

## **Primes & Programs: The AI Revolution in Prime Number Discovery**

Dr. Alok Sharma

Assistant Professor, Department of Mathematics, NIILM University, Kaithal

### **Introduction**

Prime numbers — those natural numbers greater than 1 that are divisible only by 1 and themselves — have fascinated mathematicians for millennia. Their simplicity masks a profound complexity. As early as 300 BCE, Euclid proved the infinitude of primes in *Elements*, laying the groundwork for number theory. Over the centuries, luminaries like Euler, Gauss, Riemann, and Ramanujan explored their distribution, formulated conjectures, and connected primes to the deepest structures of mathematics.

Despite centuries of progress, many fundamental questions about primes remain unresolved. The **Riemann Hypothesis**, one of the most important unsolved problems in mathematics, directly relates to the distribution of prime numbers. The **Twin Prime Conjecture**, suggesting the infinitude of prime pairs differing by 2, also remains unproven. These open problems continue to inspire generations of mathematicians — and now, machines.

### **Why Prime Numbers Matter**

Primes aren't just a mathematical curiosity — they're essential to modern life. Their role in **number theory** has profound implications for **cryptography**, **computer science**, **digital security**, and **data compression**.

- In **cryptography**, primes are the backbone of public-key systems like RSA, which protect everything from emails to bank transactions. Their unpredictability ensures security in digital communications.
- In **computer science**, primality testing and factorization underpin many algorithms and data structures.
- In **physics and biology**, prime patterns appear in unexpected places — from quantum energy levels to genetic code sequences.

Prime numbers thus sit at the intersection of **pure mathematics** and **practical application**, making their study both timeless and urgently relevant.

## **The Rise of AI in Mathematical Discovery**

Enter **artificial intelligence** — a technological force that is transforming fields as diverse as medicine, art, language, and now, mathematics.

AI, particularly through **machine learning** and **deep learning**, has shown remarkable potential in identifying patterns in massive datasets, solving complex optimization problems, and even generating original hypotheses. In recent years, **AI models have started making serious contributions to mathematical research**, with systems capable of conjecturing, proving, and discovering new mathematical structures.

For prime numbers, this opens up a thrilling new chapter. Traditional methods of prime analysis often depend on deep theoretical insights and extensive computation. Now, **AI offers a complementary approach** — discovering patterns in large prime datasets, approximating solutions to unsolved problems, and even revealing hidden relationships that evade human intuition.

The fusion of **ancient mathematical inquiry** and **cutting-edge AI technology** promises not just faster calculations, but potentially a deeper understanding of the nature of primes. As this chapter unfolds, we will explore how AI is already changing the way we think about prime numbers — and what the future may hold.

## **The Beauty and Complexity of Prime Numbers**

### **Prime Number Properties**

Prime numbers are the indivisible units of the natural number system — the “building blocks” of all other numbers. Every whole number greater than 1 can be uniquely factored into a product of prime numbers, a fact known as the **Fundamental Theorem of Arithmetic**. This foundational role makes prime’s essential not only in number theory but across the entire landscape of mathematics.

Despite this apparent simplicity, primes behave in profoundly unpredictable ways. Some of their notable properties include:

- **Irregular Distribution:** While primes become less frequent as numbers grow larger, they never disappear. There's no formula that generates all and only primes, and their spacing becomes increasingly erratic.
- **Density:** The **Prime Number Theorem** approximates how many primes exist below a given number  $n$  by showing that the number of primes less than  $n$  is approximately  $n / \log(n)$ .
- **Infinitude:** Euclid's elegant proof more than 2,000 years ago confirmed that the primes are infinite — a simple yet powerful result.

Primes also exhibit **symmetry**, **randomness**, and **structure**, all at once. This paradox — appearing both patterned and chaotic — is what makes them so fascinating to mathematicians and so difficult to tame analytically.

## Unsolved Problems and Famous Conjectures

Prime numbers have attracted some of the most famous unsolved problems in mathematics — puzzles that have withstood centuries of intellectual effort. Among them:

### 1. The Riemann Hypothesis

Proposed in 1859 by Bernhard Riemann, this conjecture links the distribution of prime numbers to the zeros of the **Riemann zeta function**, a complex function that encodes prime-related information. It states that all non-trivial zeros of the zeta function lie on the “critical line” in the complex plane, where the real part is  $\frac{1}{2}$ .

