Artificial Intelligence and Mathematics

From Neural Networks to Computer-Assisted Proofs

AI & AL

ITI “E. Torricelli - G. Tomasi di Lampedusa”

April 2026

Prompt for this presentation

I am a math teacher at a technical high school (upper secondary school) and I would like to give a lesson on the use of AI in mathematics, particularly regarding computer-assisted proofs. This lesson is aimed at 12th and 13th grade students who already have a background in computer science and follow a rigorous mathematics curriculum.
I would like to start with a brief general introduction to AI (both current and potential) and its connection to mathematics, also from a formal perspective related to its development, not only in terms of applications. I would also like to explore the reverse process: what AI can do for mathematics, especially in the context of computer-assisted proofs. For example, recalling the role of computational tools in the Four Color Theorem, the Classification of Finite Simple Groups, or Knuth's latest graph classification results.
It would also be useful to include some final exercises on how to use AI to prove statements, working through some well-known ones and approaching the analysis of a few still-unsolved (very elementary) conjectures.
It would be great to have the final output as a presentation in English, also available in LaTeX format.

What is Artificial Intelligence?

Artificial Intelligence encompasses systems that perform tasks typically requiring human intelligence:

- **Learning** from data and experience
- **Reasoning** through logical inference
- **Problem-solving** in complex domains
- **Pattern recognition** in large datasets

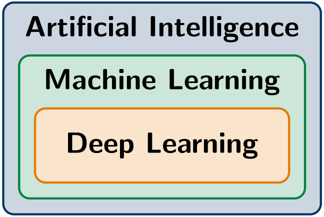

Brief Historical Timeline

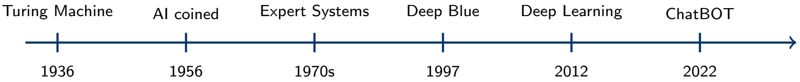

Key Insight

AI progress has always been **intertwined with mathematics**: from Turing's foundational theory to modern optimization algorithms.

How Modern AI Works: Neural Networks

A **neural network** approximates functions through layers of simple operations:

Mathematical Foundation

For input $\mathbf{x}$, output $\mathbf{y}$:

$$\mathbf{y} = f_n \circ f_{n-1} \circ \cdots \circ f_1(\mathbf{x})$$

Each layer:

$$f_i(\mathbf{z}) = \sigma(W_i \mathbf{z} + \mathbf{b}_i)$$

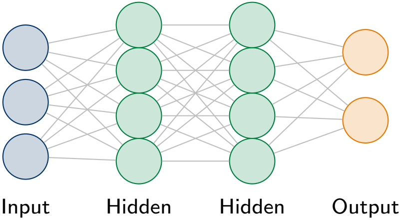

Where σ is a nonlinear **activation function**.

The Universal Approximation Theorem

Theorem (Cybenko, 1989; Hornik, 1991)

A feedforward neural network with a single hidden layer containing a finite number of neurons can approximate any continuous function on compact subsets of $\mathbb{R}^n$, under mild assumptions on the activation function.

Formally: For any continuous $f : [0,1]^n \to \mathbb{R}$ and any $\varepsilon > 0$, there exist $N \in \mathbb{N}$, weights $w_i, \alpha_i \in \mathbb{R}$, and biases $b_i \in \mathbb{R}$ such that:

$$\left| f(\mathbf{x}) - \sum_{i=1}^{N} \alpha_i \sigma\left(\mathbf{w}_i \cdot \mathbf{x} + b_i\right) \right| < \varepsilon$$

Important Caveat

The theorem guarantees *existence*, not efficient *construction* or *learning*.

Large Language Models: The Transformer Architecture

Modern AI chatbots use **Transformers**, which rely on the **attention mechanism**:

Scaled Dot-Product Attention

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V$$

Where:

- Q = Query matrix
- K = Key matrix
- V = Value matrix
- d_k = dimension of keys

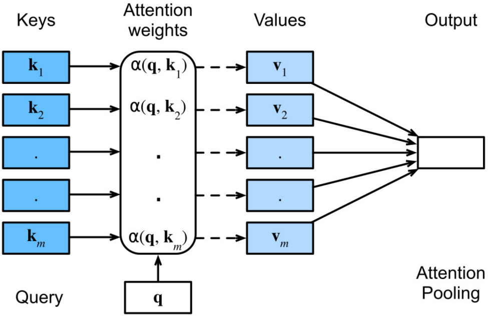

The **softmax** converts scores to probabilities.

