

Response to the AHM statement

A response to the Association for Human Mathematics Communications Working Group's [statement of October 7, 2026](#). The original statement is reproduced separately in the appendix.

Dear Professor Tao,

I would like personally to thank you for your wonderful work in disseminating a variety of perspectives on the important questions surrounding the use of artificial intelligence (AI) tools in mathematics. What I write below concerns the Association for Human Mathematics Communications Working Group's [statement of October 7, 2026](#).

I find this statement to be particularly revelatory, and I am grateful to have such a succinct representation of the views of this (rather prominent) segment of our community.

I feel it is important to offer a rebuttal. The Association for Human Mathematics (AHM) has, as of October 8, 2026, [808 members](#); its ranks boast such luminaries as Peter Scholze, Alexander Beilinson, Michael Harris, Nalini Anantharaman, and Sophie Morel. It is therefore true that any statement emanating from the AHM — even insofar as the Communications Working Group might not represent the individual views of any particular member — carries an enormous amount of force. That force is acquired not only from the intellectual power and elite status of the mathematicians of the AHM, but also from the borrowed authority of group action: the AHM takes up the posture of a professional association, labor union, or political lobby — or all three — officially representing mathematicians everywhere. Though this is no fault of the AHM, it is extremely intimidating to dissent from a group this large and powerful, especially when many members are my friends, colleagues, and mathematical idols.

But I offer this rebuttal because I believe that principled dissent is most necessary when the consensus it opposes is most intimidating — and because I am convinced that the AHM's latest stated position is so patently absurd as to verge on self-parody. I believe that to omit to dissent publicly from this position would be to fail to discourage actions and attitudes that can only undermine the credibility and well-being of the mathematical community.

I am a vastly inferior mathematician to many of the AHM's members, and it would perhaps be more effective for someone more eminent — or at least more secure in their career — to offer a statement propounding views approximating my own. Unfortunately, I have not seen such a statement. And so — and accepting whatever risk of opprobrium may come with my making these statements — I feel I must opine, out of a desire for intellectual honesty, inasmuch as I believe that stating my views frankly may inspire those who might agree with me to come forward as well. So far, I am a faction of one.

I am going to be methodical in my criticism. For practically every line contains some rhetorical flourish, some innuendo, or some fallacy that should rankle those of us who are committed to honesty and clarity of thought (which, as mathematicians, we all ought to be).

For now let me summarize the major points:

1. The AHM neither can nor should pretend to speak for all mathematicians; 2. Mathematical problems are not the exclusive property of some insular priesthood ordained to decide peremptorily if and how they may be treated by others; 3. These latest results released by OpenAI are significant mathematical discoveries, and to pretend otherwise is ridiculous; like all significant mathematical

discoveries, they should be celebrated, and we should be grateful for them; 4. Much of our current alarm over AI contributions to our discipline stems from a jealous desire to preserve our enviable status and privileges, rather than from a sincere concern for the harm they may do to “the human understanding of mathematics”; 5. Curiosity, rather than ambition, will spur and sustain future human participation in and contributions to mathematics — which is no bad thing.

The statement’s very first sentence contains the following aside:

“... OpenAI — which is currently defending lawsuits against accusations of illegal plagiarism, copyright infringement, and trademark dilution —...”

This is pure innuendo. If we are trying to offer a criticism of AI’s role *in mathematics*, then it is prejudicial for lawsuits against OpenAI to be mentioned in passing like this without showing their bearing on the mathematical work at issue. The existence of, say, the case being brought by the *New York Times* against OpenAI does not establish the case the AHM is trying to make against these manuscripts; if indeed the AHM is trying to suggest that OpenAI’s mathematical contributions are the result of plagiarism, copyright infringement, or trademark dilution, then I invite the Association to say so plainly and provide evidence rather than relying on innuendo.

Indeed, it is quite clear that mathematical publication operates under a very different economic model from journalism. The central allegation in the [New York Times complaint](#) is that OpenAI and Microsoft copied its copyrighted reporting; the *NYT*’s contention is that this permits OpenAI to generate products that substitute for its own journalism. This in turn threatens subscription, advertising, licensing, and other revenues that the *NYT* depends upon. Whether that copying and use were unlawful remains disputed — but this substitution is the central alleged economic harm (although there are others in the lawsuit).

As research mathematicians, we do not generally rest our livelihoods upon selling access to our papers. If someone, say, asks AI to explain my theorem rather than reading my paper, I have not thereby lost a subscription payment. (*This* is the part of the analogy to the *New York Times* case that I think fails to carry over.)

