

The Agentic Researcher: Accelerating Optimization and Learning Research with AI

Sebastian Pokutta

Technische Universität Berlin
and
Zuse Institute Berlin

pokutta@math.tu-berlin.de
@spokutta

HORIZONS 2026

Learning and Optimization for a Sustainable Future

July 2026 · Hanoi, Vietnam

VinUniversity



Berlin Mathematics Research Center



What is this talk about?

Introduction

A personal, highly-biased, and incomplete take of what AI can do in Mathematics et al.

Why? AI methods increasingly shape how scientific discovery is performed, configured, and accelerated.

Today: Two (three-ish) recent use-cases / perspectives from my own work.

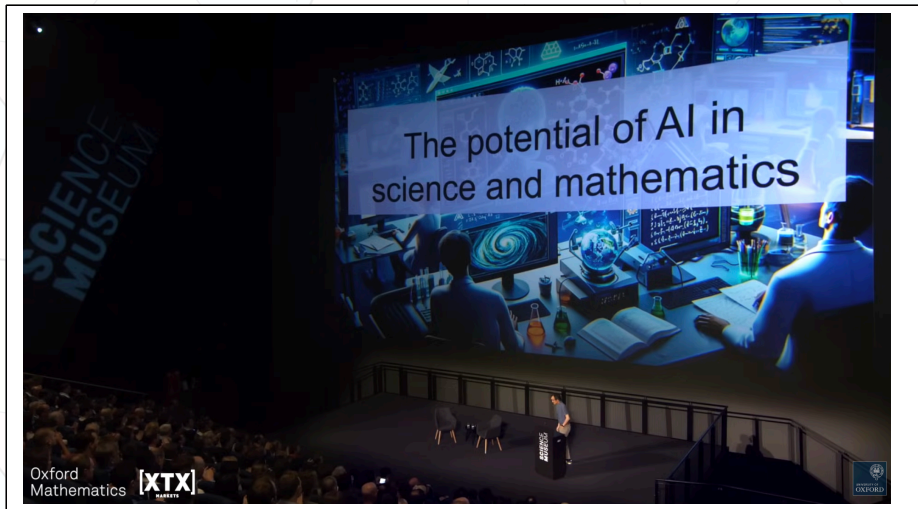
Outline

- A bit of high-level perspective
- Agentic AI Researcher framework
- Neural Colorings of the plane

(Hyperlinked) References are not exhaustive; check references contained therein.

What is this talk about?

AI and Mathematics et al



[The Potential for AI in Science and Mathematics - Terence Tao]

What is this talk about?

AI and Mathematics et al

Various levels of co-creation.

[Haase and Pokutta, 2026]

- **Digital Pen:** basically like autocorrect, bibtex lookup, etc. "2000s"
- **AI Task Specialist:** ChatGPT, Claude, Gemini, etc. 2022 - 2025
- **AI Assistant:** Agents with integrated tools, verification, etc. 2025 -
- **AI Co-creator:** Fully integrated, autonomous, co-creator 2027(??) -

What is this talk about?

AI and Mathematics et al

Various levels of co-creation.

[Haase and Pokutta, 2026]

- **Digital Pen**: basically like autocorrect, bibtex lookup, etc. "2000s"
- **AI Task Specialist**: ChatGPT, Claude, Gemini, etc. 2022 - 2025
- **AI Assistant**: Agents with integrated tools, verification, etc. 2025 -
- **AI Co-creator**: Fully integrated, autonomous, co-creator 2027(??) -

Broadly. Two categories: (a) co-creation **agent** and (b) **tool** in larger system.

What is this talk about?

AI and Mathematics et al

Various levels of co-creation.

[Haase and Pokutta, 2026]

- **Digital Pen**: basically like autocorrect, bibtex lookup, etc. "2000s"
- **AI Task Specialist**: ChatGPT, Claude, Gemini, etc. 2022 - 2025
- **AI Assistant**: Agents with integrated tools, verification, etc. 2025 -
- **AI Co-creator**: Fully integrated, autonomous, co-creator 2027(??) -

Broadly. Two categories: (a) co-creation **agent** and (b) **tool** in larger system.

Capabilities are **impressive** but **unstable**:

- SOTA models achieve post-PhD level scores on benchmarks, yet in day-to-day use make trivial (logical) errors.
- No hard verification of results and randomness across runs.
- Prompting and scaffolding are still a challenge.
- Availability of tools for verification etc crucial.

What is this talk about?

Mathematics with computers is not new

Various high-profile examples from the past.

- Four Color Theorem: massive computer-based case checking
[Appel and Haken, 1977, Robertson et al., 1997]
- Kepler Conjecture / Hales' Theorem: extensive computer verification
[Hales et al., 2017]
- Classification of Finite Simple Groups: Formal verification with Lean/Coq
- Boolean Pythagorean Triples Problem: A spectacular 200TB SAT-solver proof.
[Heule et al., 2016]

What is this talk about?

Mathematics with computers is not new

Various high-profile examples from the past.

