A LESSON ABOUT PRIME NUMBERS, FOR THE CLASSROOM OR A CLUB

What would solving the Riemann Hypothesis mean?

A 50-minute lesson for Years 7–10, with an extension for older or more curious students. It pairs with the short film *Maths in the News #1* and the student worksheet. The aim is curiosity and a clear picture, not exam technique.

What you will need

- The student worksheet (pages 1–4; keep page 5, the answers, for yourself)
- The Short, if it is available — or the spoken hook below
- Scientific calculators with an ***ln*** button, at least one between a pair, from Q5 onwards
- Optional: the lesson page at ilovemaths.app/lessons/riemann-primes

Objectives

- Use the sieve of Eratosthenes to find the primes up to 100.
- Count those primes in blocks of ten and notice that they thin out.
- Compare the estimate x / ***ln***(x) with the actual count.
- Explain, in plain words, what the Riemann Hypothesis claims, and why a proof would not break encryption.

The shape of the lesson

0–5 min — **Starter hook — play the Short if you can.** The film is about a minute long. If it is not yet on YouTube, use the spoken hook: “This week a famous puzzle about prime numbers made the news. It is called the Riemann Hypothesis. Before we find out what that name means, what do you already know about primes?” Take two or three answers. Do not correct yet. Ask what questions the film (or the hook) left them with.

5–20 min — **The sieve, together or in pairs.** Read the four steps on page 1 of the worksheet, then let students mark the 1–100 grid. Circulate. If someone crosses out a prime, ask “How did you decide that one should go?” rather than leading with “that’s wrong”. Stopping after 7 is a nice “why” moment: 11 × 11 is already past 100.

20–32 min — **Count in tens, then talk.** Students fill the table on page 2 and answer Q1–Q4. Compare a few tables. The counts should be 4, 4, 2, 2, 3, 2, 2, 3, 2, 1 — adding up to 25. If a table is off, treat the mismatch as a puzzle to check, not a failure. Draw out that the primes thin out, but not in a perfectly straight line (41–50 and 71–80 bump back up to three).

32–40 min — **Predict, then check.** Page 3. Students write a guess first, then calculate 100 / ***ln***(100). They should get about 21.7 (the film says “about 22”). The actual count is 25. Then Q8: up to 1,000 the estimate is about 145, and the actual count is 168. The rule is close, not exact, and the gap can grow.

40–47 min — **What it says, and why it matters.** Read the plain-words box together. Students attempt Q9 in their own words. Then the two important cautions: a proof would **not** break encryption, and a headline is not the same as the Clay $1 million prize being awarded. Use Q10 and Q11. Keep the November 2026 news wording careful (see below).

47–50 min — **Plenary.** In one sentence, orally or on a mini whiteboard: “The Riemann Hypothesis is a claim about…” Listen for “how evenly the primes are spread” or “the error in the estimate staying small”. Thank effort and good questions by name, specifically.

Discussion questions you can drop in

- If the primes look random in small blocks, how can there still be a rule underneath?
- Is “close” the same as “exact”? When is a rule of thumb good enough to be interesting?
- Why might mathematicians care about the *error* in an estimate, not only the estimate itself?
- If a proof is checked by a computer, what would still make you want human mathematicians to look at it?

Misconceptions to listen for

- **1 is a prime.** A prime has exactly two factors. 1 has only one. This is the most common slip on the first block.
- **Even numbers after 2 can still be prime.** Once 2 is circled, every other even number is a multiple of 2.
- **“Evenly spread as possible” means the same number of primes in every block of ten.** The table already shows that is false. The claim is about the error in the overall estimate staying small.
- **Solving the Riemann Hypothesis would break encryption or “the internet”.** It would not. RSA-style security depends on the hardness of factorising large numbers, not on this problem remaining unsolved.
- **A news headline means the $1 million prize has been won.** A computer-checked proof now looks very likely correct, but that is not the same as the prize being awarded. The Clay Mathematics Institute still lists the problem as Active, and the US$1 million prize has not been awarded.
- Some students may offer “there are only finitely many primes”. That is a different (and false) claim, proved otherwise by Euclid a very long time ago. A one-line reminder is enough; it is not the focus today.

Differentiation

If students need more support

Do the multiples of 2 and 3 together as a class, then let pairs finish 5 and 7. You can pencil a faint hint on 1, or provide the first two table rows already filled. For the logarithm, calculate $\ln(100)$ once on the board and let students do only the division. Pair work is better than a solo struggle with the calculator buttons.

If students are ready for more

The extension on page 4 (E1–E4) is for older students (Years 11–13) or a fast Year 10. E2 is the richest question: the absolute gap grows, but the relative gap tends to shrink for large x — and it is not perfectly smooth between 100 and 1,000. Praise anyone who spots that wrinkle. There is no complex analysis in this lesson, on purpose.

The November 2026 news — please keep this careful

A computer-checked proof of the Riemann Hypothesis **now looks very likely correct**. On 20 November 2026, Levent Alpöge, a mathematician at Anthropic, released a proof he worked on with the AI Claude over about six weeks. It is a Lean library of about 17 million lines, and it also covers the Generalized Riemann Hypothesis for Dirichlet L-functions. On 21 November the Lean and Mathlib communities certified the statement and the proof. Independent machine checks by Andrew Sutherland (MIT) and Michael Rubinstein (University of Waterloo) report that the proof compiles using only Lean’s standard axioms. A human-readable write-up is still being drafted. The Clay Mathematics Institute still lists the problem as Active, and the US$1 million prize has not been awarded. In class, please do not tell students that the prize has been won, and update this if the public picture changes.