A foundational aspiration of the arXiv is to grant open access to *everyone* in the community. That does not make its papers public-domain material: [copyright and reuse licenses still apply](#). Nor does it settle the legality of AI training. But each individual mathematician is encouraged to read, digest, and build upon the mathematical ideas in circulation. If the suggestion being made here is that mathematicians ordinarily receive payment whenever others learn and use their ideas, then that is not the norm under which our research community generally operates.

As mathematicians, we seek *credit* when others use our ideas; we do not *charge them* for doing so. (And, as an aside, I would like to point out that the distinction between mathematical ideas and their written expression is profound. [U.S. copyright law does not protect mathematical ideas, principles, or formulas as such](#), though it may protect the text explaining them.) It is possible that OpenAI’s output is insufficiently generous in offering credit to mathematicians. (This, indeed, is a problem in the human mathematical community as well.) But if *that* is a serious complaint, the AHM ought to state it as such — and offer evidence of the cases in which it has occurred — rather than tacitly advance that case through a tendentious reference to litigation involving OpenAI. An

omission of a citation is readily fixable; serious plagiarism is another matter. And neither should be conflated with the broader claim that a company ought not attempt mathematical research with its own models.

I also profoundly resent the rhetorical device of invoking pending civil litigation to insinuate that OpenAI is somehow nefarious. Indeed, OpenAI may ultimately be found not liable. Mentioning allegations before they become established findings as a way to prejudice a reader is unscrupulous.

The statement continues:

“Mathematicians did not ask for this work to be done.”

Pardon me, AHM, but what gives you the right to speak on behalf of (all? some? 808?) “mathematicians”? I certainly wanted this work to be done and its results released, and expressed that desire to colleagues. The AHM does not speak for all of us. And I hope OpenAI will continue to resolve mathematical mysteries and share the results.

The community does not ordinarily commission, or collectively authorize, every attempt to solve an open problem. The very nature of *being a problem* makes it something that is begging to be solved. The suggestion that a group of mathematicians should have to welcome such an attempt before it counts as legitimate research strikes me as ludicrous. Ethical conduct and sound mathematics matter; collective permission is unheard of.

“The Advisory Group on Mathematics and Artificial Intelligence, from whom OpenAI has claimed to derive its legitimacy, opened their initial advisory statement by saying that frontier AI corporations should not test advanced mathematical problems on internal models.”

In my view, the recommendations of the Advisory Group on Mathematics and Artificial Intelligence (AGMAI) are an aspect of the recent AI-in-math drama that deserves much more scrutiny. Its [opening demand](#) was specifically about *proprietary models inaccessible to the broader scientific community*. That qualification matters: AGMAI did not ask AI companies to stop all mathematical research. It also advocates wider access to models. The AHM’s opposition is broader, but it is not accurate to say that all its members reject every use of AI: abstention from AI research is a commitment of its optional [AI-free caucus](#), and its [FAQ](#) explicitly contemplates members disclosing AI use. I find it a bit suspect that the AHM is suggesting that AGMAI would provide “legitimacy” to OpenAI’s efforts (I hasten to point out that OpenAI never “claimed” to “derive its legitimacy” from anything), while the AHM and AGMAI appear to have quite distinct — and often conflicting — advice on the use and role of AI.

Still, I feel queasy about nine mathematicians — or, for that matter, any group of mathematicians — telling a private institution that it should stop using its own products for mathematical research because others cannot yet use them. AGMAI is an advisory panel, not a licensing body. Indeed, it [expressly says it does not speak for the entire mathematical community](#). Of course, its members are

entitled to offer advice — as just about anyone is. My specific objection is to treating that advice as though it were an established norm — as if AGMAI were somehow an official arbiter over the “legitimate” participation in mathematics. That still looks like gatekeeping to me. As if “advanced mathematical problems” were ours to allocate!

I also feel queasy about the possibility — though this is speculative — that AGMAI was given advance access to some of the specific results or proofs announced on October 6, 2026. I would strongly disapprove of a particular novemvirate of mathematicians — even one consisting of such distinguished personages — being granted privileged access to these results. My preference is for the *entire community* to have prompt access, rather than for disclosure to depend on a favored nine.

