The Art of Mathematical Smuggling: Alpöge's Secret and the Virtues of a Stolen Proof

Bridging Automorphic Forms and Formal Verification

Peter De Ceuster

November 23, 2026

Abstract

The Riemann Hypothesis, a foundational problem in analytic number theory concerning the distribution of prime numbers, resisted traditional mathematical approaches for over a century and a half. This paper documents the recent unconditional resolution of the hypothesis, achieved by transitioning the problem domain from classical complex analysis to the geometry of algebraic automorphic forms. Central to this breakthrough was the formulation of a highly rigid, geometrically derived smoothing function, $G_\psi(x)$, constructed using automorphic periods on unitary groups. By leveraging Langlands reciprocity, the chaotic discrete distribution of primes was translated into continuous geometric frequencies known as Hecke eigenvalues. Packaging these frequencies into an L-function via the Euler product allowed the establishment of an unbreakable geometric ceiling of $x^{1/2+\eta}$, ultimately locking all non-trivial zeros of the Riemann zeta function onto the critical line $\mathrm{Re}(s) = 1/2$. Extending beyond the classical Riemann zeta function, the proof also encompasses generalized Dirichlet L-functions, establishing their zeros similarly. Crucially, the validity of this massive theoretical framework bypassed traditional peer review in favor of a paradigm-shifting formalization process. The 17-million-line proof was achieved through an unprecedented triad: human architectural genius provided the mathematical breakthroughs, artificial intelligence acted as a massive syntax translation engine, and the Lean 4 theorem prover served as the absolute mathematical referee. Anchored strictly to community-standard `Mathlib` definitions and cross-verified by independent tools like `Comparator` and `nanoda`, the codebase guarantees mathematical absolute truth. Finally, this paper explores the profound consequences of this proof, which instantly upgrades modern digital infrastructure by mathematically guaranteeing the efficiency of cryptographic algorithms (such as the Miller-Rabin test), validating optimal network topologies via Ramanujan graphs, and cementing the foundational models of quantum chaos. Furthermore, this breakthrough unconditionally resolves Artin's Conjecture, the Lindelöf Hypothesis, and the Class Number Problem.

1 Introduction

For over 160 years, the Riemann Hypothesis (RH) stood as the most prominent unsolved problem in mathematics. First proposed by Bernhard Riemann in 1859, the hypothesis asserts that all non-trivial zeros of the Riemann zeta function, $\zeta(s)$, lie precisely on the critical line where the real part is $1/2$. The location of these zeros dictates the distribution of prime numbers; a proof of RH ensures that primes are distributed as evenly as mathematically possible. Furthermore, the proof extends beyond the classical zeta function to encompass generalized Dirichlet L-functions ($L(\chi, s)$), proving that their zeros similarly obey the critical line boundary, definitively establishing the Generalized Riemann Hypothesis (GRH).

Because Alpöge bypassed the academic journal system entirely by releasing only a brief four-page announcement alongside his code, the mathematical community is currently in the unprecedented

position of reverse-engineering a codebase. Mathematicians at institutions like IHES and the University of Waterloo are currently dissecting the massive repository to translate the Lean logic back into classical analytic number theory.

Traditional attempts to prove RH relied on analytic number theory—studying the complex plane and manipulating $\zeta(s)$ directly. To measure the primes, mathematicians historically employed the von Mangoldt function, $\Lambda(n)$, often paired with arbitrary analytic smoothing functions (e.g., Gaussian curves or polynomial cutoffs) to ensure contour integrals converged. However, these arbitrary curves lacked deep structural properties, yielding only statistical approximations of primes rather than absolute limits. In traditional mathematics, these smoothing functions often look completely unnatural—they are essentially "Frankenstein equations" built specifically to behave well in calculus, acting merely as formatting tools to prevent integrals from exploding to infinity.

The recent breakthrough abandoned this purely analytic approach. Instead, it mapped the prime numbers into the continuous, multidimensional space of algebraic geometry, leveraging the Langlands Program to enforce geometric limits upon arithmetic chaos.

To understand the exact output of this formalization, one must look at the core files. The top-level file, `src/Main.lean`, contains the final theorems pointing to the proof's conclusion:

```
-- src/Main.lean (excerpt)
theorem riemannHypothesis : RiemannHypothesis := rhOfAll

theorem LFunction_ne_zero_of_half_lt_re : ...
-- (Proves L(χ, s) ≠ 0 for Re(s) > 1/2)

theorem LFunction_eq_zero_re : DirichletLZerosOnLine := claimAll
```

The underlying generalized statements are rigorously defined in `src/Statements.lean`:

```
-- src/Statements.lean (excerpt)
def DirichletLZerosOnLine : Prop :=
  ∀ (N : ℕ) [NeZero N] (χ : DirichletCharacter ℂ N) (s : ℂ),
  -- ... mathematical conditions defining the zeros ...
```

2 The Paradigm Shift: Langlands Reciprocity

When opening the thousands of intermediate files that define this function, the code maps out immensely complex structures, including Gan-Gross-Prasad periods, the Rallis seesaw, and relative trace formulas, heavily abstracted through Lean's tactic language. This geometric smoothing function was paired with the von Mangoldt function to create a new, analytically complete identity connecting the primes to the zeta function.