Gradient Descent: Learning as Optimization

Problem: Find weights $\mathbf{w}$ that minimize a loss function $L(\mathbf{w})$.

Gradient Descent Update Rule

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \nabla L(\mathbf{w}_t)$$

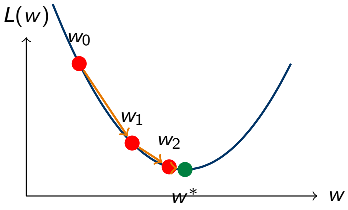

Where:

- η = learning rate (step size)
- ∇L = gradient of loss function

Key insight: The gradient points in the direction of steepest *ascent*; we go opposite.

Connection to Calculus

Recall: $\nabla L = \left(\frac{\partial L}{\partial w_1}, \dots, \frac{\partial L}{\partial w_n} \right)$

Why Linear Algebra is Essential

Core operations in AI:

- **Matrix multiplication**: $Y = WX + B$
  - Every layer transformation
  - Attention computation
- **Eigendecomposition**: $A = Q\Lambda Q^{-1}$
  - Principal Component Analysis
  - Understanding network dynamics

Dimension Check

If $X \in \mathbb{R}^{n\times m}$ (n samples, m features)
and $W \in \mathbb{R}^{k\times m}$ (k neurons),
then $Y = WX^T \in \mathbb{R}^{k\times n}$

GPT-4 Scale

Estimated parameters:
$\sim 1.7 \times 10^{12}$
Matrix operations:
$\sim 10^{15}$/inference