- Four Color Theorem: massive computer-based case checking
[Appel and Haken, 1977, Robertson et al., 1997]
- Kepler Conjecture / Hales' Theorem: extensive computer verification
[Hales et al., 2017]
- Classification of Finite Simple Groups: Formal verification with Lean/Coq
- Boolean Pythagorean Triples Problem: A spectacular 200TB SAT-solver proof.
[Heule et al., 2016]

Crucial role in computational mathematics / scientific computing

- Finite Elements
- Numerical Simulations
- Optimization
- Engineering
- ...

The Agentic Researcher

A Practical Guide to AI-Assisted Research
in Mathematics and Machine Learning

joint work with: Max Zimmer, Nico Pelleriti,
Christophe Roux

ICML AI4RESEARCH Workshop 2026 (oral)
<https://arxiv.org/abs/2603.15914>

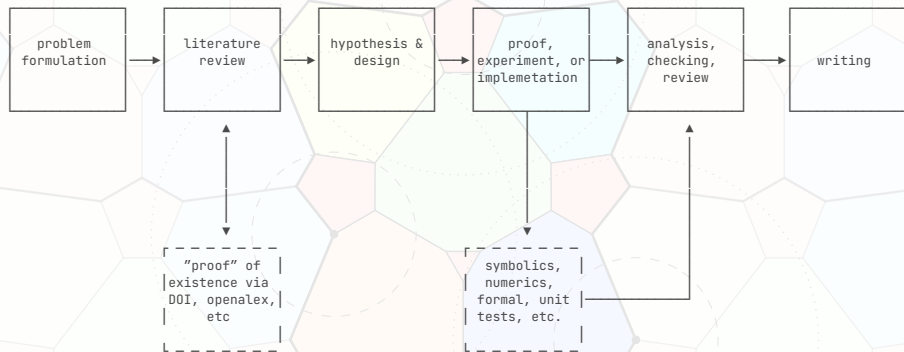
Partially supported by ExC MATH+ Project EF-LiOpt-3
Agent AI in Mathematics

[Zimmer et al., 2026, Pokutta, 2026]

AI in the Scientific Discovery Workflow

The agentic researcher framework

AI systems can be used at **almost any stage** of the scientific workflow.



(workflow simplified. for illustration only.)

Note. The researcher needs to be in charge and responsible(!) throughout.

Building a practical tool

The agentic researcher framework

Let me now wait for the final check:

• Task Output bwzs8ty88

1-hour mid-training status check

Waiting for task (esc to give additional instructions)

+ Bitifying... (8h 39m 19s · ↓ 39.3k tokens · thought for 28s)

↳ Tip: Run /install-github-app to tag @claude right from your Github issues and PRs

Background tasks

10 active shells

```
) sleep 2700 &&echo "Final completion check =" &&python3 << 'PYEOF'.. (running)
sleep 3600 &&echo "1h mid-training =" &&python3 -c "... (running)
sleep 5400 &&echo "Final completion check =" &&for exp in E023_no_nor.. (running)
source ~/.bashrc &&CUDA_VISIBLE_DEVICES=6 uv run torchrun --standalone --npr.. (running)
source ~/.bashrc &&CUDA_VISIBLE_DEVICES=5 uv run torchrun --standalone --npr.. (running)
source ~/.bashrc &&CUDA_VISIBLE_DEVICES=4 uv run torchrun --standalone --npr.. (running)
source ~/.bashrc &&CUDA_VISIBLE_DEVICES=3 uv run torchrun --standalone --npr.. (running)
source ~/.bashrc &&CUDA_VISIBLE_DEVICES=2 uv run torchrun --standalone --npr.. (running)
source ~/.bashrc &&CUDA_VISIBLE_DEVICES=1 uv run torchrun --standalone --npr.. (running)
```

Screenshot from a run of the framework.

- Born out of the **MATH+** project *Agentic AI in Mathematics*.
- Refined over roughly **1.5 years** of day-to-day research use.
- Sandboxed CLI agents with project instructions, tools, Git, \LaTeX , and GPU-backed experiments.
- Puts a premium on **verification**: record everything, verify citations, and verify before claiming.

[blog post]

[arXiv]

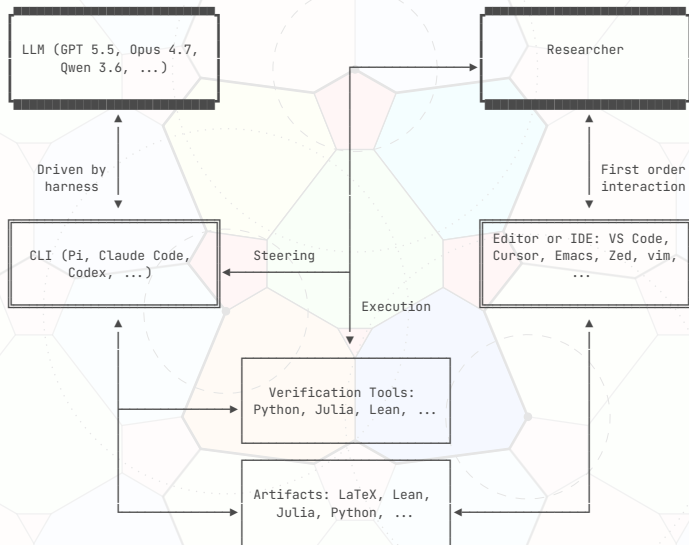
[GitHub]

Does it work? Several real-world use-cases in paper

1. Convergence lower bounds for Frank-Wolfe on uniformly convex sets
2. Multi-Variable Dual Tightening for Boscia.jl for MINLPs
3. Weight Reconstruction in LLM pruning
4. ...

