

Dr. Robert Dawson (Saint Mary's University)

Editor - CMS Notes

Things are happening fast.

As you've probably heard, OpenAI has recently (September 2026) announced a resolution of one of the Clay Millennium problems, constructing a singular solution to the Navier-Stokes equations, in which well-behaved initial conditions lead to a finite-time blowup. The claimed solution is a cylindrical vortex which "implodes" to a singularity in a finite length of time. The solution has not yet been explicitly constructed, but a machine-checkable existence proof using Lean has been released.

Physically, such a result is surprising but not unprecedented. The surprise is partially because we don't see it in real life: in most cases blowups like this in physics are prevented by the breakdown of the model. For instance, while something similar happens when wind patterns on a scale of many kilometers form a small intense tornado, a tornado does not continue to intensify until infinite wind speeds are reached! Rather, other effects that were not strong enough to prevent the tornado from forming become strong enough, relatively, to limit its intensity. Similarly, as a small bubble in a fluid collapses under surface tension, the curvature grows and the collapsing force increases faster than the outward pressure from the gas trapped inside. A simplistic model suggests that the gas should heat adiabatically to an infinite temperature, and this has been proposed as a way to achieve nuclear fusion; unfortunately, the model breaks down before this can occur, though under some circumstances not until the gas in the bubble becomes white-hot. And I once saw a solution to the five-body gravitational problem in which infinite velocities are reached in a finite time. Of course, to free this unlimited kinetic energy, some of the potentials $-M_i M_j G / r_{ij}$ must become arbitrarily large, distances r_{ij} must approach 0, and the model breaks down when the bodies collide!

But the concern among mathematicians is not that when you stir your next cup of coffee it may develop a singularity and explode. Rather, there's concern about the way in which the research has been done. There is some suspicion that the proprietary large language model that Open AI used may have been trained (in part) on publications in which mathematicians working on the same problem published their interim results. At the time of writing, this seems to be unverified, and *may* be a false alarm; what seems clear is that OpenAI, using enormous resources of a type not generally available, has horned in on an established research area, in a fashion not generally accepted in the mathematical community. The result also seems to have been announced publicly without any form of peer review, another solecism.

Now, that phrase "not generally accepted" must be read with nuance. On a small self-contained problem, it is not unknown for two (or more) mathematicians to be working on the same problem at the same time, each oblivious to the other(s). This may continue to parallel publication, or an editor may suggest that the authors combine their work in a joint publication. For a large project such as the Navier-Stokes problem, it would be unlikely for established researchers to be unaware of each other: they read the same journals, attend the same conferences, and quite likely keep each other informed before they tell the rest of the community. This is expected, but it's understood (especially in the more "accessible" areas of mathematics) that every so not-very-often an amateur or undergraduate will solve a problem that is baffling a nonempty set of experts. And that's the luck of the game. While the suggestion has occasionally been made that mathematicians should join lawyers, engineers, and geologists in setting up a closed shop and seeking the right to prosecute any outsider who tries to practice the craft from outside the tent, the idea has not caught on; and even those supporting it are not, I think, motivated by the need to prevent the occasional outsider from inadvertent trespass on an established research project.

However, OpenAI is by no means an amateur creating proofs at their kitchen table after washing the dinner dishes. It's estimated that the computing resources thrown at this frontal assault on the outstanding Clay Millennium Problems cost in the millions of dollars. Though of course OpenAI's main purpose was to showcase their software, as evidenced by the selection of problems that every science journalist has heard of, it's interesting to wonder what the result would have been if the same resources had been put into buying teaching relief for all the leading researchers in the field. Cash infusions of this scale tend to get results... as indeed was the Clay Institute's assumption when they offered the prizes.

It may be argued that it is not new for some researchers to have access to computing resources that others don't. When Appel and Haken used intensive computation to prove the four-colour theorem in 1976, not every graph theorist would have had access to the large amount of mainframe CPU time that their proof needed, even if they had known what to do with it, and the personal computers of the day were not adequate to the task. More recently, there have been periods when computer algebra systems were on some desks and not others. (As the great Piet Hein put it, "money's all wrong/ and the devil decreed it/ It doesn't belong/ to the people who need it.") Other sciences have worse inequalities of opportunity; and we generally, if grudgingly, learn to accept the situation.

But still: is it too much to ask that a nontraditional newcomer to a community, putting resources on such a scale behind their incursion, should at least find out how things are done there and try to fit in politely?

Copyright & Permissions

The Canadian Mathematical Society grants permission to individual readers of this publication to copy articles for their own personal use. Use for any other purpose is strictly prohibited. To obtain a license for anything other than copying articles for

personal use, please contact the Canadian Mathematical Society to request permissions or licensing terms.

Canadian Mathematical Society — 616 Cooper St., Ottawa, ON K1R 5J2, Canada