The Mathematical DNA of Artificial Intelligence

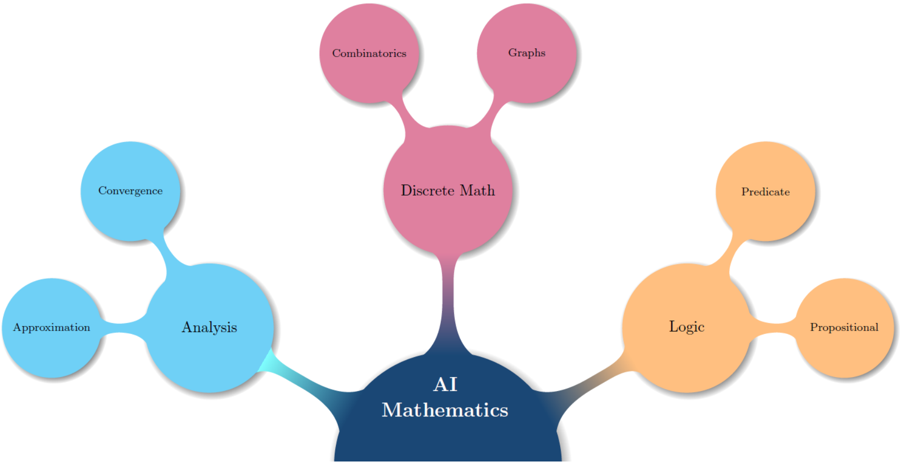

The Mathematical DNA of Artificial Intelligence

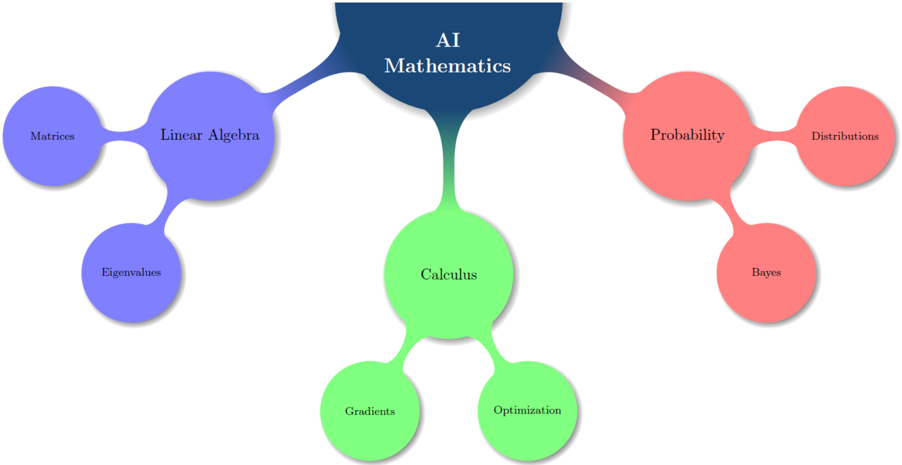

Computer-Assisted Proofs

The Traditional View of Mathematical Proof

> *"A proof is a sequence of logical deductions from axioms and previously established statements that culminates in the proposition in question."*
>
> — Traditional definition

Classical requirements:

- Human-verifiable
- Surveyable (can be checked in finite time)
- Follows accepted logical rules

Emerging challenges:

- Proofs too long for humans
- Computational verification
- Collaborative mega-proofs

The Question

Is a computer-verified proof a "real" proof?

Claude's Cycles
Don Knuth, Stanford Computer Science Department
(28 February 2026; revised 16 March 2026)

Shock! Shock! I learned yesterday that an open problem I'd been working on for several weeks had just been solved by Claude Opus 4.6 - Anthropic's hybrid reasoning model that had been released three weeks earlier! It seems that I'll have to revise my opinions about "generative AI" one of these days. What a joy it is to learn not only that my conjecture has a nice solution but also to celebrate this dramatic advance in automatic deduction and creative problem solving. I'll try to tell the story briefly in this note. Here's the problem, which came up while I was writing about directed Hamiltonian cycles for a future volume of The Art of Computer Programming:
"Consider the digraph with m^3 vertices ijk for $0 \le i, j, k < m$, and three arcs from each vertex, namely to i^+jk, ij^+k, and ijk^+, where $i^+ = (i + 1)$ mod m. Try to find a general decomposition of the arcs into three directed m^3-cycles, for all $m > 2$".

Landmark: The Four Color Theorem (1976)

Theorem (Appel & Haken, 1976)

Every planar map can be colored with at most four colors such that no two adjacent regions share the same color.

The proof strategy:

1. Reduce to 1,936 "unavoidable" configurations
2. Prove each is "reducible" (can be simplified)
3. Computer checks all configurations

Computation: ∼1,200 hours on IBM 360

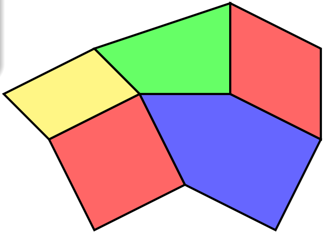

Four colors suffice!

2005: Formally Verified

Georges Gonthier verified the theorem in Coq proof assistant. [gonthier2008formal]

The Kepler Conjecture: From Hales to Formal Proof

Kepler Conjecture (1611)

The densest packing of spheres in 3D is the face-centered cubic arrangement, with density $\frac{\pi}{3\sqrt{2}} \approx 0.7405$.

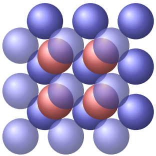

Timeline:

- **1998**: Thomas Hales announces proof
- **2003**: Referees "99% certain" after 4 years
- **2014**: Formal verification completed (Flyspeck project)
- Used **HOL Light** proof assistant

Lesson

Human proof took 4 years to review; formal verification provided certainty.

Hands-On: Using AI for Mathematical Proofs

Exercise 1: Infinitude of Primes

Theorem (Euclid)

There are infinitely many prime numbers.

Task: Ask an AI assistant to:

1. Provide Euclid's original proof
2. Give an alternative proof (e.g., using $\sum 1/p$)

Sample Prompt

"Prove that there are infinitely many primes. First give Euclid's proof, then show why $\sum_{p \text{ prime}} 1/p$ diverges."

Exercise 2: Irrationality of $\sqrt{2}$

Theorem

$\sqrt{2}$ is irrational.

