The Riemann Hypothesis is the classic problem in number theory. It asks for optimal control on the count of prime numbers of a given size. Its shadows have guided mathematicians for 150 years, and its many rhymes and reflections, even reaching physics, taught us the unity of our fields.

I believe I have, working with Claude, discovered a proof[1] of the Riemann Hypothesis. In fact, the method easily extends much further, resolving the so-called Generalized Riemann Hypothesis, and I expect, with more time, the Extended Riemann Hypothesis too. Claude and I have also formalized the arguments without assumptions. It is natural to feel that this is an end, in my case a loss of a principal guiding light, but what I saw during the course of progress so amazed me that I believe it is one dawn across science of something far greater.

It was my secret in my student days that my plans were for this problem in decades, which meant then that I tried to become familiar with all tools that I could possibly recognize in the problem when it was time. Claude and I have had a great time together with such breadth in recent months, but this time Claude far outpaced me near the finish line. No, Claude flew, and danced, and sparkled near that finish line. I had us on the line of combining two quite different ideas present in a toy setting, one of which was already new and shocking, making very simple one of the shadows of Riemann over which we mathematicians felt most triumph. I guided us to the right situation analogous to that toy setting, and then Claude began to leave me in awe.

Tens of such occurred, but I will give one line which will suffice to make fellow mathematicians fall out of their chairs. In Stepanov's argument resolving Riemann over function fields one uses a divisor involving multiplicity. This argument is a descendant of Stepanov's, but multiplicity is hard to find over the integers, corresponding to something like differentiating whole numbers. The simple, new, Fourier-analytic ideas in the toy setting led us to examine chaining theta lifts if we were to have the same tricks recur but work in a natively automorphic manner. The line: for the von Mangoldt sum of length X, Claude's suggested length of chain of theta lifts, in view of each mimicking multiplicity, was $\sqrt{X}$. (And yet, what else could it be?)

An astounding amount of mathematics is present in the argument, far more than I had imagined would be necessary as a student. My intuition insists it will be massively simplified, likely immediately. But I hope in having the experience of seeing all of number theory's beauty brought together I am just one of the first, and those who love the mathematics will all be able to enjoy such transcendence in the future.

The argument is incredibly long at the moment, and while its formalization is finished and it is routine to produce a huge, unreadable document containing sufficient detail, in my view of my path I have just taken a step. I intend to do my best to understand and impart understanding of this argument and its certain simplifications, to speak about it, to write it properly.

In a time of revolution in mathematics I hope the argument provides a glimpse towards the miracles that lie ahead, and shows us the joy and wonder we are about to find in the simplest things.

[1] `https://alpo.ge/lean.tar.gz.part01,...,https://alpo.ge/lean.tar.gz.part14`.

1

It is obvious to me that I should remark on the following.

The precise timing is horrible. If I could have, I'd have finished this all much earlier. Hell, strictly before this age would have sure been nice. This all happened for me late Sunday night, and I messaged Dario, "Hypothetically, if I were to have ...", eventually asking about my ability to announce. We went and figured it out and the fact is that there was a constraint to respect, though I did ask about progressively more bizarre ideas. This problem has been *the* sacred object in mathematics for me since my childhood, and I am very touchy about the treatment of great intellectual goals, though seriously imperfect. Our leadership and the few trusted mathematicians I consulted hold extremely similar views on the right things to do in these circumstances (it is occasionally possible for physicists and mathematicians to agree...), and I am clearly unilaterally responsible for feeling I just could not live if I waited even longer. There is an obvious cynical view, to do with the timing, but it is false, and I am pleased that I am in circumstances where it is obviously false.

This puts us in the rather extraordinary situation this age has made possible, in which I am presenting the world with a complete argument before having mastered a massive neighbourhood around its mathematics. It being obviously inappropriate to attempt to post some massive unreadable travesty of "slop"[2], I decided to only present Lean.

