Mathematics in the Age of AI

Jeremy Avigad
Department of Philosophy
Department of Mathematical Sciences
Carnegie Mellon University

Institute for Computer-Aided Reasoning in Mathematics

November 6, 2025

Overview

New technologies for mathematics

I will discuss:

- interactive theorem proving and formalization
- automated reasoning and symbolic AI
- machine learning and neural AI

In the press, these are, collectively, "AI for Mathematics."

All three come together in neurosymbolic theorem proving.

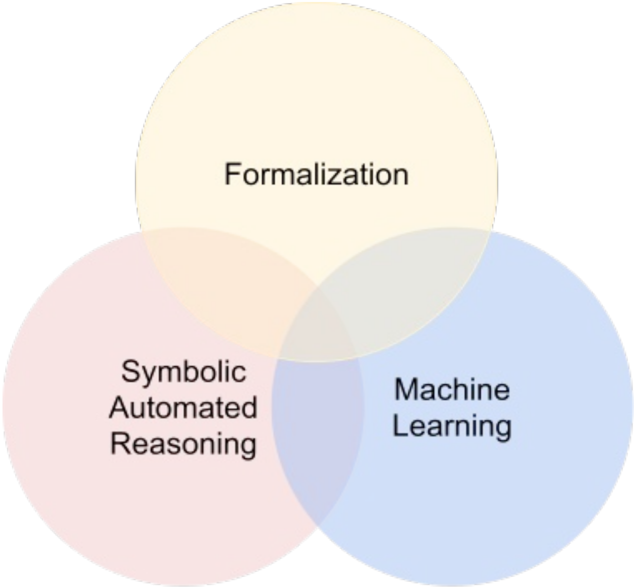

A.I. Is Coming for Mathematics, Too

For thousands of years, mathematicians have adapted to the latest advances in logic and reasoning. Are they ready for artificial intelligence?

Move Over, Mathematicians, Here Comes AlphaProof

A.I. is getting good at math — and might soon make a worthy collaborator for humans.

JUNE 8, 2024 | 12 MIN READ

AI Will Become Mathematicians' 'Co-Pilot'

Fields Medalist Terence Tao explains how proof checkers and AI programs are dramatically changing mathematics

BY CHRISTOPH DRÖSSER

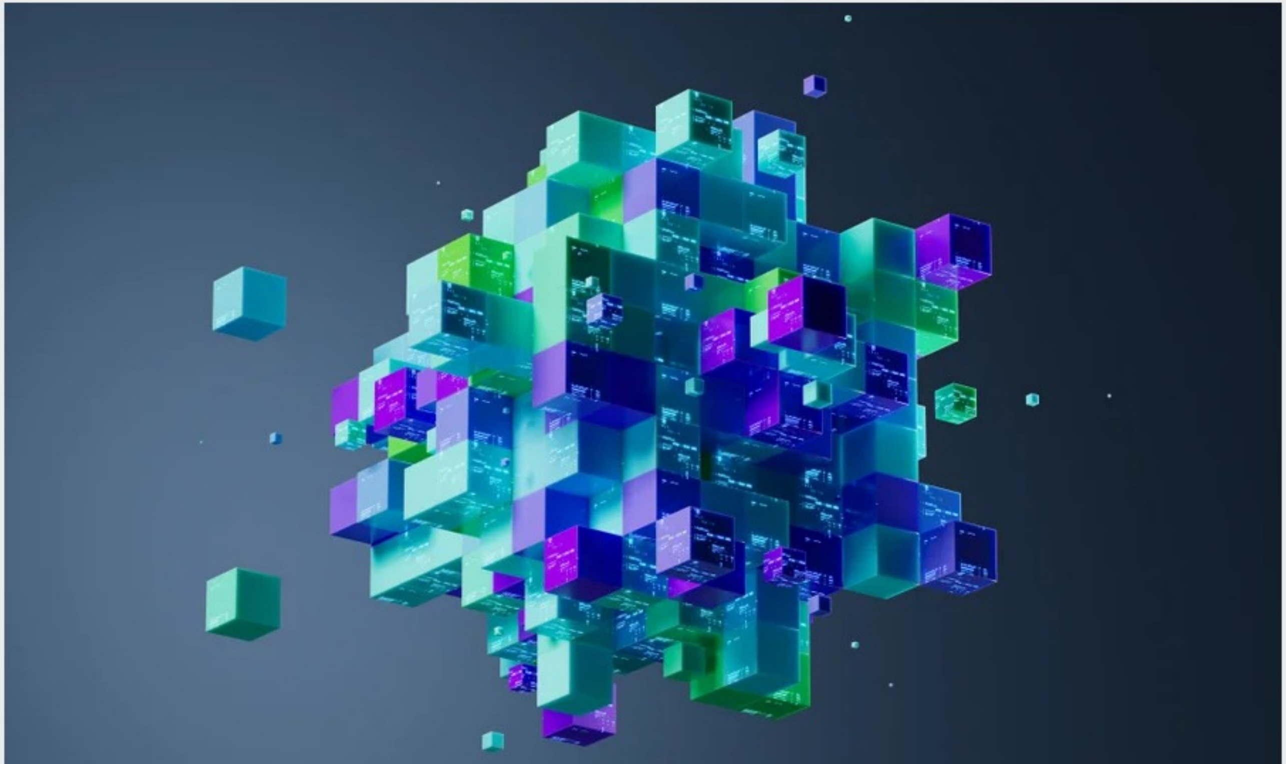

Overview

These technologies are still niche, but they are promising.

These technologies will impact mathematics:

- verification of mathematical results and mathematical computation
- communication and collaboration
- mathematical reference and search
- exploration and discovery of new mathematics
- teaching and learning

This is a survey of some of the new technologies.

Overview

Takeaway messages for mathematicians:

- The technologies hold a lot of promise for mathematics.
- They are interesting and fun.
- There are some things to worry about.
- We need to guide the next generation of mathematicians.
- Good outcomes for will require
  - collaboration between disciplines,
  - collaboration between generations, and
  - mindful attention.

Overview

Takeaway messages for computer scientists:

- Mathematics is great.
- There are lots of fun things you can do.
- There is a lot you can do to help.
- Advances in technology carry over to hardware and software verification, etc.
- Mathematical reasoning is important to the future of AI.
- Collaborating with mathematicians can be very rewarding.

New technologies for mathematics

I will discuss:

- interactive theorem proving
- automated reasoning and symbolic AI
- machine learning and neural AI
- neuro-symbolic theorem proving

I will choose examples from around 2021:

- not the earliest,
- not the most recent,
- not the most impressive to date, but
- turning points.

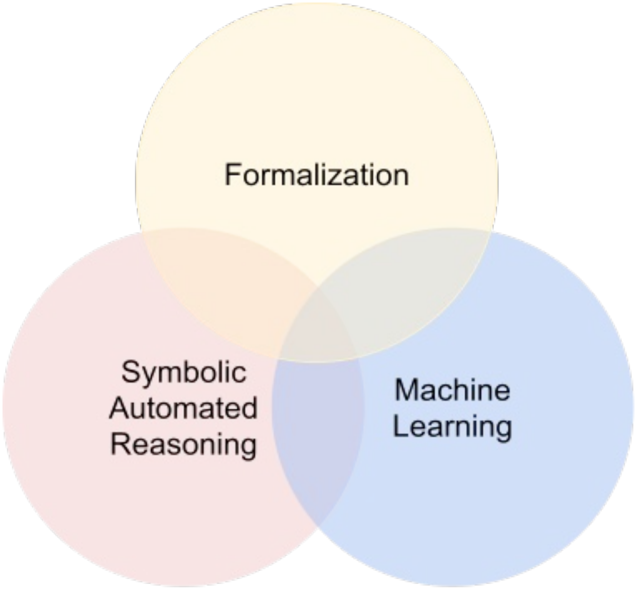

Interactive Theorem Proving

Formal methods

Formal methods are used in computer science to

- write specifications (for hardware, software, network protocols, ...), and
- verify that artifacts meet their specifications.

They rely on:

- formal languages
- formal semantics
- formal rules of inference.

Interactive theorem proving

In the early 20th century, logicians developed formal axiomatic systems for mathematics.

It soon became clear that these systems were expressive enough to formalize most mathematics, in principle.

In the early 1970s, the first proof assistants made it possible to formalize and verify proofs in practice.

Today, the practice is known as *interactive theorem proving*. Working with a proof assistant, users construct formal definitions and proofs

LEAN Community

Lean and its Mathematical Library

The Lean theorem prover is a proof assistant developed principally by Leonardo de Moura.