To be clear: I hugely revere the nine individual members of AGMAI. Consulting them may well have been a good-faith effort to gauge and assuage the mathematical community’s concerns. I also suspect that OpenAI hoped to curry favor with the community and placate some of its critics. But seeking advice created no obligation to follow it. OpenAI did not need this group’s approval to conduct mathematical research, and the public record describes consultation, not a request for permission.

So no, in my view AGMAI had no authority to confer or withhold the legitimacy of OpenAI’s mathematical work. I object to the suggestion that disregarding AGMAI’s preferred conditions for research makes that research illegitimate.

There is a broader problem here, and in my view it runs throughout the AHM’s statement: the treatment of mathematicians as some sort of special priesthood. What alarms me is the implication that these problems are somehow *ours*, and *only we* may decide who gets to try to solve them and how. *Only the clerisy gets to read and utter the holy scriptures!* I should hardly have to explain how profoundly anti-scientific this attitude is.

Next:

“In ignoring the central premise of the Advisory Group’s position, OpenAI has indicated total disregard for the norms of scientific research — norms that guarantee that mathematics remains trustworthy, ethically researched, and in the public interest.”

For starters, OpenAI had *already reported more than 100 results* when AGMAI was [publicly announced on September 21](#). It would be impossible — even for a company as extraordinary as OpenAI — to *retroactively* undo the research it had already undertaken with its new proprietary model. Whatever one thinks of AGMAI’s subsequent public demand to stop such research, work already completed cannot constitute a refusal to follow that later demand. As for dissemination, AGMAI itself recommended that significant results be responsibly released as soon as possible. Making these already-existing results publicly available seems to align rather well with that principle. So it would seem to me that it is very shaky to assert that OpenAI “[ignored] the central premise of the Advisory Group’s position.”

(Mind you: I do very much hope that, going forward, OpenAI *does* ignore AGMAI’s advice. I would like to see it solve a great many more outstanding problems; and in my view it ought to be entitled to use whatever proprietary models it develops in the future.)

I — and, I believe, the mathematical community writ large — have reason to be *tremendously grateful* that OpenAI released these results. It was under no obligation to spend resources in helping to answer questions of niche interest to us mathematicians.

In any event, declining to abide by a freshly proposed restriction on proprietary-model research in no way evinces “disregard for the norms of scientific research.” AGMAI was not elected, and it does not get to decide the norms of our field. Its recommendations should stand or fall on their own merits. It should take more than [a fortnight](#) for newfangled recommendations to settle into established scientific norms.

Then we have the suggestion that following this restriction would somehow “guarantee that mathematics remains trustworthy, ethically researched, and in the public interest” — and the concomitant implication that failing to do so imperils the trustworthiness, ethics, and public interest of its output.

“Trustworthy?” [OpenAI’s October 7 update](#) reports that the main results of 300 of the 719 remaining manuscripts have been *formalized in Lean* — about 42%. I mean, look: I trust Peter Scholze, but I *really* trust Lean — fully rigorous proofs where every step is verified by a simple computer program. For the mathematical claims actually established by those formal proofs, I regard their trustworthiness as higher than that of almost any other mathematical output in the history of our subject. And even for the rest of the release — is there any reason to believe that these outputs are untrustworthy? The fact that the model OpenAI used is internal seems to have had no adverse impact on the trustworthiness of the Lean-verified results; a good Bayesian would reckon that a system capable of generating a Lean-verified proof of the quasi-Riemann hypothesis should be quite trustworthy overall.

“Ethically researched?” If the AHM is going to allege an ethical violation — e.g., plagiarism, fabrication, or misrepresenting the correctness of a proof — I would demand specific evidence or a specific claim. Insinuation that failing to obey the dicta of AGMAI is unethical does not cut it. I do not accept that using a *proprietary model* is, in itself, an ethical violation — or likely to lead to ethical violations. The request that a private institution *stop using a mathematical tool because others lack access* is a different sort of demand than those concerning the honesty, attribution, or correctness of its results.

And to be clear, I would just as vociferously object if a similar demand came from the American Mathematical Society or the European Mathematical Society rather than AGMAI or the AHM. I do not think any mathematical organization has the authority to proscribe a private institution’s otherwise lawful research tools — and it certainly cannot do so by some loose allusion to “ethics.” Unequal access is a potentially legitimate concern — though I myself have frequently felt unequally able to access deeper mathematical truths when comparing myself to my more illustrious peers and forebears, even in the absence of their possessing frontier AI tools. For instance, I have always resented the fact that I have not had open access to the mind of Galois. But unequal access to the tools generating a result does not, in and of itself, taint a correct, honestly presented theorem. Where the proof is public and can be checked, the proprietary status of the tool that found it is not a defect in the mathematics, or of ethics. And to me there is no substantial difference between demanding that OpenAI publish access to its private models and demanding that Thurston or Gauss give public access to their private modes of thought. (At least, in the case of Thurston, an honest effort was made to publicize his private thinking. The same cannot, alas, be said for Gauss. But I do not accuse Gauss of some ethical violation because of his reluctance to do so.)