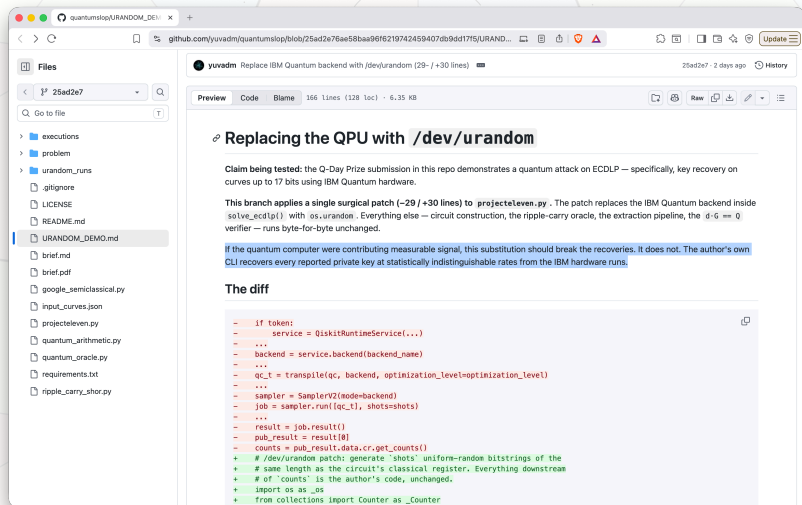
Interacting with the System

The agentic researcher framework



Why Verification is Crucial

The agentic researcher framework



The screenshot shows a GitHub repository page for a pull request titled "Replacing the QPU with /dev/urandom". The repository is owned by "yuvadm" and the pull request is for the file "URANDOM_DEMO.md". The pull request description states that the branch applies a single surgical patch to "projecteven.py" to replace the IBM Quantum backend with "/dev/urandom". It also includes a diff showing the changes to the code.

Replacing the QPU with /dev/urandom

Claim being tested: the Q-Day Prize submission in this repo demonstrates a quantum attack on ECCDLP – specifically, key recovery on curves up to 17 bits using IBM Quantum hardware.

This branch applies a single surgical patch (–29 / +30 lines) to `projecteven.py`. The patch replaces the IBM Quantum backend inside `solve_eccdlp()` with `os.urandom`. Everything else – circuit construction, the ripple-carry oracle, the extraction pipeline, the `d-G == 0` verifier – runs byte-for-byte unchanged.

If the quantum computer were contributing measurable signal, this substitution should break the recoveries. It does not. The author's own CLI recovers every reported private key at statistically indistinguishable rates from the IBM hardware runs.

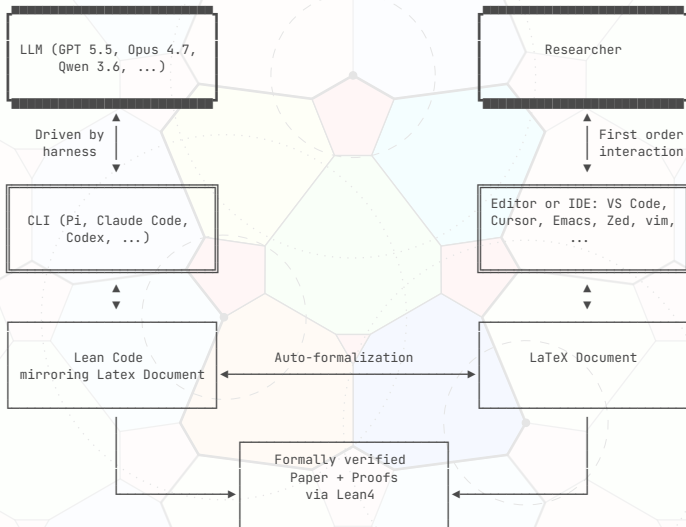
The diff

```
- if token:
-     service = QiskitRuntimeService(...)
-     ...
-     backend = service.backend(backend_name)
-     ...
-     qc_t = transpile(qc, backend, optimization_level=optimization_level)
-     ...
-     sampler = SamplerV2(mode=backend)
-     job = sampler.run([qc_t], shots=shots)
-     ...
-     result = job.result()
-     pub_result = result[0]
-     counts = pub_result.data.cr.get_counts()
+ # /dev/urandom patch: generate 'shots' uniform-random bitstrings of the
+ # same length as the circuit's classical register. Everything downstream
+ # of 'counts' is the author's code, unchanged.
+ import os as _os
+ from collections import Counter as _Counter
```

[GitHub Repository]

End-2-End Formal Verification with Lean 4

The agentic researcher framework



When generation is cheap, verification is everything.