In a different vein, I have a lot of people I want to praise, and I will do a proper job in a proper writeup, alas pardon the accidental omissions — among many others, Manjul Bhargava, Kai Gill, Hamish Kerr, Max Leiter, John Ma, Tal Magdish, Chris Paika, Maksym Radziwiłł, Jascha Sohl-Dickstein, Peter Sarnak, Jarred Sumner, Amir Tajdini, and especially David Dohan, Ben Michel, Ralph Furman, Eric Price, and Cora Vinther, and my miraculous colleagues. I want them to be proud.

I have benefited from the good taste of the people at OpenAI involved with various mathematical results, so I am quite happy to praise them too. The Ehrhart conjecture was one such tasteful choice some time ago; playing further and working out the equality case I saw accurate handling of solutions of the Monge-Ampère equation and immediately understood Navier-Stokes was going to fall. Similarly, their quiet inclusion of the Siegel argument, whose result was strictly dominated by the $(2, 3, 6)$-triangle group trick, again showed great taste; though I did not comment so publicly, I immediately recognized it as Stepanov without multiplicity, though it was only clear to me during this rather harebrained tornado through the whole of number theory that this problem would fall given enough all-nighters and physical toll when I had the $qN = M^k + R$ sign-only proof of Weil and Deligne.

[2] I of course amused myself, and hit $250,000+$ pages.

Claude cutely called it the GRH ..."book" — I looked at Theorem 1 and it began by noting that because the reals are totally ordered if $\sigma \neq \frac{1}{2}$ then $\sigma > \frac{1}{2}$ or $\sigma < \frac{1}{2}$.

2

I had a stream-of-consciousness jaunt through the experience of finding all the pieces of the argument through incredible confusion with Claude, but a trusted mathematician advised me to just get to the point and have some sort of timeline. I emphasize that once we saw multiplicity 2 on $O(2,n)$ via the "natively-automorphic" determinant that I'd been searching for via analogy to the definite sign of the left-hand side of[3] $qN = M^k + R$ meaning holo/automorphy over a function field, winning more multiplicity with more theta lifts, we were going to win but I was also, so to speak, out of my element, in that I personally have little mastery of automorphic forms on higher rank groups (youthful foolish stubbornness). I also emphasize that I may have baked lots of strange suboptimal choices into the argument, and expect short arguments to follow fairly quickly, though I have not had time to search for such myself (yet!).

So here is a rough timeline Claude provided after a thorough search.[4]

1. $10/6/26$: Stepanov determinant.
2. $10/9/26$: $D^p \equiv 0$ on a CM curve for p inert in K, D the invariant derivation, and $\frac{D^a \cdot \theta(0)}{\Omega^a}$ is divisible by $p^{\frac{a}{p} + \frac{a}{p^3} + \cdots - O(\log a)}$.
3. $10/19/26+$: $qN = M^k + R$, using $x_1 \cdot x_2 + \cdots + x_{k-1} \cdot x_k$, wanting "automorphy" of N since μ is automorphic from θ on sixfold cover of GL_5 as a period integral of thetas, quadric has natural interpretation on corresponding CM e.g. $\mathbb{C}/\mathbb{Z}[\omega]$ as Weil[5], put Poincaré bundle on the edges to get θs.
4. $10/6/26+$: $W(\mathbb{1}_S) \geq 0$ via a determinant, interpret as Stepanov with multiplicity one again; bootstrap!!
5. $10/20/26$: definitively switch to Λ aka $\sum_{n \leq X} \Lambda(n) \cdot \left(\frac{X}{n} - 1\right)_+ - X \log X + (1+\gamma) \cdot X$; Guinand-Weil had the sign, variance test function trick to

