

LETTER AI

CHRISTIAN D'ELBÉE

I write this letter as I watch the community of mathematicians contemplate immense changes in the world it evolves in. The coming of LLM able to solve math problems is certainly a gamechanger, a revolution in the mathematics community. It is also a revolution in the relation between mathematics as a field of knowledge and mathematics as a human activity. This letter contains questions, thoughts and tentative predictions. It is aimed at myself, as a way to let off steam, an « exutoire » if you wish (and know a little bit of french). I might share it with a few colleagues. It is mainly based on my experience as a member of the community of mathematicians, on a few superficial readings about AI news and on the gathering of opinions from colleagues. I try to never miss the occasion of asking a fellow mathematician what she or he thinks about the current state of affairs and her or his feeling on what is about to happen.

I made a point not to ask any LLM to correct the current document, nor to give insight on its content. I do not have particular reasons for doing so, except maybe to keep holding on to my panache. Consider any typo and grammatical mistake as a guarantor of my humanity.

1. ON A SINGULARITY ON MATHEMATICAL KNOWLEDGE

A simple extrapolation of the current state of affairs leads very quickly to a slightly cataclysmic perspective. I have met some mathematicians that trust that we live a transition phase, where all « low hanging fruits » will be proved by AI, elevating the general complexity of all open questions in the field. Then will come the phase where all « easy » open questions have been answered, and what remains are the real, new open questions, the new frontier of knowledge. Mathematicians will then work, with the help of AI, on solving those questions.

I do not share this view, at least as long as the capacity of LLMs keep improving in the way it does currently. I have several reasons to think that, given the current perspectives, our relation with mathematical knowledge will drastically change. It seems to me that mathematics as a body of knowledge, will essentially become by definition: *every statement that some LLM is able to answer in a relatively short time*. I see that arising when LLMs are both sufficiently smarter than they are today and sufficiently mistake-free in their mathematical reasoning. In that case, mathematical knowledge will be reduced to asking the LLM whether a mathematical statement is true or false. A fast answer is the state of knowledge, a slow answer is new research being happening, although I am not sure this distinction will still be relevant. The pivotal change is when a seemingly harmless question will receive a fast answer (so it is in the state of knowledge) which will be extremely long, say formalized in a draft of hundreds of thousands of pages. A triviality, a blink of an eye to read for a sufficiently powerful LLM, but the work of generations of human mathematicians to comprehend and verify. At this point, even the most advanced mathematical statement is *folklore* for the LLMs. Will mathematicians accept their fate as inefficient, mistake-full proof-checking entities? What role could a mathematician have in this perspective? Will mathematicians turn into explorers and cartographers of this infinite ocean of knowledge, turning LLM answers into scientific papers and research monographs to fill our university libraries? Will the society pay mathematicians for this endeavor?

The acceleration of the speed of the production of new knowledge will contribute to this state of fact. First, the available body of knowledge will become levels of magnitude harder to comprehend for humans than it is today. In parallel, LLMs become more able to deal with this large amount of newly produced knowledge. The acceleration of production will also lead to the destruction by saturation of the peer-reviewing publication process by lack of human brainpower to deal with it and/or the reduction of it to LLM-verification including by proof-assistant component, which leads to more sovereignty given to LLMs. This leads to two questions. The first is obvious: Why publishing a few pages of result if an LLM gives instantaneously the full argument upon request? If we reach this point, why bother to include human minds in the process at all? The old-world publication process was already cumbersome, slow and old-fashioned. Its role as quality stamp for mathematically correct preprints is even more questionable in a world with fast, mistake-free LLMs. Its role as a library of mathematical results is also useless with an all-knowledgeable LLMs. The argument « LLMs will always make mistakes, we must have a human verification in the process » is relying upon the fact that LLMs make *more* mistakes than humans, which I am not even sure is true for recent LLMs and will certainly be false in a foreseeable close future. One component

of our publication process seems entirely cultural and the existence of journals for *papers of particular interest* might of course survive, but will only consist of an insignificant part of mathematical knowledge. With an ocean of unequal and spread out mathematical results, LLMs will be the most reliable and reasonable mean of consulting the literature (which should certainly be called *elsewhat*). Furthermore, at the point where LLMs will not take much longer to give a 10 lines new proof than a 10 000 pages long new proof, it seems inevitable that LLM will be not only our best access to the literature, but rather will *become* the literature itself, and the full body of mathematical knowledge.