“In the public interest?” What is “the public interest,” and who gets to decide what it is? I don’t think it is really germane to cite the “public interest” when so much of our work is of esoteric interest to such a small group of individuals. (This is not to belittle mathematics — I view our work as profoundly important; I simply reject the frame that it is of importance because it satisfies some “public interest.”) I suppose that mathematical research proposals are indeed assessed for broader societal benefits — the [National Science Foundation’s funding criteria](#) provide us with one example. But in what world is telling a private institution that it should not pursue transformative research with its own tools “in the public interest”? In what world is the public better served by leaving a mathematical question unanswered because the available means of answering it are proprietary?

So, no. Trustworthiness, ethics, and the public interest may be laudable goals, properly defined. But bowing to the demands of AGMAI — in particular, its request that OpenAI stop attacking advanced mathematics problems with its proprietary models — would not have made OpenAI more likely to serve these interests.

“Mathematicians have a particular vision of progress that is informed by history and field-specific considerations.”

Again, this is far too broad, and the AHM has no right to speak on behalf of all mathematicians. But also, what does this mean? I personally love mathematical history; I would reckon I have spent a good deal more time studying the classic works of, e.g., Gauss, Dirichlet, Galois, and Hilbert than almost everyone in the AHM. I grant that what has been considered a breakthrough may change over time, but one thing has always been constant: resolving longstanding open conjectures has always been, and will always be, progress.

“We reject OpenAI’s assertion that this release advances our subject...”

This is an abuse of common sense and the English language.

Folks, *OpenAI has proved — and Lean-verified — the quasi-Riemann hypothesis*. OK? We now *know* that there are no zeros of the Riemann zeta function to the right of $\text{Re}(s) = 7/8$: this is explicitly covered by the [formalization’s scope](#). If that isn’t an advance, then the notion of advancement has been conceptually gerrymandered to such a degree that the word cannot mean anything.

I want to be somewhat charitable here: it is true that some solutions lead to greater understanding, new techniques, new ideas, while some solutions to problems simply end a story. It is obviously too soon to tell which class of solutions the OpenAI releases might belong to. We know, for instance, that the stories associated with three of the Millennium Prize Problems remain very much *not* complete. The full Riemann hypothesis, Hodge conjecture, and Birch and Swinnerton-Dyer conjecture [remain outstanding](#); it seems likely that some new conceptual breakthrough will be needed to resolve these rather than a virtuosic technical display. Rumors of the death of open problems on account of AI tools would appear to be much exaggerated.

It is also true that many AI outputs — and perhaps a good number in the present release — have the distinct quality of “slop.” But it simply strains credulity that there are no beautiful ideas, techniques, etc. buried in these new OpenAI documents, even if they are somewhat inhumanly written (and so getting to those beautiful core ideas might require some work). In fact, I expect almost *all* of them to have beautiful pearls, which we, as a community, need to extract, polish, and display. And, maybe, repurpose elsewhere!

I actually agree profoundly with those (like [Bill Thurston](#)) who have claimed that mathematics is about the human understanding of mathematics and that advances occur when our understanding improves. I will return to this point later. But what I fail to see is how a release initially comprising 722 manuscripts in 372 families, with many main results formalized in Lean, on some of the most extraordinary problems in our discipline, can do anything but *help* human understanding. Perhaps these manuscripts do not do so as efficiently as less “sloppy” human papers (though, to be honest, I have encountered more than my fair share of unreadable papers in my life and recently have been thrilled to be able to have an interactive discussion with AI while trying to make sense of them). But the arrow only points in one direction here, and that is the direction of *advancement*.

Continuing:

“... and we urge mathematicians and the public to view the value of this publication model with due skepticism.”

What is the publication model being offered here? Posting results and making them available for scrutiny? The primary procedural difference I see here — at least the only one I object to — was the interregnum in which OpenAI consulted with AGMAI. I do not mean that the scale of the release, the unreleased model, or the varying verification status are irrelevant. But human papers also vary in readability, and human authors also need to improve exposition and repair errors. Those are reasons for us mathematicians to roll up our sleeves and do the work of parsing these documents. The publication model here seems otherwise to differ in one way from human output: a large number of claims were publicized along with Lean verification. I regard that as an *improvement* on our usual publication model.

