# Do not forget where you were headed

Or Zamir
*Tel Aviv University*

November 22, 2026

The recent advances in AI-driven mathematics sent shockwaves through our community. It was slightly surprising that instead of joy and excitement, we saw frequent dismay or even a sense of defeat.

I'll proceed by stating the obvious: predicting the future is a notoriously foolish craft. Having failed to predict the future once is hardly a reason to believe we can predict it now. That is, the success of AI in both assisting and automating the process of proving mathematical theorems has been absolutely astonishing. We have obtained many remarkable results, some of which we sought for centuries, and many more may await. But at the same time, the vast majority of questions we care about remain unsolved.

One may try to extrapolate from the pace of success that this is only a temporary state. On the other hand, many alternative explanations are just as plausible. For example, I see not even the slightest reason to believe that the measure of a problem's difficulty is the same between humans and computers. Every undergraduate in math already sees many proofs containing a magical sequence of three technical Cauchy-Schwarzes, why wouldn't some require fifty or a hundred such unmotivated technical steps? And if a problem does have such a proof, why should we expect it to be easy for a human to discover, or hard for a computer to find? Whichever the true reason may be, as of now the vast majority of problems we consider far easier than the Riemann Hypothesis remain unsolved, including many that have resisted attempts by the very same AI companies responsible for the recent breakthroughs.

With that strong reservation in mind, let's suppose for a moment that even the wildest predictions about AI come true. Imagine a world in which a machine can solve every mathematical problem we care to pose. Even in that world, I see no reason to despair.

Much has already been said about the legibility of current machine-written proofs, and their diminished ability to influence the field for that reason. Take, for example, the color-coding paper of Alon, Yuster and Zwick. Its declared result was a fantastic breakthrough: a faster algorithm for finding any prescribed-size cycle in a graph! Imagine instead that our machine directly produced a *deterministic* algorithm for finding such cycles. Why would it bother introducing unnecessary randomization and colorings? It could simply enumerate all functions from some small and peculiar family, then run an obscure deterministic algorithm for each. Such a proof would still have been a great result, and would have been accepted at any of our top conferences or journals. At the same time, the impact of the clearly conveyed conceptual contribution of the color-coding technique on countless subsequent works would have been lost.

The lesson of that should not be to avoid using these tools, but rather to remember that a theorem's finish line is not the existence of any proof, but a good proof that we understand. We already saw instances of useful ideas distilled from initially obscure AI proofs: For example, in the $k$-Clique Hypothesis refutation, Virginia and I used a decomposition lemma we could distill out of the labs' ETH refutation and turn into a useful and self-contained tool. And perhaps future machines will be far better than us at distilling such ideas as well. That would be wonderful. It would only make the mathematics we care about more accessible to us. The more crucial question, which we must not overlook, is not *how* to obtain proofs that humans can understand, but *why we should care about understanding them in the first place?*

The answer is rather simple: because we *do* care. The general population has never heard of the Riemann Hypothesis and frankly could not care less about it. We, the mathematicians, are the ones who care about these questions and their answers. And we do not merely want answers to our questions. We want to understand them well enough to discover which questions are worth asking next. For centuries, we pursued questions whose value often seemed apparent to nobody but ourselves. We developed many mathematical fields, advanced our knowledge in them, and eventually many found irreplaceable use in the real world. Without the long history of mathematics, we would not have most of our physics, chemistry, biology, medicine, and technology. We would not have machine learning itself or the computers to run it on. Hell, we would not even know these questions were possible to ask. The most transformational effects of our work have often come not from answering questions posed to us, but from developing entirely new ideas, emerging from both the knowledge we amassed and our subjective desires as humans.

And that last part is the entire point. Our job never was and never should be being the best or quickest solvers-at-hire for the most funded company at the moment. Our profession always was and remains advancing the boundaries of *human knowledge* as we see fit, even when our objectives hold no monetary or other interest for anyone but us. But this is not merely a sport. This very social contract between academia and the external world has repeatedly proven immeasurably valuable.

We are the ones who hold that knowledge, and who decide when and how to use it to improve our lives. Those subjective choices of worthy applications are incredibly hard to predict, measure, or objectively verify. Some of these attempts disappear as quickly as their promising arrival, and some unlikely ones do end up transforming our lives. AI alignment itself offers a timely example: making sure that increasingly powerful machines serve human interests is an objective that matters to us, not one that an omnipotent machine must inherently view as worthy.

My proposition is thus modest. As long as there are so many problems we care deeply about left to solve, do not despair that solving them might become easier. And even if we someday count every last star in the sky, remember: **we hold the steering wheel not because our car is the fastest-moving object in the universe, but because we have somewhere to go**.