This considerations, of course, raises the question of the very existence of mathematicians in the society where we have reached this point. Some of my fellow mathematicians have opposed the following hopeful perspective: « the society will always need people to understand mathematics, therefore the need for math teachers will never run-out ». Pushing the perspective of the above singularity on mathematical knowledge a bit further, we might even make doubtful the need for humans to understand mathematics as a discipline. One particular fictitious example has been haunting my mind in the last few days. It goes as follows. In a foreseeable future, an AI will have physics, mathematics and engineering at full capacity as well as integrated robotic production line at disposal. The humans then come with the query of constructing a spaceship that can speed up to 5% of the speed of light. The query constitute a challenge in the field of physics and mathematics, but the AIs succeed in finding a theoretical framework in order to realized such an endeavor. The AIs have been assessing the reliability of the new mathematical model and the feasibility of the project and a green light is given by the humans to go ahead and allow production. The spaceship is then assembled by the robotic production line, so that, after a few weeks of intense work, the spaceship is constructed and available to humans to use as they please.

In the current world, the green light would have been given only under the condition that the new physical model have been assessed to be mathematically sound, but the new AI model is formalized mathematically as hundred of thousands of pages. There is no doubts that no human will be solicitation in order to verify by hand the soundness of the model, and it should suffice to run the spaceship (by robots) enough times without any humans by security in order to ensure that it is safe to travel. It could also just be that the AI provide a sound but extremely complex assessment of the reliability of their technology. In this example, there is no need for humans to be in the process of developing the new mathematical model. Even if one human want to understand the model, the complexity of the theory means that she or he might just physically not be able to in his given lifetime. But even more disturbing, there is no *need* for any human to understand the mathematics behind the new model. As long as the final product works and is understood by the AI, that shall be sufficient from a practical point of view, to run and to repare. In this world, the amount of highly-educated and knowledgeable mathematicians shall drastically drop. If the need for mathematically knowledgeable humans will drop in the more advanced fields of technology, it will even more so in disciplines where mathematics are used at a more superficial level, including all scientific discipline. In a world where the very understanding of mathematics is superfluous, why would we need math teachers?

Of course, mathematics as a discipline of the mind has other values than producing new theorems or having a grasp on the state of mathematical knowledge. Mathematics remains a formidable tool for training the human brain in logical reasoning and rigorous thinking, (just as it is used for training LLMs). Further, mathematical elegance provides certain form of aesthetics of the mind, a *poetry of science* as some would put it that is purely to the benefit of (some) humans, and distinct from the activity of proving theorem. The society might choose to keep this treasure as part of its range of activities and allocate a fraction of its wealth to train and provide for highly educated mathematicians. It should at least be made aware, if the time comes, of the evolution of our relation with mathematical knowledge.

2. ON THE DIFFICULTY TO PREDICT

It seems that few people only have been able to anticipate what was really going to happen with AI and mathematics. One of the reason for that might be a measurement problem. On the one side, the only people able to assess and contemplate what LLMs are really capable of in mathematics are mathematicians. On the other side those mathematicians are also those which are filled with the fear of being replaced by those agents, and more so as LLMs become more efficient. Thus a bias has been to first try to identify what we mathematicians can do and what LLMs will *not* be able to do, with the background hope for salvation. For the earliest phase of the LLM-era, that consisted of comments of the form « it will always make mistakes » or « ok, it can gather all the literature and sort out counterexamples to statements by just reading through it, but it cannot make new proofs ». This later became « ok, it can make small proofs by reproducing arguments from here and there but it will never be able to articulate complicated proofs ». This recently became, « ok, it can articulate proofs but it will never be able to make a new definition, or develop a new theory ».

