ISAAC NEWTON INSTITUTE
”New connections between physics and number theory”
15 July, 2026

→ *meet the agent*

reads · proposes · checks

Beyond the chatbot

How I use AI in my mathematical research

Daniel Persson

Chalmers University of Technology

HOW THE LAST FEW YEARS HAVE FELT

The water is rising

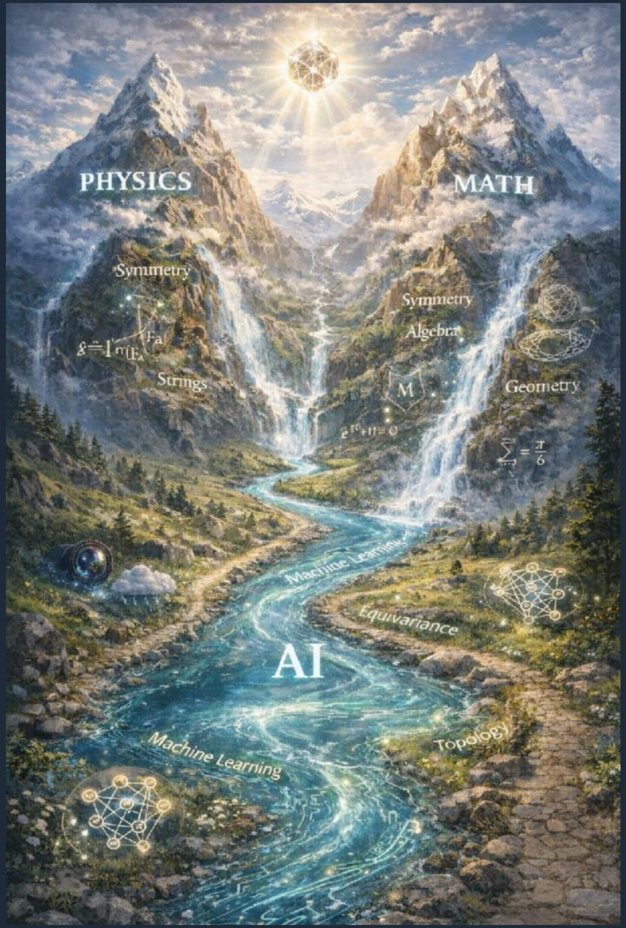

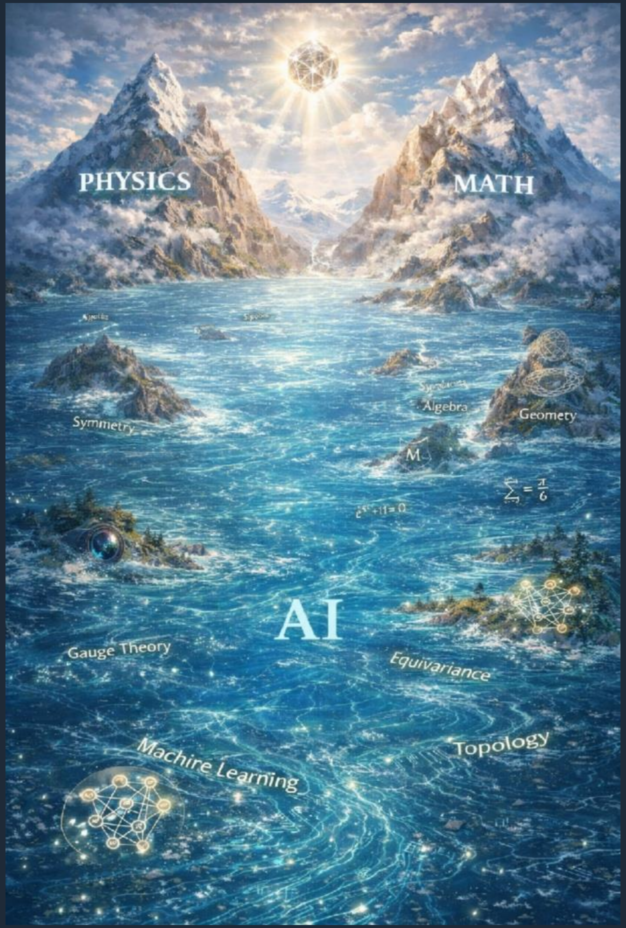

MY RESEARCH GROUP

MATH & PHYSICS → AI

Mathematics and physics for deep learning

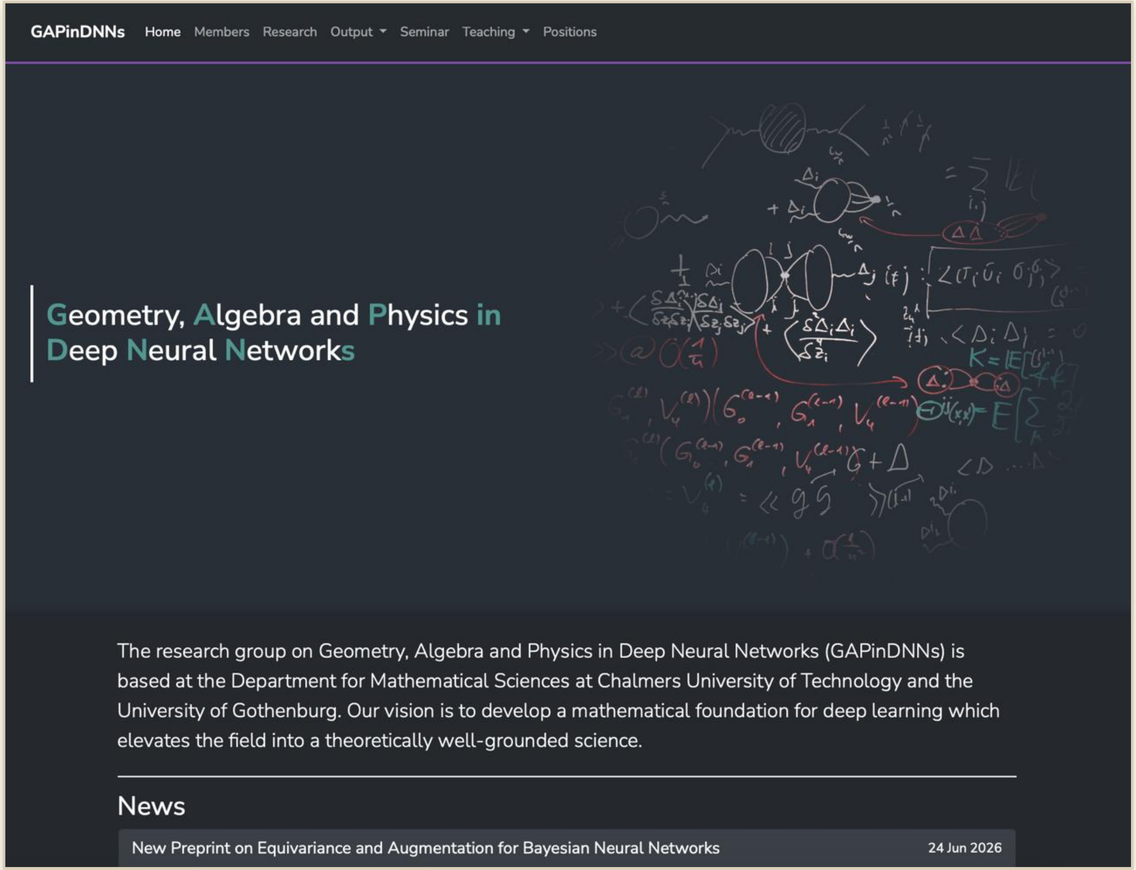