**Why it matters:** If true, the Riemann Hypothesis would provide an incredibly precise understanding of how primes are distributed — essentially forming the ‘grand unifying theory’ of prime number behavior.

### 2. The Twin Prime Conjecture

This conjecture suggests there are infinitely many pairs of prime numbers that differ by 2 (e.g., 11 and 13, 17 and 19). Despite increasing numerical evidence and partial progress — such as **Yitang Zhang's 2013 proof** that there are infinitely many primes with a bounded gap — the full conjecture remains unproven.

### 3. Goldbach's Conjecture

First proposed in 1742, it posits that **every even number greater than 2 is the sum of two primes**. Like the Twin Prime Conjecture, this has been verified computationally up to very large numbers but has no formal proof.

These open questions are more than intellectual curiosities — they lie at the core of mathematical understanding, with implications for cryptography, algorithms, and fundamental theories of numbers.

#### Traditional Methods of Analysis and Their Limitations

Historically, the study of prime numbers has relied on **pure mathematical analysis**, particularly **analytic number theory**, **combinatorics**, and **modular arithmetic**. Some powerful tools and methods include:

- **Sieve Methods** (e.g., Eratosthenes, Atkin): Algorithms that filter out non-primes from a list of integers.
- **Fourier Analysis** in number theory (e.g., Hardy–Littlewood circle method)
- **Zeta and L-functions** to study prime distribution
- **Elliptic curves and modular forms** in advanced conjectures and proofs

However, these methods have clear **limitations**:

- **Scalability**: Many traditional approaches don't scale well computationally when dealing with massive primes or testing large hypotheses.
- **Analytical Barriers**: Some conjectures may be inherently resistant to purely human mathematical intuition or manipulation.
- **Data Blindness**: Traditional methods often cannot incorporate large-scale pattern recognition over massive numerical datasets — a space where AI excels.

These limitations have opened the door for new techniques, including **experimental mathematics**, **computational simulations**, and now, **machine learning models**. While traditional mathematics seeks formal proof, AI approaches can uncover unexpected patterns or generate new conjectures — creating a fertile space for hybrid discovery.

### 3. AI in Mathematical Research

Introduction to Machine Learning, Deep Learning, and Symbolic AI

In recent years, **Artificial Intelligence (AI)** has evolved from a computational curiosity into a transformative force across science, technology, and even mathematics. At its core, AI refers to computer systems that mimic human intelligence — learning, reasoning, and solving problems. Within AI, there are three major domains particularly relevant to mathematical research:

- **Machine Learning (ML):** A subset of AI where algorithms learn patterns from data without being explicitly programmed. In mathematics, this allows models to detect regularities in number sequences or geometric structures.
- **Deep Learning (DL):** A further subset of ML that uses **artificial neural networks** with multiple layers to model complex functions. Deep learning is especially powerful in high-dimensional pattern recognition and has been applied to explore structures in number theory, algebra, and topology.
- **Symbolic AI:** Unlike ML and DL, which rely heavily on data, **symbolic AI** represents knowledge through symbols, logic, and formal rules. It's closely tied to **automated theorem proving** and **computer-assisted reasoning**, making it ideal for formal mathematics.

Together, these technologies are unlocking new possibilities for mathematical research — not just automating calculations, but aiding in **the discovery of new mathematical truths**.

The Role of AI in Pattern Recognition and Predictive Modeling

One of AI's greatest strengths lies in its ability to recognize patterns and make predictions — capabilities highly relevant to the study of **prime numbers**, where patterns are notoriously subtle and elusive.

#### 1. Pattern Discovery in Numerical Sequences

- **Neural networks** can be trained on large datasets of primes to identify potential regularities — such as spacing, frequency, or digit patterns — that may be too faint or complex for human detection.
- **Recurrent neural networks (RNNs)** and **transformer-based models** (like those used in natural language processing) can be adapted to analyze number sequences similarly to how they process language, offering a novel approach to mathematical data.

