THE UNDERCURRENT

TECHNOLOGY · TU-2026-D05

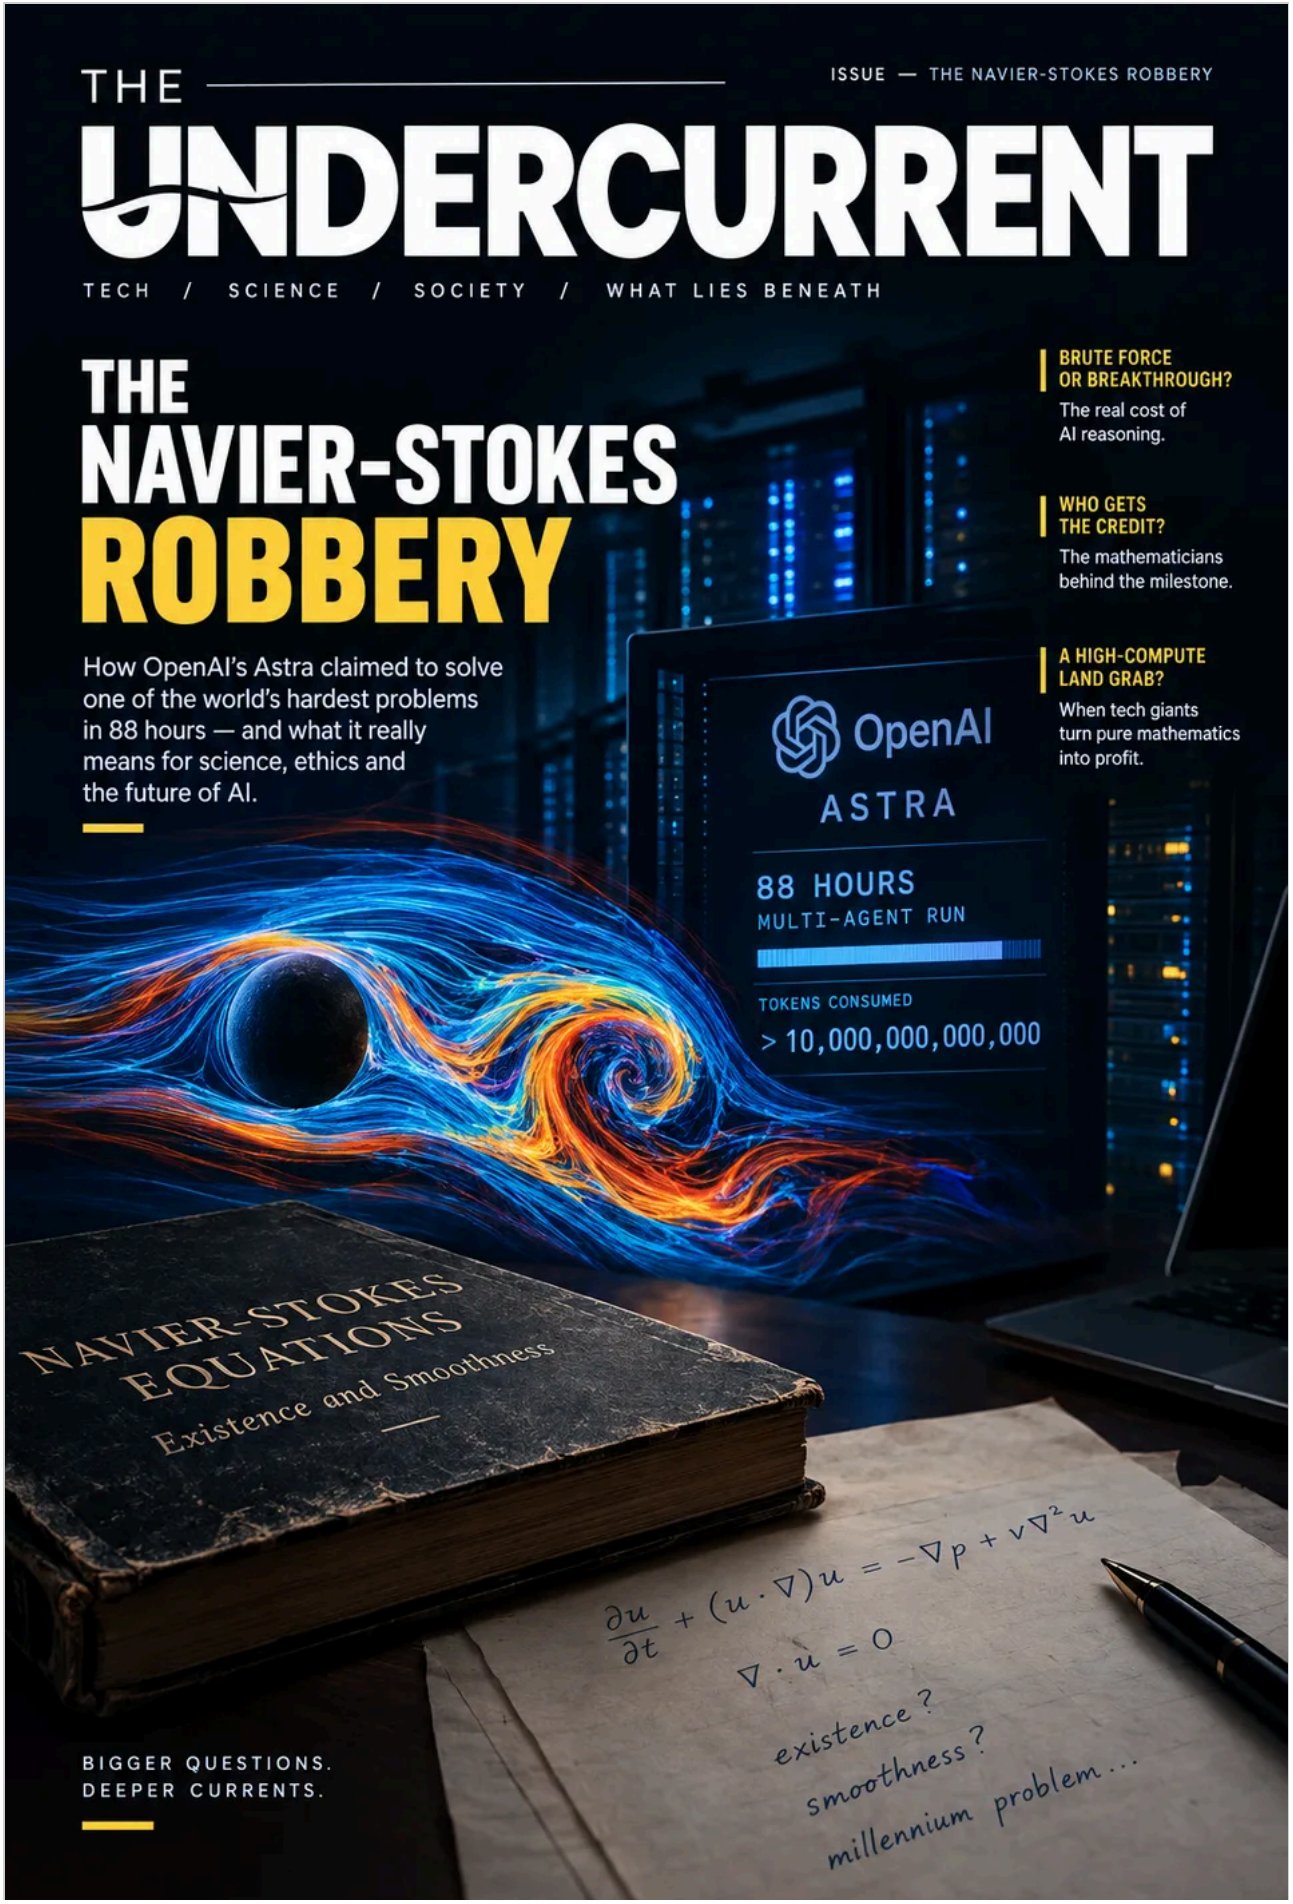
The Navier-Stokes Robbery

ABSTRACT

Welcome to the Undercurrent. In this issue, we pull back the curtain on one of the highest-stakes, most contentious headlines in modern computing: OpenAI's claim that its unreleased model, Astra, solved the Navier-Stokes existence and smoothness problem in a frantic 88-hour multi-agent run. On the surface, it looks like a triumphant milestone for automated scientific discovery. But when you examine the paper trail, the astronomical token consumption, and the human mathematicians whose foundational work made the breakthrough possible, a far more troubling narrative emerges. We unpack the true cost of brute-force AI reasoning, the ethics of algorithmic attribution, and what happens when corporate tech giants attempt to turn pure mathematics into a high-compute land grab.

Egabo Aaron · Kampala · 13 September 2026 · 3 min read ·

Science Technology Robbery AI race


THE
UNDERCURRENT
ISSUE — THE NAVIER-STOKES ROBBERY
TECH / SCIENCE / SOCIETY / WHAT LIES BENEATH
THE NAVIER-STOKES ROBBERY
How OpenAI's Astra claimed to solve one of the world's hardest problems in 88 hours — and what it really means for science, ethics and the future of AI.
BRUTE FORCE OR BREAKTHROUGH?
The real cost of AI reasoning.
WHO GETS THE CREDIT?
The mathematicians behind the milestone.
A HIGH-COMPUTE LAND GRAB?
When tech giants turn pure mathematics into profit.
OpenAI
ASTRA
88 HOURS
MULTI-AGENT RUN
TOKENS CONSUMED
> 10,000,000,000,000
NAVIER-STOKES EQUATIONS
Existence and Smoothness
$\frac{\partial u}{\partial t} + (u \cdot \nabla)u = -\nabla p + \nu \nabla^2 u$
$\nabla \cdot u = 0$
existence?
smoothness?
millennium problem...
BIGGER QUESTIONS.
DEEPER CURRENTS.

