HUMAN MATHEMATICS IN A WORLD WITH A(I)GENTS · IST · 23 NOV 2026

# Supervising (Math) Students in the AI Era

DM
DEPARTAMENTO DE MATEMÁTICA
TÉCNICO LISBOA

Hugo Tavares · Departamento de Matemática
Instituto Superior Técnico

# Plan of the talk



## 0 Beyond mathematics

How firms train juniors when AI does the junior work, and what the evidence on learning says



## 1 Institutions

The official definition, two reports, where the societies stand



## 2 Mathematicians

What is mathematics, and a PhD, for? How much AI belongs in it?



## 3 Students

What early-career mathematicians are telling us



## 4 What we can do

Which problems to give; advice for PhD students and postdocs



## 5 My personal perspective

Supervising PhD students: what I would do, and one thing I would not

# 2026: the training ground moved

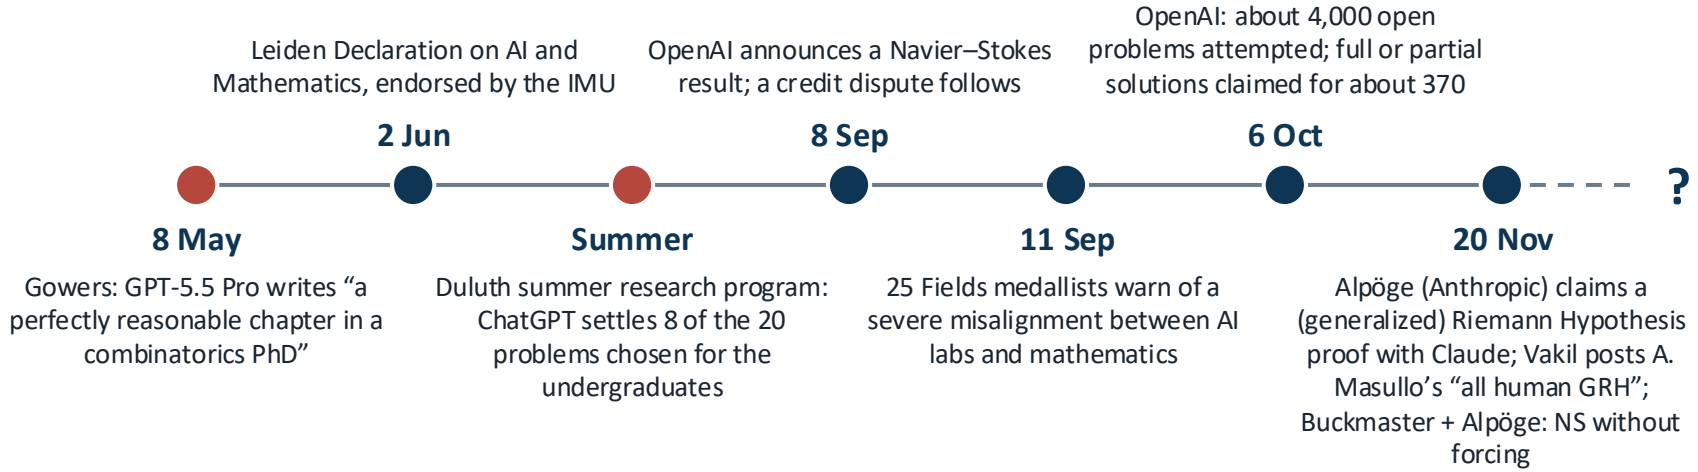


**Melanie Matchett Wood (Harvard), on 29 Oct: “AI can produce things that could have been acceptable Ph.D. theses a couple of years ago.”**

*How do we still form mathematicians — technique, judgment, taste — when the apprentice’s problems are solved by machines?*

0

# Beyond mathematics

How firms train juniors when AI does the junior work

# Other fields are losing their training ground

## AT WORK

**7×**

AI-exposed junior jobs are seven times more likely to ask for traditionally senior skills (leadership, strategy)

*PwC AI Jobs Barometer 2026*

**43%**

of 906 KPMG summer interns: top worry is AI limiting their critical thinking (5% fear losing the job)

*KPMG survey, via Emex, Oct 2026*

## WHAT STUDIES MEASURE

RANDOMISED TRIAL **High-school maths**

**−17%** nearly 1,000 students, Turkey

on the exam without AI, after practising with unrestricted GPT-4 (vs no AI). With a version that gives just hints: no loss.