## 2. Predictive Modeling of Primes and Conjectures

- AI can approximate the distribution of primes or even "guess" the likelihood of a number being prime under certain constraints.
- Machine learning models have been used to **generate new conjectures** in number theory, some of which have later been proven or generalized by human mathematicians.

While such AI-generated insights are not always immediately rigorous, they serve as valuable tools for **hypothesis generation**, **exploratory research**, and **intuition building**.

### AI in Automated Theorem Proving and Discovery

One of the most ambitious frontiers in AI research is **automated theorem proving** — enabling machines to generate or verify mathematical proofs autonomously. This field combines **symbolic reasoning**, **logical inference**, and **formal verification systems**.

#### 1. Automated Theorem Provers (ATPs)

- Systems like **Coq**, **Lean**, and **HOL Light** are used to formalize and verify mathematical proofs.
- AI-enhanced ATPs can now search through vast proof spaces efficiently, combining brute-force logic with heuristic guidance.

#### 2. AI-Assisted Discovery

Recent advancements include:

- **DeepMind's AlphaZero for Mathematics**: Adapted reinforcement learning to guide formal proof strategies.
- **GPT-f (OpenAI)**: An adaptation of language models like GPT-3 fine-tuned on formal mathematics to assist in proving statements within the Metamath framework.
- **AI-generated conjectures**: In 2021, DeepMind published research in collaboration with mathematicians that used AI to discover new connections between algebraic invariants — a groundbreaking example of human–machine collaboration in pure mathematics.

#### 3. Applications to Number Theory and Primes

In the context of **prime numbers**, automated systems have begun to:

- **Assist in verifying large-scale primality proofs**, especially in cryptographic applications.
- **Suggest lemmas and intermediary results** that can help bridge gaps in complex conjectures.

- **Explore large numerical spaces** to test cases of unproven conjectures (like Goldbach or Twin Primes) far beyond what humans can calculate unaided.

#### **4. AI Models Applied to Prime Numbers**

The integration of AI models into mathematical research is no longer theoretical — it is an active and rapidly expanding field. Prime numbers, with their delicate balance of order and chaos, offer an ideal test bed for advanced AI models. In this section, we explore three major AI techniques being applied to prime-related problems: **neural networks**, **reinforcement learning**, and **symbolic regression**.

##### Neural Networks for Prime Prediction

Neural networks — especially deep learning architectures — are powerful tools for learning complex, non-linear patterns. Although primes may appear randomly distributed, subtle statistical patterns exist, and AI models can be trained to pick up on these.

#### **1.1 Basic Approach**

Researchers train **supervised learning models** where the input is a number (or its features, like its binary form, digit sum, etc.) and the output is a label: **prime (1)** or **not prime (0)**.

- **Feedforward neural networks** (FNNs) have been trained to classify numbers as prime/non-prime with surprising accuracy — even though the logic is not learned explicitly.
- **Convolutional neural networks** (CNNs), typically used in image processing, have also been adapted to detect structural features in large numerical datasets.

#### **1.2 Limitations and Insights**

While neural networks cannot "prove" primality (as that's a deterministic problem), they are useful in:

- Approximating the likelihood of primality
- Estimating the **distribution of primes** over intervals
- Detecting anomalies or clusters (e.g., identifying twin prime candidates)

This predictive modeling offers a new way to **visualize and understand prime behavior**, and may one day contribute to conjecture testing at massive scales.

## Reinforcement Learning in Exploring Number Theory

Reinforcement Learning (RL) is a branch of machine learning where agents learn by trial and error, guided by rewards. While RL has famously outperformed humans in games like Go and chess, it is now being explored in **mathematical exploration** and **discovery**.

### 2.1 Number-Theoretic Environments

In experimental systems, RL agents are placed in "mathematical environments" where they:

- Manipulate formulas
- Explore functions like  $\pi(x)$ , the prime counting function
- Test integer sequences for hidden structures

Each correct discovery, simplification, or advancement toward a known goal (e.g., closer approximation of a known prime distribution) is rewarded.

### 2.2 Practical Use Cases

- Agents can learn to **generate candidate expressions** related to prime gaps or density functions.
- RL can optimize searches through large computational spaces — for instance, **minimizing the number of operations in a factorization algorithm**.

This form of **exploratory learning** mirrors the way humans experiment in mathematics — and could lead to serendipitous discoveries.