And what do you mean, “with due skepticism”? When it comes to AI companies posting their results, what am I supposed to be skeptical of? These papers’ accuracy? For the *Lean-verified claims*, we have an unusually strong answer. (For the others, apparently, scrutiny is warranted: OpenAI’s [October 7 update](#) withdrew three manuscripts and reported corrections to fourteen others.) Should we be skeptical of *the motivation* of the AI companies? Well, no doubt these companies are profit-seeking. But how does that necessarily misalign with our interests in truth and human understanding? Human mathematicians seek rewards in the form of prestige (a point I will return to below). Yes, OpenAI has a business interest in demonstrating its model’s power. But that creates an incentive to produce sound mathematics of interest to the mathematical community, on problems that have proved difficult for us to resolve. That seems perfectly well-aligned with our community’s interests to me.

I also do not believe that OpenAI is engaging in mathematical research solely out of raw profit-seeking selfishness. I suspect that many researchers within OpenAI and Anthropic *are* doing this out of genuine curiosity — a desire to see these conjectures resolved. And there is a serious body of

first-rate mathematical work now built up to show for it! Commercial incentives and curiosity need not exclude one another.

I go even further: by making this work publicly available, OpenAI — and other AI companies producing significant mathematics — has done an immense service to the mathematical community. That does not mean they receive no commercial or reputational benefit. It does mean that we can benefit too. And so: we should be deeply grateful.

“Releasing over 700 files at once is not a demonstration of scholarship, but a demonstration of power.”

¿Por qué no los dos?

“We urge mathematicians to discontinue their work with OpenAI and to return to a vision of science that centers human understanding.”

I find this to be presumptuous, overbearing, proscriptive, and, frankly — coming from a group containing such eminences — vaguely threatening. Please stop dictating to others how they should operate as mathematicians and whom they ought, or ought not, to associate with. Others’ use of tools you don’t like to solve problems you want to solve for yourself does not harm you.

I too value human understanding. And I have the bona fides to prove it: I have invested a great deal of time immersing myself in the study of mathematical history, and I have engaged very actively and creatively in seminars/teaching/mentoring. In my experience, these are precisely the things given *too little weight* in our community when it comes to getting a prestigious job at a top research school. We should contemplate whether our revealed preferences match our claimed values.

I do not see how making these results available, in itself, diminishes anyone’s mathematical understanding. Our stock of publicly available, verifiable mathematics has increased. The work of developing our understanding remains, and I expect it to advance as we get to work deciphering these new AI-spawned revelations.

The fact that so many are acting as if the advent of AI is apocalyptic for mathematics is horrifying to me. It will not inspire confidence in those contemplating going into mathematics.

Concluding thoughts

As I have said, it is not as if mathematicians are out of a job. We still have to unpack this new work. And doesn’t the fact that there is still much work to be done in deciphering these documents mean that we as mathematicians... have work? It seems that some would want us to be upset that AI has done our work for us — but also to be upset that AI has left work for us to do.

In fact, one can hold both positions without surrendering to incoherence, at least if one factors in the relative *prestige* associated with the two kinds of work.

What’s actually going on here is that we, as a community, have had somewhat misaligned values. I think our hypocrisy is now being exposed. If indeed we had *really* cared about the “human” understanding of mathematics, then the important work carried out by the paper referees, text-book/expository paper writers — not to mention the teaching faculty, high-quality mentors, etc. — that our community depends upon to *transmit* understanding would have a lot more prestige. But rather, we have consistently valued the point of *discovery* more highly than the points (plural) of dissemination.

It takes a huge amount of work to decipher difficult new mathematics and disseminate it to a broader mathematical audience — and also, it is worth mentioning, this is just as true about mathematics that is *not* new but must nevertheless be learned and fully absorbed by aspiring mathematicians. Yet the teacher, the expositor, and the mentor have usually been valued a lot less than the genius who solves a problem.

Now that large language models (LLMs) can solve extraordinarily difficult problems on our behalf, the value of such genius has been lessened. This is just a fact of life from now on. People *will* use AI, the AI *will* beat us to resolving hard problems, and it will become increasingly difficult for us to discern which mathematics has a purely human provenance versus which is either AI-assisted or AI-written (or even depends upon or uses mathematics which at some point was generated by an LLM). It will soon be as hard — and as absurd — to quarantine the output of LLMs from mathematics as it would be to quarantine the mathematics of Euler.