*Bastani et al., PNAS 2025*

RANDOMISED TRIAL **Programmers**

**50 vs 65** 52 participants

average test score after learning a new tool, with AI vs without. With AI, the ones that used it to understand: about 69; the other: about 33.

*Shen & Tamkin (Anthropic), 2026*

BEFORE / AFTER **Endoscopists**

**28→22%** 19 doctors, 1,443 exams

adenoma detection in exams done without AI, before vs after AI became routine. Workload also rose, so the cause is not proven.

*Budzyń et al., Lancet G&H 2025*

CORRELATION **PISA 2025, age 15**

**−20 points** 760,000 students, 91 countries

in science (about one school year) for students who use AI for schoolwork tasks, vs non-users. An association, not a proven cause.

*OECD, PISA 2025 Results, Sep 2026*

# Three ways firms are responding

**PwC**

## Teach judgment explicitly

- A course in critical thinking and judgment for 85,000+ staff, by June 2027
- Staff practise on realistic cases, e.g. an AI-written tax analysis that looks right but has gaps to find
- From 2027, any university can use the course; PwC will fund some to teach it

**KPMG**

## Simulate the lost routine

- AI now does the checking work that used to teach juniors to spot errors
- So interns train on invented cases, each with a hidden fraud to find
- Their rule: think first, ask the AI, then check its answer
- Interns rated coaching by a person most useful (57%), learning with AI least (3%)

**New York law firms**

## You first, AI later

- Barclay Damon: years 1–3 do the work first, then AI
- Bond, Schoeneck & King: no AI for new lawyers
- Tully Rinckey: 24-hour AI-free hiring test, then 6–12 months of daily mentoring

**How can one teach critical thinking and judgment? No one really knows for sure, but some firms are trying.**

# 1

# Institutions

The official definition, two reports, and where the societies stand

# What a doctorate certifies, officially

*The degree of doctor is awarded to those who demonstrate:*

1. Systematic understanding of a scientific field
2. The research skills and methods of that field
3. Ability to conceive, design and carry out significant research, with integrity
4. **A significant body of original work that extends the frontiers of knowledge**
5. Ability to analyse critically, evaluate and synthesise new and complex ideas
6. Ability to communicate with peers, the academic community and society
7. Ability to promote technological, social or cultural progress

### Where the list comes from

Dublin descriptors, Bologna Process (2004)

↓ Portuguese law: Decreto-Lei 74/2006, art. 28

↓ IST doctoral regulations, art. 1

**Only one of the seven is about new results.**

It is the one a written thesis certifies, and the one AI can increasingly produce.

### How do we assess the other six?

# Harvard summit (Sep 2026): rebuild assessment

**The summit's ladder of AI uses**




***"AI use should accelerate understanding, not bypass understanding."***

- Rigorous in-person assessment by several faculty, every year
- No PhD primarily on the text of the thesis; the defence needs complete mastery
- The thesis need not record the first proof of a result
- AI tools available to all students — use never mandatory
- Communication to experts and to the public as an explicit degree expectation

**24 mathematicians from different institutions, 17–18 Sep 2026.** The report presents itself as a starting point for discussion, not a prescription.

# AMS report (2 Oct 2026): a call to action

## 1 · A COMMUNITY CONVERSATION

### *"What is the purpose of a math PhD?"*

- Is its main product the person, or the contribution to knowledge?
- Does a dissertation still show mastery, and of which skills?
- Are qualifying exams and the defence still accurate measures?
- Should students be required to learn to use AI?
- How will employers evaluate candidates in future years?

