

October 4, 2026

AI in the University: Can We Tame the Beast?

By: Shobhit Mahajan

As AI tools colonise education and research, some of the basic building blocks of innovation are at risk: curiosity, intuition, and procedural memory.

A spectre is haunting humanity-the spectre of Artificial General Intelligence (AGI). There is hardly any field of human activity which is untouched by the prospect of AI agents replacing human beings-from customer service to coding, from drug discovery to making biological weapons, to solving centuries-old mathematical problems. The capabilities of Generative Intelligence (of which Large Language Models (LLMs) are the most well-known examples) are increasing at such a frightening pace that some of their proponents are themselves now arguing for a slowing down of development.

AI tools are already being extensively used in the field of education and research. Among students, its use is increasing rapidly. [A survey in the UK](#) has found more than 90% of the undergraduates are using AI for their coursework. In scientific research, [a 2026 study](#) has estimated that by 2025 the use of LLMs in more than 50% of the published papers. This figure would of course be much higher today. It might be fruitful, though possibly a bit late, to investigate its harmful impact on education and research and explore ways in which it could possibly be mitigated.

Almost everyone would agree that the goals of education should include making a student more aware of the physical and intellectual world, preparing the student for the job-market as well training the student to think critically. The last one is possibly the most important role of education, since critical thinking and analytical skills are required for solving intellectual as well as real-life problems.

How does the use of our current AI tools measure up to achieving these goals? Are they silver bullets that their proponents like to claim or will they prove to be disastrous for creating an educated citizenry? These are the issues which need to be carefully examined.

What Happens to Curiosity?

One can argue that AI in its current avatar of LLMs and chatbots (agents) can be very useful for information retrieval since it provides ready made, digested capsules of information about any topic. This it does because the models are trained on the accumulated knowledge available digitally in textual, visual and auditory forms. While previously a student might have had to comb through innumerable sources to extract information, LLMs provide it much more efficiently.

|| Although AI is good at providing answers to specific questions, its specificity leaves out any scope for serendipity.

The use of short cuts is hardly novel in education. In my own experience as a teacher, I have seen students graduate from guide books to YouTube videos and now to LLMs. So why should one worry about the proliferation of one more 'technology' in education? I think there are several reasons why. As they say about the stock market, this time it is different.

The primary motivation for learning about something new is curiosity. One sees this most obviously in children who are infinitely curious about their surroundings. This sense of wonder and curiosity gets conditioned by our education into a well-defined algorithm where learning is for specific goals which include scoring marks, getting admissions, or having a better chance for a job.

Although AI is good at providing answers to specific questions, its specificity leaves out any scope for serendipity. And serendipity plays an important role in learning. When one is searching for answers the old-fashioned way, one is very likely to come across things one is not looking for. For instance, looking for an entry in a dictionary or encyclopaedia one is likely to see other entries on the same page which might be interesting to explore further. The joy provided by these chance encounters with other bits of information is, of course, missing when one is delivered a packaged answer. What is also lost is the associated learning which comes from this exploration.

As an analogy, consider reading a novel as opposed to getting its gist from AI. The plot summary one gets from the AI tool might be useful to talk about at a party but can it give you the same sense of joy and satisfaction as reading the novel itself? Capsule summaries might be efficient and good for immediate gratification, but cannot provide the insight which comes from exploring and

engaging with the text.

AI and the Job Market

One of the strongest arguments given for promoting the use of AI in classrooms concerns preparing the student for the job-market. The logic is that by the time today's cohort of students graduate, most jobs will require AI and so it is best to train them to use it. There could be some merit to this since we are already seeing some jobs requiring extensive use of AI. Coding for instance is now mostly being done by Claude and similar tools. It is conceivable, indeed probable, that many of the entry-level jobs will be 'AI-ed' away over the next few years. In this scenario, students who are adept at using AI tools will have an edge over those who are not.

While making the student aware of the world and training her for the job market might be important goals, development of critical thinking and analytical skills are central to education.