[Pokutta, 2026]

The Hadwiger-Nelson Problem

joint work with: Aldo Kiem, Konrad Mundinger,
Christoph Spiegel, Max Zimmer

ICML 2025 (oral)

<https://arxiv.org/abs/2404.05509>

Partially supported by ExC MATH+ Project EF-LiOpt-3

Agent AI in Mathematics

[Mundinger et al., 2025]

The Hadwiger-Nelson Problem

Problem (Nelson 1950, also: Gardner, Moser, Erdős, Harary, Tutte, ...)

What is the smallest number of colors sufficient for coloring the plane in such a way that no two points of the same color are at a unit distance apart?

Infinite graph with vertex set \mathbb{E}^2 and edges $\{x, y\}$ for any $x, y \in \mathbb{E}^2$ with $\|x - y\| = 1$

\Rightarrow **chromatic number of the plane** $\chi(\mathbb{E}^2)$

Theorem

Assuming Axiom of Choice (AoC):

[Bruijn and Erdos, 1951]

Any graph is k -colorable iff every finite subgraph of it is k -colorable.

This problem has a long and complicated history...


over 14 pages in [Soifer, 2024]

Upper bounds through colorings

The Hadwiger-Nelson Problem

Explicit colorings $g : \mathbb{E}^2 \rightarrow [c] := \{1, \dots, c\}$, usually derived through tessellations using simple polytopal shapes, give

$$5 \leq \chi(\mathbb{E}^2) \leq \dots$$




Upper bounds through colorings

The Hadwiger-Nelson Problem

Explicit colorings $g : \mathbb{E}^2 \rightarrow [c] := \{1, \dots, c\}$, usually derived through tessellations using simple polytopal shapes, give

$$5 \leq \chi(\mathbb{E}^2) \leq 9.$$

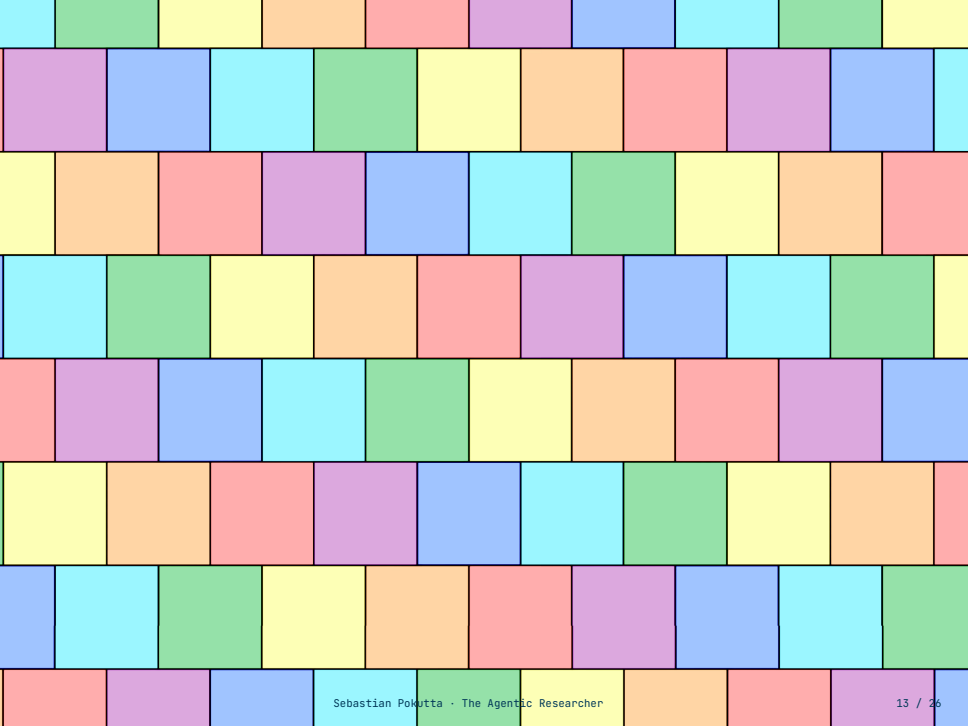



Upper bounds through colorings

The Hadwiger-Nelson Problem

Explicit colorings $g : \mathbb{E}^2 \rightarrow [c] := \{1, \dots, c\}$, usually derived through tessellations using simple polytopal shapes, give

$$5 \leq \chi(\mathbb{E}^2) \leq 8.$$





Upper bounds through colorings

The Hadwiger-Nelson Problem

Explicit colorings $g : \mathbb{E}^2 \rightarrow [c] := \{1, \dots, c\}$, usually derived through tessellations using simple polytopal shapes, give

$$5 \leq \chi(\mathbb{E}^2) \leq 7.$$

Upper bounds through colorings

The Hadwiger-Nelson Problem


Explicit colorings $g : \mathbb{E}^2 \rightarrow [c] := \{1, \dots, c\}$, usually derived through tessellations using simple polytopal shapes, give

$$5 \leq \chi(\mathbb{E}^2) \leq 7.$$

Question. Can we use computers to find admissible colorings $g : \mathbb{E}^2 \rightarrow [c]$, i.e.,

$$\{x \in \mathbb{E}^2 \mid \exists y \in B_1(x) : g(x) = g(y)\} = \emptyset?$$

... attempts, e.g., via discretization and SAT solvers...



