

Xena

Mathematicians learning Lean by doing.

-- part (a) -- delete one of these and prove the other.

theorem M1F_Sheet1_Q1a_is_T : ∀ x : ℝ, x ^ 2 - 3 * x + 2 = 0 → x = 1 := sorry

theorem M1F_Sheet1_Q1a_is_F : ¬ (∀ x : ℝ, x ^ 2 - 3 * x + 2 = 0 → x = 1) := sorry

-- part (b) etc etc.

theorem M1F_Sheet1_Q1b_is_T : ∀ x : ℝ, x = 1 → x ^ 2 - 3 * x + 2 = 0 := sorry

theorem M1F_Sheet1_Q1b_is_F : ¬ (∀ x : ℝ, x = 1 → x ^ 2 - 3 * x + 2 = 0) := sorry

Home

About Xena

Student projects

Installing Lean and mathlib

Bluesky

Useful links.

← The Annals Challenge

FLT: Anthropic has beaten me to it

Posted on September 4, 2026 by xenaproject

I guess technically it was revealed to the world by [a coffee shop in Islington](#) on Insta, but an hour later it was [officially announced by Anthropic](#): one of their internal models, using the [prove2.me](#) platform, has formalized a complete proof of Fermat’s Last Theorem (FLT) in Lean. This is the final theorem to be formalized in Freek Wiedijk’s famous [list of 100 formalization challenges](#) and thus wraps up this 20-year-old benchmark. Congratulations to Anthropic!

Mathematical details

The proof is not the modern proof which I have been formalizing myself following ideas of Khare, Taylor etc, but the Darmon–Diamond–Taylor exposition from 1995 of the Wiles–Taylor–Wiles argument, via the Langlands–Tunnell theorem and Ribet’s level-lowering theorem. Anthropic’s repository develops Fontaine theory (to study flat deformations of Galois representations) and develops enough of Mazur’s work on the Eisenstein ideal to conclude that no Frey curve can have a point of order $p \geq 17$. This means that their FLT proof only works for $p \geq 17$, however FLT was [already formalized for odd regular primes](#) by Best-Birkbeck-Brasca-Rodriguez, and the smallest irregular prime is 37, so it’s all good.

The code base

I’ve compiled [the code base](#) and run [comparator](#) on it — it checks out. It is a gigantic proof (over 13.4 million lines of code) and takes nearly 20 times as long to compile as Lean’s mathematics library (on a machine with 96 cores!). Lean can be sluggish when jumping from file to file on a repo of this size (even on a machine with 500G of ram, which Anthropic also gave me access to), but Anthropic also supplied me with some [html documents](#) which are easier in practice to explore (clone the repo and open with a web browser).

What this work is, and is not

I am currently being funded by the EPSRC to formalize a proof of Fermat’s Last Theorem, and a naive reaction to the news above is that I no longer have any work to do. This is not the case. The work certainly achieves some of the aims of the EPSRC project, and indeed it goes much further in terms of what is formalized (I only promised the EPSRC that I would reduce FLT to the 1980s; this repo proves the whole thing). But I also promised several other things to EPSRC: firstly, that I would be making pull requests to Lean’s mathematics library, adding fundamental objects from modern number theory; this is ongoing. And secondly, and perhaps most importantly, that I would be creating a dynamic document enabling humans to explore the modern proof. My guess is that it is unlikely that Anthropic are going to do this; they will feel that their job is done with the formalization (and they did not formalize the modern proof anyway).

Note that mathematically this work of anthropic tells us essentially nothing: I am on record as saying that I am 99.9% sure that the proof of FLT is OK, and most people in the number theory community are 100% sure (formalization has made me more paranoid about the mathematical literature than most). From my understanding of the argument, the formalization just faithfully follows the early literature on the proof and adds nothing.

What this work *does* tell us, however, is what is possible in the field of autoformalization. If thousands of pages of the literature can be formalized end-to-end by some kind of AI swarm in an 11 day period now, then in the future we will start to see formalization of modern research being done on the fly. We will also learn whether my paranoia about the current state of the Langlands program is justified, as machines check it and ruthlessly flag arguments which are incomplete. The ability to autoformalize hard material will ultimately make the review process for mathematics papers far less painful. It will also keep us honest — there are papers out there which assume results which are “known to the experts” and it will be interesting to see exactly what is being assumed in the proofs of various important results in my field. This is why I am so excited about the news!

I was given £1M to run my project over 5 years; Anthropic took only 11 days but I do wonder if they spent more money...

An anecdote

Thought it might be nice to finish with a personal anecdote. Wiles announced his proof of FLT at the Newton Institute in 1993 in a series of three lectures; I attended the first (I was a second year graduate student at the time) and I found it completely incomprehensible, so I skipped the next two lectures and went on holiday to Ireland with my new girlfriend instead; I was so in love that I totally forgot about the rumours, and it was only when I came back to Cambridge a week later that I heard the news that the theorem was proved. Something strangely similar happened here; when I got the email from Anthropic I was in Wales at the [Green Man music festival](#) with the same girlfriend, but with very poor phone reception; I did spot an email from someone I’d never heard of in a brief moment of 4G, with title “End-to-end Lean formalization of Fermat’s Last Theorem”, but wrote them off as a crank! It was only a week later when going through the nearly 1000 unread emails which had accrued whilst I was away, that I heard the news.

Share this:

 X

 Facebook

Reblog

Like

 2 likes

Related

Formalizing Fermat workshop May 15, 2026 In "AI"	Fermat’s Last Theorem — how it’s going December 11, 2024 In "Fermat's Last Theorem"	Does anyone “know” a proof of Fermat’s Last Theorem? September 27, 2019 In "Brian Conrad"
--	---	---

About xenaproject
The Xena Project aims to get mathematics undergraduates (at Imperial College and beyond) trained in the art of formalising mathematics on a computer. Why? Because I have this feeling that digitising mathematics will be really important one day.
[View all posts by xenaproject →](#)

This entry was posted in [Fermat's Last Theorem](#) and tagged [AI](#), [Artificial Intelligence](#), [lean](#), [math](#), [mathematics](#), [science](#). Bookmark the [permalink](#).

← The Annals Challenge

1 Response to *FLT: Anthropic has beaten me to it*



David Jao says:

September 4, 2026 at 6:57 pm

The Anthropic post claims the proof used 6 billion output tokens. At the API price, that would be \$300,000.

★ Liked by [1 person](#)

[Reply](#)

Leave a comment

Xena

Create a free website or blog at WordPress.com.

Privacy & Cookies

We, WordPress.com, and our advertising partners store and/or access information on your device and also process personal data, like unique identifiers, browsing activity, and other standard information sent by your device including your IP address. This information is collected over time and used for personalised advertising, advertising measurement, audience research and services development specific to our ads program. If this sounds good to you, select "I Agree!" below. Otherwise, you can get more information, customize your consent preferences for 100 different ad vendors, or decline consent by selecting "Learn More". Note that your preferences apply to all websites in the [WordPress.com network](#), and if you change your mind in the future you can withdraw your consent anytime by visiting the Privacy link displayed under each ad or by using the "Privacy" option in the Action Bar located at the bottom-right corner of the screen. One last thing, our partners may process some of your data based on legitimate interests instead of consent but you can object to that by choosing "Learn More" and then disabling the