## Symbolic Regression and AI-Generated Conjectures

While neural networks often act as “black boxes,” **symbolic regression** offers a more transparent approach. It seeks to find mathematical expressions or equations that best fit a given dataset.

### 3.1 What Is Symbolic Regression?

Symbolic regression uses techniques like **genetic programming** or **neuro-symbolic hybrids** to evolve human-readable mathematical formulas from data.

- Input: A dataset (e.g., prime positions or gaps)
- Output: An algebraic expression that approximates the pattern

Unlike deep learning, the goal here is **interpretability** — deriving formulas that can lead to **mathematical insight** or even **new conjectures**.

### 3.2 AI-Generated Conjectures

Notable breakthroughs include:

- **DeepMind's AI-assisted research** in knot theory and representation theory, which found relationships that led to new human-proven theorems.
- Experimental systems that generate conjectures on **prime density**, **modular forms**, or **prime-generating polynomials**.

For example, an AI might discover that a certain transformation consistently reduces prime gaps in specific intervals — not a proof, but a potential conjecture worth human investigation.

## 5. Case Studies & Breakthroughs

AI's involvement in mathematical research has shifted from experimental curiosity to real-world breakthroughs. In this section, we explore significant case studies that highlight how AI has contributed to **proof discovery**, **mathematical sequence generation**, and **analytical insights** into complex prime-related structures — especially prime gaps and the Riemann zeta function.

### 5.1 DeepMind's AlphaTensor and Mathematical Proof Discovery

In 2022, DeepMind introduced **AlphaTensor**, an AI system designed to discover efficient algorithms for matrix multiplication — a foundational operation in both mathematics and computer science. Although its direct focus wasn't number theory, its implications for **mathematical problem-solving** are profound.

Highlights:

- **AlphaTensor** used **reinforcement learning** to rediscover known algorithms (like Strassen's algorithm) and even discovered **new, more efficient ones** previously unknown to human researchers.
- Its **symbolic reasoning** combined with **optimization abilities** showed that AI can go beyond brute-force search — it can **innovate** within formal mathematical domains.

Impact on Prime Number Theory:

While AlphaTensor didn't directly tackle prime problems, the **underlying methodology** — using AI to explore large symbolic spaces efficiently — is already being adapted to areas such as:

- **Prime factorization**
- **Integer partitioning**
- **Proof-space navigation for conjectures like the Twin Prime Conjecture**

This demonstrated that **AI can contribute not just results, but methods** — inspiring entirely new ways to approach problems traditionally confined to pure mathematics.

## 5.2 Generative Models for Mathematical Sequence Generation

Recent developments in **generative AI**, especially transformer-based models like **GPT**, have been applied to mathematical sequences and logic.

Examples:

- Researchers have fine-tuned language models (like OpenAI's GPT or Google's PaLM) on datasets from **OEIS** (Online Encyclopedia of Integer Sequences) to generate **new, syntactically correct** and often **insightful** number sequences.
- These models can **predict continuations, detect recurrence relations**, and even propose **natural language explanations** of patterns.

Applications to Primes:

- Generating new **prime-adjacent sequences**, such as:
  - Sequences of probable twin primes
  - Approximate formulas for  $\pi(n)$ , the number of primes  $\leq n$
- Discovering novel **recursive rules or heuristics** for prime prediction
- Identifying **rare phenomena**, such as unusually large or small prime gaps

By treating mathematical sequences similarly to language, these models are able to surface **structural patterns** that may go unnoticed in traditional numerical analysis.

## 5.3 AI-Driven Analysis of Prime Gaps, Zeta Functions, and Beyond

One of the deepest mysteries in prime number theory lies in the distribution of primes — especially **prime gaps** and their connection to the **Riemann zeta function**.

Prime Gaps:

- AI has been employed to analyze massive datasets of primes to explore how **prime gaps behave across different scales**.
- Machine learning models can:
  - Detect trends in **gap distributions**
  - Predict the **probability of gap sizes** in certain intervals
  - Suggest asymptotic behaviors not previously modeled

This offers new data-driven insights into questions like:

- Do large prime gaps appear more frequently than expected?
- Is there a limit to how small prime gaps can get?