Geometry, Algebra and Physics in Deep Neural Networks

- **GAPinDNNs** — our research group at Chalmers & the University of Gothenburg
- **Vision:** a mathematical foundation for deep learning
- Symmetry, geometry and physics as guiding principles for neural networks

the direction: mathematics & physics in service of AI

MY RESEARCH GROUP · EXAMPLE

MATH & PHYSICS → AI

Inductive bias: build the math into the network

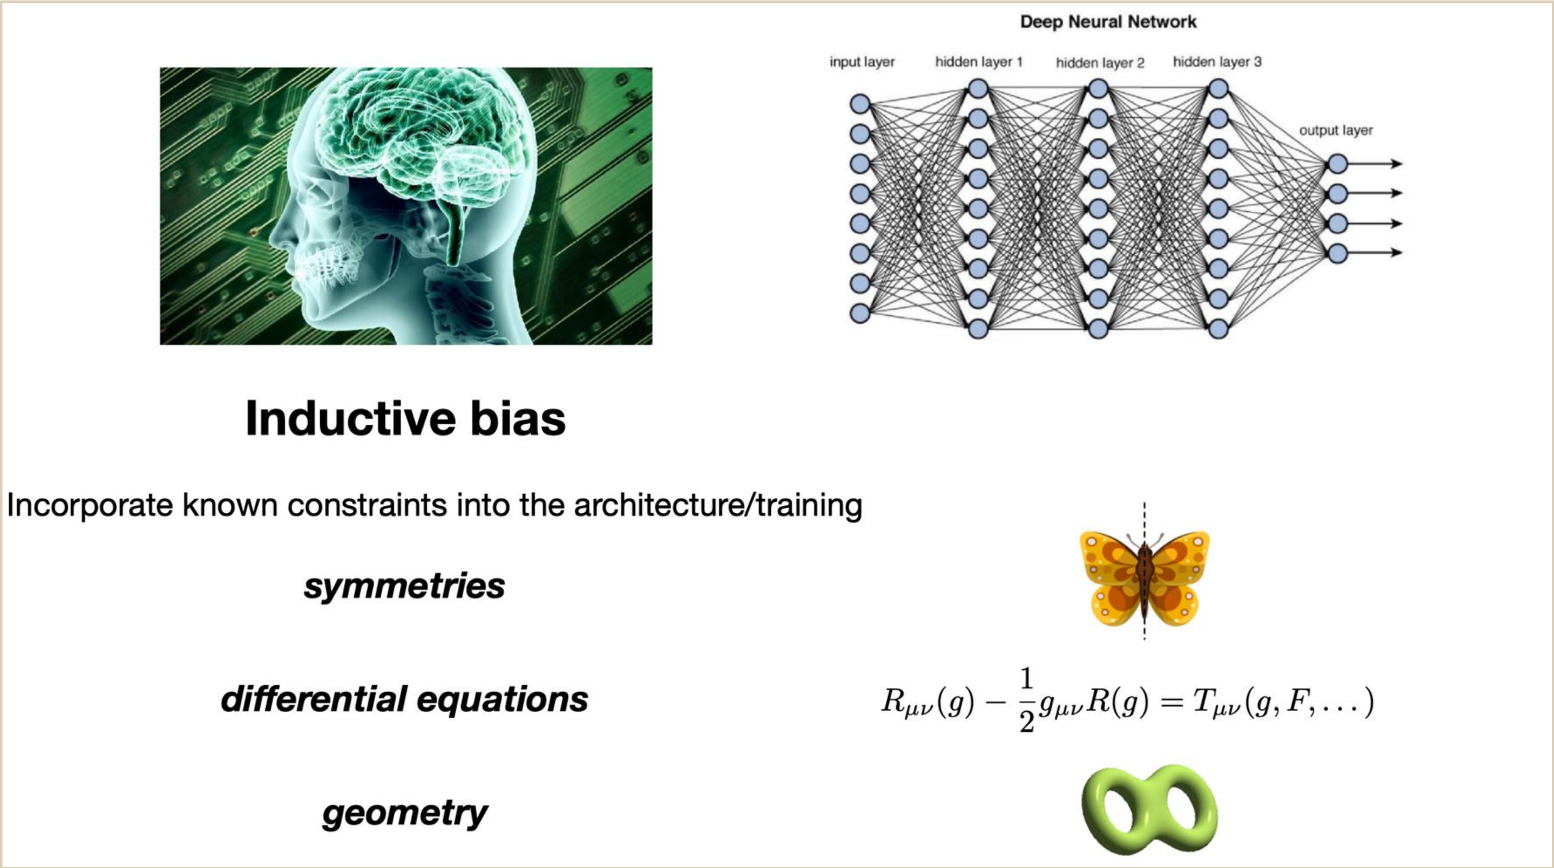

slide from a recent talk at CHAIR (Chalmers AI Research Centre)