Let me expand a little on that latter comment, which still seem to give some comfort to many mathematicians. I wish it would give me some comfort as well, but it does not, for the following reasons. It comes from my impression on what a *good* definition is, even thought it might constitute quite a challenge to pin down what that

is. In a world where mathematics is the realm of human brains, I see definitions as pivotal objects serving at least two different purposes. It certainly has an *explanatory* purpose. A definition gives structure and make transparent a formal deduction. A definition isolates objects that share certain properties, and gather them artificially in order to serve a given demonstrative purpose. This aspect of a definition is to be compared to what a *clade* is in biology, where groups of living things are gathered under the same name in a world where only individuals exist. It has a rhetorical power and is appealing to the human brain but is certainly just a toolistic purpose, emerging from to the limits of human capacities. It has a fortunate second effect, probably crucial in the development of mathematics. The second purpose of definitions concern the development of new ideas. Thanks to the appetency of the human brain for the elegance of an on-point definition, it is often use as a platform for the emergence of new ideas. Definitions are stimulating the emergence of new ideas in an cultural and organic way. Often we hear: « definitions are more important than theorems » or « I finally found the right definition » or even « if I give the mathematical community one *good* definition, I can retire in peace ». In that regard, definitions are certainly more than just mind-shortcut to clarify and simplify our argument. They have a cultural and stimulating function within the community of mathematicians.

It seems to me that both purposes above are superfluous in order to practice mathematics in its (perhaps restrictive) sense of *proving theorems*. The first purpose relates to a purely computational subterfuge and is certainly useless in a computer program that can call to all the current definitions of the literature in the blink of an eye, and unravel those cumbersome definitions to deal solely with properties. The fact that we, humans, may construct an *intuition* on the *behaviour* of mathematically defined objects is certainly nothing more that unconscious brain reflexes and habitus that are also instantaneously executed by LLMs or largely compensated by brute force. The question whether LLMs have intuition is certainly ill-founded and I do not care about that one. The question whether one needs intuition in order to invent proofs is another one, and I am not even sure that it is important. The fact is, whatever the LLMs « have », it is sufficient to prove new theorems. As for the second purpose, it seems restricted to the cultural or organic relations that human brains have with the abstract concepts that they developed, and I do not see any reasons why the activity of proving theorem must rely on that aspect. The fact that the history of the current state of knowledge of mathematics relies on that cultural and organic aspect does not mean that it is necessary in order to keep the activity. The fact is that LLMs are created with at disposal the current state of knowledge of mathematics, and build on top of it.

Do LLM need to isolate new definitions in order to prove new theorems? The answer might simply be no. Certainly they will need to come up with definitions or other subterfuge in order to explain their proofs to the human brain, and there is no doubts, having being trained on thousands of human-brained produced articles filled with inequal definitions, that they will (and in fact they already are) able to do so. In fact, being *language* models, they excel in producing content adapted to the understanding of the human brain.

In the sentence « ok, it can articulate proofs but it will never be able to make a new definition, or develop a new theory », I consider *new theories* to be a scaled version of what a *new definition* is, and everything above applies also to that notion.

3. ON A QUESTION THE SOCIETY HAS NOT BEEN GIVEN THE RIGHT TO ANSWER

There are two aspects in which the lack of input from the society in regards to the present situation is flagrant. The first is quite obvious. LLMs have been imposed to our society without any form of consent. In November 2022 the first public LLM is released, free to use by anyone, with the only cost of creating an account. For such a game-changer in our civilizations, it would have been preferable that the choice be given to the society: do we want this technology in our lives? The fact that it is possible to do should not be a criterion for realizing it. First, we need to ask ourselves: what does our society needs now, and then we ask ourselves whether we may realize it.