Upper bounds through colorings

The Hadwiger-Nelson Problem

Explicit colorings $g : \mathbb{E}^2 \rightarrow [c] := \{1, \dots, c\}$, usually derived through tessellations using simple polytopal shapes, give

$$5 \leq \chi(\mathbb{E}^2) \leq 7.$$

Question. Can we use computers to find admissible colorings $g : \mathbb{E}^2 \rightarrow [c]$, i.e.,

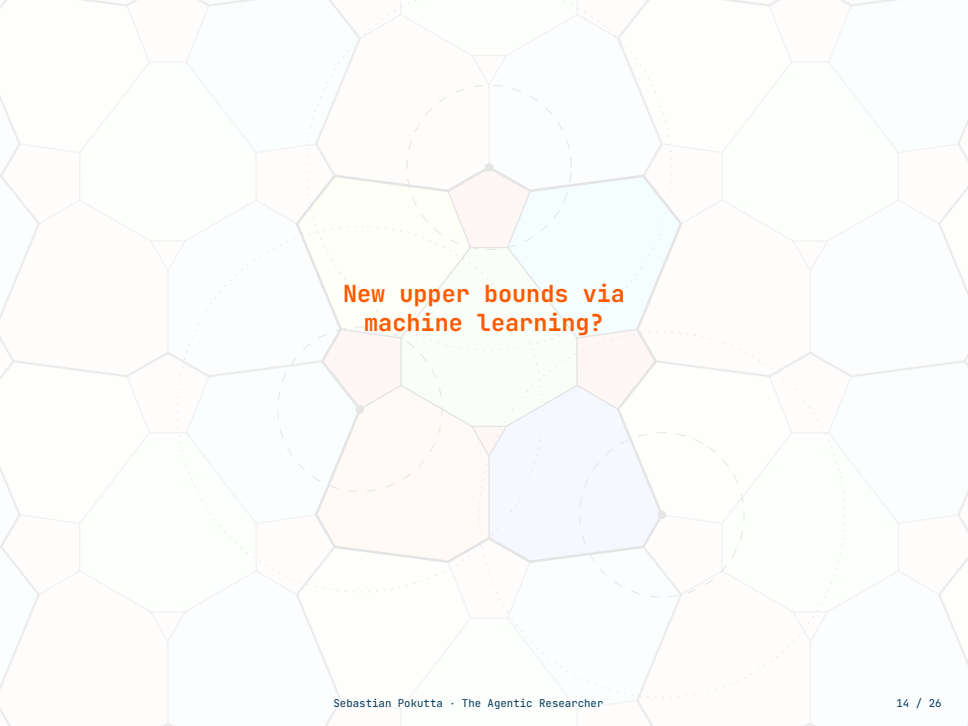
$$\{x \in \mathbb{E}^2 \mid \exists y \in B_1(x) : g(x) = g(y)\} = \emptyset?$$

... attempts, e.g., via discretization and SAT solvers...

Idea. Use a parameterized and easily differentiable family $g_\theta : \mathbb{E}^2 \rightarrow \Delta_c$ and find

$$\operatorname{argmin}_{\theta} \mathbb{E} \left[\int_{B_1(x)} g_\theta(x) \cdot g_\theta(y) dy \mid x \in \mathbb{E}^2 \right].$$

Key Point. Approach is continuous in nature.



**New upper bounds via
machine learning?**

Can we improve the upper bound?

Neural Networks as Colorings

Idea. Use gradient descent to train a feedforward network g_θ to minimize

$$L(\theta) = \int_{[-b,b] \times [-b,b]} \int_{B_1(x)} g_\theta(x) \cdot g_\theta(y) dy dx$$

for some reasonable $b \in \mathbb{R}$?

Can we improve the upper bound?

Neural Networks as Colorings

Idea. Use gradient descent to train a feedforward network g_θ to minimize

$$L(\theta) = \int_{[-b,b] \times [-b,b]} \int_{B_1(x)} g_\theta(x) \cdot g_\theta(y) dy dx$$

for some reasonable $b \in \mathbb{R}$?

Problem. Still a continuous problem. How to compute?

Can we improve the upper bound?

Neural Networks as Colorings

Idea. Use gradient descent to train a feedforward network g_θ to minimize

$$L(\theta) = \int_{[-b,b] \times [-b,b]} \int_{B_1(x)} g_\theta(x) \cdot g_\theta(y) dy dx$$

for some reasonable $b \in \mathbb{R}$?

