INSTITUTE OF
MATHEMATICS &
ITS APPLICATIONS

# Foundational Mathematics for Trustworthy AI

## Abstract of a Discussion Paper on Research Investment, International Position and National Mathematical Capability

**Submitted by:** *The Chair, IMA Special Interest Group on Foundational Mathematics for AI*

**Occasion:** *2nd IMA Artificial Intelligence Congress, AI Unlocked, Hyatt Regency Birmingham, 17 to 18 September 2026*

**Status:** Abstract of a discussion paper issued to inform debate. It does not represent an agreed position of the Institute of Mathematics and its Applications, and it advances no commercial interest. The full paper is fully referenced and states the policy position as at 24 July 2026.


**The position**

Over eighteen months the United Kingdom has built the infrastructure, the capital and the institutions for artificial intelligence at national scale. Public compute has grown roughly tenfold, sixty-eight billion pounds of investment has been pledged, and thirty-eight of the fifty commitments in the *AI Opportunities Action Plan* are recorded as met. Artificial intelligence is now inside the machinery of the state: around a third of NHS chest X-rays are AI-assisted, and Making Tax Digital for Income Tax began on 6 April 2026 and will draw millions of taxpayers into an automated, software-mediated relationship with HMRC by 2028.

**The argument**

**The country is deploying at the speed of capability and checking at the speed of anecdote.**

The evidence on which these systems are approved is *observational*. It consists of test scores, pilot results, supplier accuracy claims and confidence figures produced by the systems themselves.
Observational evidence establishes what a system did on the examples it was shown. It does not establish what it will do tomorrow, why it did what it did, where it will go wrong, how badly, or whether anyone will notice when it does. For a film recommendation that is sufficient. For a system that sets a citizen's tax liability, reads a chest X-ray, cues a weapon, sequences an approach path or brakes a vehicle, it is not.

Three questions follow, and no regulator in the United Kingdom can currently answer any of them for a deployed system: *what does it promise, how does it fail, and can it show its working?* More testing will not resolve them, because the difficulty is not a shortage of tests but the absence of a measuring instrument. That instrument is mathematical. The paper identifies seven areas in which it has to be built: the geometry of the space in which machines hold meaning; enforcement of a subject's own exact rules by construction rather than by training; honest uncertainty in place of self-reported confidence; the limits of what the available data can support; the sensitivity of an answer to a small change in its input; guarantees for systems assembled out of other systems; and the characterisation, rather than the counting, of failures.

**The international position**

Every major jurisdiction has now discovered that the instrument does not exist. The European Union has deferred its high-risk obligations to December 2027 and August 2028 because the conformity standards on which they depend did not materialise. The United States is pursuing a measurement science for artificial intelligence through NIST. China has paired law with technical standards, at speed. France has committed seventy-three million euros to a national programme with an explicit mathematical-foundations pillar. Germany has run a national priority programme on the theoretical foundations of deep learning since 2020 and has funded fifteen further independent research groups, expressly to reverse brain drain, including researchers recruited from the United Kingdom. Comparable British investment consists of a small number of excellent but individually modest grants totalling single-digit millions. Britain is competing in this race with world-class individuals and no team.

The country that builds the measuring instrument will set the terms on which artificial intelligence is bought and sold in every regulated market in the world, and the window closes with the European deadlines.

**National mathematical capability**

Research alone is insufficient: a guarantee nobody can read is not a guarantee. England's national curriculum is being rewritten now. The new curriculum is due in spring 2027 for first teaching from September 2028, with statutory consultation on draft content running through 2026. **This is the window, and it will not reopen for a decade.**

The paper sets out a single progression from primary school to professional practice. The objects underlying artificial intelligence met early and without machinery, at ages five to eleven. Formalised at Key Stage 3. A compulsory strand at GCSE across all tiers, including why a machine's confidence score is not the probability that it is right. Linear algebra, iteration and statistical inference brought into A-level Mathematics rather than reserved for Further Mathematics. Quantitative literacy to eighteen on every route, academic and technical alike, including the further education and apprenticeship routes from which the assurance workforce will largely be recruited. A compulsory, credit-bearing, discipline-appropriate module in artificial intelligence and quantitative analysis in every undergraduate degree, in every faculty. A reformed core in mathematics, statistics and computer science degrees themselves, accredited by the professional bodies. And demonstrated competence as a condition of chartered status wherever a professional signs a decision to which a machine contributed.

Two constraints bind all of it: mathematics teacher supply, and initial teacher education. The workforce capable of teaching the new content has to exist before September 2028.

Artificial intelligence education should be compulsory in four specific senses: statutory content in schools, an entitlement that funding follows at sixteen to nineteen, a condition of degree award, and a condition of chartered status. It should mean the mathematics that allows a person to check a machine, not training in the use of this year's products.

**Ethics, and an exportable template**

Britain's distinctive contribution to artificial intelligence has been neither capital nor compute but vocabulary: Turing's reframing of the question as one about observable evidence, Onora O'Neill's distinction between trust and trustworthiness, the first cross-sector code, the first regulator-backed

guidance on explaining an automated decision, the first international declaration, and the first state institute mandated to evaluate frontier models. That contribution is unfinished, because its terms have never been given quantities, and **an ethical commitment that cannot be measured cannot be audited.** Supplying the quantities is work that only a mathematical community can do.

British qualifications are examined in more than twelve thousand schools across more than one hundred and sixty countries. A change to the mathematics curriculum in England is therefore not only a domestic act. The paper proposes that the reform be written once and published under an open licence as a reference curriculum, on the argument that the country which defines what it means to understand artificial intelligence defines the vocabulary in which conformity is later argued.

**What the paper asks for**

Twenty-eight recommendations, in four groups.

- **A ten-year national programme** in foundational mathematics for trustworthy AI: cross-council, UKRI-hosted, each research area paired with the regulator whose evidential needs define success, and scaled against comparable European investment rather than against the current British total.
- **A Mathematical Conformance Standard for artificial intelligence in regulated decisions**, made a condition of public procurement. This is the central lever. It costs nothing and it reorganises the market.
- **Mathematics curriculum reform submitted inside the 2026 statutory consultation window**, from primary school to university, with teacher supply treated as the binding constraint it is.
- **Compulsory artificial intelligence education**, an *Ethics-to-Mathematics Concordance* giving each principle in British guidance a quantity and a test, and an openly published reference curriculum distributed through the awarding bodies, the British Council, the Commonwealth and UNESCO.

**The cost of inaction**

If the country continues to deploy at the speed of capability and check at the speed of anecdote, the failure will not announce itself. There will be a slow accumulation of confident, plausible, undetected errors distributed through the tax system, the health service, the transport network and the defence enterprise: individually deniable, collectively enormous, and invisible to every metric currently collected. It will emerge in a tribunal, in an inquest, or in a public inquiry several years and several billion pounds later.

Sixty-eight billion pounds has been pledged to build these systems. A ten-year national programme to establish that they can be trusted would cost a fraction of one per cent of that, and it is the only part of the expenditure that protects the rest.

The United Kingdom has the tradition to do this: the tradition of Turing, of the numerical analysts, the statisticians and the coding theorists, of a mathematical community that has always understood that the interesting question is not whether a thing works, but *why*, and *when it will not*. That community exists. It is ready. It asks to be used.

*The full discussion paper, with sources, is available from the author.*