



SuperDataScience

SDS PODCAST

EPISODE 1038:

AI'S CONTENTIOUS

MATH

BREAKTHROUGHS:

THE FIELDS

MEDALLISTS PUSH

BACK



Jon Krohn:	00:00	This is episode number 1038 on AI's contentious math breakthroughs. Welcome back to the Super Data Science Podcast. I'm your host, Jon Krohn. Today's episode covers the remarkable and remarkably contentious run of mathematical results that AI systems have produced over the past few weeks. The headline is OpenAI's claim to have resolved one of the Clay Mathematics Institute's millennium prize problems. I'm going to cover what was claimed and why it's exciting. I'm also going to cover why many of the world's top mathematicians are alarmed. Some background first. To mark the turn of the century, the Clay Institute selected seven of the most important unsolved problems in mathematics and attached a million dollar prize to each. In the 26 years since, humans have solved only one of them, the Poincare conjecture, hopefully I'm pronouncing that correctly, back in 2000, Alen Khan, a winner of a prestigious fields medal, which is roughly analogous to a Nobel Prize in terms of prestige, but specific to mathematics, that guy Alen Khan said these problems were totally inaccessible to computers, but that was 26 years ago.
	01:15	AI has changed a fair bit. Last month, OpenAI announced that about 10,000 autonomous AI agents running on an unreleased internal model had found a singularity in the three-dimensional Navier Stokes equations. The Navier Stokes equations describe how fluids move such as air flowing over an aircraft wing, oil moving through a pipe or plasma circulating inside a star. Engineers rely on these equations every day. However, mathematicians have never been able to prove whether solutions in three dimensions always stay well-behaved or whether they can blow up by producing a point where a quantity like the fluid speed becomes infinite in a finite amount of time. That is the singularity. Not like the future, the singularity when artificial superintelligence happens, but just the mathematical term singularity. So OpenAI says its agents found such a blow up, such a singularity, a swirling



vortex whose velocity grows without bound. The agents started working on September 1st.

02:16 After 50 hours, a team of about a hundred agents had answered a closely related question, so OpenAI's researchers redirected the entire swarm toward Navier Stokes. After 88 hours and nearly three million messages passed between agents, a team of roughly 10,000 agents had a solution. How much would something like that cost? By The Economist Magazine's estimate, the effort used at least six and a half million dollars of compute. Wow. OpenAI then spent another 17 hours formalizing it and verifying the proof in a programming language called lean in which every step of a proof can be checked by machine, and they did that all using GPT-6 Astra. OpenAI has made the writeup, the paper and the lean code public so anyone can machine check the results. This is a staggering achievement on its face, but a balanced picture requires three big caveats. The first caveat concerns what exactly has been proved.

03:14 The millennium problem has four official variants. OpenAI's fluid has a smooth external force applied to it the whole way from rest until the singularity forms. On that basis, OpenAI says it has established statements C and D of the four official variants. The unforced variants, options A and B, remain unsolved. Many mathematicians consider those unforced variants at the heart of the problem. Tellingly, OpenAI has said it does not intend to claim the prize. The Clay Institute has acknowledged that the problem has apparently been settled, but says its verification process is intentionally slow and it now lists the problem status as active. Clay won't even consider a proposed solution until at least two years after publication, so we will be waiting a while for that official result. So that's the first caveat. The second caveat is a messy credit dispute. Tristan Buckmaster of New York University and Lavont Alpugy of Anthropic, I'm probably



butchering Levant's name, the two of them had been using AI to hunt for blowups in related fluid equations.

04:18 So in mid-August, they found blow up results for simpler cousins of Navier Stokes, including in something called the three-dimensional Euler equations with forcing. OpenAI says it launched its own effort after hearing a rumor about their work. Buckmaster, the NYU guy, says OpenAI proposed that he alone write up the Navia Stokes result while acknowledging that OpenAI's model had resolved it. The proposal excluded Alpugy, the anthropic guy, citing a conflict of interest because he works at Anthropic. Buckmaster declined. OpenAI denies drawing on the pairs on published work. However, the two mathematicians, Buckmaster and Alpugi, had used OpenAI's products and OpenAI told the economist it cannot rule out that their drafts ended up in its training data. OpenAI has since credited Buckmaster and Alpugy for their concurrent work and offered a joint priority announcement. The Navia Stokes claim turned out to be the opening act. On October 6th, OpenAI posted 722 manuscripts to a public GitHub repository grouped into 372 result families.