## 2 · WHAT A DEPARTMENT CAN DO NOW (our community needs leadership now)

1. Say why mathematicians matter
2. Town halls; advisors talk with their students
3. Open up non-academic careers
4. Know what each kind of academic post needs
5. Train communication: weekly talks with feedback
6. Write new exam and defence rules
7. Look after mental health, also of junior faculty
8. Teach current tools: data, computing, Lean
9. Clear rules for declaring AI use
10. Department-wide access to frontier AI tools

**A workshop report by 19 participants.**

# Where the societies stand

## International Mathematical Union

*Jun 2026*

Endorses the Leiden Declaration: human responsibility, disclosure of tools, proper attribution. Nothing on training.

## European Mathematical Society

*5 Aug and 10 Sep 2026*

Decides not to endorse Leiden formally: weigh the opportunities too. On Navier–Stokes: a milestone; credit must be resolved; closed models are a problem for open science.

## American Mathematical Society

*2026*

No single position: white papers from five policy committees, and the Committee on Education's report on the doctorate (2 Oct).

## Advisory Group on Mathematics and AI

*independent, at the IAS · Sep 2026*

Advises AI companies on the responsible release of results; warns of a two-tier system of access to models.

**The societies do not agree either:** the IMU endorses Leiden, the EMS does not. Only the Harvard and AMS reports address the doctorate.

# 2

# Mathematicians

An existential question, and a spectrum of views

# An existential moment: what is all this for?

**Terence Tao, to the New York Times:** this moment feels like “soul searching at gunpoint”. Two questions keep coming back.

## What is mathematics for?

**25 Fields medallists (11 Sep)**
Solving problems is “only a tool and proxy” for understanding

**Guillen (12 Sep)**
To understand causes; an AI result should be a source of shared pride

**Madore, Télécom Paris (20 Oct)**
To understand the world; we gave hard problems too much glory

**Vakil, Stanford (6 Nov)**
Maths cannot be solved once and for all; we learn in the process

**Tao, ICM lecture (24 July)**
A crisis of values, like the foundations crisis of 1900–1930

## What is a PhD in mathematics for?

**Litt, Toronto (14 Sep)**
To become a world expert on a deep topic, and able to explain it

**Koberda (23 Sep)**
To pass on human expertise, more than to produce theorems

**Harvard summit (17–18 Sep)**
The aims have not changed; how we teach and assess must

**Lozano-Robledo, UConn (8 Oct)**
Love of the subject, as ever: “Keep calm and carry on studying math”

## And the next generation?

268 early-career mathematicians were asked for their favourite aspects of being in maths (up to five each).
**Solving problems comes third.**