The community recently switched from using Lean 3 to using Lean 4. This website is still being updated, and some pages have outdated information about Lean 3 (these pages are marked with a prominent banner). The old Lean 3 community website has been archived.

The Lean mathematical library, *mathlib*, is a community-driven effort to build a unified library of mathematics formalized in the Lean proof assistant. The library also contains definitions useful for programming. This project is very active, with many regular contributors and daily activity.

You can get a bird's-eye view of what is in the mathlib library by reading the library overview, and read about recent additions on our blog. The design and community organization of mathlib are described in the 2020 article The Lean mathematical library, although the library has grown by an order of magnitude since that article appeared. You can also have a look at our repository statistics to see how the library grows and who contributes to it.

Try it!

You can try Lean in your web browser, install it in an isolated folder, or go for the full install. Lean is free, open source

Learn to Lean!

You can learn by playing a game, following tutorials, or reading books.

Meet the community!

Lean has very diverse and active community. It gathers mostly on a Zulip chat and on GitHub. You

```
File  Edit  Selection  View  Go  Run  ···        mathlib4
RieszLemma.lean ×                                   Lean InfoView   Lean InfoView ×
Mathlib > Analysis > NormedSpace > RieszLemma.lean > riesz_lemma
```

```lean
variable {𝕜 : Type*} [NormedField 𝕜]
variable {E : Type*} [NormedAddCommGroup E] [NormedSpace 𝕜 E]
variable {F : Type*} [SeminormedAddCommGroup F] [NormedSpace ℝ F]

/-- Riesz's lemma, which usually states that it is possible to find a
vector with norm 1 whose distance to a closed proper subspace is
arbitrarily close to 1. The statement here is in terms of multiples of
norms, since in general the existence of an element of norm exactly 1
is not guaranteed. For a variant giving an element with norm in `[1, R]`, see
`riesz_lemma_of_norm_lt`. -/
theorem riesz_lemma {F : Subspace 𝕜 E} (hFc : IsClosed (F : Set E)) (hF : ∃ x : E, x ∉ F) {r : ℝ}
    (hr : r < 1) : ∃ x₀ : E, x₀ ∉ F ∧ ∀ y ∈ F, r * ‖x₀‖ ≤ ‖x₀ - y‖ := by
  classical
    obtain ⟨x, hx⟩ : ∃ x : E, x ∉ F := hF
    let d := Metric.infDist x F
    have hFn : (F : Set E).Nonempty := ⟨_, F.zero_mem⟩
    have hdp : 0 < d :=
      lt_of_le_of_ne Metric.infDist_nonneg fun heq =>
        hx ((hFc.mem_iff_infDist_zero hFn).2 heq.symm)
    let r' := max r 2⁻¹
    have hr' : r' < 1 := by
      simp only [r', max_lt_iff, hr, true_and]
      norm_num
    have hlt : 0 < r' := lt_of_lt_of_le (by simp) (le_max_right r 2⁻¹)
    have hdlt : d < d / r' := (lt_div_iff₀ hlt).mpr ((mul_lt_iff_lt_one_right hdp).2 hr')
    obtain ⟨y₀, hy₀F, hxy₀⟩ : ∃ y ∈ F, dist x y < d / r' := (Metric.infDist_lt_iff hFn).mp hdlt
    have x_ne_y₀ : x - y₀ ∉ F := by
      by_contra h
      have : x - y₀ + y₀ ∈ F := F.add_mem h hy₀F
      simp only [neg_add_cancel_right, sub_eq_add_neg] at this
      exact hx this
    refine ⟨x - y₀, x_ne_y₀, fun y hy => le_of_lt ?_⟩
    have hy₀y : y₀ + y ∈ F := F.add_mem hy₀F hy
    calc
```

```
▼ RieszLemma.lean:55:4
▼ Tactic state
1 goal
𝕜 : Type u_1
inst✝² : NormedField 𝕜
E : Type u_2
inst✝¹ : NormedAddCommGroup E
inst✝ : NormedSpace 𝕜 E
F : Subspace 𝕜 E
hFc : IsClosed ↑F
r : ℝ
hr : r < 1
x : E
hx : x ∉ F
d : ℝ := infDist x ↑F
hFn : (↑F).Nonempty
hdp : 0 < d
r' : ℝ := max r 2⁻¹
hr' : r' < 1
hlt : 0 < r'
hdlt : d < d / r'
y₀ : E
hy₀F : y₀ ∈ F
hxy₀ : dist x y₀ < d / r'
⊢ ∃ x₀ ∉ F, ∀ y ∈ F, r * ‖x₀‖ ≤ ‖x₀
- y‖
▶ Messages (1)
▶ All Messages (0)
Restart File

master  ⊗ 0 ⚠ 0        Ln 55, Col 5   Spaces: 2   UTF-8   LF   {} lean4
```

FOUNDATIONS OF MATHEMATICS

Building the Mathematical Library of the Future

A small community of mathematicians is using a software program called Lean to build a new digital repository. They hope it represents the future of their field.

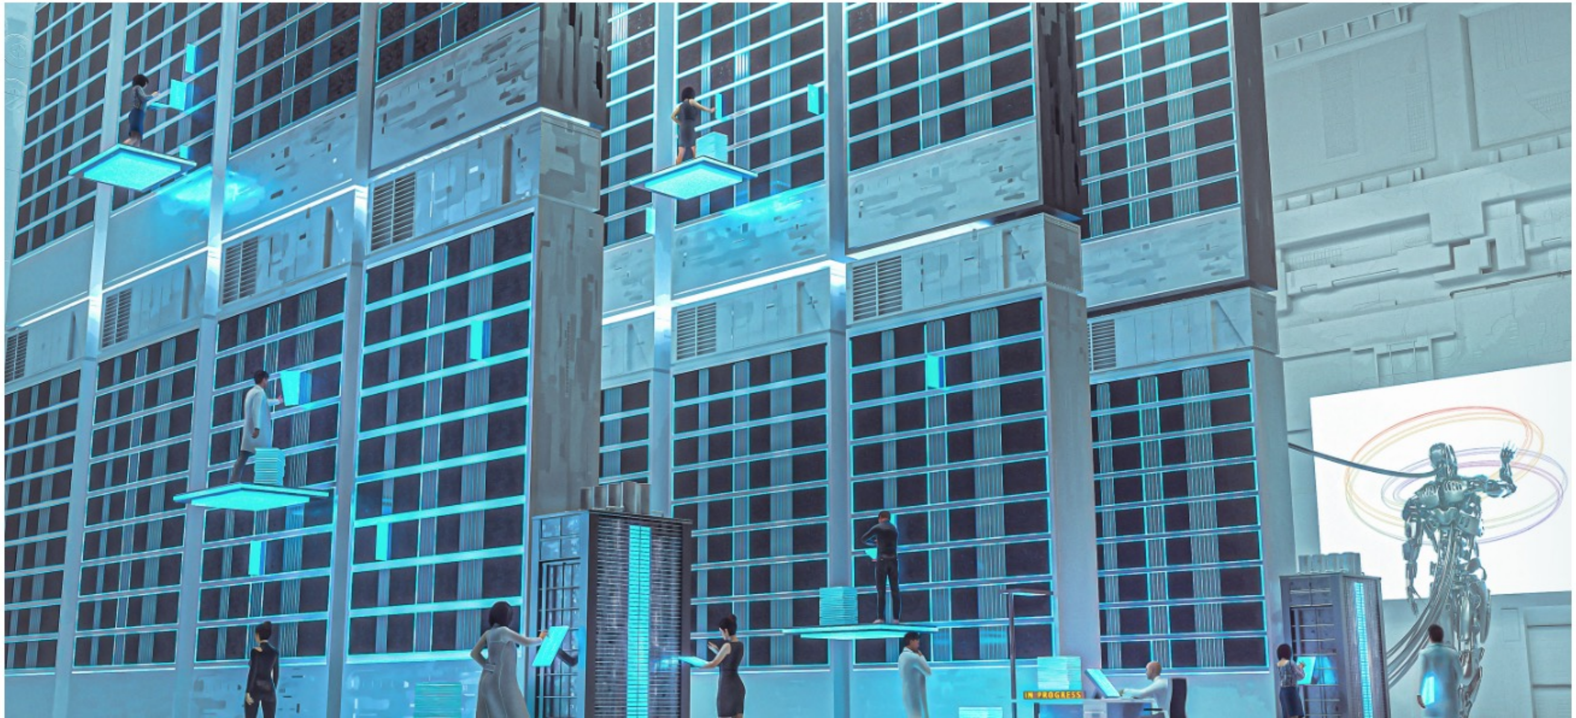

Liquid tensor experiment

Posted on December 5, 2020 by xenaproject

This is a guest post, written by Peter Scholze, explaining a liquid real vector space mathematical formalisation challenge. For a pdf version of the challenge, see here. For comments about formalisation, see section 6. Now over to Peter.