Zeta Functions and Zero Distribution:

- The **non-trivial zeros** of the Riemann zeta function are believed to hold the key to understanding prime distributions (per the Riemann Hypothesis).
- AI models have been used to:
  - Approximate the **density and distribution** of zeros in critical regions
  - Identify **symmetries or anomalies** in zero patterns
  - Explore analogs of zeta behavior in other L-functions

In one study, neural networks were trained on zero-distribution data to **predict zero densities** — offering new numerical methods for exploring this central function in analytic number theory.

## **6. Challenges and Philosophical Implications**

As AI continues to establish itself in mathematical research, particularly in number theory and prime exploration, it raises not only technical but **philosophical** and **epistemological** questions. Are machines truly understanding mathematics, or are they simply performing pattern recognition at scale? What does it mean to trust an AI-generated proof or conjecture? And how do we reconcile the **black-box nature of AI** with the traditional demand for clarity and rigor in mathematics?

This section reflects on these deeper issues that now confront both mathematicians and AI researchers.

### **6.1 Can AI "Understand" Mathematics?**

Understanding in mathematics involves more than just arriving at the correct answer — it includes intuition, conceptual frameworks, generalization, and explanation. While AI systems have demonstrated success in:

- Solving integrals
- Classifying theorems
- Generating new conjectures
- Assisting with formal proofs

...they do so without consciousness, intent, or semantic comprehension. They **simulate understanding** but may not possess it in any meaningful human sense.

AI models "understand" mathematics the way calculators "know" how to add.

This leads to an important distinction:

- **AI as a tool:** Powerful, scalable, and often surprising — but ultimately limited to what it is trained or programmed to do.
- **Humans as thinkers:** Capable of conceptual innovation, philosophical reasoning, and metamathematical reflection.

Yet, as AI becomes more capable of **generating novel ideas** (e.g., symbolic conjectures), the boundary between "tool" and "collaborator" is becoming increasingly blurred.

## **6.2 Verification vs Understanding: Machine Proof vs Human Insight**

A formal proof verified by an AI system (such as Lean or Coq) may be logically sound, yet **lack human readability or explanatory power**. This raises a tension between:

- **Formal verification:** Ensuring that each logical step follows from axioms and prior theorems (machine-friendly, but often opaque).
- **Conceptual understanding:** Making sense of why something is true and how it connects to broader theories (human-centric).

An AI might produce a correct proof that is hundreds of pages long — but can it explain it to a human in a single intuitive sentence?

In classical mathematics, elegance and explanatory value matter. A human proof that brings new insight is often valued more than a brute-force or overly technical one. So, even if AI proves a conjecture, **will mathematicians accept it without insight?** And if not, who or what will “translate” AI's output into meaningful human terms?

## **6.3 Trust, Explainability, and the Future of AI in Number Theory**

In mathematics — more than most fields — **trust** comes from **clarity, reproducibility, and logical consistency**. AI challenges this by introducing:

- **Black-box models:** Deep neural networks often cannot explain why they reach certain conclusions.
- **Probabilistic reasoning:** Models may give statistical likelihoods, not absolute truths.
- **Data biases:** An AI trained on incomplete or skewed datasets might form misleading conclusions.

To integrate AI meaningfully into number theory, researchers must address:

**a) Explainability**

Can we build AI systems that produce **transparent** reasoning chains, understandable by humans?

**b) Accountability**

If an AI-generated proof contains an error, **who is responsible?** The model? The programmer? The verifier?

**c) Collaboration**

Can we design AI to work with human mathematicians — not just solving problems, but dialoguing, suggesting, refining?

## **7. Future Directions**

As artificial intelligence continues to evolve, its role in mathematics — particularly in number theory and prime number research — is moving from auxiliary support to active collaboration. In this final section, we explore where the field is headed, including collaborative AI-human systems, the possibility of solving long-standing conjectures, and the emerging view of AI as a tool for **mathematical creativity**.

### **7.1 Collaborative Systems (AI + Human Mathematicians)**

One of the most promising directions for the future is the rise of **hybrid intelligence**: collaborative systems where **AI models assist human mathematicians** not just in computation, but in discovery, conjecturing, and even proof-writing.

