



Neural Discovery in Mathematics: Do Machines Dream of Colored Planes?

E. Balalaxminarayana¹, Mullagiri Venkatakrishna²

¹Lecture in Computer Science, S. R. Arts & Science Government Degree College, Kothagudem

²Lecture in Computer Science, Government Degree College (A) Paloncha

Abstract- This paper reviews the 2025 research work titled "Neural Discovery in Mathematics: Do Machines Dream of Colored Planes?" by Konrad Mundinger, Max Zimmer, Aldo Kiem, Christoph Spiegel, and Sebastian Pokutta. The study presents a novel approach that uses neural networks to drive mathematical discovery. Focusing on the Hadwiger-Nelson problem—a long-standing open question in discrete geometry and combinatorics—the researchers reformulated a complex geometric coloring challenge into an optimization task that neural networks could solve. Through this method, they discovered two novel six-colorings of the Euclidean plane, achieving the first improvement in thirty years to a variant of the original problem. This paper provides a conceptual, equation-free explanation of the research, its methodology, key findings, and broader implications for the use of artificial intelligence in mathematical discovery.

Keywords: Artificial Intelligence (AI), Neural Networks, Mathematical Discovery, Machine Learning, Hadwiger-Nelson Problem, Euclidean Plane Coloring.

I. INTRODUCTION

Mathematics is often viewed as a discipline built on rigid logic and absolute proof. A statement is either true or false; a theorem is either proven or remains open. Yet, the path to a proof is rarely straightforward. It is often messy, relying on intuition, visualization, and trial and error. In recent years, an intriguing question has emerged: can artificial intelligence serve as an engine of intuition for pure mathematics?

The 2025 paper by Mundinger and colleagues offers a compelling affirmative answer. Presented at the 42nd International Conference on Machine Learning (ICML) and published in the Proceedings of Machine Learning Research (PMLR), the work demonstrates how neural networks can actively contribute to mathematical discovery, not merely by verifying known results but by generating genuinely new ones. The paper's playful title—"Do Machines Dream of Colored Planes?"—evokes Philip K. Dick's novel "Do Androids Dream of Electric Sheep?" and invites readers to consider whether machines can engage in the creative, exploratory aspects of mathematical thinking.

II. THE HADWIGER-NELSON PROBLEM: A GEOMETRIC PUZZLE

To understand the contribution of this research, one must first grasp the mathematical problem at its heart: the Hadwiger-Nelson problem. This problem, named after mathematicians Hugo Hadwiger and Edward Nelson, asks a seemingly simple question: what is the minimum number of colors required to color the entire Euclidean plane such that no two points exactly one unit apart share the same color?



The problem was first conceived by Edward Nelson in 1950 and later popularized by Martin Gardner in Scientific American magazine in 1960. Despite its straightforward formulation, it has resisted a complete solution for over seven decades. For many years, it was known that the answer lies between four and seven. In 2018, the mathematician Aubrey de Grey improved the lower bound, proving that at least five colors are needed. However, the exact chromatic number of the plane—whether it is five, six, or seven—remains unknown to this day.

The version of the problem that Munding and his team tackled is a variant known as the "off-diagonal" Hadwiger-Nelson problem. In this variant, different colors are allowed to avoid different distances. Specifically, the researchers focused on a six-coloring scenario: five colors must avoid having any two points at unit distance with the same color, while the sixth color must avoid a different, specified distance. This variant had not seen significant progress for thirty years prior to this work.

III. THE NEURAL APPROACH: REFORMULATING GEOMETRY AS OPTIMIZATION

The key innovation of the research lies in how the team reformulated the problem. Traditional mathematical approaches to the Hadwiger-Nelson problem involve discrete, hard constraints: a point either receives a color or it does not. This either-or nature makes the problem difficult to solve using conventional continuous optimization tools.

The researchers introduced a breakthrough by converting this discrete coloring problem into a continuous optimization task. Instead of assigning each point a fixed color, they assigned a probabilistic coloring function—a distribution of probabilities over the available colors for every point on the plane. In other words, rather than saying "this point is definitely red," they allowed the neural network to express, "this point has a 90% chance of being red, a 5% chance of being blue, and so on."

This probabilistic representation made the problem differentiable, meaning that the researchers could define a loss function that measured how well a given coloring avoided forbidden monochromatic distances. A loss function is a standard tool in machine learning that quantifies the error of a model's predictions; the lower the loss, the better the solution. By making the loss function differentiable, the team enabled gradient-based exploration—a technique where the neural network repeatedly adjusts its parameters in the direction that most reduces the loss, gradually moving toward a valid coloring.

Effectively, the neural network was tasked with "painting" the plane in a way that minimized conflicts. Over many iterations, the network learned patterns and strategies that satisfied the coloring constraints.

This approach is fundamentally different from earlier attempts to apply AI to the Hadwiger-Nelson problem. Rather than forcing the network to find a perfect coloring from the start, the method allowed it to propose almost-perfect colorings. Human mathematicians could then examine these near-solutions, identify the small regions where conflicts remained, and manually adjust them to achieve a fully valid coloring.