This argument is, of course, a valid one since students ultimately need to enter the workforce and education needs to prepare them for it. However, it is important to keep in mind the caveat about the broader role that education plays in a society. Are we seriously advocating the vocationalisation of education? Training for the job market is an important part of education but is it something which can override other concerns? And in any case, even for things like coding which could be done with AI tools, one would need to formulate problems which require more skills than just familiarity with AI. It would require knowledge of the field as well as training in problem solving even though the problems themselves are more easily solved by these tools.

It is also a fact that the pace of technological change in the world today makes it very difficult to actually predict what kinds of skills might be in demand in a few years. Did we, for instance, even 10 years ago, envisage some of the jobs which exist now? Or imagine that some of the jobs 10 years ago would be obsolete today?

While making the student aware of the world and training her for the job market might be important goals, development of critical thinking and analytical skills are central to education. Developing these comes with an engagement with the subject which usually takes many forms like being taught by a teacher who has been trained in the subject, self-study, problem solving, and writing essays. The cognitive processes which come into play when the student struggles with a subject are precisely what train her to think. The time that one spends with the problem trains the mind's infrastructure that will eventually lead one to do original work.

Engaging with a subject also builds the intuition which is essential for the development of knowledge. Intuition, based on the accumulated knowledge and trained by critically thinking about a problem is an important component of the process of discovery and enhancement of any domain of knowledge.

Atrophy of the Mind

Extensive reliance on AI could lead to the atrophy of the mind and cognitive abilities. The technological dependence which comes with it can be disastrous and has been called cognitive surrender. Just like our limbs and other body parts can become atrophied by extended disuse, the mind can also lose its cognitive abilities if these are not used optimally. [A recent report from the Brookings Institution's Center for Universal Education](#) discusses the dependence on AI in students and its impact. Students are using AI extensively and in the process losing their capacity for creativity, problem solving and reasoning as well as the ability to focus on a problem. Given that these are precisely the things which should develop during one's education, this should worry us.

The atrophy of any ability from disuse is most apparent in the use of language. Students nowadays are by and large linguistically challenged when it comes to writing in any language-not just in a second language but also in their mother tongue. Lack of reading, extensive use of the visual as opposed to the textual media and writing only short messages and that too in a particular kind of abbreviated language are some of the reasons for this. Many studies have shown that the process of writing itself enhances mental skills. The process of formulating one's thoughts in writing clarifies and deepens the understanding of the subject. This linguistic deficit which is already apparent is likely to become widespread in the future. Students will simply lose the capacity to formulate ideas and thoughts cogently in any language.

Students are using AI extensively and in the process losing their capacity for creativity, problem solving and reasoning as well as the ability to focus on a problem.

We have seen this happen with mathematical skills. Till about a generation ago, simple arithmetic operations were performed easily mentally by most people. This was true whether one had gone through any schooling or not. After all, the street vendor could easily figure out the amount of change that he needs to return to his customer. With the proliferation of calculators and then mobile phones, we seem to have lost this ability. It is not that everyone needs to remember the multiplication table of 19-it is just that one should have some idea of the logic of simple mathematical processes. This is needed more now than before since our world is so technology dependent.

I had this experience with a student once which demonstrated this. An experiment involved the multiplication of two double-digit numbers. As expected, she had punched them into her mobile phone and came up with a few million as the answer. However much I tried to convince her that this simply cannot be true, her belief in technology was so strong that she would not agree. I then told her to multiply the highest 2-digit number with itself and check the result. She again used her phone and got the expected answer. What had happened earlier was that she had accidentally inputted wrong numbers and yet she didn't realise that the result couldn't be correct because she had no inherent understanding about what multiplication meant.

Interestingly, the famous physicist, Erwin Schrodinger had been prescient enough to foresee precisely this danger. [Speaking in 1956 on the dangers of excessive mechanisation](#), he said "I believe that the increasing mechanization and 'stupidization' of most manufacturing processes involve the serious danger of a general degeneration of our organ of intelligence" (*Mind and Matter*, Cambridge University Press, 1959).