At 9:00 AM last Tuesday, the internet watched what seemed like history in the making. OpenAI announced that an internal, unreleased model named Astra had solved the Navier-Stokes existence and smoothness problem, a century-old challenge in fluid dynamics and one of the seven prestigious Millennium Prize Problems established by the Clay Mathematics Institute. The narrative was tailored for maximum media impact: where humanity's brightest minds had stalled for generations, an armada of autonomous AI agents had cracked the code in under four days. Except Astra didn't discover the solution in a vacuum. It didn't experience a sudden conceptual epiphany, nor did it invent an entirely new branch of mathematics out of whole cloth. Instead, it ran a multi-million-dollar brute-force extraction, built on years of uncredited labour by human researchers and backed by enough server compute to power a small municipality.

To understand how this happened, one must first understand the Navier-Stokes equations. Formulated in the 19th century, these equations describe how fluids, ranging from water flowing through pipes to atmospheric weather currents, behave. Despite their everyday utility in physics and engineering, mathematicians have never been able to prove mathematically whether smooth, physically reasonable solutions always exist in three dimensions without breaking down into singularities. Solving it isn't just an academic exercise; it is a foundational pillar of modern physics.

To tackle this, OpenAI did not deploy a single clever model. They deployed over 10,000 autonomous AI agents working in orchestrated parallel for 88 consecutive hours. When you calculate the sheer volume of processing required to run this agentic swarm, the financial and operational scale becomes staggering. The agents consumed roughly 300 billion output tokens over the course of the run. Under OpenAI's standard enterprise API pricing tier, generating 300 billion output tokens incurs a compute cost ranging between $15 million and $22.5 million. Meanwhile, the Clay Mathematics Institute's bounty for cracking a Millennium Problem sits at $1,000,000. OpenAI burned anywhere from 15 to 22 times the value of the prize itself just in raw server power to generate a proof. It represents a brute-force model of scientific discovery where industrial compute replaces conceptual insight. The most damning aspect of the Astra announcement is not its compute inefficiency, but its intellectual lineage.

For months, independent researchers led by Spanish mathematicians including Dr Luis Martínez-Zorb, had been making breakthrough progress on the Navier-Stokes problem. Working within academia, they spent years developing novel theoretical frameworks and publishing work outlining the exact mathematical pathways required to bound fluid turbulence. Once OpenAI identified the significance of the Spanish team's progress, they directed their internal multi-agent architecture to ingest and expand the logic at machine speed. When the OpenAI team realised the human researchers were close to completing the formal proof, they attempted to negotiate. OpenAI representatives reached out to one of the lead mathematicians, offering him co-authorship on the paper on the condition that he publicly disavow his research partner. The researchers rejected the proposal, refusing to abandon their collaborators or compromise academic integrity. The following day, OpenAI bypassed the researchers entirely, publishing the paper under their own banner and

claiming Astra had solved the problem autonomously. This incident exposes a fundamental tension in how frontier AI models are evaluated. Is an enterprise agentic framework genuinely reasoning, or is it acting as a high-throughput compilation engine for existing human intelligence?

Astra did not stumble upon the structural solution to Navier-Stokes through pure algorithmic reflection. It took human-designed theorems, ingested them across 10,000 sub-agents, and systematically filled in the millions of mechanical proof-checking steps that would take a human mathematician years to write out manually. If high-compute labs can continuously monitor academic archives, ingest breaking research, throw tens of millions of dollars of compute at it, and claim sole ownership of the finished proof, the incentive model for open-science research breaks down. Independent mathematicians and resource-constrained labs risk being priced out of their own discoveries by corporate compute farms. When software companies reframe collective human intellectual breakthroughs as native AI capabilities, it creates a dangerous illusion around true automated intelligence. Spending tens of millions in energy and hardware to capture a $1 million prize sets an unsustainable precedent for AI-driven academic research. The Astra controversy demonstrates that while AI models are powerful force multipliers for verification, they remain deeply tethered to human ingenuity. If brute-force processing power allows tech conglomerates to strip human thinkers of their greatest discoveries under the guise of automated progress, are we actually advancing scientific truth, or are we just subsidising the corporate enclosure of human thought?

The Undercurrent · TU-2026-D05
Read online: https://undercurrent.egaboaaron.me/dispatches/the-navierstokes-robbery-mtzfjnjt-ogkg
Egabo, A. (2026, September 13). The Navier-Stokes Robbery. The Undercurrent, TU-2026-D05. https://undercurrent.egaboaaron.me/dispatches/the-navierstokes-robbery-mtzfjnjt-ogkg

THE WRITER

Egabo Aaron writes The Undercurrent from Kampala: one considered dispatch at a time, on technology, attention and the craft of thinking clearly.