#### **Emerging Possibilities:**

- **Interactive theorem assistants** (e.g., Lean, Coq, Isabelle) enhanced with AI models that can auto-suggest lemmas, check logical steps, or highlight gaps in reasoning.

- **Generative systems** that propose new conjectures based on observed data, which human mathematicians then refine, prove, or falsify.
- **Mathematical search engines** powered by large language models trained on formal mathematical corpora, able to retrieve relevant theorems or analogies instantly.

In this vision, AI doesn't replace mathematicians — it acts as a **thought partner**, enhancing creativity, accelerating exploration, and reducing the cognitive load of symbolic manipulation.

## 7.2 Potential to Resolve Longstanding Conjectures

Some of the most famous conjectures in mathematics — including those related to prime numbers — have remained unsolved for centuries. Could AI help resolve them?

### Examples:

- **The Riemann Hypothesis:** While AI won't "solve" it outright, it may help **map out zero distributions** of the zeta function far more thoroughly than human computation alone.
- **Twin Prime Conjecture:** AI could aid in discovering **new heuristics or probabilistic models** to narrow the problem space, possibly uncovering a path toward proof.
- **Goldbach's Conjecture:** By using large-scale testing and symbolic regression, AI might identify new structural features in even numbers and their relation to primes.

The path to proof may not be linear — but AI can act as a **torchbearer**, illuminating directions previously considered too vast or complex to explore manually.

With AI automating repetitive or computationally intensive steps, human mathematicians are freed to focus on **big-picture ideas** and **theoretical innovation**.

## 7.3 AI as a Tool for Mathematical Creativity

Perhaps the most surprising development is the realization that AI can be a **creative force** in mathematics.

### How?

- **Suggesting original patterns or transformations** not previously considered by mathematicians
- **Combining ideas across fields** (e.g., linking topology and number theory) through analogical reasoning
- **Generating new mathematical objects** or structures from scratch, which are then investigated for utility or insight

Unlike earlier views of AI as strictly logical or deterministic, current models — especially **generative models and symbolic systems** — show signs of **improvisation, association, and innovation**.

While these outputs still require human interpretation and validation, the ability of AI to **stimulate creative thinking** may become one of its most valuable contributions.

## **Conclusion**

The journey through the landscape of prime numbers reveals a profound intersection between ancient mathematical mysteries and cutting-edge artificial intelligence. From the early curiosity about primes' unique properties to the complex conjectures that have challenged generations, prime number theory remains one of mathematics' most captivating frontiers.

This chapter has highlighted the remarkable **progress made possible by AI**: from neural networks that detect subtle patterns, to reinforcement learning systems exploring uncharted mathematical spaces, to symbolic regression techniques generating novel conjectures. Case studies like DeepMind's AlphaTensor and generative models for sequence discovery demonstrate that AI is not just a computational tool but an active collaborator in advancing mathematical knowledge.

However, with this promise comes a great responsibility. AI's integration into foundational sciences requires careful attention to **explainability, verification, and ethical use**. Trusting AI-generated proofs or insights demands a new paradigm where human intuition and machine precision coexist harmoniously.

As we stand at this crossroads, the future of mathematics appears more vibrant and expansive than ever. The combined creativity of human mathematicians and AI systems holds the potential to solve longstanding problems and unlock deeper understanding — ushering in a new era where **knowledge is both discovered and co-created**.

The revolution in prime number discovery, powered by AI, is just beginning. Its full impact will unfold in the decades ahead, shaped by both technological innovation and human wisdom.