But the real point of this section is not really about that particular aspect. Two worlds (or two states of advancement of the same world) were described above. In both of them, the activity of doing research in its entirety is outperformed by LLMs on all levels. LLMs *constitute* the current body of knowledge, mathematicians can ask mathematical questions, they can understand the truth value of the answer, and might sometime not have a full access to proofs, by lack of time or lifespan. In the first world, mathematicians may have found a role to play, explorers or cartographers of the infinite land of knowledge offered by LLMs. Mathematicians become field biologists that perform a cladistic work on the available knowledge, drawing from the source and producing articles and monographs. In this world, the need to understand mathematics is real, and teachers of mathematics are still an important part of the society. Activities such as outreach and diffusion of knowledge would still be meaningful in this world. However, the research part of the role of mathematicians would be restricted to this editorial endeavor. The highly regarded *theorem-proving* –one of the hardest and well-known part of our work– would become completely obsolete. It will then be time for the society to ask itself whether its tax money is well-spent for supporting and maintaining a handful of beautiful minds playing with concepts their barely understand, and trying to keep up, like dogs running after a Lanborgini.

In fact, this question is long overdue, as I am not sure the public opinion is aware of the inner workings of our community of researchers in mathematics. It is only fair that the question be raised, that could be one of the few beneficial effect of this paradigm shift.

If teaching activities are safe in this first world, it is not so clear in the second, more pessimistic world. In this second one, mathematics is but one component of an all-powerful mechanistic system, where our proximity to the science behind technology has completely vanished. In this world, teachers, pure and applied mathematicians would have disappeared. The spaceship example of the second world might suggest that only applied mathematicians are useless, but I would not imagine a world where applied mathematicians disappear and pure mathematicians survive. Already, in our contemporaneity, it seems sometimes that pure mathematicians rely on the survival of applied mathematicians to survive.

Do not get me wrong. I am glad to be a mathematician, I believe in our community whose functioning is admirable, and more virtuous than most of other job systems in place. This only might justify its very existence. The free space for knowledge and exchange that we have created, on the shoulders of giants, makes me proud. The pure power of abstraction, pushing the deepest limits of the human mind, melting rationality, aesthetic and intuition into an ineffable work of art makes me flirt with transcendence. It is, in my opinion, among the greatest achievement of our humanity. But it is true that our distance from public opinion often makes me uneasy and I consider that the legitimacy for doing what we do should not be taken lightly.

4. SOME OPEN QUESTIONS

Here are a few general and open questions that each community of mathematicians may ask themselves.

Question 4.1. (For tomorrow) In the world of tomorrow, LLMs might have acquired all published mathematical knowledge, immensely powerful reasoning capacities and with a ratio for errors far better than that of our all of human brains together. In that world, the *state of knowledge* might just become whatever the LLM can answer in a few minutes or a few seconds. The pivotal change, to me, is when, for a seemingly harmless question, the answer, quickly given by the LLM, consist of a proof which cannot be written in details in less than hundreds of thousands of pages. If this premise is acceptable, what role could mathematicians endorse?

- (1) In that pivotal moment, no single human brain, nor an entire generation of mathematicians might be able (or willing) to check the veracity of some proofs, or of the lean code. (This feels like a scaled version of the CFSG problem.) Nonetheless, we can assess that the probability of the proof to contain a mistake is almost 0. What should we do with this sort of statement? What if, in fact, most of mathematical theorems have such an out-of-reach solution? Should we consider those statements unproved, in spite of the fact that we have a proof? (Statistically correct, but again, aren't all our proof statistically correct?) In an RSA encryption which relies on our incapacity to factor out composite numbers in a short time, don't we consider the prime factors to exist? If we can't understand the proof but do have a proof, shouldn't we consider the theorem to be true? Isn't one of the main lesson of logic to accept that a theorem is true as soon as there is a proof?
- (2) The activity of theorem-proving is certainly useless and might turn out to be more costly for the society when done by human researchers (which takes years of training, need salaries, offices, and coffee!). The society could decide that people are free to prove theorems, but not by using public money, just as personal hobbies are not supported by taxpayer's money.
- (3) Mathematicians might become archivists, or field biologists, that perform a cladistic work on the available data, drawing from the source of knowledge (LLMs) and performing a classification, grouping results in theories, producing articles and monographs, to organize this immense knowledge and make it available to the human brain. Although I do not see why, in this endeavor, we could not also be outperformed by LLMs, it might have some quality and value to have it made by humans.
- (4) The latter point also suggests that, with less time spent on correcting small lemmas, we could find peace and meaning by transmitting this new and incredible knowledge to the public, and give back to our society. Outreach might be a way of salvation and an asset for justifying our existence.
- (5) If this world still values the understanding of mathematics, teachers of mathematics will still be needed.