Problem. Still a continuous problem. How to compute?

Stochastic (Batch) Gradient Descent. Sample point $x^{(i)} \in [-b,b] \times [-b,b]$ and $y^{(i)} \in B_1(x)$ for $i = 1, \dots, m$

Can we improve the upper bound?

Neural Networks as Colorings

Idea. Use gradient descent to train a feedforward network g_θ to minimize

$$L(\theta) = \int_{[-b,b] \times [-b,b]} \int_{B_1(x)} g_\theta(x) \cdot g_\theta(y) dy dx$$

for some reasonable $b \in \mathbb{R}$?

Problem. Still a continuous problem. How to compute?

Stochastic (Batch) Gradient Descent. Sample point $x^{(i)} \in [-b,b] \times [-b,b]$ and $y^{(i)} \in B_1(x)$ for $i = 1, \dots, m$ and use that

$$\nabla_\theta L(\theta) \approx \hat{\nabla}_\theta L(\theta) \doteq \frac{1}{m} \sum_{i=1}^m \nabla_\theta g_\theta(x^{(i)}) \cdot g_\theta(y^{(i)}),$$

Can we improve the upper bound?

Neural Networks as Colorings

Idea. Use gradient descent to train a feedforward network g_θ to minimize

$$L(\theta) = \int_{[-b,b] \times [-b,b]} \int_{B_1(x)} g_\theta(x) \cdot g_\theta(y) dy dx$$

for some reasonable $b \in \mathbb{R}$?

Problem. Still a continuous problem. How to compute?

Stochastic (Batch) Gradient Descent. Sample point $x^{(i)} \in [-b,b] \times [-b,b]$ and $y^{(i)} \in B_1(x)$ for $i = 1, \dots, m$ and use that

$$\nabla_\theta L(\theta) \approx \hat{\nabla}_\theta L(\theta) \doteq \frac{1}{m} \sum_{i=1}^m \nabla_\theta g_\theta(x^{(i)}) \cdot g_\theta(y^{(i)}),$$

where $\nabla_\theta g_\theta(x^{(i)}) \cdot g_\theta(y^{(i)})$ is easily computed through backpropagation,

Can we improve the upper bound?

Neural Networks as Colorings

Idea. Use gradient descent to train a feedforward network g_θ to minimize

$$L(\theta) = \int_{[-b,b] \times [-b,b]} \int_{B_1(x)} g_\theta(x) \cdot g_\theta(y) dy dx$$

for some reasonable $b \in \mathbb{R}$?

Problem. Still a continuous problem. How to compute?

Stochastic (Batch) Gradient Descent. Sample point $x^{(i)} \in [-b,b] \times [-b,b]$ and $y^{(i)} \in B_1(x)$ for $i = 1, \dots, m$ and use that

$$\nabla_\theta L(\theta) \approx \hat{\nabla}_\theta L(\theta) \doteq \frac{1}{m} \sum_{i=1}^m \nabla_\theta g_\theta(x^{(i)}) \cdot g_\theta(y^{(i)}),$$

where $\nabla_\theta g_\theta(x^{(i)}) \cdot g_\theta(y^{(i)})$ is easily computed through backpropagation, to adjust the parameters θ with an appropriate step size α_k through

$$\theta_{k+1} = \theta_k - \alpha_k \hat{\nabla}_\theta L(\theta).$$


⇒ **Very flexible approach “Deep Annealing”**

(also: tropicalization of loss function aka softmax... “minimize the max”)



Unfortunately this coloring was already known...

Neural Networks as Colorings



Theorem

99.985% of the plane can be colored with 6 colors such that no two points of the same color are a unit distance apart.

[Pritikin, 1998, Parts, 2020]

Corollary

Any unit distance graph with chromatic number 7 must have at least 6992 vertices.

⇒ While coloring was known already maybe on the right track?

A Voronoi diagram with a central cell highlighted in green. The diagram consists of several colored cells: a central green cell, a light blue cell to its right, a light orange cell below it, and a light blue cell to its left. The background is filled with other cells in various colors (light green, light blue, light orange, light pink). Two dashed circles are centered on the vertices of the central green cell, and two dotted circles are centered on the vertices of the light blue cell to its right. The text "Off-diagonal variant" is written in orange in the center of the diagram.

Off-diagonal variant

Going off-diagonal

Neural Networks as Colorings

If we cannot solve the original problem, we study variants of it:

Going off-diagonal

Neural Networks as Colorings

If we cannot solve the original problem, we study variants of it:

A c -coloring of the plane has **type** (d_1, \dots, d_c)
if color i does not contain any points at distance d_i .

Going off-diagonal

Neural Networks as Colorings

If we cannot solve the original problem, we study variants of it:

A c -coloring of the plane has **type** (d_1, \dots, d_c)
if color i does not contain any points at distance d_i .