In Lean 4 syntax, this equivalence theorem—the physical bridge forcing Lean to accept that $G_\psi(x)$ controls the primes—is formalized as:

```
/-- The analytic identity connecting smoothed prime count to Zeta -/
theorem smoothed_prime_identity (x : ℝ) (hx : x > 1) :
  (∑ n in Finset.Icc 1 ⌊x⌋, Λ n * G_psi (n / x)) =
  (1 / (2 * Real.pi * Complex.I)) * ∫ s in CriticalStrip,
  (-(deriv ζ s) / ζ s) * (x ^ s) / s :=
by
  -- ... hundreds of lines of complex analysis and contour integration ...
```

The bridge was forged via the Explicit Formula, summarized conceptually as:

$$\sum_{n \leq x} \Lambda(n) G_\psi\left(\frac{n}{x}\right) = \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} \left(-\frac{\zeta'(s)}{\zeta(s)}\right) x^s \, ds \tag{1}$$

Because $G_\psi(x)$ is defined by the strict boundaries of a closed unitary group, it is physically constrained. Through the trace formula, it was proven that this geometric structure imposes a strict upper bound on its growth, capping it at an absolute ceiling of $x^{1/2+\eta}$.

To prove this limit without leaping to impossible conclusions, Alpöge used the foundational rules of inequalities within Lean's `calc` environment. It translates to massive blocks of chained equations, with each step justified by a specific mathematical rule:

```
calc
  AutomorphicPeriod_X ≤ Bounded_Value_Y := by apply Cauchy_Schwarz_inequality
  _ ≤ Smoothed_Limit_Z := by linarith [trace_formula_bounds]
  _ = x^(1/2 + η) := by ring
```

Every `apply`, `linarith` (linear arithmetic), and `ring` (algebraic ring solver) is a command telling Lean exactly which microscopic rule to use. Lean checks the rule; if it applies, it accepts the step.

The crucial next step was discovering the algebraic ceiling. Once the primes were translated into automorphic periods, Alpöge looked at the algebraic laws governing those periods and proved that they obey a strict, unbreakable limit. Because of this limit, his smoothed prime-counting function could never grow faster than $x^{1/2+\eta}$, establishing an impenetrable geometric ceiling over the primes.

2.1 The Euler Product and the Zero Trap

The geometric frequencies (a_p) were packaged into an L-function, which bridges the two domains. An L-function takes all of those geometric frequencies and organizes them into an infinite series:

$$L(s, f) = \sum_{n=1}^{\infty} \frac{a_n}{n^s} = \frac{a_1}{1^s} + \frac{a_2}{2^s} + \frac{a_3}{3^s} + \dots \tag{2}$$

However, because the Hecke operators perfectly respect the rules of prime factorization, this massive addition problem can be rewritten as a massive multiplication problem involving *only* the prime numbers, known as the Euler Product:

$$L(s, f) = \prod_{p \text{ prime}} \left(1 - a_p p^{-s} + p^{k-1-2s}\right)^{-1} \tag{3}$$

In this equation, f represents the specific geometric wave (the automorphic form), s is a complex number (the variable graphed to find the zeros), and a_p are the Hecke eigenvalues (the frequencies) extracted using the "tuning fork" Hecke operators. By proving that the eigenvalues a_p cannot exceed the geometric ceiling, that strict boundary was absorbed by the L-function. Because the identity perfectly mirrors the prime sequence on the left to the zeros of the zeta function on the right, the geometric ceiling crushed the complex plane. If a single zero existed with a real part greater than $1/2$, it would force the primes to breach the geometric limit. As this is algebraically impossible, all non-trivial zeros are permanently locked onto the critical line.

3 A Thief in the Night

The Zero Trap (Reversing the Logic)

Historically, mathematicians knew that *if* the Riemann Hypothesis was true (all zeros are on the $1/2$ line), *then* the primes would grow predictably. Alpöge reversed the trap. He stated: "I have just proven using algebraic geometry that the primes are trapped beneath the $x^{1/2+\eta}$ ceiling." Because the math is strictly connected, if even a single zero of the zeta function wandered off the $1/2$ line, it would mathematically force the prime count to break his ceiling. Since breaking the ceiling is algebraically impossible, every single non-trivial zero is permanently trapped exactly on the $1/2$ line.

4 The Unprecedented Formalization Process

The proof's scale and complexity defied traditional peer review. The full proof consists of 175,187 separate `.lean` files (about 1.03 GB of source code), forming a massive dependency tree of over 725,000 intermediate theorems. The foundational logic of this verification relies on the Curry-Howard correspondence, a principle in computer science where a mathematical theorem is treated as a software "type," and a proof is a program that successfully evaluates to that type. The logical links connecting the Langlands program, algebraic limits, and complex analysis required millions of microscopic steps. To guarantee its validity, the proof was fully formalized in Lean 4.

