A Prime in Every Gap: Shaving 22/25 Down to 0.8512

2026·11·23 12 min read mathematics · primes · riemann-hypothesis · lean

ABSTRACT

Assuming the Riemann Hypothesis, every interval $[x, x + 0.8512 \cdot \sqrt{x} \cdot \log x]$ with $x \geq 4$ contains a prime. The previous constant was 22/25. How I got there with Fourier optimization, certified interval arithmetic and a Lean 4 formalization, and what is still unchecked.

Three Days After

On 20 November 2026 a proof of the Riemann Hypothesis was announced, by L. Alpöge together with Claude. It is written in Lean on top of Mathlib (about seventeen million lines, according to the announcement). There is no human-readable version yet, and the mathematical community has not formally accepted it.

So I asked a simple question: **if RH is true, what can we actually do with it today?**

RH says something about the zeros of the zeta function. What I care about is the primes, and the zeros control how evenly the primes are spread. One of the oldest consequences is about gaps. About a hundred years ago Harald Cramér showed that, under RH, the gap after a prime p is at most a constant times $\sqrt{p} \log p$. The question that has been open since then is: **which constant?**

This post is about making that constant smaller.

> *Everything below is conditional.* *Every theorem in this post starts with "assuming RH". RH is used as a black box, meaning only the statement that all non-trivial zeros lie on* $\mathrm{Re}\, s = \frac{1}{2}$. *Nothing here depends on the contents of the announced proof. If that proof falls apart, the results stay true as conditional statements, the same way they would have been in 2019.*

The State of the Art

There are two kinds of results, and the difference matters.

An **asymptotic** result says that for *large enough* x there is a prime in $[x,\ x + c\sqrt{x}\log x]$. "Large enough" might mean beyond 10^{1000}. It tells you nothing about any number you could actually write down.

An **explicit** result says this holds *for every* $x \geq 4$, with no exceptions and no hidden threshold. Explicit results are much harder. Every error term that the asymptotic proof can ignore has to be bounded by an actual number.

The best results I could find:

result	constant c	source
explicit	$4/\pi \approx 1.27$	Dudek, 2015 (prime in $(x - c\sqrt{x}\log x,\ x]$, $x \geq 2$)
explicit	$1 + 4/\log x$	Dudek, Grenié, Molteni, 2016
explicit	$22/25 = 0.88$	Carneiro, Milinovich, Soundararajan, 2019 (CMS)
asymptotic	$21/25 = 0.84$	CMS, 2019
asymptotic	0.8358	Chirre, Pereira Júnior, de Laat, 2021 (CPdL)

The Dudek-Grenié-Molteni result is stated for the symmetric interval $(x - c\sqrt{x}\log x,\ x + c\sqrt{x}\log x)$ with $c = \frac{1}{2} + 2/\log x$, so its total length is $(1 + 4/\log x)\sqrt{x}\log x$. They also note that $c = 0.55$ works for all $x \geq 5$, which is a length of $1.1\sqrt{x}\log x$. All lengths in the table are comparable in that sense.

The CMS paper, *Fourier optimization and prime gaps*, is beautiful. It turns the prime gap question into a pure optimization problem over functions. Pick an even, integrable function F and look at its Fourier transform $\hat{F}$. Then this number:

$$J(F) = \frac{F(0) - A\int_{|t|>1} \max(\hat{F}(t),\ 0)\,dt}{\|F\|_1}$$

gives a prime gap bound of $1/J(F)$. A bigger J means a smaller constant. You want F to be tall at the origin and light overall, and you are penalised whenever its Fourier transform goes positive outside $[-1, 1]$. The weight A comes from a sieve bound on how many primes can fit in a short interval.

The gap between "explicit" and "asymptotic" in that table has a simple cause. For the explicit 22/25, CMS did not use their best function. They used a simple, hand-picked one ($\cos 2\pi x/(1 - 16x^2)$), because every norm and every error term has to be bounded by hand, and a simple function makes that bearable.

That is where I went in.

• A Better Test Function

CPdL built their function as a polynomial times a single Gaussian, $p(x^2)\,e^{-\pi x^2}$, and optimised it with semidefinite programming, the same machinery Cohn and Elkies used for sphere packing. Instead I took a **sum of Gaussians of different widths**:

$$F(y) = \sum_j c_j\, e^{-\pi b_j y^2}, \qquad \hat{F}(t) = \sum_j \frac{c_j}{\sqrt{b_j}}\, e^{-\pi t^2/b_j}$$

The Fourier transform of a Gaussian is a Gaussian, so $\hat{F}$ is explicit and every integral is a combination of error functions. If you fix the widths b_j (geometrically spaced), the problem is convex in the coefficients c_j, and after discretisation it is a linear program. Thirty Gaussians, widths from 0.15 to 20, solved with HiGHS.