Links

- This lesson on I Love Maths: **https://ilovemaths.app/lessons/riemann-primes**
- The Short, once it is published, and Jane’s channel: youtube.com/@getsussed-y4h
- The student tutor app (enjoyment and confidence, not exam marks): https://ilovemaths.app
- Clay Mathematics Institute, Millennium Prize Problems (RH still listed as Active): https://www.claymath.org/millennium-problems/
- Lean and Mathlib communities, 21 November 2026: https://leanprover-community.github.io/blog/posts/riemann-hypothesis/
- Scientific American, updated 21 November 2026: https://www.scientificamerican.com/article/the-ai-riemann-hypothesis-proof-is-almost-certainly-true-independent-mathematicians-attest/
- Michael Rubinstein, University of Waterloo: https://math.uwaterloo.ca/~mrubinst/alpoge_claude_grh/
- The Guardian, 21 November 2026: https://www.theguardian.com/technology/2026/nov/21/anthropic-riemann-hypothesis-solved-claude-ai

For teachers: how students learn best with AI

The evidence is still early, and it is mixed. These points are here to be useful and honest — not to oversell I Love Maths, or AI in general. Links were checked in November 2026.

1. **AI is more likely to help when it supports understanding than when it does the work.** A November 2026 review by Stanford's SCALE Initiative (the AI Hub for Education) looked at 52 cause-and-effect studies. It found that using AI to understand new material is more promising than using it to produce the work — although neither use consistently improves learning. Using AI to offload the task can reduce what students later know.
2. **Treat sweeping "AI improves learning" claims with care.** An Education Week report (2 November 2026) covered an audit of 14 research reviews that all claimed positive effects. The audit found they failed basic research standards: they mixed very different studies, and often lumped together engagement, motivation or "how much students felt they had learned" with actual achievement.
3. **A few careful trials are encouraging, but they are still small.** In one randomised trial in Ghana, students who used Rising Academies' WhatsApp maths tutor Rori for two 30-minute sessions a week, over eight months, gained about 0.37 standard deviations in maths compared with a control group (around 1,000 students in grades 3–9, across 11 schools). That is a meaningful gain in education research, and the authors asked that these first-year results be interpreted carefully.
4. **Good design coaches step by step, rather than handing over the answer.** Google's Guided Learning mode in Gemini is built to ask questions and walk through the steps, rather than simply giving the answer. That is a design choice, in line with the idea above. It is not, by itself, independent proof that the tool improves results.
5. **When AI is used to mark work, human checks still matter.** NZQA uses automated text scoring for digitally submitted NCEA Literacy writing, and human markers check scripts around the Achieved / Not Achieved boundary (about 40% of assessments, in NZQA's published description). If the human mark differs, the human mark is used. This is about assessment marking, not tutoring — it is a reminder that AI works best as a support, with people still in the loop.
6. **The older findings about how people learn still apply, with or without AI.** Students generally learn more when they have to pull ideas out of memory (retrieval practice), explain their thinking (self-explanation), work through a genuine challenge rather than being rescued too quickly (productive struggle), and get feedback on their method, not only on the final answer.

How I Love Maths uses this

I Love Maths is built to support understanding, not to do the work for the student. The tutor asks questions before it gives answers. To reach "Mastered" on a topic, a student needs to type an explanation in their own words. If they send a photo of their working, it is marked for method — and the photo is not stored.

Sources (checked November 2026):

1. Stanford SCALE, "The Evidence Base on AI in K-12: A November 2026 Review" — https://scale.stanford.edu/ai/evidence-base-review/evidence-base-ai-k12-november-review
2. Stephen Sawchuk, "The Research on AI in Education Is a Mess", *Education Week*, 2 November 2026 — https://www.edweek.org/technology/the-research-on-ai-in-education-is-a-mess/2026/11
3. Owen Henkel, Hannah Horne-Robinson, Nessie Kozhakhmetova and Amanda Lee, "Effective and Scalable Math Support: Evidence on the Impact of an AI-Tutor on Math Achievement in Ghana" (arXiv, May 2024 version) — https://arxiv.org/abs/2402.09809
4. Google, "How Gemini's Guided Learning can help you study more effectively", 23 September 2025 — https://blog.google/products-and-platforms/products/gemini/guided-learning-google-gemini/
5. NZQA, "How we mark co-requisite assessments" — https://www2.nzqa.govt.nz/ncea/subjects/litnum/marking/
6. Retrieval practice: https://www.retrievalpractice.org/ · Michelene Chi and colleagues, "Eliciting Self-Explanations Improves Understanding", *Cognitive Science*, 1994 — https://onlinelibrary.wiley.com/doi/10.1207/s15516709cog1803_3 · Manu Kapur, "Productive Failure in Learning Math", *Cognitive Science*, 2014 — https://onlinelibrary.wiley.com/doi/full/10.1111/cogs.12107 · John Hattie and Helen Timperley, "The Power of Feedback", *Review of Educational Research*, 2007 — https://journals.sagepub.com/doi/10.3102/003465430298487

Jane, an experienced maths teacher
I Love Maths — a friendly maths tutor for GCSE and NCEA students who want maths to make sense again.
https://ilovemaths.app · This lesson: https://ilovemaths.app/lessons/riemann-primes