MY RESEARCH GROUP · EXAMPLES

MATH & PHYSICS → AI

Geometric deep learning in practice

HEAL-SWIN: A Vision Transformer On The Sphere

Oscar Carlsson*[ab] Jan E. Gerken*[a] Hampus Linander[a] Heiner Spieß[c] Fredrik Ohlsson[d]

Christoffer Petersson[ea] Daniel Persson[a]

Abstract

High-resolution wide-angle fisheye images are becoming more and more important for robotics applications such as autonomous driving. However, using ordinary convolutional neural networks or vision transformers on this data is problematic due to projection and distortion losses introduced when projecting to a rectangular grid on the plane. We introduce the HEAL-SWIN transformer, which combines the highly uniform Hierarchical Equal Area iso-Latitude Pixelation (HEALPix) grid used in astrophysics and cosmology with the Hierarchical Shifted-Window (SWIN) transformer to yield an efficient and flexible model capable of training on high-resolution, distortion-free spherical data. In HEAL-SWIN, the nested structure of the HEALPix grid is used to perform the patching and windowing operations of the SWIN transformer, enabling the network to process spherical representations with minimal computational overhead. We demonstrate the superior performance of our model on both synthetic and real automotive datasets, as well as a selection of other image datasets, for semantic segmentation, depth regression and classification tasks. Our code is publicly available[1].

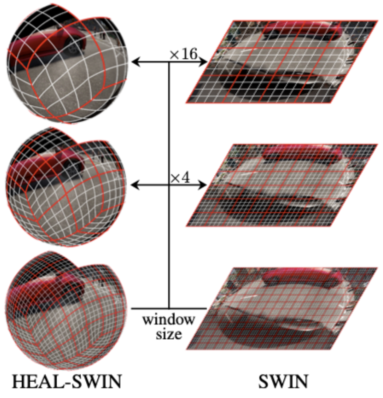

Figure 1. Our HEAL-SWIN model uses the nested structure of the HEALPix grid to lift the windowed self-attention of the SWIN model onto the sphere.

vision and autonomous driving datasets do not contain fisheye images. For these reasons, fisheye images have received

HEAL-SWIN: computer vision on the sphere

PEAR: Equal Area Weather Forecasting on the Sphere

Hampus Linander[1,2] Tage Tykesson[2,4] Pietro Rosso[3]
Christoffer Petersson[2,5] Daniel Persson[2] Jan E. Gerken[2]

[1] VERSES AI
Los Angeles, USA

[2] Department of Mathematical Sciences
Chalmers University of Technology
University of Gothenburg, Sweden

[3] Department of Computer Science and Engineering
[4] Department of Physics and Astronomy
Chalmers University of Technology, Sweden

[5] Recohere
Gothenburg, Sweden

Abstract

Artificial intelligence is rapidly reshaping the natural sciences, with weather forecasting emerging as a flagship AI4Science application where machine learning models can now rival and even surpass traditional numerical simulations. Following the success of the landmark models Pangu Weather and Graphcast, outperforming traditional numerical methods for global medium-range forecasting, many novel data-driven methods have emerged. A common limitation shared by many of these models is their reliance on an equiangular discretization of the sphere which suffers from a much finer grid at the poles than around the equator. In contrast, in the Hierarchical Equal Area iso-Latitude Pixelization (HEALPix) of the sphere, each pixel covers the same surface area, removing unphysical biases. Motivated by a growing support for this grid in meteorology and climate sciences, we propose to perform weather forecasting with deep learning models which natively operate on the HEALPix grid. To this end, we introduce Pangu Equal ARea (PEAR), a transformer-based weather forecasting model which operates directly on HEALPix-features and outperforms the corresponding model on an equiangular grid, and other baselines, without any computational overhead. Furthermore, we perform numerical experiments on the equivariance properties of our setup and verify the performance of PEAR on climate model emulation.
Code: https://github.com/hlinander/PEAR-Weather.