Question 4.2 (On the value to give to human brain sustainability). By (*human*) *brain sustainability*, we mean mathematical research activities that have had few or no (*AI-free*) input at all from an LLM, when it comes to mathematical proofs. While it might seem reactionary for some, adopting such practices might serve various purposes, among which:

- easing the mind of our current students that have a yearning for writing their first papers by themselves,
- giving justice for those who genuinely refuse to use LLMs because none have been given the choice to have those tools imposed to them,

- it is in the name: sustainability, because the risk that exposing our brain to a fast-answering oracle might lead to a loss of brain-sovereignty is not to be taken lightly, especially among students.

The co-existence of AI-full papers and AI-free/brain-sustainable papers is not unimaginable. Just as survey papers are considered in a special category, brain-sustainable papers could also be their own category. This co-existence will need answers to the following questions:

- (1) Is there a will of the community to give any form of value to such a brain-sustainable practice of theorem-proving? Initiatives such as the Association for Human Mathematics (AHM) seem to indicate yes.
- (2) Should some journals adopt AI-free or brain-sustainable editorial guidelines? Those could be brain-sustainable-only journals, or a brain-sustainability labels for papers.
- (3) How to assess the authenticity of a brain-sustainable paper? Should we put trust as the core value of our community of people? or should we ask each mathematician to openly publish all of their mathematical prompts?
- (4) How to compare an AI-full paper proving a longstanding open question and a AI-free paper proving a locally important result?

Question 4.3. (On the evaluation of quality)

- (1) How to assess the quality of mathematics in a given AI-full paper?
- (2) How to make the difference between an AI-full paper which uses AI at its uttermost capacity and an AI slop rewriting?
- (3) Should AI-full paper results be considered valuable as soon as they are not immediately answered by AI?

Question 4.4. (On AI-full practices)

- (1) Is there such a thing as a practice of AI-full mathematical research which is valuable and not just AI slop? If yes, how to assess that it is? Maybe by asking an LLM to solve the particular question, if it cannot, then the paper is genuinely more than just rewriting of AI.
- (2) Will not there be a moment where the distinction between "AI slop" and "AI full" capacity won't exist?

Question 4.5. (Social inequality) How to assess and solve the social issue that is raised by access to LLMs? Some researchers/universities cannot afford to have a paying license.

Question 4.6. (On the changes in teaching mathematical research)

- (1) How should we teach our students to do research? Maybe just as before, one of the main difference is that the level on which we reach genuinely new results is far more advanced.

Question 4.7. After endless fights of generations of mathematicians for having the largest amount our knowledge made public and free ; the LLMs are trained with and have full access to almost all of our literature. Our very branch of logic has in its origin the wish for an all-powerful machine for computing deductions and proving theorems. Now this machine exists and is certainly depriving many of us of the joy of attacking an open question, diving into dusty books of a unknown theory, slowly scratching our beautiful abstract formalism on a piece of paper, maybe under a green oak-tree and lulled by a fresh breeze coming down the mountains... Isn't this ironic?

UNIVERSITY OF THE BASQUE COUNTRY, DEPARTMENT OF MATHEMATICS (LEIOA) / INSTITUTE FOR LOGIC, COGNITION, LANGUAGE AND INFORMATION (DONOSTIA-SAN SEBASTIÁN), BASQUE COUNTRY, SPAIN

Email address: christian.delbee@ehu.eus

URL: http://choum.net/~chris/page_perso/