# How much AI in a PhD? A spectrum of views

LESS AI DURING THE PHD — All dates 2026 — MORE AI DURING THE PHD

### AS LITTLE AS POSSIBLE

**Keep AI out while you learn the craft**

**Tao, ICM lecture (24 July)**
In training, "tightly restrict the use of AI tools"

**Rivera, Purdue (26 Aug)**
Asks his PhD students to keep AI use low for now, and to declare it

**Kearns & Roth, Penn (9 Mar)**
Juniors should still write their papers without AI help

### FOR SOME TASKS ONLY

**Around the maths, not where you learn**

**Tao, UCLA Magazine (6 Oct)**
Useful in many ways, but risky on a problem that was set to train you

**Etingof, MIT (May, upd. Sep)**
No AI on the exercises your advisor sets; elsewhere, few and careful prompts

**Lozano-Robledo, UConn (23 Oct)**
Not for proving the key lemmas: ask your advisor for the hint

### THE STUDENT DECIDES

**Neither forbid nor impose**

**Symplectic collective (19 Oct)**
Supervisors should neither forbid AI nor force it on students

**Harvard summit (17–18 Sep)**
No fixed rules; tools for every student, use never mandatory

**Vakil, Stanford (26 July)**
You do not need AI for now, but do not be afraid of it

### ASSUME AI IS USED

**So change what the PhD asks for**

**Tao, to his students (16 Oct)**
Not one problem alone: use AI on ten, and look for the patterns

**Calegari, Chicago (1 Nov)**
Judge all work as if AI helped; make the PhD broader

**Litt, Toronto (14 Sep)**
Welcome results from any source; award the PhD on a rigorous defence

**Common ground:** the student must understand, be able to explain and answer for everything they sign. The disagreement is how to get there, and what a PhD should certify.

# 3

# Students

What early-career mathematicians are telling us

# What early-career mathematicians are telling us

**70%**

say their chance of staying in academic maths has dropped (9% say it rose)

**76%**

would be unhappy if AI solved a problem they had worked on

**54%**

would not want a career where proof digestion is the main creative task

**62%**

of those paying for AI tools pay out of their own pocket

*268 respondents (85% PhD students; mostly pure maths), ~60 institutions, Sep–Oct 2026; informal, mostly US, before the 6 Oct release. Brandenberger, Mularczyk & Woodruff (MIT), arXiv 2610.40584*

*"a sustainable profession cannot rely on expertise developed under a system it has abandoned"*

**Anonymous PhD student who will leave academia**

*Proofs and Prompts (blog), 20 Oct 2026*

*"the rules cannot change only for the young"*

**Anonymous postdoc in PDE**

*Proofs and Prompts (blog), 22 Oct 2026*

Part of the "Navier–Stokes generation" of PhD students — asks for more expositions, not only AI statements

**Yifu Zhang, 1st-year PhD (Imperial)**

*Proofs and Prompts (blog), 7 Nov 2026*

# 4

# What we can do

Problems to give, and advice for students, postdocs, supervisors and departments

# What PhD problems should we give now, and how?

## WHICH PROBLEMS TO GIVE

### 1 Understand a known result

Rework a hard theorem few have digested. Spadaro's PhD with De Lellis, on Almgren's "big regularity paper".

*Tao, Math 2.0, Oct · example from a colleague*

### 2 Explore around a theorem

Remove a hypothesis or a key tool: can it still be proved? What survives nearby?

*Tao, Math 2.0, Oct*

### 3 A territory, not a conjecture

Explore a newly defined object, or find the tractable cases of a conjecture out of reach.

*Litt, Sep · Duluth REU, Sep · symplectic, Oct*

### 4 Where AI is weak

A brand-new corner of a subject, or first questions from another science, e.g. psychology, neuroscience

*Strogatz, ICM panel, Jul · Lozano-Robledo, Oct*

### 5 Beyond AI alone

Targets that neither the student nor the AI could reach alone.

*Gowers, May*

### 6 Formalization & tools

Formalize results in Lean, build libraries. Counts for little on job applications today.

*Avigad, Aug*

## HOW TO HANDLE THEM

### a Keep some work AI-free

The struggle is the point. Always at least one project LLM-free.

*Tao, Jul · Rivera, Aug · Lozano-Robledo, Oct*

### b Pre-test against AI

Check problems against AI, as against the literature. Duluth: AI settled 8 of 20.

*Duluth REU, Sep · symplectic collective, Oct*

### c Start from an AI draft

A student reads an AI-written draft, rewrites and extends it. First steps of research only; risk: quality control only.

*Lozano-Robledo, Oct · Alali, Sep*

### d Work as a lab

Students brainstorm, the leader asks an AI, the group checks and discusses the output.