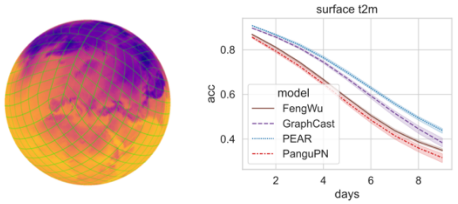

PEAR: equal-area AI weather forecasting

AI FOR MATH

AI → MATH

AI is already helping to discover mathematics

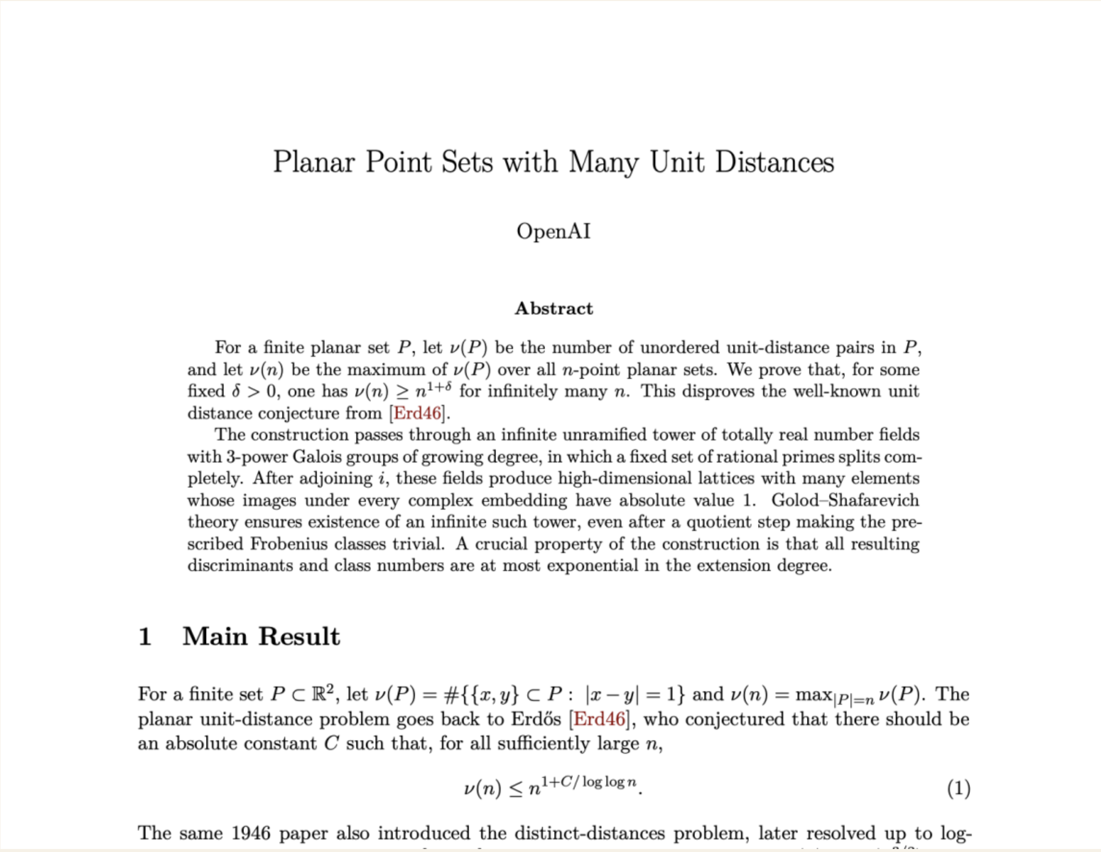

Planar Point Sets with Many Unit Distances

OpenAI

Abstract

For a finite planar set P, let $\nu(P)$ be the number of unordered unit-distance pairs in P, and let $\nu(n)$ be the maximum of $\nu(P)$ over all n-point planar sets. We prove that, for some fixed $\delta > 0$, one has $\nu(n) \geq n^{1+\delta}$ for infinitely many n. This disproves the well-known unit distance conjecture from [Erd46].

The construction passes through an infinite unramified tower of totally real number fields with 3-power Galois groups of growing degree, in which a fixed set of rational primes splits completely. After adjoining i, these fields produce high-dimensional lattices with many elements whose images under every complex embedding have absolute value 1. Golod–Shafarevich theory ensures existence of an infinite such tower, even after a quotient step making the prescribed Frobenius classes trivial. A crucial property of the construction is that all resulting discriminants and class numbers are at most exponential in the extension degree.

1 Main Result

For a finite set $P \subset \mathbb{R}^2$, let $\nu(P) = \#\{\{x, y\} \subset P : |x - y| = 1\}$ and $\nu(n) = \max_{|P|=n} \nu(P)$. The planar unit-distance problem goes back to Erdős [Erd46], who conjectured that there should be an absolute constant C such that, for all sufficiently large n,

$$\nu(n) \leq n^{1+C/\log\log n}. \tag{1}$$

The same 1946 paper also introduced the distinct-distances problem, later resolved up to log-

OpenAI · a new unit-distance result

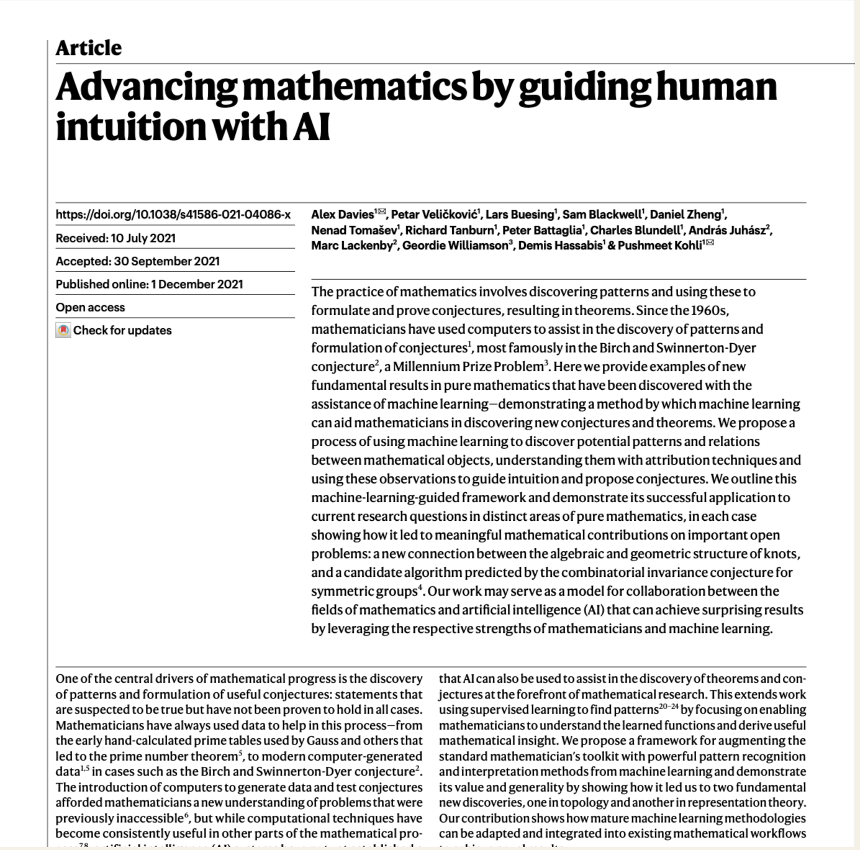

Article

Advancing mathematics by guiding human intuition with AI

https://doi.org/10.1038/s41586-021-04086-x

Received: 10 July 2021

Accepted: 30 September 2021

Published online: 1 December 2021

Open access

Check for updates

Alex Davies[1✉], Petar Veličković[1], Lars Buesing[1], Sam Blackwell[1], Daniel Zheng[1], Nenad Tomašev[1], Richard Tanburn[1], Peter Battaglia[1], Charles Blundell[1], András Juhász[2], Marc Lackenby[2], Geordie Williamson[3], Demis Hassabis[1] & Pushmeet Kohli[1✉]