IV. KEY FINDINGS: TWO NOVEL SIX-COLORINGS

The neural network did not disappoint. After extensive training and optimization, the system produced two novel six-colorings of the Euclidean plane that satisfied the constraints of the off-diagonal variant. These were not minor variations of known colorings. They represented the first improvement in thirty years to this variant of the problem. The neural network consistently suggested patterns that, after careful mathematical analysis and formalization, led to valid geometric constructions that human mathematicians had not previously discovered.



The researchers emphasized that while the neural network identified promising patterns, formalizing these discovered patterns into rigorous mathematical constructions still required substantial human effort. The machine did not simply output a final, verified coloring. Instead, it acted as a creative collaborator—generating candidate solutions and revealing structural regularities that human researchers could then analyze, generalize, and prove correct.

In addition to the two main findings, the team demonstrated the broader applicability of their method through further numerical experiments, showing that the same neural framework could be adapted to other geometric coloring problems.

V. BROADER IMPLICATIONS FOR MATHEMATICAL DISCOVERY

This research is significant for reasons that extend far beyond the Hadwiger-Nelson problem itself. First, it provides a concrete example of how machine learning can contribute to pure mathematics—a field not typically associated with empirical or data-driven methods. While AI has made impressive strides in applied domains such as image recognition, natural language processing, and drug discovery, its role in advancing fundamental mathematical knowledge has been less explored. This study shows that neural networks can generate genuinely new mathematical insights, not merely rediscover known ones.

Second, the work establishes a generalizable methodology. The technique of reformulating discrete geometric problems as continuous optimization tasks using probabilistic colorings and differentiable loss functions can be applied to other combinatorial and geometric challenges. This opens the door to using neural networks as discovery engines across a wide range of mathematical domains.

Third, the research highlights a collaborative paradigm for human-AI interaction in mathematics. The machine does not replace the mathematician; rather, it augments human intuition. The neural network explores vast solution spaces far more quickly than any human could, identifying promising patterns and counterexamples. The mathematician then brings rigorous analysis and formal proof to bear, transforming machine-generated suggestions into verified mathematical truths. This synergistic relationship could accelerate progress on many open problems that have resisted solution through traditional methods alone.

The paper's presentation as an "oral paper" at ICML 2025—one of the most prestigious conferences in machine learning—attests to the significance and novelty of this approach within the research community.

VI. CONCLUSION

"Neural Discovery in Mathematics: Do Machines Dream of Colored Planes?" represents a landmark at the intersection of artificial intelligence and pure mathematics. By applying neural networks to the Hadwiger-Nelson problem, the researchers achieved the first progress in three decades on a variant of this long-standing open question. The discovery of two novel six-colorings of the Euclidean plane demonstrates that machines can indeed "dream" in mathematical terms—generating creative, non-obvious solutions that human mathematicians had overlooked.

Yet the dream is not autonomous. The neural network needed human guidance to formalize its findings into rigorous proofs. The true power of this approach lies in collaboration: machines exploring vast possibility spaces with superhuman speed, and humans applying deep insight and formal reasoning to transform those explorations into lasting mathematical knowledge. As this methodology is extended to



other problems, the partnership between human intuition and machine computation promises to unlock new frontiers in mathematical discovery.

REFERENCES

1. Mundinger, K., Zimmer, M., Kiem, A., Spiegel, C., & Pokutta, S. (2025). Neural Discovery in Mathematics: Do Machines Dream of Colored Planes? Proceedings of the 42nd International Conference on Machine Learning, PMLR 267. <https://arxiv.org/abs/2501.18527>
2. Mundinger, K., Zimmer, M., Kiem, A., Spiegel, C., & Pokutta, S. (2025). Neural Discovery in Mathematics: Do Machines Dream of Colored Planes? – Max Zimmer's blog post. <https://maxzimmer.org/2025/05/29/neural-discovery-in-mathematics-do-machines-dream-of-colored-planes/>
3. Zuse Institute Berlin (ZIB). (2025). Neural Networks Guide Mathematical Discovery (Research highlight). <https://www.zib.de/news/neural-networks-guide-mathematical-discovery>
4. Mundinger, K., Pokutta, S., Spiegel, C., & Zimmer, M. (2024). Extending the Continuum of Six-Colorings (Informal summary). <https://www.pokutta.com/blog/2024/07/28/extending-the-continuum-of-six-colorings/>
5. Hadwiger–Nelson problem. Wikipedia. https://en.wikipedia.org/wiki/Hadwiger%E2%80%93Nelson_problem
6. Hadwiger–Nelson Problem. Wolfram MathWorld. <https://mathworld.wolfram.com/Hadwiger-NelsonProblem.html>
7. Bytez.com research summary. (2025). Neural Discovery in Mathematics: Do Machines Dream of Colored Planes? <https://bytez.com/papers/2501.18527>