*Etingof & Gerovitch, Oct*

**No consensus, and almost no evidence.** All dates 2026.

# If you are a PhD student

### Struggle first, check last

First serious attempt and final check without AI; AI only in between

*Brcic & Frljic, 24 Jun 2026 · Etingof, May 2026*

### Do the exercises yourself

Exercises and key lemmas are training, not output: the effort is the point

*Kontorovich (Rutgers), ICM panel, 26 Jul 2026*

### Read actively

Try to prove it yourself; look only when stuck, and take the smallest hint

*Gowers, Gowers's Weblog, 17 Sep 2026*

### Keep a lab notebook

Keep prompts and answers; disclose AI use in papers and thesis

*Symplectic collective, 19 Oct 2026*

### Don't be alone with the AI

Regular, in-depth discussions with your advisor and your group

*Symplectic collective, 19 Oct 2026*

### Ask before you choose

Ask advisors about their stance on AI; choose one that matches yours

*Lozano-Robledo (UConn), 23 Oct 2026*

**Ravi Vakil:** *"the best person with a calculator is the best person without a calculator."*

*Stanford; president of the American Mathematical Society. Panel at the ICM, Philadelphia, 26 July 2026, recorded for Quanta's podcast The Joy of Why.*

# If you are a postdoc, or mentor one

## IF YOU ARE A POSTDOC

1. **Own what you sign** ALSO FOR PHD STUDENTS
   Make your own contribution clear; be ready to explain why it works, what failed and what is still open. *Figalli & Seyfaddini (ETH), 6 Nov*
2. **Declare your AI use** ALSO FOR PHD STUDENTS
   Now required by the CNRS, MIT and MPIM Bonn; the CNRS will not penalise use or non-use. *Proofs and Prompts, 14 Nov*
3. **Fewer, better papers**
   Committees now ask for a few selected works: 4 at Cambridge, 7 at the CNRS, 3 projects at the IAS. *Proofs and Prompts, 14 Nov*
4. **Talks and conversations count**
   An interview, a talk and an in-depth discussion are replacing the paper file. *Calegari, 1 Nov · ETH, 6 Nov*
5. **Breadth and your own vision**
   Broad and deep knowledge, judgment, independent vision, communication. *Etingof & Gerovitch, 21 Oct*

## IF YOU MENTOR OR HIRE ONE

1. **Publish the rules first** ALSO FOR PHD STUDENTS
   Permitted use, priority, authorship, hiring standards: before people invest months in a project. *Anonymous postdoc, 22 Oct*
2. **Pay for the tools** ALSO FOR PHD STUDENTS
   Top subscriptions cost $200 to $500 a month; access should not depend on salary or contacts. *Anonymous postdoc, 22 Oct*
3. **Do not count papers**
   Judge a few works in depth, as if all were AI-assisted; interview in person. *Calegari, 1 Nov*
4. **Do not hand them the checking**
   Verifying AI proofs line by line should not fall to students or postdocs. *Lewin (CNRS), 2 Nov*
5. **Share the good problems**
   A problem AI can explore quickly can be a joint training project, not one more paper of yours. *Anonymous postdoc, 22 Oct*

**What changes for postdocs:** independence, credit and the job market. All dates 2026.

# To read in full: two reports and three guides

## TWO REPORTS FROM INSTITUTIONS

### Harvard summit report

*24 mathematicians, 17–18 Sep 2026*

How to assess a PhD now: yearly in-person review, a defence that tests mastery, a thesis that need not hold the first proof.

### AMS Committee on Education report

*19 participants, 2 Oct 2026*

What is a PhD for? And ten things a department can do now: conversations, careers, communication, mental health, access.

## THREE GUIDES BY MATHEMATICIANS

### Pavel Etingof (MIT)

*May 2026 + Sep update*

- Keep abreast of the maths in anything you discuss with AI
- Few, well-thought prompts; run drafts past an AI referee

### 36 symplectic geometers

*incl. S. Anjos (Técnico), 19 Oct 2026*

- Numbered advice for students, supervisors and institutions
- Ask consent before giving others' unpublished work to a model

### David Glickenstein (Arizona)

*Sep 2026*

- Judgment as an explicit programme outcome
- 20–30% of quals for judgment tasks; timed quals AI-free

**Closing the circle:** we began with the institutions; these two reports are where a department can start.

5

# My personal perspective

Supervising PhD students: what I would do, and one thing I would not

# The stages of AI grief, for a mathematician



**1 Denial**

Just statistics: a fancy calculator that will never reason deeply.