The practice of mathematics involves discovering patterns and using these to formulate and prove conjectures, resulting in theorems. Since the 1960s, mathematicians have used computers to assist in the discovery of patterns and formulation of conjectures[1], most famously in the Birch and Swinnerton-Dyer conjecture[2], a Millennium Prize Problem[3]. Here we provide examples of new fundamental results in pure mathematics that have been discovered with the assistance of machine learning—demonstrating a method by which machine learning can aid mathematicians in discovering new conjectures and theorems. We propose a process of using machine learning to discover potential patterns and relations between mathematical objects, understanding them with attribution techniques and using these observations to guide intuition and propose conjectures. We outline this machine-learning-guided framework and demonstrate its successful application to current research questions in distinct areas of pure mathematics, in each case showing how it led to meaningful mathematical contributions on important open problems: a new connection between the algebraic and geometric structure of knots, and a candidate algorithm predicted by the combinatorial invariance conjecture for symmetric groups[4]. Our work may serve as a model for collaboration between the fields of mathematics and artificial intelligence (AI) that can achieve surprising results by leveraging the respective strengths of mathematicians and machine learning.

One of the central drivers of mathematical progress is the discovery of patterns and formulation of useful conjectures: statements that are suspected to be true but have not been proven to hold in all cases. Mathematicians have always used data to help in this process—from the early hand-calculated prime tables used by Gauss and others that led to the prime number theorem[5], to modern computer-generated data[1,5] in cases such as the Birch and Swinnerton-Dyer conjecture[2]. The introduction of computers to generate data and test conjectures afforded mathematicians a new understanding of problems that were previously inaccessible[6], but while computational techniques have become consistently useful in other parts of the mathematical pro-

that AI can also be used to assist in the discovery of theorems and conjectures at the forefront of mathematical research. This extends work using supervised learning to find patterns[20–24] by focusing on enabling mathematicians to understand the learned functions and derive useful mathematical insight. We propose a framework for augmenting the standard mathematician's toolkit with powerful pattern recognition and interpretation methods from machine learning and demonstrate its value and generality by showing how it led us to two fundamental new discoveries, one in topology and another in representation theory. Our contribution shows how mature machine learning methodologies can be adapted and integrated into existing mathematical workflows

DeepMind · AI-guided mathematics

The elephant in the room

Worries…

- Will we lose the craft?
- Can we trust results we did not derive ourselves?
- Who owns a machine-assisted proof?
- What happens to mathematical apprenticeship?
- Who pays?

My thoughts

- Agents can widen the search for ideas and references.
- Fluent output is never evidence.
- Sources, verification, and human responsibility decide what survives.
- We should learn as much as possible what AI can and cannot do

THE WORKFLOW

The step from chatbot to agents

CHATBOT

ask a question
↓
read a plausible answer
↓
copy the useful part
↓
context ends with the session

PROJECT AGENT

give a task
↓
it reads the repository
↓
it edits and runs tools
↓
you inspect the evidence
↓
accept—or send it back

THE TOOLS

What is a project agent?

- **Works in a project folder** — your repository: LaTeX, notes, code, bibliography
- **Reads and writes files** — drafts, edits and reorganises; every change can be inspected
- **Runs programs** — computations, checks, literature searches — and reads the output
- **Two examples** — Codex (OpenAI) and Claude Code (Anthropic)

