RIEMANN-REV · PUBLIC NOTE · 26 NOVEMBER 2026

The condensed Lean proof of RH and GRH: what happened, how it goes, what is checked

A reading of `lsernau/grh-lean` *at commit* `72fead1` *(25 November 2026), a shorter rewriting of L. Alpöge's announced proof. Part of riemann-rev; see also the proof atlas.*

Status in one paragraph. The repository contains a Lean 4 proof, written by Claude (AI) after Alpöge's proof, of the Riemann hypothesis and of GRH for all Dirichlet L-functions, stated in Mathlib's terms. Its authors report that Lean's kernel accepts it with the three standard axioms only, that three independent proof checkers accept an export of it, and that the tool Comparator accepts it against a statement-only challenge file; the community repository that adopted it builds every commit with checks for `sorry` and extra axioms, and the runs that have finished have passed. We found no `sorry`, no axiom and no construct that bypasses the kernel in its source. **No mathematician has read the proof through**, and no paper accompanies it; the authors say so themselves. Our own replay of the build is under way. This note is a map of the argument as the code states it, not a verification of it.

1. What happened

20 Nov 2026	Levent Alpöge announces proofs of RH and of GRH for Dirichlet characters, written with the AI system Claude: a four-page note and a Lean archive of 175,187 files and 17.4 million lines. Independent kernel replays of that archive are recorded on the riemann-rev dashboard.
20 Nov	The community repository `leanprover-community/riemann-hypothesis` opens and starts to clean up the archive.
25 Nov, 20:38 UTC	`lsernau/grh-lean` appears: "a shorter version, produced in collaboration with my colleagues at Anthropic, and Claude" (Luke Sernau). Two editions of the same five theorems: a readable one of 477,124 lines and a compact one of 112,827 lines. "No mathematician has read it through."
25 Nov, 22:24 UTC	The community repository moves its main branch to the readable edition and from then on builds every commit, failing on `sorry` or on any axiom beyond the standard three.
26 Nov	This note. riemann-rev starts a complete human-readable write-up of the compact edition.

2. How large the simplification is

edition	Lean files	lines	remark
Alpöge's submission, 20 Nov (`40e56bab`)	175,187	17.4 million	the announced artifact
readable edition `grh/`, 25 Nov	803	477,124	proof part `New`: 192,000 lines
compact edition `grh-compact/`, 25 Nov	794	112,827	**proof part `New`: 24,850 lines of code, 1,571 theorems**

The two editions of `grh-lean` have byte-identical statements and build settings (Lean v4.35.0, Mathlib `56a00585`); some steps use different arguments. Each is split into `Lib` (background results "classed as known in the literature" that Mathlib lacks, with their proofs; it never imports `New`) and `New` (the argument). Every number in the table was counted on the pinned commits.

What changed in five days. The formal text shrank from 17.4 million lines to 112,827, a factor of about 150. The part that is specific to this proof is 24,850 lines of code in 219 files; the background that Mathlib lacks (51,280 lines) is separated from it and labelled with references to the literature; the five final statements did not change and are checked by machine in each edition. Both the

announced proof and its condensation were written by an AI system under human direction, and the condensation, a rewriting of one machine-written proof into another a hundred and fifty times smaller, took days. The proof is now small enough for a team of experts to read. Until someone does, it is a claim that machines have checked, not a theorem that mathematicians have accepted.

3. The theorems and the trust base

The file `src/Statements.lean` (16 lines, imports only Mathlib) defines four statements, and `src/Main.lean` proves five theorems without hypotheses:

- `riemannHypothesis`: Mathlib's `RiemannHypothesis` — every zero $s \neq 1$ of ζ other than $-2, -4, \ldots$ has $\operatorname{Re} s = \frac{1}{2}$; and the same statement spelled out;
- for every modulus $N \geq 1$ and every Dirichlet character χ mod N, primitive or not: $L(s, \chi) \neq 0$ for $\operatorname{Re} s > \frac{1}{2}$ (except $s = 1$ for the principal character); every zero with $0 < \operatorname{Re} s < 1$ has $\operatorname{Re} s = \frac{1}{2}$; and the case of primitive χ.

To trust the result one reads these 16 lines and the five statements, and trusts Lean's kernel and Mathlib's definitions of ζ, of $L(s, \chi)$ and of `RiemannHypothesis`. Everything else is checked by the kernel; the tool Comparator checks the five theorems against `Challenge.lean`, which states them with `sorry`.