1. The challenge

I want to propose a challenge: Formalize the proof of the following theorem.

Theorem 1.1 *(Clausen-S.) Let* $0 < p' < p \leq 1$ *be real numbers, let* S *be a profinite set, and let* V *be a* p*-Banach space. Let* $\mathcal{M}_{p'}(S)$ *be the space of* p'*-measures on* S*. Then*

$$\mathrm{Ext}^i_{\mathrm{Cond(Ab)}}(\mathcal{M}_{p'}(S), V) = 0$$

for $i \geq 1$.

Comment

Half a year of the Liquid Tensor Experiment: Amazing developments

Posted on June 5, 2021 by xenaproject

[This is a guest post by Peter Scholze.]

Exactly half a year ago I wrote the Liquid Tensor Experiment blog post, challenging the formalization of a difficult foundational theorem from my Analytic Geometry lecture notes on joint work with Dustin Clausen. While this challenge has not been completed yet, I am excited to announce that the Experiment has verified the entire part of the argument that I was unsure about. I find it absolutely insane that interactive proof assistants are now at the level that within a very reasonable time span they can formally verify difficult original research. Congratulations to everyone involved in the formalization!!

1.1 Breen–Deligne data

The goal of this subsection is to a give a precise statement of a variant of the Breen–Deligne resolution. This variant is not actually a resolution, but it is sufficient for our purposes, and is much easier to state and prove.

We first recall the original statement of the Breen–Deligne resolution.

Theorem(Breen–Deligne)

For an abelian group A, there is a resolution, functorial in A, of the form

$$\dots \longrightarrow \bigoplus_{i=1}^{n_i} \mathbb{Z}[A^{r_{ij}}] \longrightarrow \dots \longrightarrow \mathbb{Z}[A^3] \oplus \mathbb{Z}[A^2] \longrightarrow \mathbb{Z}[A^2] \longrightarrow \mathbb{Z}[A] \longrightarrow A \longrightarrow 0.$$

What does a homomorphism $f\colon \mathbb{Z}[A^m] \to \mathbb{Z}[A^n]$ that is functorial in A look like? We should perhaps say more precisely what we mean by this. The idea is that m and n are fixed, and for each abelian group A we have a group homomorphism $f_A\colon \mathbb{Z}[A^m] \to \mathbb{Z}[A^n]$ such that if $\phi\colon A \to B$ is a group homomorphism inducing $\phi_i\colon \mathbb{Z}[A^i] \to \mathbb{Z}[B^i]$ for each natural number i then the obvious square commutes: $\phi_n \circ f_A = f_B \circ \phi_m$.

The map f_A is specified by what it does to the generators $(a_1, a_2, a_3, \dots, a_m) \in A^m$. It can send such an element to an arbitrary element of $\mathbb{Z}[A^n]$, but one can check that universality implies that f_A will be a $\mathbb{Z}$-linear combination of “basic universal maps”, where a “basic universal map” is one that sends $(a_1, a_2, \dots, a_m)$ to $(t_1, \dots, t_n)$, where t_i is a $\mathbb{Z}$-linear combination $c_{i,1} \cdot a_1 + \dots + c_{i,m} \cdot a_m$. So a “basic universal map” is specified by the $n \times m$-matrix c.

Definition 1.1.1 ✓

A *basic universal map* from exponent m to n, is an $n \times m$-matrix with coefficients in $\mathbb{Z}$
.