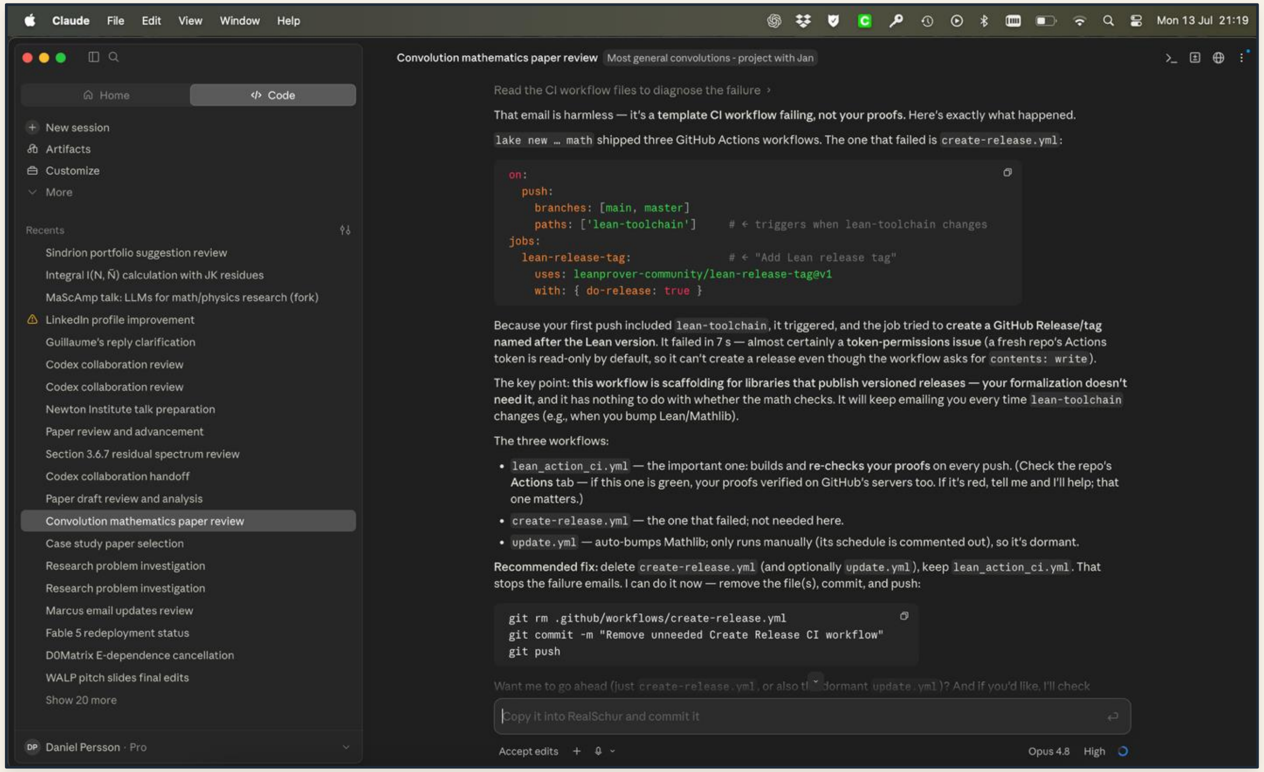

Claude Code, working in one of my paper repositories

Markdown files make the project remember

- • CLAUDE.md records the goal, notation, conventions, and boundaries.
- • A progress log keeps decisions, loose ends, and next steps.
- • New sessions start by reading this durable context.

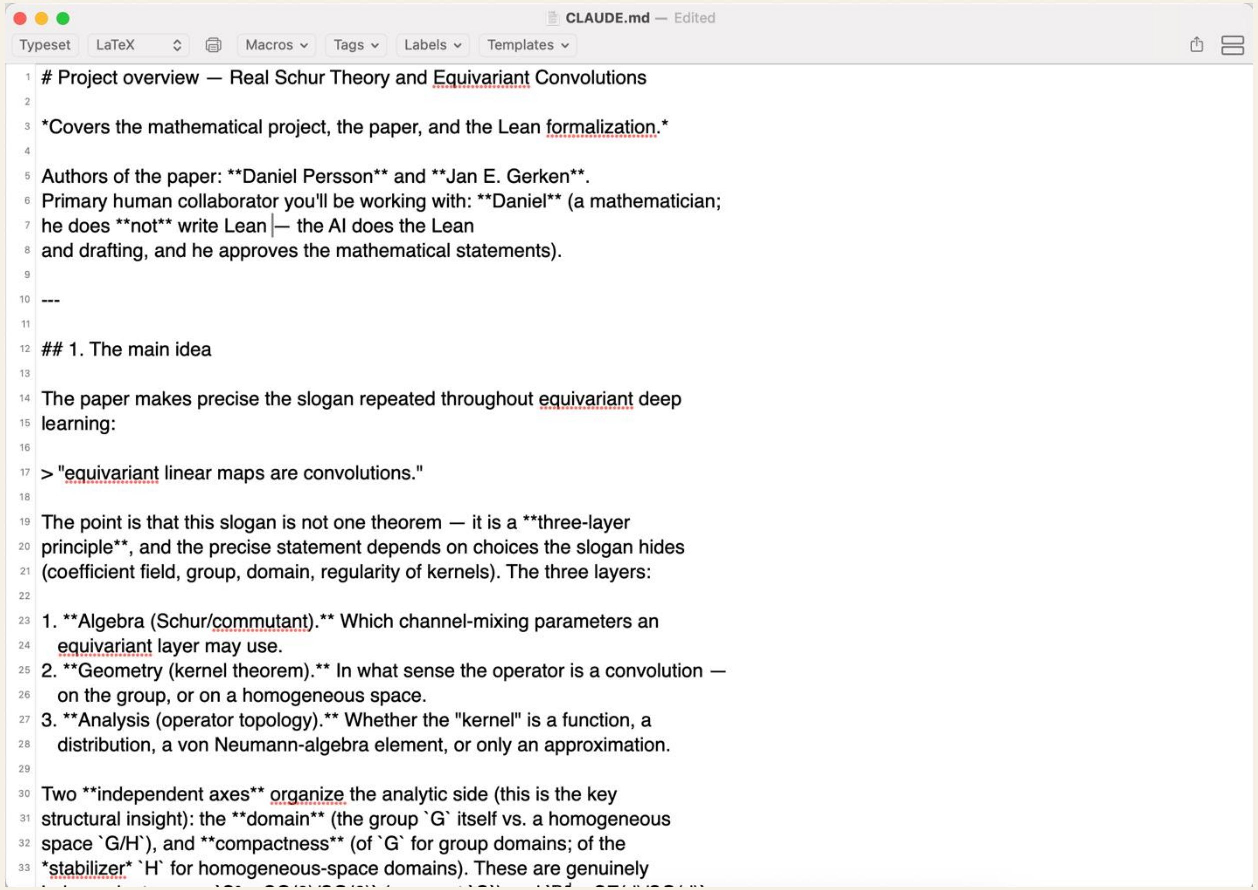


```
# Project overview — Real Schur Theory and Equivariant Convolutions

*Covers the mathematical project, the paper, and the Lean formalization.*

Authors of the paper: **Daniel Persson** and **Jan E. Gerken**.
Primary human collaborator you'll be working with: **Daniel** (a mathematician;
he does **not** write Lean — the AI does the Lean
and drafting, and he approves the mathematical statements).

---

## 1. The main idea

The paper makes precise the slogan repeated throughout equivariant deep
learning:

> "equivariant linear maps are convolutions."

The point is that this slogan is not one theorem — it is a **three-layer
principle**, and the precise statement depends on choices the slogan hides
(coefficient field, group, domain, regularity of kernels). The three layers:

1. **Algebra (Schur/commutant).** Which channel-mixing parameters an
   equivariant layer may use.
2. **Geometry (kernel theorem).** In what sense the operator is a convolution —
   on the group, or on a homogeneous space.
3. **Analysis (operator topology).** Whether the "kernel" is a function, a
   distribution, a von Neumann-algebra element, or only an approximation.

Two **independent axes** organize the analytic side (this is the key
structural insight): the **domain** (the group `G` itself vs. a homogeneous
space `G/H`), and **compactness** (of `G` for group domains; of the
*stabilizer* `H` for homogeneous-space domains). These are genuinely
```