Yes, the mathematical problem-solving task has been mechanized. But that does not diminish the dignity of mathematics or mathematicians any more than a tractor diminishes the dignity of agriculture or farmers. While it seems only human for individuals who are phenomenal problem solvers to feel their intelligence is diminished by AI, it is not rational.

Humans are not in competition with their tools, but with each other. Do we still award a Fields Medal to someone who used AI to solve a great math problem? Do individual mathematicians get the kudos of publishing in “top” journals if they seriously utilized AI? Whom do the top math departments hire, if most of what we consider to be “genius” is generated by AI? It is understandably human questions like these that seem to me to be driving much of our community’s angst about AI, rather than concerns about how AI is affecting human understanding.

We have, in my view, been far too blasé about fostering human understanding heretofore. We should require ourselves to be honest about how much of our newfound concern is for “human understanding,” and how much is anxiety about our own place in the profession.

The concerns listed above — publication credit, jobs, Fields Medals — have a common thread: prestige and, more fundamentally, ego. I suspect that much of what is troubling us concerns who gets to publish in top journals, who gets the prestige jobs, and who wins the Fields Medal. In other words, the hierarchy and relative respect we have towards one another. We ought to be honest about that.

And here’s my suggestion — not my public urging, not my demand, but my honest suggestion: maybe we should actually follow our precepts. Our community should start to value, study, and revere our history. We should not treat great mathematics of the past as if it were only of peculiar antiquarian interest. We should start to esteem more highly those who help in the internal mechanism through

which our knowledge and understanding — tacit and explicit — are communicated. We should start working out how students can use AI — to explore, to be mentored, to help clarify, and maybe to assist — in such a way that their own understanding does not get stunted. We should throw ourselves more passionately into the conferences, seminars, and department teas through which so much of our own human understanding is conveyed.

Ego and ambition will matter less and less to the math community. We should cease to imagine that our telos *qua* mathematicians is to win glory through resolving some famously hard question. *Rather than being driven by ambition, we must start to be driven by curiosity.* And once we clarify that it is our curiosity that needs to be scratched — rather than our ego stroked — the AI systems, in their extraordinary capacity to help us arrive at the truth, will be seen to be our allies, not our foes.

My fellow numerati: if we continue along these lines, we will embarrass ourselves utterly before the world. We must accept that we do not get to freeze in amber how mathematics is done or who/what may participate in it. We may not be as quick as AI to answer our questions, but it would be very unintelligent of us to slow it down in a vain effort to catch up. Comments like those of the AHM are going to profoundly undermine our credibility at a time when we should be putting forward our most confident optimism. And this is indeed such an exciting time! World-historic problems are being solved — and a whole new tool is being introduced to the mathematics community, one which can so readily help us to scratch that curiosity which is, ultimately, the great motive behind mathematics.

Sincerely,

Aaron Slipper

Phillip Griffiths Assistant Research Professor

Duke University

Original AHM statement

Association for Human Mathematics Communications Working Group, October 7, 2026. Source: [AHM statements page](#); reposted by [Terence Tao](#). The statement below is reproduced as supplied, without copyediting. “Yesterday” refers to October 6, 2026.

Yesterday, on October 6th, 2026, OpenAI — which is currently defending lawsuits against accusations of illegal plagiarism, copyright infringement, and trademark dilution — released a repository of manuscripts purporting to contain solutions to a number of high-profile problems in mathematics.

Mathematicians did not ask for this work to be done. The Advisory Group on Mathematics and Artificial Intelligence, from whom OpenAI has claimed to derive its legitimacy, opened their initial advisory statement by saying that frontier AI corporations should not test advanced mathematical problems on internal models. In ignoring the central premise of the Advisory Group’s position, OpenAI has indicated total disregard for the norms of scientific research — norms that guarantee that mathematics remains trustworthy, ethically researched, and in the public interest.

Mathematicians have a particular vision of progress that is informed by history and field-specific considerations. We reject OpenAI’s assertion that this release advances our subject, and we urge mathematicians and the public to view the value of this publication model with due skepticism.

Releasing over 700 files at once is not a demonstration of scholarship, but a demonstration of power. We urge mathematicians to discontinue their work with OpenAI and to return to a vision of science that centers human understanding.

Association for Human Mathematics Communications Working Group