05:22 The claimed results span famous open problems such as the unique games conjecture, the exponent of matrix multiplication, Hadwiger's conjecture, and the chromatic number of the plane. Yes, I am just reading those four famous open problems. I was not personally familiar with them, but I do have a citation here. A academic named Alvaro Lozano Robledo of the University of Connecticut called this posting of more than 700 manuscripts on all of these famous open problems. He called it possibly the single most important day in the history of mathematics thus far. Wow. But this is where the third caveat comes in, verification. The very next day, OpenAI withdrew three manuscripts after a sign error, so positive or negative, broke one proof and the two papers that depended on it leaving 719. Over the following week, OpenAI revised 14



more manuscripts and told a organization called Retraction Watch that about half of the release results had been published without full verification.

06:30 OpenAI says 42% of the top line results are formalized in lean, that programming language I was talking about earlier, but it doesn't define what counts as formalized and independent recounts put the figure between 24% and 36%. So yeah, not much verification and a bunch of recounts. So how have mathematicians reacted to all of this? Many are distressed. On September 11th, 25 fields medalists, again, that's kind of like the equivalent to a Nobel Prize in mathematics in terms of prestige. These 25 fields metalists published a joint declaration warning that the goals of AI companies and the goals of mathematicians are severely misaligned. More than 8,000 researchers have since endorsed that declaration. Their core argument is that solving a problem is a means to an end, not the end itself. The end is human understanding and the long struggle toward a solution is what produces new tools.

07:26 Imaginary numbers, for example, were dismissed as useless in the 16th century and now underpin electrical engineering and quantum physics. The fields medalist Hugo Dominil Kopin specifically wrote that I am paralyzed. Columbia's Michael Harris told the journal nature he considers the bulk release illegitimate and won't be reading the papers. Dana Moshkowitz, an expert on the unique games conjecture said OpenAI's pre-print on that problem was so poorly written that it can't be read without AI assistance. Harvard's Melanie Wood pointed out that when a model solves a problem, nobody understands the solution yet at release, so that's when the human work begins. A newly formed group called the Association for Human Mathematics described the release as a demonstration of power. So yeah, largely negative feedback, but other reactions are more positive. Al Pagee himself, the guy that works at Anthropic, despite his own



dispute with OpenAI, called the release a historic moment.

08:27 Ying Shieh at Kennesaw State University, apologies to Ying if I'm pronouncing their name incorrectly, is treating the release as a mine of new lemmas, L-E-M-M-A-S, Lemas, which are constructions and proof techniques to extract and study. That work cuts against the claim that the papers contain no new ideas, though the novelty assessments are preliminary. There's also an earlier precedent for how this can go well when an OpenAI model disproved something called the AirDos unit distance conjecture in May, five mathematicians posted a digested human verified version of the counterexample that same day. In that case, the machine found the result and humans turned it into understanding and very quickly. So yeah, that's kind of the critical feedback and some of the positive feedback on this big OpenAI release. But all of this is specific to mathematics. Beyond mathematics, a University of Chicago analysis of 41 million scientific papers found that researchers who use AI published more and are cited more.

09:30 However, AI use also narrowed the range of topics studied by almost 5% and cut engagement with other scientists' work by 22%. Theoretical physics, a closer relative of mathematics is probably next in line for this kind of disruption. For those of us who build AI systems, I see a lesson here that extends well beyond mathematics. When a model can generate hundreds of plausible proofs or thousands of lines of code in hours, generation stops being the bottleneck, of course. Verification is what becomes the bottleneck. The results from this period that will hold up best are the ones paired with machine checkable evidence such as a lean proof or a strong test suite and with humans who take the time to understand them. Results released in a rush without that evidence and understanding will draw much of the skepticism.



These are wild times for sure and for many mathematicians, probably unsettling times indeed.

- 10:24 The capability leap here seems to be real. Whether that leap ends up enriching human understanding or bypassing it will depend on the norms we build around verification, around credit and around explanation. And yeah, we might need to start building those norms over the coming months. Capabilities are moving quickly and math is surely just the beginning here. Theoretical physics I already mentioned may be next, but computer science and perhaps scientific discovery in general may be automated by AI to a large extent in the coming years. Lost a process, but hopefully this largely ends up being for the best for humankind. For more on how AI-enabled breakthroughs could herald a positive new era for all of us, check out the preceding episode of this podcast, number 1037 with Dr. Richard Socher, as well as you can go back to episode number 975 from earlier this year with Zach Cass.
- 11:18 Those are both very positive, optimistic episodes on how AI breakthroughs could mean lots of positive outcomes for us. I, for one, will continue to maintain as much optimism as I can. All right, that's the end of today's episode. If you enjoyed it or know someone who might consider sharing this episode with them, leave a review of the show on your favorite podcasting platform or YouTube. Tag me in a LinkedIn post with your thoughts. I will respond to all of those. And if you aren't already, obviously, subscribe to the show. Most importantly, however, we hope you'll just keep on listening. Until next time, keep on rocking it out there and I'm looking forward to enjoying another round of the Super Data Science podcast with you very soon.