores three key advancements bridging statistical signal processing and deep learning. First, we present cross-modal visual foundation models for Synthetic Aperture Radar (SAR) target classification, trained on corresponding SAR and optical image patches to overcome feature interpretability challenges. Next, we review the recent concept of transformer-based architectures for multi-target tracking, highlighting their ability to model complex spatio-temporal dynamics across dense clutter. Finally, we examine PiVoT (Poisson Measurements-based Variational Multi-object Detection and Tracking), a real-time training-free Bayesian multi-object detection and tracking algorithm for radar perception. PiVoT is built for the hard cases: dense clutter, Doppler point clouds, and large, fast-changing object populations, while working directly from noisy measurements without external clustering or detector stages. Its performance is with Graph Neural Network architectures.

Together, these methodologies demonstrate how combining foundation models, transformers, and principled Bayesian inference unlocks robust next-generation sensing capabilities across challenging operational environments.

This talk will also give updates about how this work aligns directly with the research mission of the EPSRC/MoD Centre for Doctoral Training in Sensing, Processing, and AI for Defence and Security (SPADS CDT), and will give an overview of the broad activities taking place in SPADS.

Understanding The Confidence with which a Machine Learning Model Classifies an Input Datum

Daniel Sumler

Understanding the confidence with which a machine learning model classifies an input datum is an important, and perhaps under-investigated, concept. In this presentation, I will talk about a new calibration metric, the Entropic Calibration Difference (ECD).

Based on existing research in the field of state estimation, specifically target tracking (TT), I will show how ECD may be applied to binary and multi-class classification machine learning models.

This will include the importance of relative under- and over-confidence and how, specifically, the ECD metric distinguishes between them.

Examples of ECD in practise, as well as how the metric may work best in combination with other metrics, such as the Expected Calibration Error (ECE) will also be explored.

Trajectory-Free Parameter Estimation in Interacting Particle Systems via Topological-Frequencies

Mu'izz Asif

Particle Tracking Velocimetry (PTV) contains an expensive computational bottleneck in its frame-to-frame matching step once targets are identified. This process is typically used to attain statistical physics metrics or infer system parameters, such as interaction law exponents. Persistent homology, a tool of Topological Data Analysis is used to capture high-level information of connectivity and voids in data at multiple spatial scales. Over time, the captured multiscale profiles can present structural evolution in the form of CROCKER vectorisations. This research offers an alternative topologically-driven methodology to estimate and compare system parameters and physical metrics without requiring complete particle trajectories, tested on electric-like synthetic data with a wide parameter mesh. This frequency analysis approach on the (Betti) CROCKER vectorisations, enables understanding of the different rates of topological variations of a confined system, directly embedding temporal information compared to treating it as a tuneable hyperparameter. This inherits stability from the frequency transform to large frame rate differences. We present a progression of interpretable metrics (from spectral peak characteristics to whole-spectrum) on this topological-frequency spectra that are able to characterise the system at multiple time and spatial scales, shown by using low-compute regression models to be able to estimate particular properties to a high accuracy (some with $<7\%$ mean absolute percentage error), as well as be able to predict relative differences between systems (some $>90\%$ accuracy). These results show the capacity for TDA to interpret highly-coupled dynamical systems without resolving individual particle trajectories, offering a powerful and efficient alternative to traditional tracking methods.

Scientific Machine Learning

Prof Chris Budd OBE, University of Bath, email: mascjb@bath.ac.uk

Artificial Intelligence (AI) is transforming fields in many disciplines fuelling innovation and advancing our ability to solve complex problems. This transformation is built in part on a strong mathematical foundation that makes such AI systems more reliable, interpretable, and robust. A strongly emerging field within AI is Scientific machine learning (SciML), which combines AI-driven insights and architectures, with traditional mathematical models (such as numerical analysis, differential equations, operator theory, inverse problems) to tackle complex, multi-scale scientific challenges varying from medical imaging to climate change and from power generation to material science. The impacts of SciML are far-reaching. In mathematics it can be used for function approximation and classification, solution of differential equations (ODEs and PDEs), efficient simulation of differential operators, uncertainty quantification, solution and regularization of inverse problems, equation discovery, prediction and time series analysis, classification and generation of high-resolution images. In physics, it enables the simulation of complex phenomena; in medicine, it improves diagnostics, personalised treatments, and drug discovery. SciML's ability to produce high-quality results from limited data makes it indispensable in fields like climate science and materials research. SciML also provides a unique advantage for industry, models that balance accuracy, speed, and complexity, enabling product and process optimization on an unprecedented scale.

In this talk I will give an overview of this exciting field, talking about approximation methods, physics informed neural networks, variational machine learning, and neural operators. I will also describe their applications to environmental modelling. In particular I will consider both the strengths and weaknesses of the methods, and the very important need for a careful mathematical approach to ensure that they are robust, transparent, and reliable.