Problem (Soifer in Nash and Rassias' *Open Problems in Mathematics*)

Determine the continuum of six-colorings

$$X_6 = \{d \mid (1, 1, 1, 1, 1, d) \text{ can be realized}\}.$$

[Soifer, 1994a, Nash and Rassias, 2016]

Going off-diagonal

Neural Networks as Colorings

If we cannot solve the original problem, we study variants of it:

A c -coloring of the plane has **type** (d_1, \dots, d_c)
if color i does not contain any points at distance d_i .

Problem (Soifer in Nash and Rassias' *Open Problems in Mathematics*)

Determine the continuum of six-colorings

$$X_6 = \{d \mid (1, 1, 1, 1, 1, d) \text{ can be realized}\}.$$

[Soifer, 1994a, Nash and Rassias, 2016]

Status. Six-colorings exist for:

1. $d = 1/\sqrt{5}$

[Soifer, 1992]

2. $d = \sqrt{2} - 1$

[Hoffman and Soifer, 1993, 1996]

3. Family with $0.414 \approx \sqrt{2} - 1 \leq d \leq 1/\sqrt{5} \approx 0.447$

[Hoffman and Soifer, 1996, Soifer, 1994b, 2009]

Going off-diagonal

Neural Networks as Colorings

If we cannot solve the original problem, we study variants of it:

A c -coloring of the plane has **type** (d_1, \dots, d_c)
if color i does not contain any points at distance d_i .

Problem (Soifer in Nash and Rassias' *Open Problems in Mathematics*)

Determine the continuum of six-colorings

$$X_6 = \{d \mid (1, 1, 1, 1, 1, d) \text{ can be realized}\}.$$

[Soifer, 1994a, Nash and Rassias, 2016]

Status. Six-colorings exist for:

1. $d = 1/\sqrt{5}$

[Soifer, 1992]

2. $d = \sqrt{2} - 1$

[Hoffman and Soifer, 1993, 1996]

3. Family with $0.414 \approx \sqrt{2} - 1 \leq d \leq 1/\sqrt{5} \approx 0.447$

[Hoffman and Soifer, 1996, Soifer, 1994b, 2009]

Deep Annealing approach provides two new colorings leading to...


Theorem

X_6 contains the closed interval $[0.354, 0.657]$.

[Mundinger et al., 2024, 2025]


First coloring: $0.354 \leq d \leq 0.553$

Neural Networks as Colorings



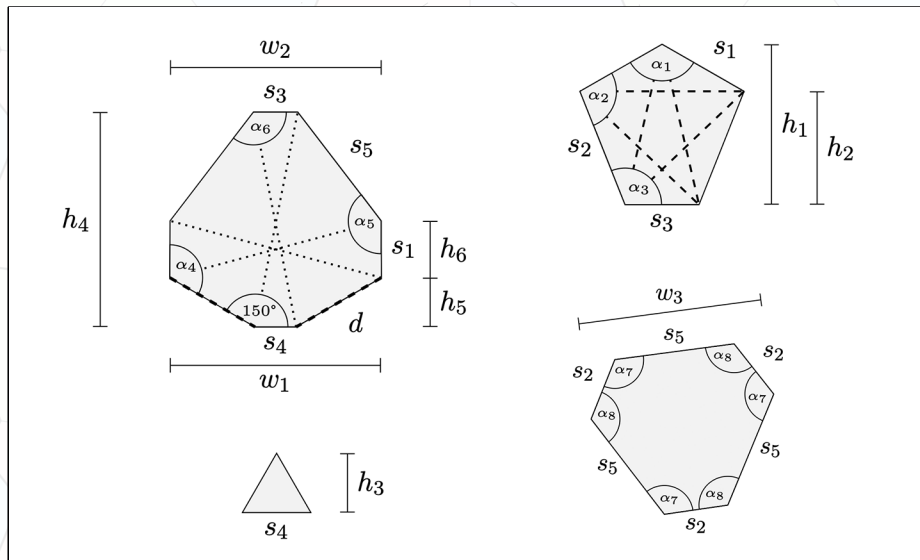
Second coloring: $0.418 \leq d \leq 0.657$

Neural Networks as Colorings







First coloring: exact components

Neural Networks as Colorings



Second coloring: exact components

Neural Networks as Colorings



Final Remarks

1. AI can be used in **various ways** in modern mathematical research workflows (actual discovery, verification, etc); beyond simple black-box prompting
2. **Fully-automatic discovery** of new mathematics might be possible in the future but relies on *strong* verification approaches
3. The agentic harness seems to be key: how does the agent receive feedback on its work, how is it guided, and which tools are available?
4. Empirically: the human-in-the-loop is crucial to guide the search

The promises of AI4MATH are great but need to go beyond simple black-box prompting “your favorite Erdős problem” (which then turns out having a solution that is already known)...



Thank you!

Mundinger, K., Pokutta, S., Spiegel, C., and Zimmer, M. (2024). Extending the Continuum of Six-Colorings. *Geombinatorics Quarterly*. Available at <https://arxiv.org/abs/2404.05509>.

