AI News Digest

10/12/2026 - 10/18/2026

AI's Rapid Leap into Core Mathematical Research

Recent work reveals that large language models are not truly random when asked to pick a number, exposing systematic biases in their output distributions. A forward-looking study predicts that between mid-2024 and mid-2026 AI will dramatically expand its capacity to solve theoretical problems across probability, economics, high-energy physics, cosmology and quantum physics. Meanwhile, OpenAI disclosed a trove of 719 AI-generated math papers, spanning 372 result families and over 34,000 pages, underscoring the scale at which generative models are now producing scholarly content. Together, these signals point to an accelerating integration of AI into the heart of mathematical and scientific inquiry.

KEY FINDINGS

✓ Language models exhibit a measurable bias toward certain integers when asked to generate a "random" number, with extracted logits showing deviations of several percentage points from a uniform 1% chance per value.

✓ Forecasts indicate AI's ability to solve theoretical problems in probability, statistics, economics, high-energy physics, gravitation, cosmology, and quantum physics will increase markedly between mid-2024 and mid-2026.

✓ OpenAI released 719 AI-generated mathematics papers, organized into 372 result families across 17 disciplines, totaling 34,669 pages, with theoretical computer science accounting for 40 families.

IMPLICATIONS

→ Researchers will increasingly rely on AI for hypothesis generation, proof assistance, and rapid exploration of complex scientific problems.

→ The academic community will need new standards for evaluating, attributing, and mitigating bias in AI-produced research outputs.

→ The sheer volume of AI-generated literature could reshape publishing models, funding allocations, and the pace of scientific discovery.

> Key Metrics: Bias in random-number generation: certain numbers appear up to ~2-3x more often than the uniform 1% baseline.; Projected AI capability growth window: mid-2024 to mid-2026.; OpenAI math corpus: 719 papers, 372 result families, 34,669 pages.; Disciplinary coverage: 17 fields; theoretical computer science dominates with 40 families.

1. How Many Guesses Does It Take to Crack an AI's "Random" Number?

Category: Research
Departments: Computer Science, Mathematical Sciences
Topics: AI Ethics & Safety, Number Bias, RLHF Amplification, Randomness Quantification, LLM Security

Impact Score: 25.0
Rubric Scores

Research Relevance: 4	Educational Value: 4	Innovation / Novelty: 4
Practical Impact: 4	Interdisciplinary Potential: 4	Ethical / Policy Implications: 5

The article quantifies the bias of language models toward specific numbers when asked to pick a random integer between 1 and 100 by extracting exact probability distributions from model logits, avoiding sampling noise. It compares base models with their instruction-tuned counterparts across five open model families and two prompt formats, finding that a 1.5-billion-parameter instruction-tuned model selects the number 47 in roughly 75 % of cases, and still does so over one-third of the time with a reworded prompt. By converting the probability vectors into bits of randomness and expected guess counts, the study shows that these biases dramatically reduce the security of secret selections. The work also highlights the reproducibility advantages of open checkpoints and documents how post-training stages, particularly RLHF, amplify the concentration of probability on favorite numbers.

https://towardsdatascience.com/how-many-guesses-does-it-take-to-crack-an-ais-random-number/

2. AI Came for Math. Now It's Chasing the Rest of Science

Category: Research
Departments: Computer Science, Mathematical Sciences, Physics, Economics & Finance
Topics: Science & Research, Theoretical Problem Solving, Probabilistic Reasoning, Quantum Physics Modeling

Impact Score: 22.0
Rubric Scores

Research Relevance: 4	Educational Value: 3	Innovation / Novelty: 4
Practical Impact: 4	Interdisciplinary Potential: 5	Ethical / Policy Implications: 2

Thompson's study predicts AI will dramatically increase its ability to solve theoretical problems in probability, statistics, economics, high-energy physics, gravitation, cosmology, and quantum physics. Between mid-2024 and mid-2026, the proportion of such problems correctly solved by leading OpenAI models rose from roughly 2 % to about 35 %, with no indication of slowdown. The evaluation used problems from papers published after the models' training cutoffs (Dec 2025-May 2026), blocked internet access, and required the models to generate solutions independently. The results imply that the disruptive impact currently observed in mathematics is likely to extend to other theory-intensive scientific domains.

https://www.wired.com/story/ai-is-coming-for-science-next-watch-out/

3. What OpenAI's 719 Math Papers Mean for Data Science

Category: Research
Departments: Computer Science, Mathematical Sciences
Topics: Science & Research, Randomized Edit Distance, Automated Theorem Generation, AI-Generated Research Catalog

Impact Score: 15.0
Rubric Scores

Research Relevance: 2	Educational Value: 2	Innovation / Novelty: 3
Practical Impact: 2	Interdisciplinary Potential: 4	Ethical / Policy Implications: 2

On 6 October 2026 OpenAI published a catalogue of 719 papers generated by an internal model, grouped into 372 result families across 17 disciplines and totaling 34,669 pages. The collection is dominated by theoretical computer science (40 families) and probability/statistical mechanics (21 % of pages), with most families containing one or two papers. The article highlights a new edit-distance algorithm from the catalogue that, using randomness, estimates distance within a fixed error percentage in time proportional to the combined string length, contrasting with the quadratic time of exact computation. It notes that practical deployment will depend on choices of accuracy, distance metrics, and randomness, and on bridging the gap between the proved theorems and usable implementations.

https://towardsdatascience.com/what-openais-719-math-papers-mean-for-data-science/