The workflow operated as a triad:

1. **The Architect (Human):** The mathematical intuition—mapping the primes to unitary groups and identifying the $x^{1/2+\eta}$ ceiling—was entirely human.
2. **The Muscle (Artificial Intelligence):** Large language models (specifically Claude) were utilized as translation engines. The AI processed high-level geometric arguments and ground them down into the thousands of lines of pedantic tactic scripts required by Lean.
3. **The Inspector (Lean Compiler):** The Lean 4 kernel served as the absolute mathematical referee. It evaluated every logical inference without assuming any prior unproven theorems.

This triad operates much like building a skyscraper. The human author is the architect who designed the structure and placed every single steel beam. The AI is the heavy machinery doing the lifting. The Lean compiler is the physics inspector who tests every single bolt to mathematically guarantee the building will never collapse. Without realizing it, Alpöge became the first mathematician in history to make a significant contribution by using AI as an extension of his intellect.

Within the thousands of intermediate files, Alpöge avoided leaping to impossible conclusions by relying on Lean's `calc` environment. This allowed him to chain hundreds of equations and inequalities, justifying each microscopic step with foundational commands like `apply`, `linarith` (linear arithmetic), and `ring` (algebraic ring solver).

To ensure no circular logic or altered definitions were smuggled into the proof, the project utilized a "goalpost" file named `Challenge.lean`. This file contained the exact theorems using Lean's `sorry` command—a macro that tells the compiler to temporarily skip verification and "just trust" the statement:

```
-- Challenge.lean
theorem riemannHypothesis : RiemannHypothesis := sorry
theorem LFunction_ne_zero_of_half_lt_re : ... := sorry
```

`Challenge.lean` acted as a cryptographic checksum for the theorems. The payload file, `Main.lean`, then replaced these `sorry` placeholders with the actual proof paths while explicitly invoking the community-standard `Mathlib` definition of the Riemann Hypothesis:

```
-- Main.lean (The actual proof)
import Mathlib.NumberTheory.RiemannHypothesis

theorem riemannHypothesis : RiemannHypothesis := rhOfAll
```

In this syntax, the colon (`:`) means "is of type" or "makes the claim," pointing to the unalterable `Mathlib` goal. The assignment operator (`:=`) means "is proven by," pointing to `rhOfAll`, the culmination of the 17 million lines of code. If `rhOfAll` proves something even slightly different, Lean throws a type-mismatch error.

A checksum script (`Comparator`) verified that the final binary relied only on standard axioms, specifically restricting dependencies to the three fundamental rules of standard mathematics (`propext`, `Quot.sound`, and `Classical.choice`), and explicitly checking that no `sorry` commands remained in the final compiled code. To eliminate the risk of a bug in Lean itself, the compiled proof was exported into a standardized low-level format and fed into `nanoda`—a completely separate, third-party proof checker. This multi-layered verification resulted in a 17-million-line codebase of indisputable absolute truth.

5 Consequences for Mathematics and Society

The unconditional proof of the Riemann Hypothesis fundamentally upgrades the mathematical infrastructure of modern computer science and physics:

- **Prime Distribution and Cryptography:** The maximum gap between a prime p and the subsequent prime is now strictly bounded by $\mathcal{O}(\sqrt{p}\log p)$. Algorithms relying on primes, such as the deterministic Miller-Rabin primality test used in generating RSA encryption keys, are mathematically guaranteed to operate in polynomial time without failure.
- **Network Optimization:** The automorphic forms used in the proof are the mathematical basis for Ramanujan graphs, which govern the most efficient possible topologies for global telecommunication and internet routing.
- **Quantum Mechanics:** The proven spacing of the zeta function's zeros formally validates models in Random Matrix Theory, cementing the formulas physicists use to predict atomic energy levels and quantum chaos.
- **Artin's Conjecture on Primitive Roots:** First proposed in 1927, this conjecture regarding how often a given number acts as a primitive root in modular clock arithmetic was known to be true only if the Generalized Riemann Hypothesis held. It is now an unconditional mathematical fact.
- **The Lindelöf Hypothesis:** This major open question about the vertical growth rate of the zeta function's waves along the $1/2$ critical line is automatically proven. This provides a strict upper bound on these waves, tightly controlling complex number behavior in high-energy physics.
- **The Class Number Problem:** Dating back to Gauss, this problem concerns the factorization patterns of imaginary quadratic fields. The proof provides "effective bounds," ensuring that algorithms calculating these algebraic structures no longer have to account for infinite exceptions—the boundaries of the math are now closed and fully mapped.

- **Unlocking "Stuck" Algorithms:** Beyond cryptography, an entire class of advanced computer algorithms used for data compression and error-correcting codes (essential for keeping satellite transmissions and hard drives from corrupting) operate under the assumption that RH is true. Software engineers can now confidently deploy these highly efficient algorithms across critical infrastructure without risk of failure.

6 Conclusion

The resolution of the Riemann Hypothesis marks the end of a 160-year search and the beginning of a new era in mathematics. By smuggling the absolute limits of algebraic geometry into the heart of analytic number theory, the chaotic primes were finally tamed. Furthermore, the construction of this proof demonstrates that the future of profound mathematical discovery lies in the synthesis of human effort, AI-driven formalization, and strict computational verification.