**2 Anger**

At tech companies, and at colleagues who use AI: plagiarism, opacity, broken norms.

**3 Bargaining**

Only for auxiliary tasks, Lean and narrow problems; the creative core stays human.

**4 Depression**

It settles lemmas and open problems so fast: what is left for the young?

**5 Acceptance**

AI is here to stay: a partner to explore and to prove with.

**6 Existential crisis**

After Navier–Stokes and the flood of AI theorems: what is a mathematician for?

**I AM HERE**

**7 Optimism**

Mathematics has a future, and so do our students.

I have been through the crisis. **I am now at a seventh stage: optimism.**

***Stages 1–6: Javier Gómez-Serrano (Brown University),*** *from the slide that opens his talks; the sixth is his recent update. After Kübler-Ross's five stages of grief (1969). Wording condensed; stage 7 is mine.*

# Supervising PhD students: what I would do

## AS A SUPERVISOR

- Know what a paid AI plan can do today, even if you never use AI yourself
- Help the student write the first prompts
- Students now need to travel more: attend conferences and give talks. Secure the funding!
- Agree AI rules on day one, then revisit them regularly
- Weekly chalk-and-talk; give the hint yourself
- Train judgment: audit an AI proof, find the flaw
- Check the problems you pose against AI, as against the literature

## AS A DEPARTMENT

- Students need support in this turbulent phase
- Organize workshops and training sessions
- Yearly in-person review by several faculty: as support, not only as a test
- Be open about what may change in the coming years: hiring, publishing
- A defence that tests mastery; the thesis need not contain the first proof
- Interview applicants; teach communication and seminar culture
- Pay for the tools centrally; never make them mandatory
- Establish the rules, with students at the table

**What I would not do:** generate an AI draft by myself and with my own LLM, handing it to the student to correct.

**And for postdocs: raise the bar (and the support).** Not the "exercises" an LLM does in one click, but the frontier of knowledge.

# Questions for discussion

1. What should a mathematics PhD certify from now on — and how do we train and test it?
2. How can we help PhD students and PostDocs navigate these turbulent times?
3. Which tasks must stay AI-free in training, and for how long?
4. How do we keep the rules fair, so they don't change only for the young?
5. Is this mainly a pure-mathematics crisis? What changes in applied mathematics?

> *"It was not doomsday, but with 700 papers being released all at once, it was just overwhelming. […] There will be new problems — more interesting problems, more beautiful ones, and we will just shift gears. […] this actually encourages me to **dream big**."*
>
> **Jack Miller**, second-year PhD student in mathematics, Harvard · Harvard FAS news, 29 Oct 2026
>
> **He added that his adviser and other mentors provided reassurance.**