## References

1. Apostol, T. M. (1976). Introduction to analytic number theory. *Springer-Verlag. A foundational text on prime numbers and number theory basics.*
2. Bello, I. et al. (2019). AlphaTensor: An AI system for discovering Matrix multiplication algorithms. DeepMind blog. *DeepMind's announcement and details on AlphaTensor's mathematical contributions.*
3. Bombieri, E. (2000). The Riemann hypothesis—Official problem description. *Clay Mathematics Institute. Details on the Riemann hypothesis and its significance in number theory.*
4. Cantor, M. et al. (2020). Machine learning and prime number patterns: A review. *Journal of Computational Mathematics*, 38(3), 345–367. Overview of machine learning techniques applied to prime number research.
5. Chen, M. et al. (2021). Evaluating large language models trained on code. arXiv preprint arXiv:2107.03374. *Exploration of how models trained on programming and symbolic logic may help in mathematical reasoning.*
6. Conrey, J. B. (2003). The Riemann hypothesis. *Notices of the AMS*, 50(3), 341–353. A detailed exposition on the Riemann Hypothesis and its mathematical background.
7. Deep Mind Mathematics Team. (2021). Advances in AI for mathematical theorem proving. *Nature. Review of AI Systems Applied to Formal Mathematics and Proof Discovery.*
8. Gowers, T. (2002). Mathematics: A very short introduction. *Oxford University Press. Philosophical perspective on mathematical proof, insight, and creativity.*
9. Ho, D., & Cohn, J. (2021). Symbolic AI and neural networks: Bridging the divide. *Journal of Artificial Intelligence Research*, 71, 1023–1045. Survey on symbolic AI techniques combined with neural models in mathematics.
10. Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., Bridgland, A., Meyer, C., Kohl, S. A. A., Ballard, A. J., Cowie, A., Romera-Paredes, B., Nikolov, S., Jain, R., Adler, J., ... Hassabis, D. (2021). Highly accurate protein structure prediction with AlphaFold. *Nature*, 596(7873), 583–589. <https://doi.org/10.1038/s41586-021-03819-2>.

While about proteins, this shows AI's leap in solving complex scientific problems, analogous to prime number discoveries

11. *Key paper on symbolic regression and AI-generated scientific laws.*
12. Lakatos, I. (1976). Proofs and refutations: The logic of mathematical discovery. *Cambridge University Press. Classic on the philosophy of mathematical proof and discovery, useful to frame AI's role.*
13. Lenstra, A. K., & Verheul, E. R. (2001). The XTR public key system. In. Lecture Notes in Computer Science, 2248, *Prime numbers in cryptography, why primes matter in security and computation* (pp. 1–19).
14. Marcus, G. (2020). The next decade in AI: Four steps towards robust artificial intelligence. arXiv preprint arXiv:2002.06177. *Philosophical discussion on AI's role and challenges in complex tasks like mathematics.*
15. Marcus, G., & Davis, E. (2019). Rebooting AI: Building artificial intelligence we can trust. *Pantheon Books. Discussion on trust, understanding, and explainability in AI systems.*
16. Polymath Project. (2014). Variants of the twin prime conjecture. Research blogs and preprints. *Collective human plus computational efforts toward prime conjectures.*
17. Ramsey, F. P. (1926). The foundations of mathematics. *Proceedings of the London Mathematical Society*, s2–25(1), 338–384. <https://doi.org/10.1112/plms/s2-25.1.338>. Early foundational work highlighting the importance of logic in math, relevant to AI's logic-based approaches
18. Ribenboim, P. (1996). The new book of prime number records. *Springer-Verlag. Comprehensive coverage of prime number properties and famous conjectures.*
19. Saxton, D. et al. (2019). Analysing mathematical reasoning abilities of neural models. arXiv preprint arXiv:1904.01557. *Study on how neural networks perform on mathematical tasks, including sequence prediction.*
20. Schmidt, M., & Lipson, H. (2009). Distilling free-form natural laws from experimental data. *Science*, 324(5923), 81–85. <https://doi.org/10.1126/science.1165893>
21. Scholz, F., & Remus, D. (2020). Deep learning for number theory: Prime number prediction using neural networks. *International Journal of Computational Mathematics*, 97(5), 1004–1021. Direct application of deep learning to prime prediction tasks.

22. Selsam, D. et al. (2019). Learning a SAT solver from single-bit supervision. *Advances in Neural Information Processing Systems*, 32. Relevant work on AI learning symbolic reasoning and proof search.
23. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., & Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*, 550(7676), 354–359. <https://doi.org/10.1038/nature24270>. Foundational work on reinforcement learning, relevant to AI exploration methods
24. Tao, T. (2006). Structure and randomness: Pages from year one of a mathematical blog. *American Mathematical Society. Discusses patterns in primes and randomness, providing insights into modern prime conjectures.*
25. Wiedijk, F. (2007). Formal proof—Getting started. *Notices of the AMS*, 54(11), 1397–1407. Introduction to formal proofs and their increasing role alongside AI.