

DeepMind's AlphaGeometry solves 25 of 30 olympiad geometry problems

Source copy kept by wheresthe.ai for the story "DeepMind's AlphaGeometry solves 25 of 30 olympiad geometry problems" (<https://wheresthe.ai/d/deepmind-alphageometry-2024-01/>).

Source: <https://deepmind.google/blog/alphageometry-an-olympiad-level-ai-system-for-geometry/>

Text saved 18 Oct 2026 (the page as published).

SHA-256 of the saved text: fac71052f2a9988da2a3cca439e4e8ce319eb642dc8cc9b9e2521e8c470936e2

January 17, 2024 Science AlphaGeometry: An Olympiad-level AI system for geometry Trieu Trinh and Thang Luong

Share

Our AI system surpasses the state-of-the-art approach for geometry problems, advancing AI reasoning in mathematics. Reflecting the Olympic spirit of ancient Greece, the International Mathematical Olympiad is a modern-day arena for the world's brightest high-school mathematicians. The competition not only showcases young talent, but has emerged as a testing ground for advanced AI systems in math and reasoning. In a paper published today in *Nature*, we introduce AlphaGeometry, an AI system that solves complex geometry problems at a level approaching a human Olympiad gold medalist - a breakthrough in AI performance. In a benchmarking test of 30 Olympiad geometry problems, AlphaGeometry solved 25 within the standard Olympiad time limit. For comparison, the previous state-of-the-art system solved 10 of these geometry problems, and the average human gold medalist solved 25.9 problems.

In our benchmarking set of 30 Olympiad geometry problems (IMO-AG-30), compiled from the Olympiads from 2000 to 2022, AlphaGeometry solved 25 problems under competition time limits. This is approaching the average score of human gold medalists on these same problems. The previous state-of-the-art approach, known as "Wu's method", solved 10.

AI systems often struggle with complex problems in geometry and mathematics due to a lack of reasoning skills and training data. AlphaGeometry's system combines the predictive power of a neural language model with a rule-bound deduction engine, which work in tandem to find solutions. And by developing a method to generate a vast pool of synthetic training data - 100 million unique examples - we can train AlphaGeometry without any human demonstrations, sidestepping the data bottleneck. With AlphaGeometry, we demonstrate AI's growing ability to reason logically, and to discover and verify new knowledge. Solving Olympiad-level geometry problems is an important milestone in developing deep mathematical reasoning on the path towards more advanced and general AI systems. We are open-sourcing the AlphaGeometry code and model, and hope that together with other tools and approaches in synthetic data generation and training, it helps open up new possibilities across mathematics, science, and AI.

It makes perfect sense to me now that researchers in AI are trying their hands on the IMO geometry problems first because finding solutions for them works a little bit like chess in the sense that we have a rather small number of sensible moves at every step. But I still find it stunning that they could make it work. It's an impressive achievement.

Ngô Bảo Châu Fields Medalist and IMO Gold Medalist

AlphaGeometry adopts a neuro-symbolic approach. AlphaGeometry is a neuro-symbolic system made up of a neural language model and a symbolic deduction engine, which work together to find proofs for complex geometry theorems. Akin to the idea of "thinking, fast and slow", one system provides fast, "intuitive" ideas, and the other, more deliberate, rational decision-making. Because language models excel at identifying general patterns and relationships in data, they can quickly predict potentially useful constructs, but often lack the ability to reason rigorously or explain their decisions. Symbolic deduction engines, on the other hand, are based on formal logic and use clear rules to arrive at conclusions. They are rational and explainable, but they can be "slow" and inflexible - especially when dealing with large, complex problems on their own. AlphaGeometry's language model guides its symbolic deduction engine towards likely solutions to geometry problems. Olympiad geometry problems are based on diagrams that need new geometric constructs to be added before they can be solved, such as

points, lines or circles. AlphaGeometry's language model predicts which new constructs would be most useful to add, from an infinite number of possibilities. These clues help fill in the gaps and allow the symbolic engine to make further deductions about the diagram and close in on the solution.

AlphaGeometry solving a simple problem: Given the problem diagram and its theorem premises (left), AlphaGeometry (middle) first uses its symbolic engine to deduce new statements about the diagram until the solution is found or new statements are exhausted. If no solution is found, AlphaGeometry's language model adds one potentially useful construct (blue), opening new paths of deduction for the symbolic engine. This loop continues until a solution is found (right). In this example, just one construct is required.