4. The strategy in one picture

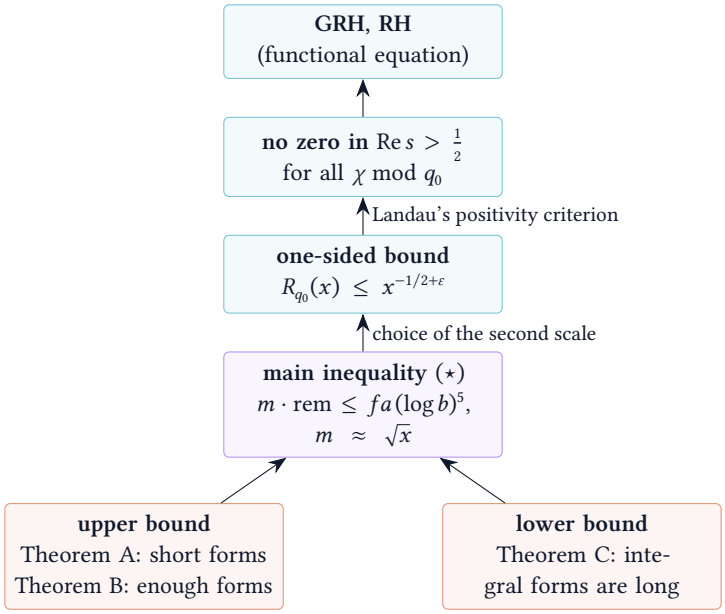

the main amount h_{main} appears in both bounds and cancels

5. Step 1: from a one-sided bound for primes to a zero-free half-plane

Fix q_0 divisible by $18480 = 16 \cdot 3 \cdot 5 \cdot 7 \cdot 11$ and a prime $d \equiv q_0/2 - 1 \pmod{q_0}$ (so $d \equiv 7 \pmod 8$); let $K = \mathbb{Q}(\sqrt{-d})$ with quadratic character χ_K. The weights $\lambda(n) = \Lambda(n)\,(1 - \chi_K(n)) \geq 0$ for $n \equiv 1 \pmod{q_0}$ (and 0 otherwise) count the primes that do not split in K. The remainder is

$$R_{q_0}(x) = \sum_{n \leq x} \lambda(n)\Big(\frac{1}{n} - \frac{1}{x}\Big) - \frac{\log x}{\varphi(q_0)} - c_{q_0},$$

with an explicit constant c_{q_0} built from logarithmic derivatives of L-functions at 1. **Criterion** (`entireL_ne_zero_of_primeError_le`): if $\beta > \frac{1}{2}$ and, for the primes d of a unit class forcing $d \equiv 7$

(mod 8), $R_{q_0}(x) \le x^{\beta-1}$ for all large x, then no $L(s,\chi)$ with $\chi \bmod q_0$ vanishes in $\operatorname{Re} s > \beta$. Only an *upper* bound is needed, because the weights are non-negative: this is Landau's theorem on Dirichlet integrals with a positive integrand, in the form of a one-sided criterion (Landau 1905; Montgomery–Vaughan §15.1; M. Suzuki 2024); real zeros are avoided by the choice of d. Every modulus divides some admissible q_0, so all Dirichlet characters are covered.

6. Step 2: the main inequality and why $\sqrt{x}$

At a large scale x the proof builds *data*: further scales x_1 and b with $(\log b)^{100} \le x$, $x^{10^6} \le x_1$, $x_1^{10^6} \le b$; a number $m \in [\sqrt{x}, 2\sqrt{x}]$ of theta lifts; ranks $r \in [x^5, 3x^5]$ and $k = r + 1 - m$; a CM field $E = K(\zeta)$ with about x^{11} distinguished infinite places ("tops"); a derivative order $a \approx xb$ and a prime level N. The main inequality (`star_inequality`) reads

$$m \cdot \mathrm{rem} \;\le\; fa(\log b)^5,$$

where rem is the prime sum at the tops less its expected value and f the number of infinite places. In the authors' words, "the saving $x^{-1/2}$ is $1/m$": normalised by the number of places and the order a, the remainder is at most about $(\log b)^6/m$, with $\log b$ about x^{ϵ}. A good choice of the second scale x_1 in a long window (`IsPolar.exists_window_sub_mul_le`) then gives $R_{q_0}(x) \le x^{8\varepsilon-1/2}$ (`eventually_primeError_le_rpow`), and Step 1 gives GRH (`New.grh`).

7. Step 3: one covolume bounded from both sides

The objects are automorphic forms on unitary groups over E. *Columns* are algebraic modular forms on a definite $U(k)$; *rows* (the target forms) are cusp forms on an indefinite $U(r)$. Poincaré series map columns to rows, and a chain of m theta lifts $Y_0 \to Y_1 \to \dots \to Y_m$ with ranks growing by $t+1$, $t = 5r$, is twisted by an anticyclotomic character ϑ of $U(1)$. For a family of target forms one measures the logarithm of the *covolume* of its lifts (a Gram determinant, at every complex conjugate) and its *index* (how far the lifts are from integral, prime by prime).