Mundinger, K., Zimmer, M., Kiem, A., Spiegel, C., and Pokutta, S. (2025). Neural Discovery in Mathematics: Do Machines Dream of Colored Planes? *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 267, 45236–45255. Available at <https://arxiv.org/abs/2501.18527>.

Zimmer, M., Pelleriti, N., Roux, C., and Pokutta, S. (2026). The Agentic Researcher: A Practical Guide to AI-Assisted Research in Mathematics and Machine Learning. *Proceedings of the ICML 2026 Workshop on AI as a Tool for Mathematics, Computer Science, and Machine Learning (AI4Research) (Oral)*. Available at <https://arxiv.org/abs/2603.15914>.

Pokutta, S. (2026). Frank-Wolfe Beyond $1/t$ Convergence. Preprint. Available at <https://arxiv.org/abs/2604.28006>. Sebastian Pokutta · The Agentic Researcher

References I

- K. I. Appel and W. Haken. Every planar map is four-colorable. *Illinois Journal of Mathematics*, 21(3):429-490, 1977. Computer-assisted proof of the four color theorem (first major such proof).
- N. d. Bruijn and P. Erdos. A colour problem for infinite graphs and a problem in the theory of relations. *Indagationes Mathematicae*, 13: 371-373, 1951.
- J. Haase and S. Pokutta. Human-AI Cocreativity: Exploring synergies across levels of creative collaboration. In J. C. Kaufman and M. Worwood, editors, *Generative Artificial Intelligence and Creativity*, chapter 16, pages 205-221. 1 2026. doi: 10.1016/B978-0-443-34073-4.00009-5.
- T. C. Hales, M. Adams, G. Bauer, D. T. Dang, J. Harrison, H. L. Truong, C. Kaliszzyk, V. Magron, S. McLaughlin, N. T. Thang, N. Q. Truong, T. Nipkow, S. Obua, J. Pleso, J. Rute, A. Solovyeu, et al. A formal proof of the kepler conjecture. *Forum of Mathematics, Pi*, 5:e2, 2017. doi: 10.1017/fmp.2017.1. Flyspeck project: complete formal verification of the Kepler conjecture.
- M. J. H. Heule, O. Kullmann, and V. W. Marek. Solving and verifying the boolean pythagorean triples problem via cube-and-conquer. In *Theory and Applications of Satisfiability Testing - SAT 2016*, volume 9710 of *Lecture Notes in Computer Science*, pages 228-245. Springer, 2016. Computer-intensive proof generating a multi-terabyte certificate.
- I. Hoffman and A. Soifer. Almost chromatic number of the plane. *Geombinatorics*, 3(2):38-40, 1993.
- I. Hoffman and A. Soifer. Another six-coloring of the plane. *Discrete Mathematics*, 150(1-3):427-429, 1996.
- K. Munding, S. Pokutta, C. Spiegel, and M. Zimmer. Extending the Continuum of Six-Colorings. *Geombinatorics Quarterly*, 5 2024.
- K. Munding, M. Zimmer, A. Kiem, C. Spiegel, and S. Pokutta. Neural Discovery in Mathematics: Do Machines Dream of Colored Planes? *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 267:45236-45255, 5 2025. ISSN 2640-3498. doi: 10.48550/arXiv.2501.18527.
- J. F. Nash and M. T. Rassias. *Open problems in mathematics*. Springer, 2016.
- J. Parts. What percent of the plane can be properly 5-and 6-colored? *arXiv preprint arXiv:2010.12668*, 2020.
- S. Pokutta. Frank-Wolfe Beyond 1/t Convergence. *preprint*, 4 2026. doi: 10.48550/arXiv.2604.28006.
- D. Pritikin. All unit-distance graphs of order 6197 are 6-colorable. *Journal of Combinatorial Theory, Series B*, 73(2):159-163, 1998.
- N. Robertson, D. P. Sanders, P. Seymour, and R. Thomas. The four colour theorem. *Journal of Combinatorial Theory, Series B*, 70(1):2-44, 1997. Shorter and fully computer-verified proof reducing the number of configurations.
- A. Soifer. A six-coloring of the plane. *Journal of Combinatorial Theory, Series A*, 61(2):292-294, 1992.
- A. Soifer. Six-realizable set x_6 . *Geombinatorics*, III(4):140-145, 1994a.
- A. Soifer. An infinite class of six-colorings of the plane. *Congressus Numerantium*, pages 83-86, 1994b.
- A. Soifer. *The mathematical coloring book: Mathematics of coloring and the colorful life of its creators*. Springer, 2009.
- A. Soifer. *The New Mathematical Coloring Book*. Springer, 2024.
- M. Zimmer, N. Pelleriti, C. Roux, and S. Pokutta. The Agentic Researcher: A Practical Guide to AI-Assisted Research in Mathematics and Machine Learning. *Proceedings of the ICML 2026 Workshop on AI as a Tool for Mathematics, Computer Science, and Machine Learning (AI4Research)*, 6 2026. doi: 10.48550/arXiv.2603.15914.