AlphaGeometry solving an Olympiad problem: Problem 3 of the 2015 International Mathematics Olympiad (left) and a condensed version of AlphaGeometry's solution (right). The blue elements are added constructs. AlphaGeometry's solution has 109 logical steps.

See the full solution

Generating 100 million synthetic data examples Geometry relies on understanding of space, distance, shape, and relative positions, and is fundamental to art, architecture, engineering and many other fields. Humans can learn geometry using a pen and paper, examining diagrams and using existing knowledge to uncover new, more sophisticated geometric properties and relationships. Our synthetic data generation approach emulates this knowledge-building process at scale, allowing us to train AlphaGeometry from scratch, without any human demonstrations. Using highly parallelized computing, the system started by generating one billion random diagrams of geometric objects and exhaustively derived all the relationships between the points and lines in each diagram. AlphaGeometry found all the proofs contained in each diagram, then worked backwards to find out what additional constructs, if any, were needed to arrive at those proofs. We call this process "symbolic deduction and traceback".

Visual representations of the synthetic data generated by AlphaGeometry

That huge data pool was filtered to exclude similar examples, resulting in a final training dataset of 100 million unique examples of varying difficulty, of which nine million featured added constructs. With so many examples of how these constructs led to proofs, AlphaGeometry's language model is able to make good suggestions for new constructs when presented with Olympiad geometry problems. Pioneering mathematical reasoning with AI The solution to every Olympiad problem provided by AlphaGeometry was checked and verified by computer. We also compared its results with previous AI methods, and with human performance at the Olympiad. In addition, Evan Chen, a math coach and former Olympiad gold-medalist, evaluated a selection of AlphaGeometry's solutions for us. Chen said: "AlphaGeometry's output is impressive because it's both verifiable and clean. Past AI solutions to proof-based competition problems have sometimes been hit-or-miss (outputs are only correct sometimes and need human checks). AlphaGeometry doesn't have this weakness: its solutions have machine-verifiable structure. Yet despite this, its output is still human-readable. One could have imagined a computer program that solved geometry problems by brute-force coordinate systems: think pages and pages of tedious algebra calculation. AlphaGeometry is not that. It uses classical geometry rules with angles and similar triangles just as students do."

AlphaGeometry's output is impressive because it's both verifiable and clean...It uses classical geometry rules with angles and similar triangles just as students do.

Evan Chen Math Coach and Olympiad Gold Medalist

As each Olympiad features six problems, only two of which are typically focused on geometry, AlphaGeometry can only be applied to one-third of the problems at a given Olympiad. Nevertheless, its geometry capability alone makes it the first AI model in the world capable of passing the bronze medal threshold of the IMO in 2000 and 2015. In geometry, our system approaches the standard of an IMO gold-medalist, but we have our eye on an even bigger prize: advancing reasoning for next-generation AI systems. Given the wider potential of training AI systems from scratch with large-scale synthetic data, this approach could shape how the AI systems of the future discover new knowledge, in math and beyond. AlphaGeometry builds on Google DeepMind and Google Research's work to pioneer mathematical reasoning with AI – from exploring the beauty of pure mathematics to solving mathematical and scientific problems with language models . And most recently, we introduced FunSearch , which made the first discoveries in open problems in mathematical sciences using Large Language Models. Our long-term goal remains to build AI systems that can generalize across mathematical fields, developing the sophisticated problem-solving and reasoning that general AI systems will depend on, all the while extending the

frontiers of human knowledge.

[Learn more about AlphaGeometry](#)

[Read our paper in Nature](#) [Access the AlphaGeometry Github](#)

Acknowledgements This project is a collaboration between the Google DeepMind team and the Computer Science Department of New York University. The authors of this work include Trieu Trinh, Yuhuai Wu, Quoc Le, He He, and Thang Luong. We thank Rif A. Saurous, Denny Zhou, Christian Szegedy, Delesley Hutchins, Thomas Kipf, Hieu Pham, Petar Veličković, Edward Lockhart, Debidatta Dwibedi, Kyunghyun Cho, Lerrel Pinto, Alfredo Canziani, Thomas Wies, He He's research group, Evan Chen, Mirek Olsak, Patrik Bak for their help and support. We would also like to thank Google DeepMind leadership for the support, especially Ed Chi, Koray Kavukcuoglu, Pushmeet Kohli, and Demis Hassabis.

Related posts

[Exploring the beauty of pure mathematics in novel ways](#) December 2021 Science [Learn more](#)

[FunSearch: Making new discoveries in mathematical sciences using Large Language Models](#) December 2023 Science [Learn more](#)