Tool for Language Skills?

An argument in favour of the use of AI which can be given especially in our country is about equalisation of disparate initial conditions. With increasing Gross Enrolment Ratio (GER), students from very diverse backgrounds are entering higher education. The socio-economic profile of students has changed significantly in the last couple of decades; there are many more first generation learners now. The socio-economic disparity is not just about unequal school education but also about language. A lot of students find themselves at a huge disadvantage when it comes to English comprehension and writing. And since higher education in our country mostly uses English as a medium of instruction, this deficit escalates into learning shortfalls.

One can argue that using AI can mitigate this issue to an extent. That might be true, but would that impart language skills to the student? And once again, the ability to express one's thoughts cogently comes from repeated practice and engagement, not by taking short cuts. Technology can help. I recall the case of a research student who could not write a single grammatically correct sentence in English. I encouraged him to write the first draft of a paper by himself. The result was, as expected, totally incomprehensible. After several iterations of my corrections and his rewriting, he was able to produce a reasonable paper. He then polished the final draft with the help of Grammarly, a grammar tool. In the process of his writing the several drafts by himself, he developed some confidence in himself which would have been missing if he had used an AI tool to being with.

There is also the issue of homogeneity. LLMs are trained on a finite set of textual data. They are also probabilistic and not deterministic. There is a good chance that given similar set of prompts, they would generate almost identical text. (Parenthetically, given that almost all the text that could be used to train the models has been used, they would in future be trained on their own outputs, giving a self-reinforcing loop.) We would thus have what has been called the McDonaldisation of writing. This has been documented in many studies. For instance, narrative fiction generated by AI is far more homogenous than that which is written by humans. [There is also a definite Western bias in the models](#) since most of their training is on material from the West.

Impact on Scientific Research

The use of AI tools in scientific research as well as writing papers is by now well known. We have many examples of how they have made a dent in previously intractable problems like protein folding, climate science and even solving mathematical problems the solutions to which have eluded us for centuries. Despite these successes, there are some issues which need to be highlighted.

Take the example of Astrophysics. In some sense, quantitative analysis has its origins in astronomical data. Ancient civilisations developed mathematical tools to make sense of astronomical data. Today, the amount of data has increased exponentially because of our ability to probe the universe using not just light but almost all wavelengths of the spectrum. AI tools are being used to analyse the data and find patterns. This has led to a huge increase in the number of scientific publications. This quantitative increase in research output is seen with an increase in homogenisation-the content and style is becoming more and more similar. There are also concerns that AI tools are generating non-existent references. Add to this the well-known phenomenon of hallucinations in these models and things are

not as rosy as they seem.

The use of AI is so extensive that [some respected figures in astrophysics claim](#) that one really doesn't need graduate students-the ones who hitherto did most of the painstaking analysis for research-anymore. This is a scary scenario. If this actually comes to pass, imagine what will happen in a few years: there will be no astrophysicists since we didn't train any. The chain of trainer-trainee-trainer or teacher-student-teacher which has been unbroken for several millennia will no longer exist. Who then would be asking the relevant and interesting questions to the AI tools if there are no people who actually understand the subject after having worked with it for decades?

Who then would be asking the relevant and interesting questions to the AI tools if there are no people who actually understand the subject after having worked with it for decades?

What about truly path breaking discoveries? Some time ago, it was felt that if one trained an AI model on data from before a particular discovery was made, it will come up with the said discovery. The idea was that whoever made that discovery had access only to the data prior to the date and so AI model should be able to use the same reasoning to discover the theory.

[This was tested with General Relativity \(GR\)](#) which Einstein started working on in 1911. A model was trained on relevant material from before that date. The model came up with all kinds of bizarre theories but not GR. In another experiment, models were trained on the material Newton had available for formulating his law of gravitation. Instead of coming up with a universal law of gravitation, the model came up with laws of gravity which were different for each of the planets .