Dependencies

Legend ≡

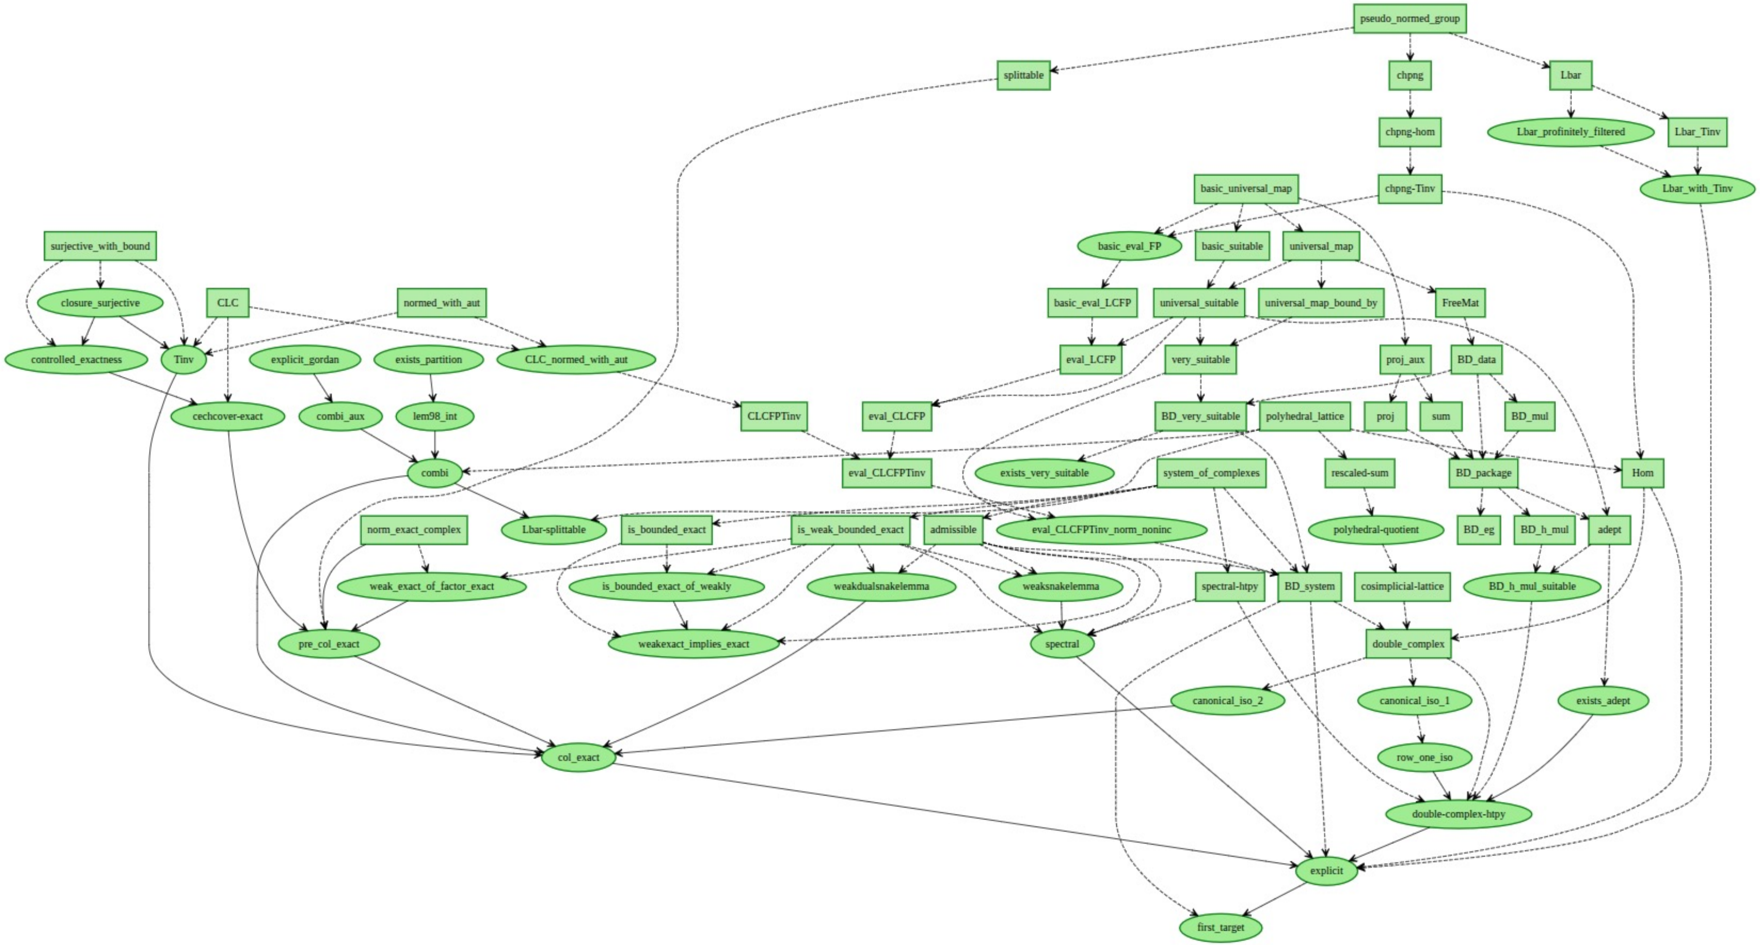

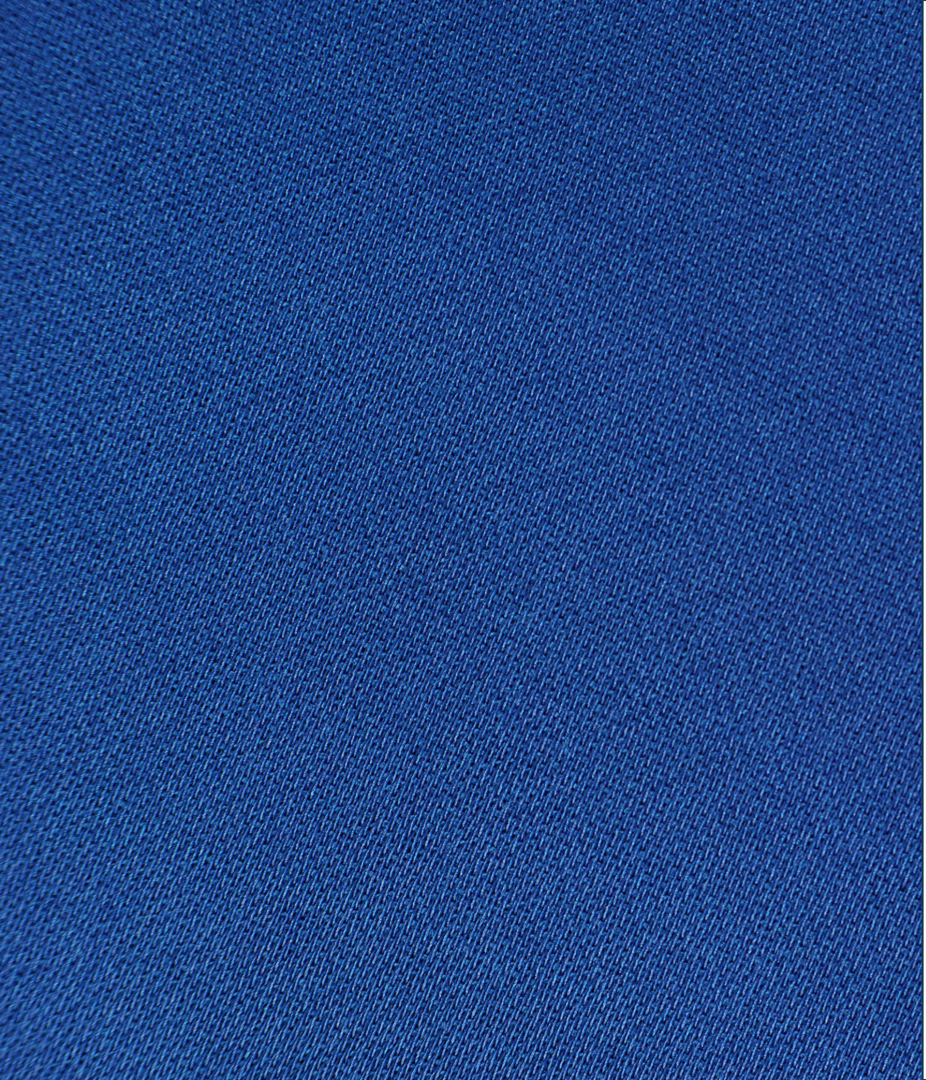

Automated Reasoning

Automated reasoning

Even before computers were invented, logicians were interested in algorithmic procedures to

- decide the truth of mathematical statements, and
- search for proofs.

The first automated provers appeared in the 1950s and 1960s.

Now we have

- first-order provers,
- SAT solvers, and
- SMT solvers.

SAT solvers

A formula in propositional logic (like $P \vee Q \rightarrow Q \wedge R$) is true or false depending on the truth assignments to the variables.

A satisfiability solver determines whether a formula has a satisfying assignment.

Modern SAT solvers can decide industrial formulas with tens of millions of variables and hundreds of millions of clauses, often in minutes.

Recipe for mathematics:

- Encode / reduce a problem to a SAT problem.
- Use a SAT solver.

Keller's conjecture

In 1930, Ott-Heinrich Keller conjectured that for every *n*, any tiling of *n*-dimensional space with unit *n*-dimensional cubes must have at least two cubes that fully share a face.

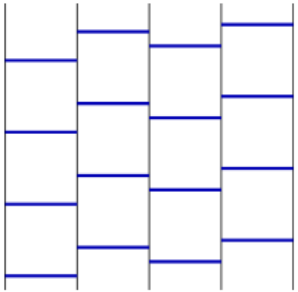

Figure 1: Two-dimensional tiling

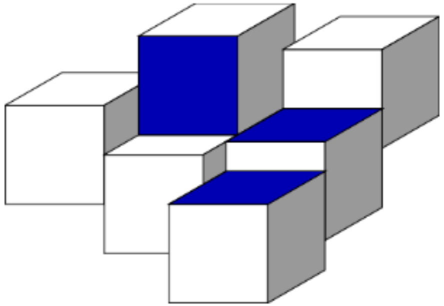

Figure 2: Three-dimensional tiling

In 1940, Perron showed that the conjecture is **true** up to dimension 6.

Keller's conjecture

In 1990, Corradi and Szabo showed that the conjecture is true if and only if there are no cliques of a certain size in certain associated graphs, now known as *Keller graphs*.

In 1992, Lagarias and Shor showed that the conjecture is false in dimensions 10 and up.

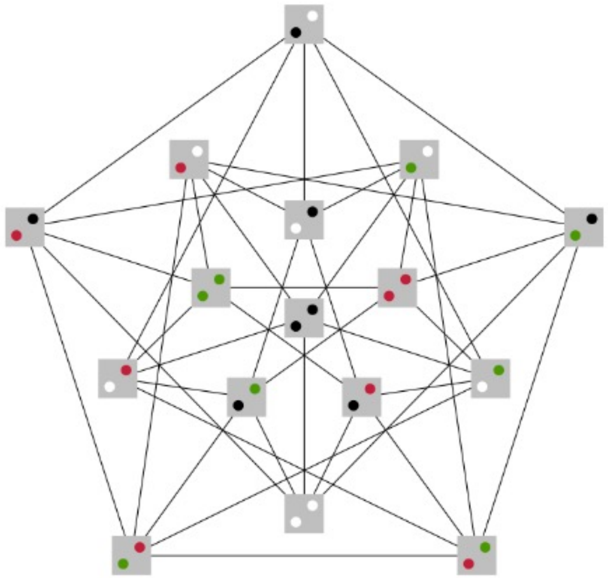

Figure 3: a Keller graph

Keller's conjecture

It may be a difficult matter to determine exactly the critical dimension. Exhaustive search for Szabó-type counterexamples already seems infeasible for G_7^*; the maximum clique problem is a well-known NP-complete problem, which is also computationally hard in practice. The authors ruled out the existence of any 2^7-clique in G_7^* that is invariant under a cyclic permutation of coordinates by computer search. It is conceivable that there exist Szabó-type counterexamples in dimension 7, 8, or 9, which are all so structureless that they will be hard to find. In any case we have so far found no variant of the constructions of Theorem A that work in these dimensions.

In 2002, Mackey showed that it is false in dimensions 8 and 9.

Keller's conjecture

In 2020, Brakensiek, Heule, Mackey, and Narváez showed that there is no counterexample in dimension 7.

They used a SAT solver.

The search space is huge; they used additional reductions and symmetry breaking, to rule out the result.

The images are taken from their web page.

Keller's conjecture

Should we trust the result?

Contemporary SAT solvers produce proofs of unsatisfiability that can be checked independently.

- Joshua Clune (and then James Gallicchio) verified the reduction to a combinatorial problem in Lean.
- James Gallicchio verified the encoding as a propositional formula.
- The UNSAT result was checked by an independent checker.
- The correctness of the checker was verified by Cayden Codel.