One trap along the way: with too wide a range of widths and truncated integrals, the LP "cheated" and reported values above 2, which the theory says is impossible. The fix was choosing truncation points with tails below e^{-45}. After that, **every** candidate was checked three independent ways:

A	LP	mpmath (closed form)	numpy (quadrature)	Arb (rigorous lower bound)
36/11	1.199002	1.199005647	1.199005645	≥ 1.19900564740
4	1.174857	1.174877488	1.174877482	≥ 1.17487748806

The last column is the one that counts. It comes from Arb[1], a library for interval arithmetic: every number is a ball guaranteed to contain the true value. The coefficients are stored as exact dyadic rationals. Integrals are in closed form, sign changes of F and $\hat{F}$ are isolated rigorously, and the tails are bounded with erfc. The output is a *proof* that the number is at least what it says.

Two things follow directly:

1. The asymptotic constant. CMS state their asymptotic theorem in terms of the supremum of J over *all* admissible functions, at $A = 36/11$ (Iwaniec's sieve constant). Any single function gives a lower bound on that supremum. So with no new analysis:

$$\limsup_{n \to \infty} \frac{p_{n+1} - p_n}{\sqrt{p_n} \log p_n} \leq 0.834025 \qquad \text{(assuming RH; was 0.8358)}$$

2. A by-product. At $A = 4$ we get 1.17488, while CPdL report 1.17233. That improves their constant for primes in arithmetic progressions (under GRH) from about $0.8531\,\varphi(q)$ to $0.8512\,\varphi(q)$.

But the main target was the explicit result.

• Making It Explicit

The explicit version uses $A = 4$. That comes from the Montgomery-Vaughan inequality $\pi(u + y) - \pi(u) < 2y/\log y$, a fully explicit sieve bound. My function gives $J \geq 1.1749$ there, so in principle $1/1.1749 \approx 0.8512$ instead of 0.88.

"In principle" took most of the work. The proof is by contradiction. Suppose there is no prime in $[x,\ x + c\sqrt{x}\log x]$. Plug F into the Guinand-Weil explicit formula, which links a sum over zeta zeros to a sum over primes. RH says the zero side is under control, and the missing primes make the prime side too small. For large x the two sides cannot be balanced.

The CMS skeleton survives, but replacing their F with thirty Gaussians means re-deriving almost every estimate for a general F:

- the sum over zeros, using Trudgian's explicit bound for $N(T)$, the number of zeros up to height T;
- the Gamma-factor term, which CMS could treat as $O(1)$ and I could not;
- the prime sum near the gap, split by a chain of short intervals $(u,\ u + \sqrt{u}]$ with Montgomery-Vaughan on each link. Near the gap this costs the factor $A = 4$, and far away a total-variation bound on $\hat{F}$;
- prime powers, and super-exponentially small tails (around $10^{-5.9 \cdot 10^{14}}$, but they still have to be bounded);
- every $x \geq 4 \cdot 10^{18}$, not just one point. The small error terms depend on x, so I cover the range in balls up to $\log x = 1000$ and use monotone crude bounds beyond that.

Below $4 \cdot 10^{18}$ there is a separate argument. Oliveira e Silva, Herzog and Pardi verified by computer in 2014 that all prime gaps up to that point are smaller than $\log^2 p$. That is far below $c\sqrt{p}\log p$, apart from a few tiny cases handled by hand.

Every one of these steps is written out in full in proof.org[2]. The final numbers, all from Arb at 256 bits:

quantity	bound
J (worst case over all $x \geq 4 \cdot 10^{18}$)	≥ 1.17487747839
sum of small error terms (max)	≤ 0.0027020
final condition $K/\lVert F \rVert_1$	≤ 1.30487 (need < 1.99883)
$c = 1/J$, **rounded up**	**0.8511525827**

> *Theorem (assuming RH).* *For every $x \geq 4$ the interval $[x,\ x + 0.8511525827\sqrt{x}\log x]$ contains a prime.*

The previous explicit constant was 0.88.

• The Review

I don't know any number theorists I could just send this to. So the first referee was a Claude session, given the write-up and told to find mistakes. It did:

- one inequality in the prime-power estimate was simply **false**. The correct version changes the result by $2 \cdot 10^{-5}$ against a safety margin of about 0.69;

- the argument that the worst case over x is covered by the finite computation had a gap. I closed it by switching to the "balls plus crude tail bounds" scheme described above;
- several cosmetic issues.

No critical error. The constant survived. Yet?

• Lean

If the RH proof itself is a Lean artefact, why should a corollary of it be checked any less strictly?

