All the math we will not see

A. Della Corte

Imagine that in Babylon, 600 BC, they had today's AI technology: deep learning, transformers, GPUs. The training data would consist only of pre-Greek mathematics, like ancient Egyptian and Babylonian mathematical traditions. I doubt that anyone, given only those ingredients, would believe that this pre-Greek-AI could autonomously produce Archimedes' results on the stability of equilibrium of paraboloids or the solution to the Brachistochrone problem.

Now shift the experiment to London, 1550 AD. Again, I suspect very few people would expect a pre-Newton-AI, trained only on the mathematics available then, to produce, say, C*-algebras, ∞-categories, or set-theoretic forcing.

In these cases, the task would have been not (just) “too difficult” for AI. It would have been enormously improbable because mathematics is not an extrapolation from prior state, it is an unpredictable cultural/historical phenomenon.

But shift the experiment to today, and suddenly many people are confident that AI could one day replace human mathematical work altogether.

To me, this confidence simply reflects our ignorance of what future mathematics will look like and our over-confidence. The thought experiment suggests that, whenever we already know the future trajectory of mathematics, we become much more cautious about believing that it could have been extrapolated from the past. Yet when we stand at the frontier ourselves, we are tempted to assume that the distance from today's math to tomorrow's is smaller than it actually is.

Part of the confusion may come from a rather systematic misevaluation of what drives mathematical evolution. Working mathematicians spend a very large chunk of their time studying other people's papers. It is natural to believe that what one does most of the time is what matters most. But the daily practice of reading and building on existing work makes it easy to forget the big picture of historical and cultural conditions that guide the development. So let us try to separate (approximately, just enough for our purposes) what actually goes into mathematical work.

Mathematics is produced using, basically, three ingredients:

A: pre-existing mathematics: the mathematics that already exists, has been recorded and is available and known to the subject producing mathematics.

B: external stimuli from the real world: sensory experience, spatial intuition, practical needs (cartography for differential geometry, heat power for Fourier analysis, etc.), technological novelties (think of the math born out of QM experimental discoveries), particular phenomena one wants to understand or control (population dynamics, weather forecasting, transmission of information, etc.).

C: some internal labor of the subject that produces it.

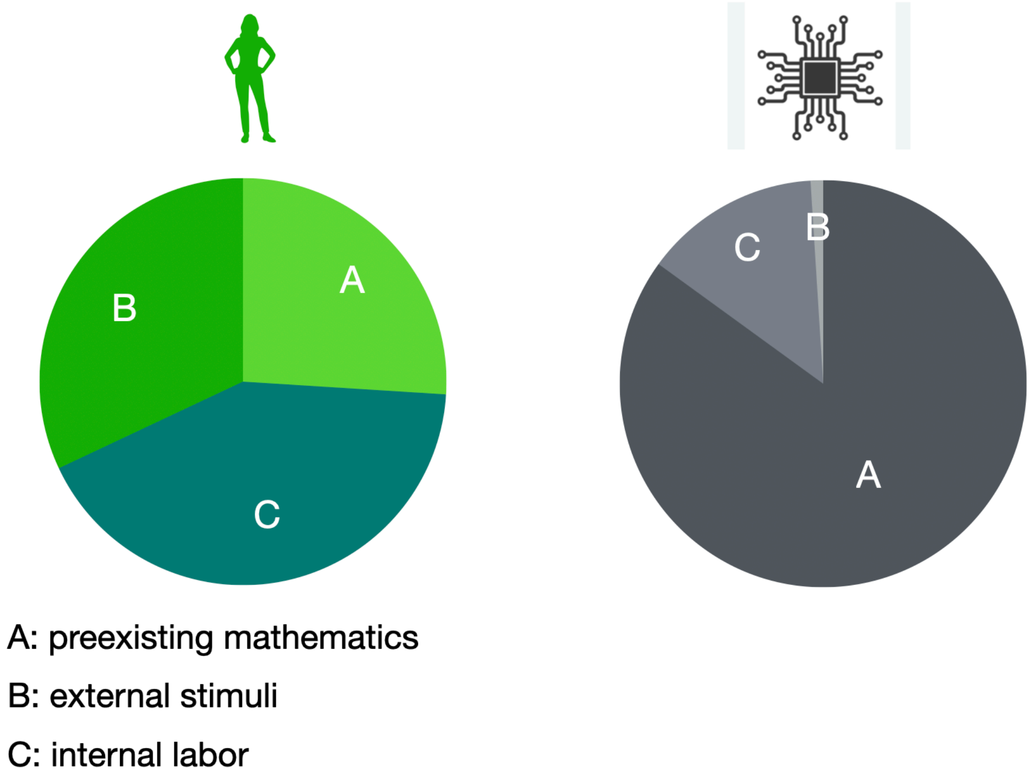

One plausible division of the ingredients of mathematics for people and AI tools

It's difficult to compare directly how "large" are A, B or C for mathematician Y in a particular piece of work. But it's easier and more sensible to make a comparison between A, B and C, on average, for mathematicians and AI tools. A mathematician knows only an infinitesimal fraction of A. In the case of AI tools, instead, A grows disproportionately, while B is almost nonexistent. It’s worth noticing that, in either case, we know very little about C. We have only a very vague idea of what the internal labor involved in mathematical creation actually consists of. But if the proportions of A, B and C are indeed so different in humans and in current AI, then C must be doing something very different in the two cases. Moreover, whoever has worked in mathematics knows that, for mathematicians, C is not a disembodied logic engine. When people do mathematical work, C is constantly being shocked, redirected, fertilized by ingredients from B. This is true at the level of the individual mathematician, and all the more so when we consider mathematics across longer time scales and as a collective enterprise.

So perhaps one important question about AI and mathematics (besides well-documented problems like deskilling, dilution of good content, dependence on enormous private companies, environmental impact) is what kind of mathematical developments we will miss if we adopt too quickly, too pervasively, too universally the same kind of tools in every stage of education and research. The main long-term effect of this adoption might be to significantly change the general shape of mathematics (and science at large), from a diversified, unpredictable, historically contingent ecology of ideas and practices to a more homogeneous and path-dependent one. This would be perfectly compatible with the periodic, regular appearance of AI-aided results that, locally, are of the highest technical difficulty and even the finest mathematical quality. And even if one day AI goes beyond within-paradigm technical tasks, and propose meaningful novel paradigms, the mathematics it creates will be its own, not a substitute for the one we would have produced, and that we may well never see.