Machine Learning

Machine Learning

Key approaches:

- **Supervised learning:** the system is presented with (input, output) pairs, and learns a rule connecting them.
- **Unsupervised learning:** the system is presented with data, and learns some sort of structure.
- **Reinforcement learning:** the system acts in a space and is rewarded accordingly; it learns to maximize rewards.

Models can be very simple (linear regression, decision trees) to very complex (neural networks).

Reinforcement learning

Before the explosion of LLMs, Google DeepMind achieved landmark successes on training neural networks using reinforcement learning.

- By December 2013, they had a system that could learn to play Atari 2600 games and surpass human performance on three of them.
- In March 2016, AlphaGo beat go champion Lee Sedol.
- In October 2017, they published an article in *Nature* on AlphaGo Zero, which was trained without using data from human games.
- Soon after, AlphaZero was able to master chess, shogi, and go trained entirely with self-play.

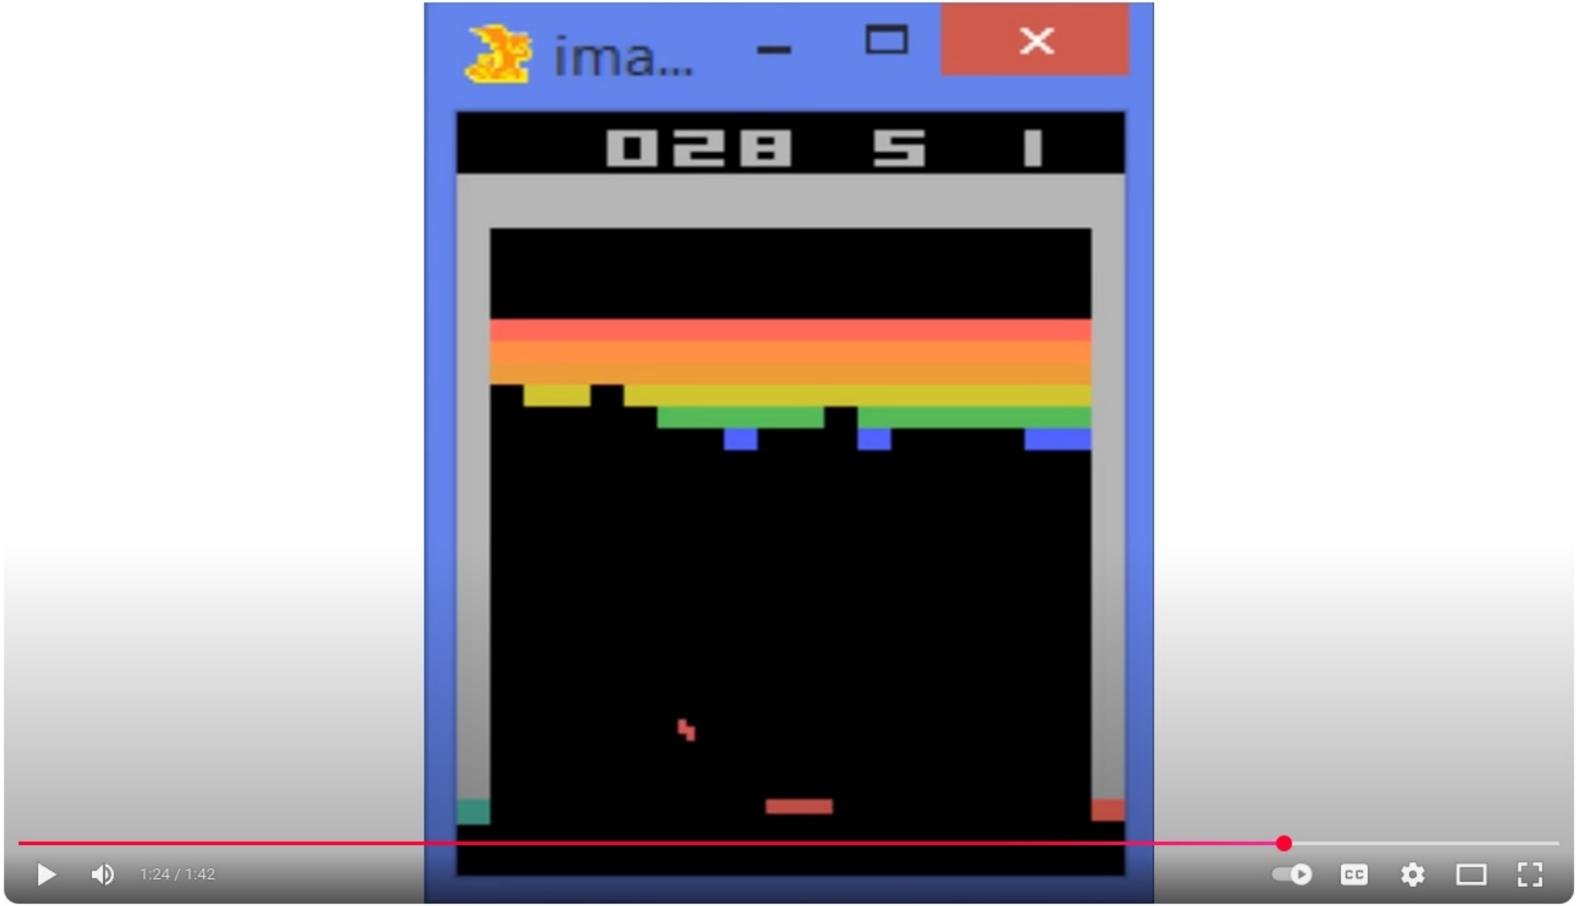

Google DeepMind's Deep Q-learning playing Atari Breakout!

Two Minute Papers
1.71M subscribers

Join Subscribe

7K Share Download Thanks

1.4M views 10 years ago

The game of Go is the most challenging of classic games. Despite decades of effort, prior methods had only achieved amateur level performance. We developed a deep RL algorithm that learns both a value network (which predicts the winner) and a policy network (which selects actions) through games of self-play. Our program AlphaGo combined these deep neural networks with a state-of-the-art tree search. In October 2015, AlphaGo became the first program to defeat a professional human player. In March 2016, AlphaGo defeated Lee Sedol (the strongest player of the last decade with an incredible 18 world titles) by 4 games to 1, in a match that was watched by an estimated 200 million viewers.

Applications to mathematics

In 2021, Adam Zsolt Wagner published a paper in which he used reinforcement learning to find counterexamples to several graph-theoretic conjectures.

The method:

- Cook up a suitable reward function.
- Ask a network to generate graphs and select the ones with the highest score.
- Update the network to nudge it in the direction of moves that generated these graphs.
- Iterate.

Constructions in combinatorics via neural networks

Adam Zsolt Wagner*

Abstract

We demonstrate how by using a reinforcement learning algorithm, the deep cross-entropy method, one can find explicit constructions and counterexamples to several open conjectures in extremal combinatorics and graph theory. Amongst the conjectures we refute are a question of Brualdi and Cao about maximizing permanents of pattern avoiding matrices, and several problems related to the adjacency and distance eigenvalues of graphs.

1 Introduction

Computer-assisted proofs have a long history in mathematics, including breakthrough results such as the proof of the four color theorem in 1976 by Appel and Haken [7], and the proof of the Kepler conjecture in 1998 by Hales [29]. Recently, significant progress has been made in the area of machine learning algorithms, and they have have quickly become some of the most exciting tools in a scientist's toolbox. In particular, recent advances in the field of reinforcement learning have led computers to

- In Section 2 we present our results we obtained via the deep cross-entropy method.
  - In §2.1 we give a short introduction to the cross-entropy method and describe how we will use it to produce constructions to extremal combinatorics problems.
  - In §2.2 we illustrate the method by finding counterexamples to a conjecture about the sum of the largest eigenvalue and matching number of graphs, which was proposed in [4].
  - In §2.3 we refute a similar conjecture of Aouchiche–Hansen [6] about the distance spectrum and proximity of graphs.
  - In §2.4 we refute an old conjecture of Collins [18] by showing that the peaks of the coefficient sequences of the adjacency and distance polynomials of trees can be far apart.
  - In §2.5 we show that transmission regularity of graphs is not preserved under cospectrality of the distance Laplacian, answering a question of Hogben and Reinhart [31].
  - In §2.6 we address a problem of Brualdi and Cao [15] about maximizing the permanent of an $n \times n$, 312-pattern avoiding binary matrix. Among others, we find that the best possible answers for $n \leq 8$ are given by the rather remarkable sequence

$$1, \quad 2, \quad 4, \quad 8, \quad 16, \quad 32, \quad 64, \quad 120.$$

- In Section 3 we present two constructions obtained via LP solvers.
  - In §3.1 we refute a conjecture of Aaronson–Groenland–Grzesik–Kielak–Johnston [2] about a problem of covering certain subsets of the hypercube with few hyperplanes.
  - In §3.2 we answer a problem of Király–Nagy–Pálvölgyi–Visontai [32] about the maximum

An Example

Conjecture. Let G be a connected graph on $n \geq 3$ vertices, with largest eigenvalue λ_1 and matching number μ. Then $\lambda_1 + \mu \geq \sqrt{n-1} + 1$.