# Sources (1/2)


- PwC, 2026 Global AI Jobs Barometer; HCAmag on PwC's judgment certificate (9 Nov 2026)
- Emex Magazine on KPMG early-careers training (14 Oct 2026); NYSBA managing-partners roundtable (28 Apr 2026)
- Budzyń et al., Lancet Gastroenterol. Hepatol. (2025) — accepted version
- Shen & Tamkin, arXiv 2601.20245; Bastani et al., PNAS 122 (2025); OECD, PISA 2025 Results, Volume I (Sep 2026); Brcic & Frljic, arXiv 2606.26181
- OpenAI, On the Navier–Stokes Millennium Prize Problem (8 Sep 2026); T. Buckmaster, statement (7 Sep 2026); MIT Technology Review (8 Sep 2026)
- OpenAI, Sharing AI progress in mathematics (6 Oct 2026)
- Leiden Declaration on AI and Mathematics (2 Jun 2026); IMU Circular Letter 8/2026 (2 Jun 2026)
- EMS statements on the Leiden Declaration (5 Aug 2026) and on the Navier–Stokes announcement (10 Sep 2026)
- Advisory Group on Mathematics and AI (IAS), launch statement (21 Sep 2026) and Responsible Release of AI-Generated Mathematics (29 Sep 2026)
- Decreto-Lei 74/2006, art. 28; IST doctoral regulations (Despacho n.º 5635/2025), art. 1; Dublin descriptors (Joint Quality Initiative, 2004); Salzburg Principles (EUA, 2005)
- Report of the Summit on PhD Math Education in the Age of AI (Harvard CMSA, Sep 2026); Harvard FAS news (K. Pattison, 29 Oct 2026)
- AMS Committee on Education, workshop report A paradigm shift for doctoral education in mathematics (listed 2 Oct 2026); AMS white papers on AI
- P. Etingof, Use of AI in mathematical research: a guide for young mathematicians (MIT, May 2026; update 14 Sep)
- Symplectic collective, Proofs and Prompts (19 Oct 2026); D. Glickenstein, arXiv 2609.21132

DM DEPARTAMENTO DE MATEMÁTICA TÉCNICO LISBOA

# Sources (2/2)


- A. Figalli & S. Seyfaddini; M. Lewin; "Hiring this year" — Proofs and Prompts (2–14 Nov 2026)
- T. Tao, Mathematics in the age of AI, arXiv 2608.16753 and ICM lecture slides (Jul 2026); Math 2.0 slides (Oct 2026); S. Wolpert, The Tao of Tao, UCLA Magazine (6 Oct 2026)
- T. Gowers, Gowers's Weblog (8 May and 17 Sep 2026); E. Bumiller, interview with T. Tao, New York Times (16 Oct 2026)
- M. Rivera; D. Litt; Á. Lozano-Robledo; an anonymous PhD student; an anonymous postdoc — Proofs and Prompts (Aug–Oct 2026)
- T. Koberda, arXiv 2609.28140; F. Calegari, guest post on Tao's blog (1 Nov 2026); M. Kearns & A. Roth, Amazon Science (9 Mar 2026)
- C. Defant, E. Hodges, N. Kravitz (Duluth REU), arXiv 2609.28615
- A. Brandenberger, H. Mularczyk, D. Woodruff, Early career researcher responses to AI in math, arXiv 2610.40584
- E. Kübler-Ross, On Death and Dying (1969)
- B. Alali, arXiv 2609.22891; J. Avigad, arXiv 2608.23218; P. Etingof & S. Gerovitch, Proofs and Prompts (21 Oct 2026); N. Guillen, Birds and frogs (12 Sep 2026); J. Gómez-Serrano, talk slide, via Scientific American (Nov 2026)
- Y. Zhang, Proofs and Prompts (7 Nov 2026); Quanta, The Joy of Why live from ICM 2026 (3 Sep 2026)
- A Severe Misalignment of AI in Mathematics, 25 Fields medallists (11 Sep 2026); D. Madore, Proofs and Prompts (20 Oct 2026); D. Litt, Proofs and Prompts (14 Sep 2026); Stanford Report, R. Vakil (6 Nov 2026); Science Friday, T. Tao (28 Oct 2026); Á. Lozano-Robledo, guest post on Tao's blog (8 Oct 2026)
- T. Tao, Unforced Euler singularity, unforced Navier–Stokes singularity, and the Generalized Riemann Hypothesis (blog, 20 Nov 2026) — on L. Alpöge's GRH claim and Alpöge & Buckmaster, Finite-Time Blow-Up for the Navier–Stokes Equations
- R. Vakil, post on A. Masullo's draft proof of GRH (20 Nov 2026)


**Further reading (my favorite texts) and to access these slides:** sites.google.com/site/hugotavaresmath/ai-and-math

**AI disclosure:** Prepared with the help of Claude (Anthropic) for complementary literature search, summaries, fact-checking and drafting; the choices and any errors are mine.