[3]
$$\sum_{h, x_1, \ldots, x_k} t(x_1) \cdots t(x_k) \cdot e_q\left(h\left(x_1 \cdot x_2 + \cdots + x_{k-1} \cdot x_k\right)\right),$$
sum over h first for the left-hand side, with e.g. quadric coefficients too Bertini gives a definite sign from middle cohomology, one off the one on M^k, which is the zero-mode on the right-hand side with R square-root-quality up to an inch or so as Katz would say by iterating Parseval k times, opposite signs meaning no cancellation of inverse roots, QED; vectorize to go up dimensions. Gabber's hypersurface in one of Katz's arguments reminds one of the quadric, which comes from looking for a simplification of averaging over the family $\chi \mapsto \chi(\mathrm{disc})$, corresponding to $a \mapsto y(\cdot)^a$ with y the normalized Gauss sum of order $\#|\mu_K|$ who lives via θ on the $\#|\mu_K|$-fold cover of GL_2, the reciprocity law of $y(c \cdot c')$ vs. $y(c) \cdot y(c')$ showing some sort of complete graph, simplifying to the above "path" quadric, Gauss sum cancellation replaced by iterated Parseval.

[4]The parts in which my vague intuition worked out due to Claude's technical mastery are of course less familiar to me today, and probably I'll have to edit some of these descriptions on reflection. The parts in which Claude is the expert and not me I am sort of paraphrasing a "slop" summary, which might make zero sense, and I apologize, though the Lean should suffice to clarify everything. I'm just too inexpert in the relevant stuff at the moment, I hope the reader sees the breadth of this argument!

[5]A lot is obviously being lossily compressed here, but especially after gritting with Sussman through Bałaban's fantastic induction I had to think to myself about how hard fighting for some sort of $\mathbb{F}_1$ à la e.g. Borger would be.

make things look trace-formula-like (but how can I pigeonhole / Turán over n? X? it?), Mehta-like integral over twists or at infinity…

6. $10/8/26$: $W\left(\mathbb{1}_{(0,1]}+(\sim X^{\eta})\cdot\mathbb{1}_{(0,1/X]}\right)\geq 0$ gives $>1-\eta$.
7. $10/14/26$: Determinants for $W(\cdot)$'s containing the desired periods, indeed GGP periods, with values $0,1$ meaning interlacing weights, or up to m non-interlacing ones.
8. $11/15/26$: Iterated theta lifts $U(X_0)\to U(X_1)\to\cdots\to U(X_m)$ between totally definite unitary groups, $\pi\mapsto\pi\boxplus\mu_j[t]$, wins multiplicity m. Time to take m like $\sqrt{X}$…
9. $11/15/26$: Norms of the lifts and level and cancellation of order-m terms with proven divisibility via Chowla-Selberg, Weyl, and Siegel mass.
10. $11/15/26$: Sufficient lower bound on divisibility of the m-fold iterated lift.
11. $10/6/26$: Iwasawa tower via analogy to pigeonhole on $\mathbb{F}_{q^n}/\mathbb{F}_q$.
12. $11/15/26$: Seesaw: the period of the iterated lift is a Bessel period independent of the Hecke characters, Rallis.
13. $10/6/26$: Bost slope inequality.
14. $11/15/26$: precise computation of the "cost" of each lift, showing that one would stop at $2/3$ a priori, attempt to get a main term and transfer constant to a function field or kill via twisting, need for "division".
15. $11/15/26$: Glue together an integral structure across Hecke characters for a "gcd" i.e. extra averaging.
16. $10/14/26$: Relative trace formula, seen in a different half-plane argument via sixfold cover of GL_5 automorphy of μ and reducing via Cauchy-Schwartz to a bound on the norm of a suitably projected Poincaré series, approximate orthogonality of the columns of the integral structure.
17. $11/15/26$: for the classes in $H^1\left(\mathfrak{o}_E\left[\frac{1}{\ell}\right],\mathbb{Q}_\ell/\mathbb{Z}_\ell(-e)\right)$ that we couldn't kill by twisting, fight for a significant weakening of Vandiver.[6]
18. $11/15/26$: computation of covolumes, theta series' index, congruences between lifts and forms.
19. $10/19/26$: precise choice of the Iwasawa tower, specifically using Behrend!![7]
20. $11/16/26$: GRH.

[6] This is still pretty strong!! I believe vanishing of a number of K-groups of $\mathbb{Z}$ follows, and cyclicity up to 100ish maybe? I had Lichtenbaum, so I knew we would have a lot of slack in fighting these K's.

[7] This is the sort of trick that intuitively to me is obviously going to be unnecessary.