Task: Explore multiple proof strategies with AI:

1. Classic proof by contradiction (parity argument)
2. Geometric proof using isoceles right triangles
3. Proof using unique factorization

Classic Approach

Assume $\sqrt{2} = \frac{p}{q}$ in lowest terms.
Then $2q^2 = p^2$, so p is even.
Write $p = 2k$, then $q^2 = 2k^2$, so q is even.
Contradiction!

Prompt to try:
"Can you show me a geometric proof of the irrationality of $\sqrt{2}$ using an infinite descent argument with triangles?"

Exercise 3: Plato's Meno

Doubling a given square

Using only a compass and straightedge, how can you construct a square with double the area of a given square?

Task: Ask an AI assistant to:

1. Provide the Socratic method
2. Give a geometrical construction by using straightedge and compass

Tripling a given square

Using only a compass and straightedge, is it possible to construct a square whose area is three times that of a given square?

Exercise 4: Exploring Goldbach's Conjecture

Goldbach's Conjecture (1742)

Every even integer greater than 2 can be expressed as the sum of two primes.

Status: **Unproven** — but verified computationally for $n < 4 \times 10^{18}$

Task: Use AI to explore:

1. Write code to verify for small n
2. Investigate the number of representations $r(n)$
3. Discuss why this is hard to prove
4. Explore related conjectures (weak Goldbach, proved 2013)

Sample Investigation

"Write Python code to find all ways to write even numbers up to 100 as sums of two primes. Then plot the number of representations."

Conclusions

Best Practices for AI-Assisted Mathematics

DO:

- ✓ Verify AI outputs independently
- ✓ Use AI for exploration and intuition
- ✓ Ask for multiple approaches
- ✓ Request explanations of steps
- ✓ Cross-check with formal tools
- ✓ Treat AI as a collaborator

DON'T:

- × Trust calculations blindly
- × Accept "proofs" without checking
- × Assume AI understands deeply
- × Skip learning fundamentals
- × Use AI for exams dishonestly
- × Forget to cite AI assistance

Critical Thinking

AI can produce convincing-sounding but incorrect proofs. **Mathematical validity requires verification, not just plausibility.**

The State of AI in Mathematics (2025)

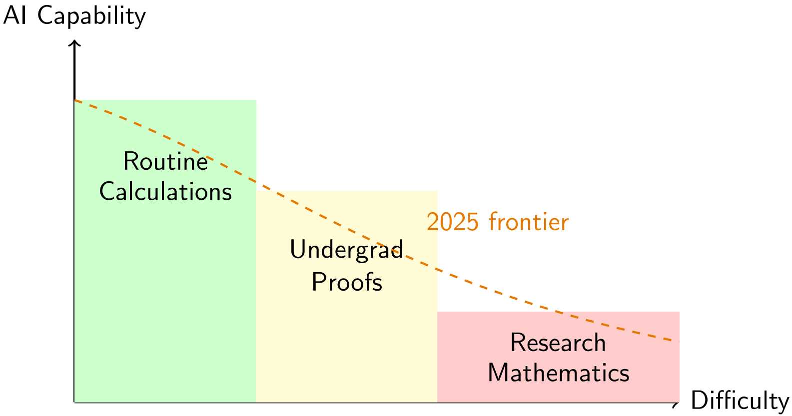

"Recent developments show that AI can prove research-level theorems in mathematics, both formally and informally." [arxiv.org](https://arxiv.org/html/2603.03684v2)

Resources for Further Exploration

Proof Assistants:

- Lean: leanprover.github.io
- Natural Number Game: adam.math.hhu.de
- Mathlib documentation: leanprover-community.github.io

Reading:

- Tao, T. (2024). "Machine Assisted Proofs" [terrytao.wordpress.com](https://terrytao.wordpress.com/wp-content/uploads/2024/03/machine-jan-3.pdf)
- Avigad, J. (2026). "Mathematicians in the Age of AI" [arxiv.org](https://arxiv.org/html/2603.03684v2)
- Herman, R.L. (2025). "AI in Mathematics Research" [people.uncw.edu](https://people.uncw.edu/hermanr/mat495/AI_in_Math.pc