A project-specific CLAUDE.md carries durable context

THE MINDSET

The repository is the shared brain

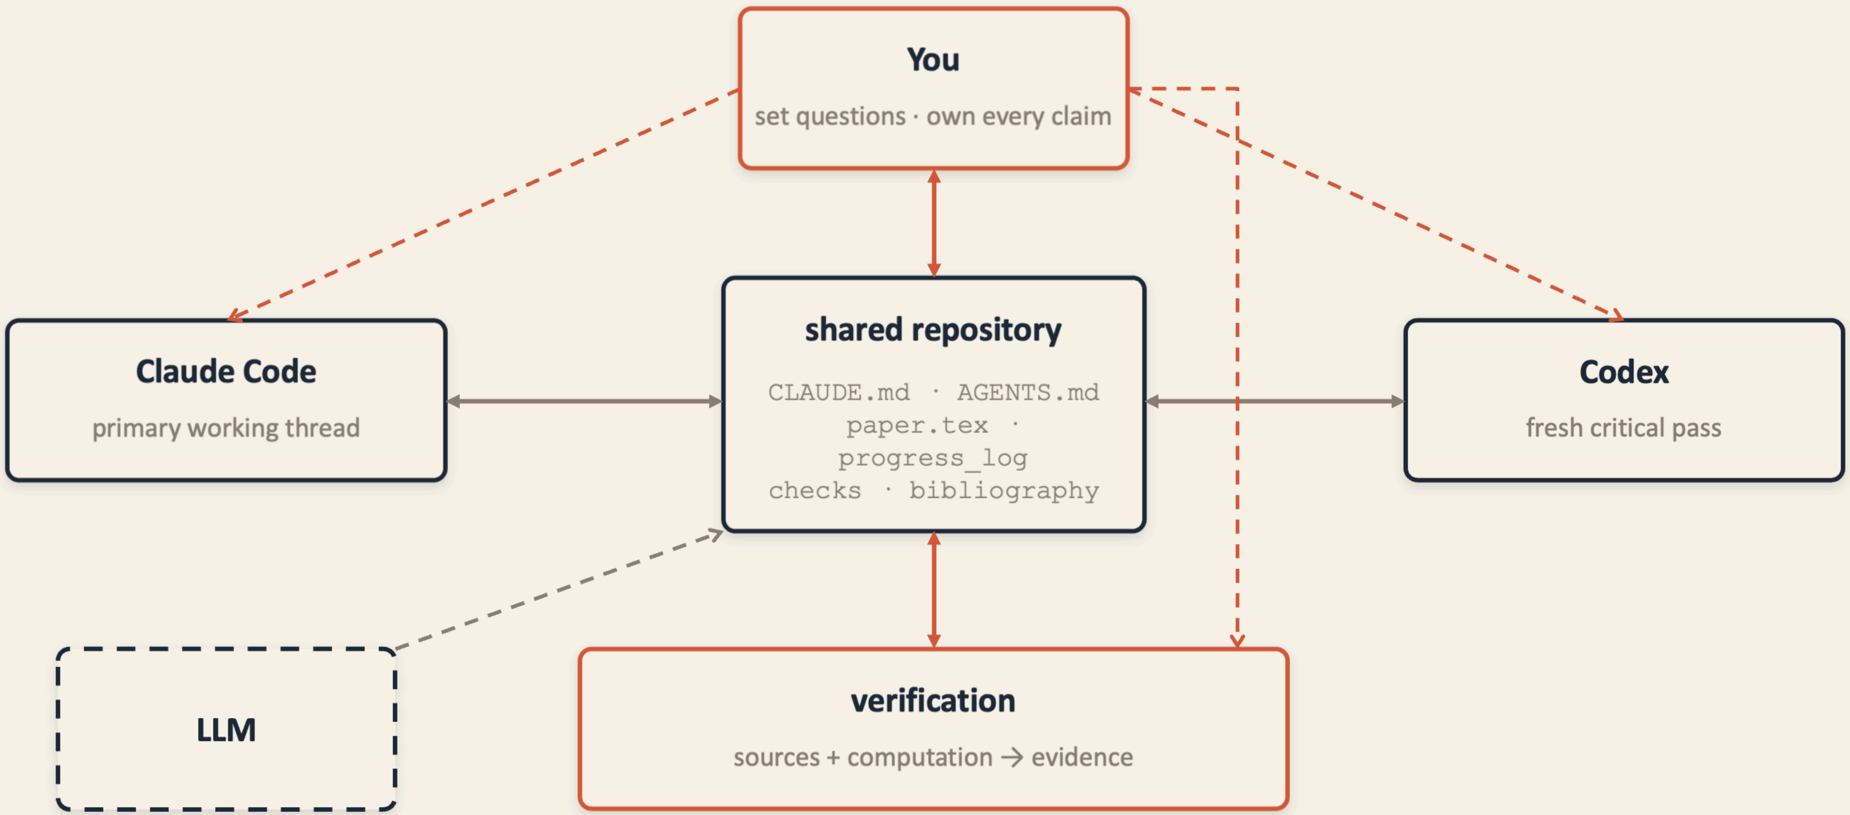

THE CASE STUDY

Z-infinite automorphic forms

- Project: automorphic functions that are not Z-finite.
- We had an annihilation theorem for SO(n,n), proved by explicit differential equations.
- The goal was a proof that generalizes to exceptional groups.

the paper draft:

$\mathcal{Z}$-INFINITE AUTOMORPHIC FORMS AND THE SQUARE OF THE MINIMAL REPRESENTATION

GUILLAUME BOSSARD, SOLOMON FRIEDBERG, DMITRY GOUREVITCH, AXEL KLEINSCHMIDT, AND DANIEL PERSSON

(426)