So the proof went into Lean 4 with Mathlib. The main theorem reads:

```
theorem RhGaps.main (hRH : RiemannHypothesis) (x : ℝ) (hx : 4 ≤ x) :
∃ p : ℕ, p.Prime ∧ x ≤ p ∧ (p : ℝ) ≤ x + cstar * √x * log x
-- cstar = 8511525827 / 10^10
```

The project is in research/rh-prime-gaps/RhGaps[3]. `RiemannHypothesis` is Mathlib's own definition, not mine. RH is a *hypothesis* of the theorem, not an axiom. The project is 1345 lines in eight files. It compiles with no `sorry`, and `#print axioms` shows exactly what is assumed:

axiom	what it says	where it comes from
`oliveira_e_silva`	prime gaps below $4 \cdot 10^{18}$ are $< \log^2 p$	computer search, 2014
`montgomery_vaughan`	$\pi(u+y) - \pi(u) < 2y/\log y$	Montgomery-Vaughan, 1973
`explicit_formula_core`	the explicit formula, the zero sum, and the small error terms	Guinand-Weil, Trudgian, Arb

axiom	what it says	where it comes from
`cert_L1` … `cert_TV`	five numbers: norms and integrals of this specific F	Arb certificates

Everything else is proved. That includes the whole small-x argument, the parameter algebra, the lower bound for the left-hand side using the exact thirty rational coefficients, the full chain argument over primes on both sides of the gap (with the Montgomery-Vaughan and total-variation bookkeeping), and the final assembly, where `nlinarith` checks the inequalities on exact rationals.

The core axiom is still large. Yeah. Mathlib has no explicit formula and no zero-counting function $N(T)$, and building them is a project of its own. But the axiom does not hide anything vague. Its prime sums are concrete sums over `Nat.Prime` of a concrete function built from concrete rationals. The five certificates are single numbers that any interval-arithmetic package can recheck.

The best part: **Lean found a bug that a human reader would very likely have missed.** The constant was originally 0.8511525826, which is 1/J rounded to ten digits.

Lean refused to accept $c \cdot J \geq 1$, because rounding to nearest had gone *down*, and $c \cdot J$ fell short of 1 by $3 \cdot 10^{-11}$. The theorem now uses 0.8511525827, rounded up. Lean also flagged an inequality in my first draft ($2\pi\Delta \geq L/(4.004\sqrt{x})$) that was false at the edge of the range. The correct version, with 4.01, does not change the result, but the draft was wrong.

• What This Is and What It Isn't

To be precise about the claims:

- **Proved, assuming RH:** every $[x,\ x + 0.8511525827\sqrt{x}\log x]$, $x \geq 4$, contains a prime. The previous explicit constant was $22/25 = 0.88$.
- **Proved, assuming RH:** asymptotically, the constant 0.834025, improving 0.8358 (CPdL). This one is a direct consequence of the CMS theorem plus a certified function.
- **Not peer reviewed** by a human expert. One AI review, one Lean formalisation with the axioms listed above.
- **Literature search was automated.** As of today I found nothing better than 22/25 explicitly or 0.8358 asymptotically. Quesada-Herrera (2024) and Visser (2026) still cite 22/25 as the current explicit constant. That does not prove there is nothing newer. At least I'm honest about it right?
- **How far this method can go:** CMS show it cannot go below about 0.75 at $A = 36/11$. With $A = 4$ the bar is even higher. Getting much further needs a different idea, not a better function. But I cannot think of another one.

None of this changes the world unfortunately. It moves a constant from `0.88` to `0.8512`. But it is a real, checkable improvement, and three days ago it was not possible to state it with "assuming RH" as anything more than a hypothetical.

If you work in analytic number theory and this caught your eye: please try to break it. Everything is in research/rh-prime-gaps[4] in this site's repository: the full proof with every estimate and how to rerun it (proof.org[2]), the LP search and Arb certificates (code/[5]), the exact coefficients (results/[6]), and the Lean project (RhGaps/[3]). Write to me.

Credits

The mathematics, code and Lean proofs in this project were written by me with Claude, working with over a long sessions. The approach rests entirely on Carneiro, Milinovich and Soundararajan's *Fourier optimization and*

prime gaps (Comment. Math. Helv. 94, 2019) and on Chirre, Pereira Júnior and de Laat's *Primes in arithmetic progressions and semidefinite programming* (2021).

colophon

words	2323
written in	Org mode, GNU Emacs
set in	Literata & Typus Mono 95
mended	1× · last 2026·11·23
source	a-prime-in-every-gap.org ↓ [7] view
pdf	a-prime-in-every-gap.pdf ↓ [8]



Links

1. arblib.org
2. github.com/szymonwilczek/szymon-wilczek.me/blob/main/research/rh-prime-gaps/proof.org
3. github.com/szymonwilczek/szymon-wilczek.me/tree/main/research/rh-prime-gaps/RhGaps
4. github.com/szymonwilczek/szymon-wilczek.me/tree/main/research/rh-prime-gaps
5. github.com/szymonwilczek/szymon-wilczek.me/tree/main/research/rh-prime-gaps/code
6. github.com/szymonwilczek/szymon-wilczek.me/tree/main/research/rh-prime-gaps/results
7. 127.0.0.1:41381/writings/a-prime-in-every-gap.org
8. 127.0.0.1:41381/writings/a-prime-in-every-gap.pdf