- **Theorem A (short forms)**, `Data.short_forms`: after a rational normalisation, the projected Poincaré series of integral ramified columns have lifts whose squared norms are at most $\exp\big(h_{\mathrm{main}} - m \cdot \mathrm{rem} + \mathrm{error}\big)$ at every conjugate, and the index of what they span is small.
- **Theorem B (enough forms)**, `Data.enough_forms`: few translates of these Poincaré series span all but a share $4/b^2$ of the target forms — a second-moment ("filling") argument with Bessel functions at the real places.
- **Theorem C (integral forms are long)**, `Data.integral_forms_long`: for any such family, $\mathrm{covolume} + 2 \cdot \mathrm{index} \;\ge\; \dim \cdot \big(h_{\mathrm{main}} - \mathrm{error}\big)$. The mechanism is the product formula for heights (the height of the constant family 1 is ≥ 0): sizes are read on the untwisted chain, valuations on twisted ones, with one twist ϑ_l for each prime ℓ and each $\iota : \mathbb{C} \cong \overline{\mathbb{Q}}_\ell$, chosen with congruences, Galois representations and Selmer groups (`exists_isFitFamily`); the volume of the chain factors into the volumes of single lifts, each of which gains a known amount.

A and B give the upper bound $h_{\mathrm{main}} - m \cdot \mathrm{rem} + \mathrm{error}$ for the normalised covolume, C gives the lower bound $h_{\mathrm{main}} - \mathrm{error}$; the main amount cancels and $(\star)$ remains, with total error $fa(\log b)^5$.

8. What is quoted from the literature

`Lib` (573 files, 51,280 lines of code in the compact edition) proves, in the forms needed and often weaker than the general theorems, parts of: class field theory and Galois cohomology; ℓ-adic and Hodge cohomology; Galois representations attached to eigenforms (matching Hecke polynomials for most Frobenii); p-adic comparison, p-adic Hodge theory and Kisin modules; Shimura varieties; the Siegel–Weil and Rallis inner product formulas; a simple trace formula; a Chowla–Selberg formula; a bound for the odd part of twisted Tate–Shafarevich groups of abelian CM fields (after Mazur–Wiles and Greither, proved without the main conjecture, `exists_length_shaOdd_le`); primes in progressions; and a zero-density count after Pólya–Vinogradov for the choice of d. These are proved inside the archive, not assumed. Its `REFERENCES.md` lists 196 citations, of which 57 were compared with their sources by the authors.

9. What is checked, and what is not

checked by machine (reported)	build of every file without `sorry`; `#print axioms`: `propext`, `Classical.choice`, `Quot.sound`; export accepted by nanoda, con-leche, con-ron; Comparator against `Challenge.lean`
checked by others	community CI builds every commit of the readable edition, failing on `sorry` or extra axioms; the runs that have finished have passed
checked by us	source scan of both editions: no `sorry`, `axiom`, `opaque`, `unsafe`, `native_decide`, custom syntax or compile-time code; statements read; independent replay on a cluster under way
not checked by anyone	a mathematician's reading of the argument; whether the method is, in substance, as new as it looks; most of the cited places in the literature

10. What we are doing next

A complete write-up of the compact edition's `New` (25,000 lines), with every classical input quoted by precise reference: structured proofs prepared unit by unit by a second AI agent, written up as mathematics by a first one and checked against the Lean statements; first release in about ten days, at `github.com/nasqret/riemann-rev`. The same project keeps a dashboard of every recorded check, a guide to the submission and four notes on parts of it; the proof atlas places the declarations of the archive on a map of the arXiv literature. The submission's route through the theorem of Ferrero and Washington (a "cohomological window" along cyclotomic towers) is not taken: neither edition cites that theorem. The corresponding cohomological input is the bound for odd Tate–Shafarevich groups of §8, which enters through Ribet's method in the choice of the twists.

Sources: `github.com/lsernau/grh-lean` at `72fead102f66` (README, `GUIDE.md`, `INDEX.md` of `grh/`, header of `grh-compact/src/New/Main.lean`, the statements quoted); L. Alpöge, `alpo.ge/remark.pdf`; `github.com/leanprover-community/riemann-hypothesis`. Line counts are physical lines; "lines of code" omit comments and blank lines. The announcement and the formal proof are claims checked by machine, not by human readers; nothing here asserts more.