To refute the conjecture:

- Make the reward function $\lambda_1 + \mu$.
- Fix n, and ask the system to generate graphs on n vertices.
- Let the system learn to minimize it.

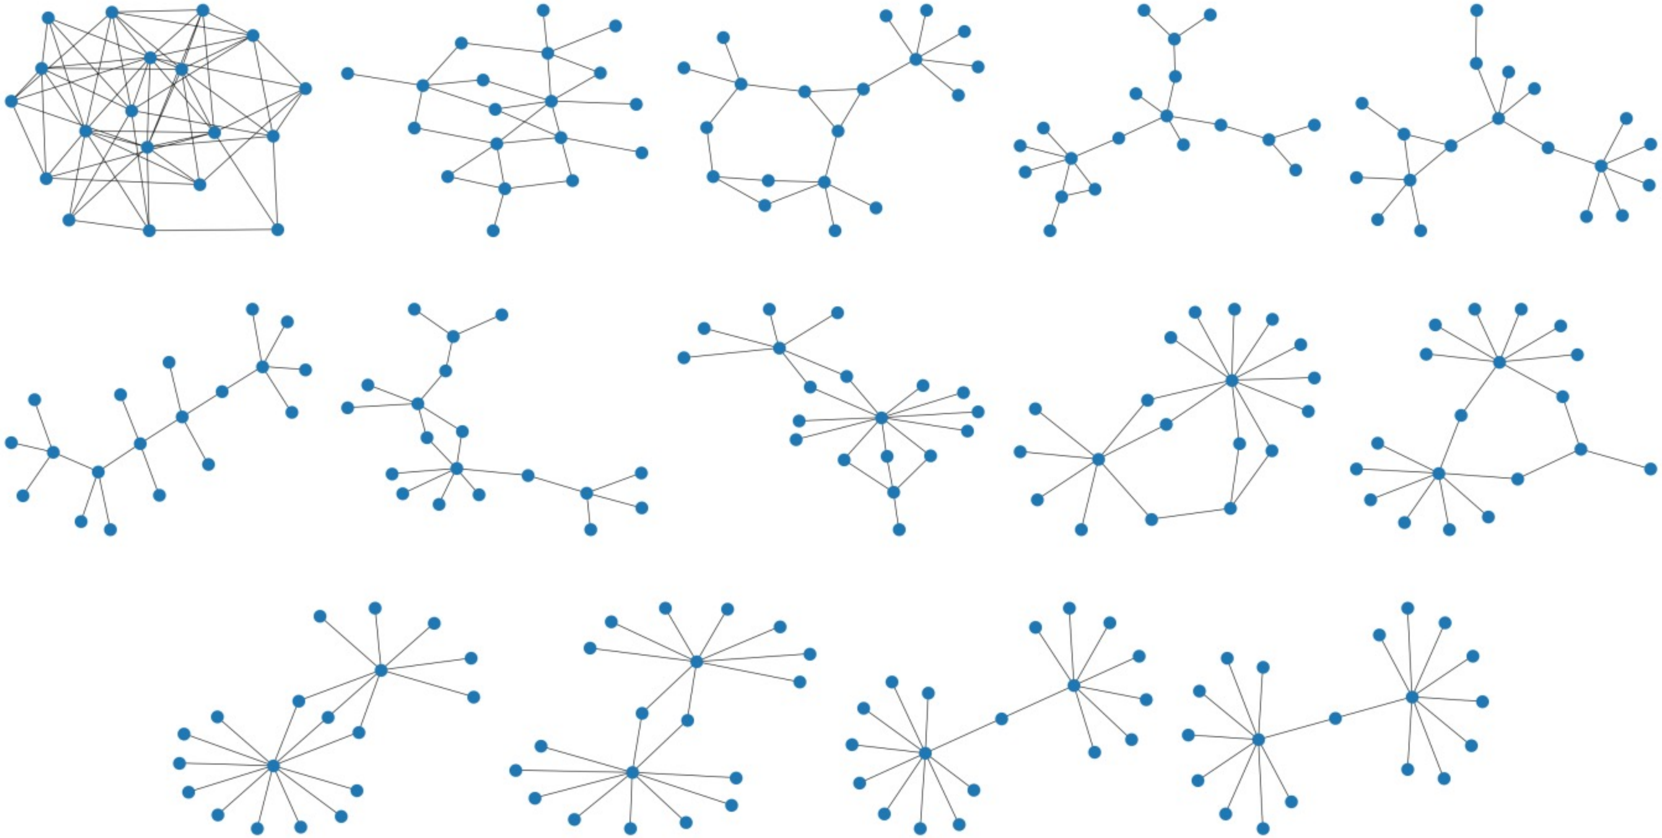

Figure 3: The evolution of the best construction over time. The network quickly realizes that sparse graphs are best, and eventually the "balanced double star" structure emerges.

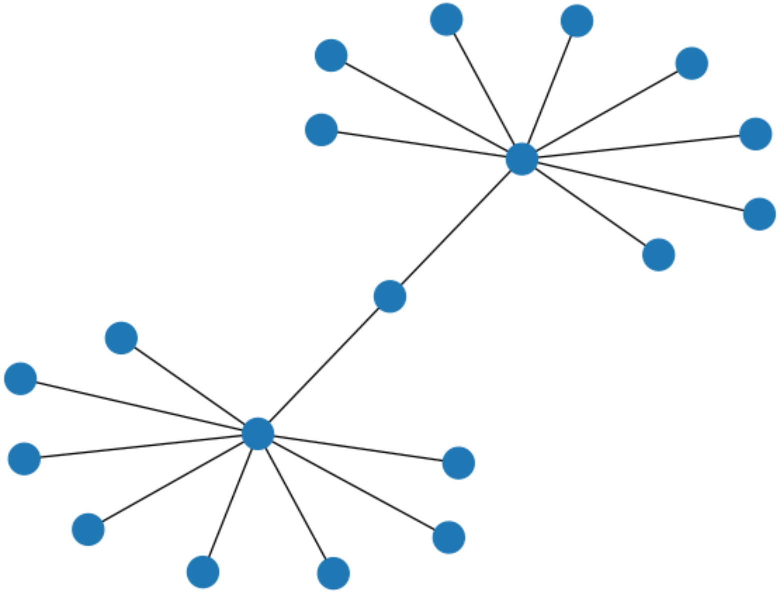

Figure 4: A graph on 19 vertices satisfying $\lambda_1 + \mu < \sqrt{n-1} + 1$.

Applications to mathematics

Adam’s online talks are highly recommended.

Since then, more sophisticated methods have been developed:

- PatternBoost
- FunSearch
- AlphaEvolve
- OpenEvolve
- ShinkaEvolve

Interlude

Skepticism

Keller's conjecture and graph theory deal with finite combinatorial objects.

Real mathematics deals with infinitary objects and spaces.

Does AI help us do *real* mathematics?

Optimism

Don't devalue the finite. In the long run we are all dead.

But even for a fixed dimension, Keller’s conjecture is about infinite configurations.

Finite objects and expressions, like invariants, can tell us things about the infinite.

Thanks to formalization:

- Any mathematical object can be described with a finite expression.
- Any mathematical claim is a finite expression.
- A mathematical proof is a finite expression.