Despite the fact that the archival data might have been incomplete, some researchers feel that the problem is more with reasoning. Models rely on inductive reasoning: they generalise a principle from lots of data points. Humans, on the other hand, rely on what is called abductive reasoning. This entails a combination of guesswork and intuition and finding the simplest explanation. We take the most likely explanation and if that cannot be generalised, we iterate to find a better one etc.

There is also the question of outliers in the data. Most discoveries have come about when people have been troubled by outliers, that is, data points which are far away from the general trend. The anomalous data points are what usually lead the researcher to question the assumptions and could in fact lead to a paradigm shift in the field. AI models being probabilistic in nature, would typically ignore outliers in the data and thus miss out on truly revolutionary breakthroughs in our understanding. The key here is human intuition, built up by working and engaging with a subject consistently over a long time.

Some researchers have pointed out that the problem is that the models lack any internal representation of the real world and no real understanding of, say, the physics behind the phenomena. Even a state of the art model like Alpha Fold does not enhance our understanding of how proteins fold but just gives us the most probable shapes.

Interestingly, the famous satirist Jonathan Swift in his *Gulliver's Travels* talks about a world where machines have replaced the human intellect . On one of his voyages, Gulliver is shown a wooden contraption which allows "even the most ignorant person[to] write books in philosophy, poetry, politics, law, mathematics, and theology, without the least assistance from genius or study." [As the Canadian academic Donal Gill notes](#), "Swift, who was deeply learned, was envisaging a nightmare world in which reading and writing-the essential foundation of wit and wisdom-are replaced by a flashy simulation of human knowledge ."

Risks and Prospects of Regulation

The risks of the proliferation of these models for the society at large are well documented. These include environmental (from the enormous amounts of power and water that the data centres use for training as well as for inference), rogue agents wreaking havoc, their use for creating physical and digital malware etc. Their impact on human interaction is equally serious. With the increasing use of social media, many people, especially youngsters, already lead an atomised existence with minimal social interaction. This trend will get exacerbated when people start treating chatbots as their friends and confidants. The chatbot is trained to reflect and reaffirm one's own views. The echo chamber which would result from this could be disastrous. There is also the tendency to think of the chatbot as an omniscient being whose pronouncements cannot be wrong. After all, instinctively, like the student with the calculator, we feel that technology must always be right. This can have dangerous consequences for our already hyperpolarised public discourse.

AI models being probabilistic in nature, would typically ignore outliers in the data and thus miss out on truly revolutionary breakthroughs in our understanding.

Can we think of regulating the use of these tools in education? Our education and assessment system which incentivise acing exams makes this difficult. In research, the incentive is to publish more and do it fast since one's career advancement depends solely on this metric. This is a major deterrent in regulating its use. Although policing is being tried in various forms like inserting Trojan Horse text in the test questions and using AI detection tools (which themselves are unreliable and what is more, have spawned the growth of new tools which can fool the AI detection tools), these cannot be replicated on a large scale.

Educators are trying to respond in different ways to the omnipresence of AI tools in the classroom. Some universities are now allowing its use to generate papers but asking the students to do a critical analysis of the output. Some institutions are banning its use in class and switching to technology-free in-class written or oral examinations. Others encourage the use of different models and then ask the students to do a comparative analysis of the results. While these and similar initiatives might work on the margins, replicating them on scale seems to be impractical. So maybe we are indeed fighting a losing battle.

Our modern industrial age is characterised by a fetish for efficiency. On this benchmark, the AI tools are of course unmatched. We would indeed save a lot of time as they proliferate in our daily lives. More than a century ago, a bearded philosopher had foreseen that technological progress will leave humanity with a lot of free time for developing its creative potential. This is definitely the case with AI tools, except that the time saved is more likely to be used for watching endless AI generated reels than anything genuinely creative.

Shobhit Mahajan used to teach physics at the University of Delhi.