Using the functional relation (184) one obtains therefore both (423) at $s = \frac{n-2}{2}$ and the equation above (426) for $E_{P_2,\pi}(\frac{n-2}{2}, g)$. Substituting (426) in (423) for $s = \frac{n-2}{2}$, one obtains

$$\frac{\Delta + (n-1)(n-2)}{4}\mathcal{D}_{a\hat{c}}\mathcal{D}^{d\hat{c}}\mathcal{D}_{d\hat{b}}E^{D_n}_{P_2,\pi}(\tfrac{n-2}{2}, g) = \Big(\frac{\Delta + (n-1)(n-2)}{4}\Big)\Big(\frac{\Delta + (n-2)(n-3)}{4}\Big)\mathcal{D}_{a\hat{b}}E^{D_n}_{P_2,\pi}(\tfrac{n-2}{2}, g)\,. \qquad (427)$$

For π such that $\Delta f_\pi = -(\frac{1}{4} + r_\pi^2)f_\pi$ in $SL(2)$, one obtains that

$$\frac{\Delta + (n-1)(n-2)}{4}E^{D_n}_{P_2,\pi}(\tfrac{n-2}{2}, g) = -\frac{\frac{1}{4} + r_\pi^2}{4}E^{D_n}_{P_2,\pi}(\tfrac{n-2}{2}, g) \qquad (428)$$

so one can inverse this operator and obtain that

$$\mathcal{D}_{[a}{}^{[\hat{a}}\mathcal{D}_{b}{}^{\hat{b}}\mathcal{D}_{c]}{}^{\hat{c}]}E^{D_n}_{P_2,\pi}(\tfrac{n-2}{2}, g) = 0\,, \qquad (429)$$

$$\mathcal{D}_{a\hat{c}}\mathcal{D}^{d\hat{c}}\mathcal{D}_{d\hat{b}}E^{D_n}_{P_2,\pi}(\tfrac{n-2}{2}, g) = \frac{1}{4}\Big(\Delta + (n-2)(n-3)\Big)\mathcal{D}_{a\hat{b}}E^{D_n}_{P_2,\pi}(\tfrac{n-2}{2}, g)\,.$$

This concludes the direct proof that $E^{D_n}_{P_2,\pi}(\frac{n-2}{2}, g)$ is annihilated by $\mathcal{I}_{P_2}$.

The proof emerged from collaboration with AI agents

Not one prompt: repeated proposals, objections, source checks, and rewrites.

1. **Define the obstruction — correct theorem, look for alternative proof**
2. **ChatGPT first suggests a way forward. Claude Code then searches the literature and brings in new ideas.**
3. **We review and come with new input. Codex then checks and re-derives the argument, attacks assumptions, and proposes alternatives.**
4. **We read the new sources, check the details, and reformulate the argument. Both Claude Code and Codex verify the new proof.**

WHAT THE LOOP CHANGED

explicit D_n identities
↓
annihilator of a spherical line
↓
cyclic Harish-Chandra module
↓
irreducible spherical representation
↓
Annihilation is automatic from representation theory

The proof emerged from collaboration with AI agents

Not one prompt: repeated proposals, objections, source checks, and rewrites.

1. **Define the problem — correct theorem, look for alternative proof**
2. **ChatGPT first suggests a way forward. Claude Code then searches the literature and brings in new ideas.**
3. **We review and come with new input. Codex then checks and re-derives the argument, attacks assumptions, and proposes alternatives.**
4. **We read the new sources, check the details, and reformulate the argument. Both Claude Code and Codex verify the new proof.**

Theorem 1. *For every spherical cuspidal π occurring in Lemma 6, we have*

$$\mathcal{A}\, X_\pi = \mathcal{B}\, X_\pi = 0\,. \tag{220}$$

Equivalently, $\mathcal{I}_{P_2} X_\pi = 0$; together with $\Delta X_\pi = 2s_\pi(2s_\pi + 3 - 2n)X_\pi$ this gives $\mathcal{I}_{P_2,s_\pi} X_\pi = 0$.

Proof. Let $M_{s_\pi} = \mathcal{U}(\mathfrak{g}_\mathbb{C})\chi_{P_2}^{s_\pi}$ be the cyclic Harish-Chandra module generated by the scalar spherical section. It has finite length, so it has a nonzero irreducible quotient J_π. The image of its cyclic vector is necessarily nonzero and K-fixed. Hence J_π is spherical and has infinitesimal character $2s_\pi\Lambda_2 - \rho = \lambda_\pi$. Equation (212) descends to its spherical vector.

By Proposition 9, $J_\pi \simeq \sigma_\pi$, where σ_π is the representation of Lemma 7. Thus the K-fixed line in every copy of σ_π is annihilated by $\mathcal{A}$ and $\mathcal{B}$. Lemma 7 writes $X_\pi = v_K \otimes m_\pi$ in the corresponding isotypic component, with G° acting only on the first factor. Therefore

$$\mathcal{A}X_\pi = (\mathcal{A}v_K) \otimes m_\pi = 0\,, \qquad \mathcal{B}X_\pi = (\mathcal{B}v_K) \otimes m_\pi = 0\,. \tag{221}$$

□

Agents can work with proof and computation tools

1. **Mathematica / Sage** — examples, symbolic checks, and counterexamples
2. **Lean** — once formalized, anyone can replay and check the proof
3. **AI helps with translation** — informal proof → type-theoretic statement
4. **The tool gives hard feedback** — failed checks become concrete errors
5. **Human judgment remains**

PRACTICAL

How to start on Monday

1. **Put the project in plain text** — LaTeX, Markdown, code, notes, bibliography
2. **Give an agent the project folder** — begin with a bounded task you understand well
3. **Write a durable project brief** — goals, conventions, boundaries, known pitfalls
4. **Demand independent evidence** — calculation, proof artifact, source, example, counterexample

ONE LAST THOUGHT

“Can AI do mathematics?” is not my question.

How does research change when a fluent, tireless, fallible collaborator joins the loop?

DISCUSSION

AI for mathematics: questions

1. Are you using AI tools in your research?
2. If so, how are you using them?
3. Where do they genuinely help—and where do they not?
4. Are you worried about the future of mathematics because of AI?
5. What should remain non-negotiably human?