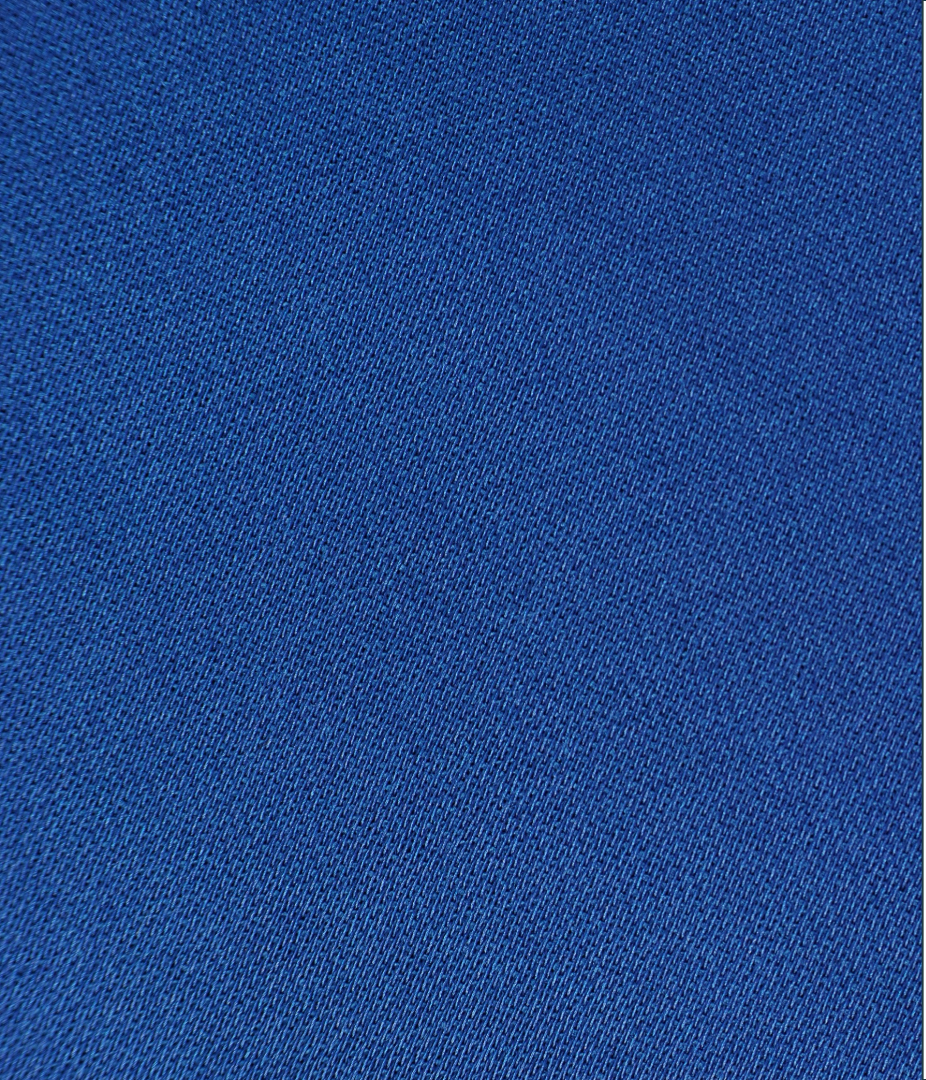

Neurosymbolic Theorem Proving

IMO Grand Challenge

The International Mathematical Olympiad (IMO) is perhaps the most celebrated mental competition in the world and as such is among the ultimate grand challenges for Artificial Intelligence (AI).

The challenge: build an AI that can win a gold medal in the competition.

To remove ambiguity about the scoring rules, we propose the formal-to-formal (F2F) variant of the IMO: the AI receives a *formal* representation of the problem (in the Lean Theorem Prover), and is required to emit a *formal* (i.e. machine-checkable) proof. We are working on a proposal for encoding IMO problems in Lean and will seek broad consensus on the protocol.

Other proposed rules:

Credit. Each proof certificate that the AI produces must be checkable by the Lean kernel in 10 minutes (which is approximately the amount of time it takes a human judge to judge a human's solution). Unlike human competitors, the AI has no opportunity for partial credit.

ABSTRACTIONS BLOG

At the Math Olympiad, Computers Prepare to Go for the Gold

15

Computer scientists are trying to build an AI system that can win a gold medal at the world's premier math competition.

SCIENCE

AI achieves silver-medal standard solving International Mathematical Olympiad problems

25 JULY 2024

AlphaProof and AlphaGeometry teams

Share

Score on IMO 2024 problems

Graph showing performance of our AI system relative to human competitors at IMO 2024. We earned 28 out of 42 total points, achieving the same level as a silver medalist in the competition.

AlphaProof: a formal approach to reasoning

The IMO Grand Challenge today

After the 2025 IMO, four groups claimed gold medal performance:

- Harmonic AI (formal)
- ByteDance (formal)
- OpenAI (informal)
- Google DeepMind (informal)

ByteDance's SeedProver solves 78.1% of formalized past IMO problems, and more than 50% on PutnamBench.

On September 26, a group at Apple and UC San Diego claimed 70% on PutnamBench with its publicly available Hilbert prover.

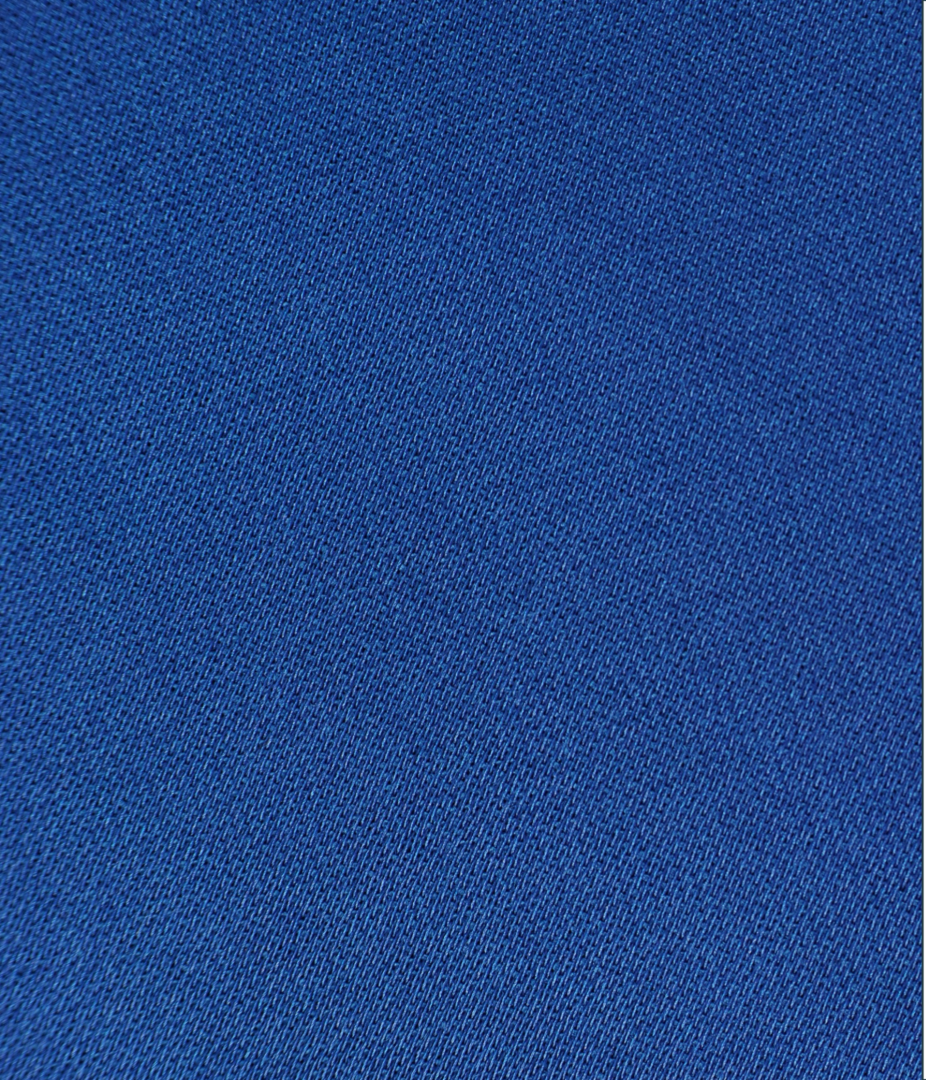

Concerns

What do I worry about?

- AI in society
    - Concerns about reliability, explainability, alignment, access
    - Economic and social concerns

  Keeping mathematics in the loop is part of the solution.

- AI in mathematics
    - Changes to the discipline: as AI gets better at discovering patterns and proving theorems, what's left for us?
    - Ceding mathematical thought to big tech
    - Access to resources for mathematicians
    - Access to mathematics for the public at large

Institute for Computer-Aided Reasoning in Mathematics

NSF

NSF News

NSF invests over $74 million in 6 mathematical sciences research institutes

From improving medical care to detecting planets in other solar systems, the institutes will explore mathematical sciences with a broad range of applications

August 4, 2025

The U.S. National Science Foundation is investing over $74 million in six research institutes focused on the mathematical sciences and their broad applications in all fields of science, technology and many industries.

For over 40 years, NSF has funded Mathematical Sciences Research Institutes to serve as catalysts for U.S. research in mathematics and statistics and to produce mathematical innovations to rapidly address new and emerging challenges and opportunities. The institutes collectively investigate a wide range of mathematical research areas with potential impacts, including better patient outcomes in hospital emergency rooms, enhanced safety of semiautonomous vehicles, and detection of exoplanets using quantum physics. Previous research conducted at the institutes has had broad impacts, such as improved speed and accuracy of MRI imaging and the development of mathematical foundations of artificial intelligence-based technologies.

Mission

The mission of the *Institute for Computer-Aided Reasoning in Mathematics* is to:

- empower mathematicians to take advantage of new technologies for mathematical reasoning and keep mathematics central to everything we do;
- unite mathematicians of all kinds, computer scientists, students, and researchers, to achieve that goal; and
- ensure that mathematics and the new technologies are accessible to everyone.

Motivation

In 2023, a workshop by the *National Academies for Science, Engineering, and Medicine* explored the promise of these technologies and the challenges that lie ahead.

The Institute for Computer-Aided Reasoning in Mathematics is designed to meet the challenges.

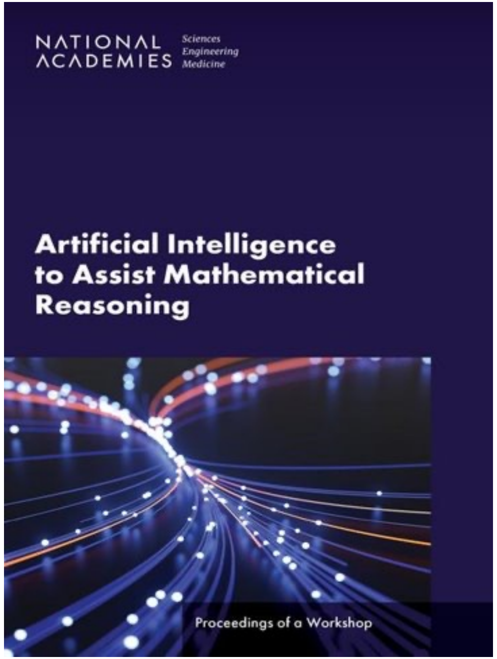

Empowering Mathematicians

Some challenges:

- Existing tools aren't designed for mathematicians.
- Documentation isn't written for mathematicians.
- Mathematicians don't have the relevant expertise.
- Mathematicians don't have time to learn to use AI.
- Collaborations are needed between computer scientists and mathematicians.
- Nobody "owns" AI for mathematics; it falls through the cracks.
- Some of the work is tedious, doesn't yield academic credit.
- The mathematics community doesn't know how to support/assess mathematicians using AI.

Empowering Mathematicians

We will maintain a staff of innovation engineers who will:

- help mathematicians learn to use the technologies
- answer questions and provide technical support
- maintain documentation, tools, infrastructure, and other community resources
- serve as liaisons to computer science and industry
- carry out essential tasks that academics don’t have time or incentives to do
- be community leaders in the use of technology
- gather resources and coordinate efforts.

Bringing us together

We will also provide:

- workshops
- summer schools
- collaborative visits
- an annual conference

These will build a community of students, researchers, mathematicians, computer scientists, engineers, and others to address the challenges together.

We need a combination of perspectives and expertise.

Improving access

AI and the digitization of mathematics can lead to greater democratization but it can also lead to greater inequities.

A central goal of ICARM is to ensure that all communities have the resources they need to participate in mathematics and take advantage of the new technologies.

Our original proposal included a summer school for college students, a workshop for graduate students, and an after-school program for high school students to address this challenge head on.

Current status

The institute has been launched as a three-year pilot:

- 2-3 administrative staff
- 3 innovation engineers
- At least two workshops each year
- At least one summer school each year
- A conference in the second year
- Collaborative visits

Current status

We are:

- Setting up administrative and financial infrastructure within CMU
- Constituting our governing boards
- Setting up our space
- Setting up our web pages and computing infrastructure
- Hiring staff and innovation engineers
- Starting to plan our first activities and events
- Collaborating with the other institutes

At the Joint Mathematics Meetings, we will hold tutorials, participate in the institutes’ reception, and have a booth.

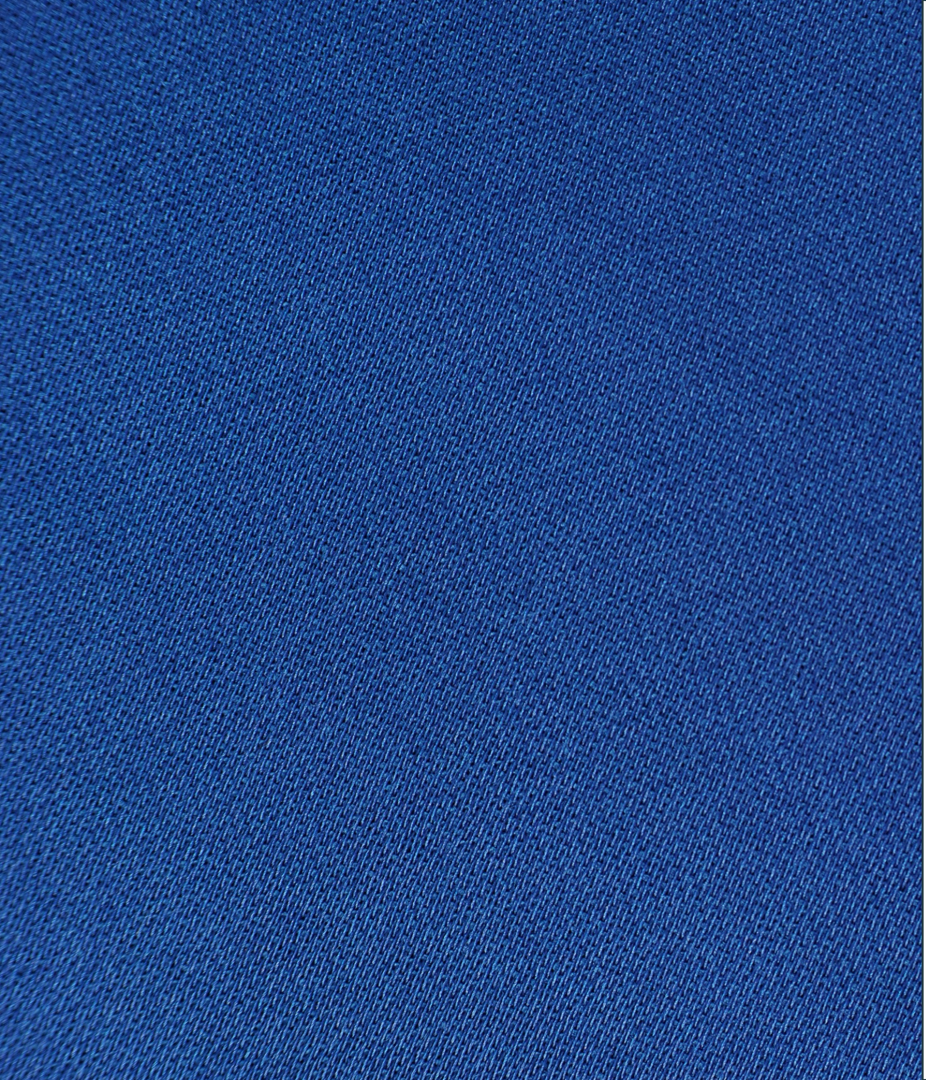

Conclusions

Recap

We have considered examples of various technologies for mathematics:

- interactive theorem proving and formalization
- automated reasoning and symbolic AI
- machine learning and neural AI

These technologies are still niche, but they are promising.

All three come together in neurosymbolic systems that conjecture and prove theorems.

Recap

These technologies will impact mathematics:

- verification of mathematical results and mathematical computation
- communication and collaboration
- mathematical reference and search
- exploration and discovery of new mathematics
- teaching and learning

We need to help the next generation of mathematicians navigate the changes.

Final thoughts

“Today we serve technology. We need to reverse the machine-centered point of view and turn it into a person-centered point of view: Technology should serve us.”

From *Things That Make Us Smart: Defending Human Attributes in the Age of the Machine*, by Donald A. Norman (1994)

The question is not “how can mathematicians use the technology?” but rather “what can technology do for mathematicians?”

Institute for Computer-Aided Reasoning in Mathematics

This is an exciting time for mathematics. Let’s make the most of it.