

The Digitally Diluted Mind: A Crisis of Cognition in Classrooms

Pranav Srinivas¹, Asha G R²

¹Final-year Undergraduate Student, Department of Computer Science and Engineering,
B.M.S. College of Engineering, P.O. Box No.: 1908, Bull Temple Road, Bengaluru — 560 019, Karnataka, India.
Email: [pranavsrinivasannayyappa\[at\]gmail.com](mailto:pranavsrinivasannayyappa[at]gmail.com)

²Professor, Department of Computer Science and Engineering,
B.M.S. College of Engineering, P.O. Box No.: 1908, Bull Temple Road, Bengaluru — 560 019, Karnataka, India.
Email: [asha.cse\[at\]bmsce.ac.in](mailto:asha.cse[at]bmsce.ac.in)

Abstract: *This general article explains the accelerating incorporation of artificial intelligence (AI) in educational settings and explores its far-reaching effects on pedagogy, ethics, and student development. AI provides efficiency gains through automated grading, personalised feedback, and fast content creation. However, these benefits carry significant risks that threaten the core values of education. The article argues that AI systems lack ethical reasoning, contextual understanding, and human judgment, which can compromise academic integrity and stifle creativity and critical thinking. Learners are increasingly relying on AI, which means they skip the important challenges that foster deep learning and intellectual growth. Additionally, this paper points out the hidden financial, psychological, and social costs tied to using AI, including growing educational inequality and the commercialisation of learning. As AI tools become common, the focus in education shifts from active, student-centred learning to passive content consumption. This change undermines students' sense of agency and ownership. This work advocates for a human-centred educational paradigm that preserves the indispensable role of teachers, prioritises learning processes over outcomes, and cultivates ethical, imaginative, and resilient learners. Without this recalibration, widespread reliance on AI in education may yield a generation efficient in task completion yet deficient in original thought and moral discernment.*

Keywords: Ethical Implications of AI, Academic Integrity, Educational Equity, Human-Centred Learning

1. Introduction

For as long as one can remember, humans have ruled the Earth by virtue of unrivalled intelligence. Up until now, we have had complete dominance over the planet. We are on the verge of an irrevocable change, one that code and computation, not textbooks or teachers, will announce. Artificial intelligence (AI) is a new force that is quietly usurping education, not just using it as a tool. Under the pretence of efficiency and innovation, what once seemed like speculative fiction is now becoming a reality and is being incorporated into curricula and classrooms. However, there is a danger hidden beneath its elegant appearance. In the same way that the mythic Prometheus brought fire to mankind, we are on the verge of sparking a fire of our own — a machine mind.

To understand what might come, let us start at the beginning. What is intelligence? In simple terms, it is the ability to learn from experience, to reason, to acquire knowledge and skills, and to solve problems. Intelligence has always been life's currency of power. Every niche in nature is won through some form of intelligence, however simple or specialised. Humans have exploited this power more than any other species on Earth. We leveraged our brains to break free from nature's constraints and build civilizations. Yet our journey to this point was not easy or straightforward. We forgot how much struggle it took to arrive here. Over millennia, tiny innovations led to bigger breakthroughs, which then begot even bigger ones. But the pace at which knowledge built upon knowledge has exploded only in recent history. And now, with the dawn of advanced AI, we may once again be on the brink of changing everything. In a way, we stand at the foot

of history's greatest pyramid: we have already constructed so much, but the next layer – the creation of an artificial mind, greater than our own – could be both the crowning glory and Pandora's box. Intelligence may seem invaluable, but it is also costly. Brains guzzle energy. In fact, the human brain, which is only about 2% of our body weight, consumes 20% of our resting metabolic energy. For most creatures, brains beyond a certain size simply are not worth it. The energy to power a big, complex brain must come from somewhere, usually extensive feeding or other metabolic sacrifices. Consequently, for hundreds of millions of years, most life forms do with primitive or narrowly focused nervous systems. About 500 million years ago, the first rudimentary brains appeared. Think of simple flatworms with tiny clusters of neurons barely enough to control a few reflexes. As life exploded into new forms, a few lineages took the gamble on bigger brains. When competition heated up, natural selection sometimes rewarded cleverness. In those fortunate branches of the tree of life, brains grew more complex. Cephalopods like octopuses, some birds like crows, parrots, and especially mammals saw leaps in cognitive ability. These smarter animals gained advantages as they could navigate complex environments, remember useful information, and even communicate in sophisticated ways. However, even among these species, their intelligence was often limited to specific tasks. Hawks excelled at hunting, dolphins at echolocation and social bonding, humans at everything. For the vast majority of species, there was no need for a universal problem-solver. A rabbit does not require the invention of the wheel to ensure its survival. Only a few evolutionary experiments like ours ventured into general intelligence. Those experiments succeeded, but only at a substantial metabolic price. It was a careful trade-off in the biosphere:

most life sticks to specialised instincts, while a few adventurers, like the humans, risk the cost of a big brain for all its benefits.

The Rise of Humans

About 7 million years ago, a substantial change started in the hominin branch of evolution. Our early hominin ancestors began to separate from other apes. At first, their brains were only a little bigger than a chimpanzee's, but they grew over time. People still argue about the reasons — diet, climate, or social change had something to do with it — but one important change happened: intelligence changed from being good at solving problems to being more general. Our minds became more flexible, just like a knife that turned into a Swiss Army knife. *Homo Erectus* had started to see the world differently by 2 million years ago. They learnt how to control fire, make tools, and show early signs of culture. These notable events needed careful planning and foresight. *Homo sapiens* appeared about 250,000 years ago. Their brains were 30% bigger than those of *Homo erectus*. With abstract thought and language, intelligence became collective. Shared ideas could evolve across generations, accelerating progress. As language and culture flourished, knowledge built exponentially. The Neolithic revolution gave way to agriculture. Then nearly every field exploded. The industrial era ensued, and now, merely 40 years into the digital revolution, global connectivity has reached unprecedented heights.

The Rise of AI

AI is the idea that computers can perform tasks that normally require human-like mental processes. It is a strange switch: we use circuits and silicon to mimic what in other animals is done by neurons and biology. During its initial stages, AI merely consisted of a few lines of code, serving as a proof of concept to demonstrate that a machine could indeed perform certain tasks. In the 1960s, for instance, a primitive chatbot was created and a simple molecule-sorting algorithm. These early systems could perform only one specific task in a highly controlled environment, and they were incapable of learning a language or understanding a novel. In effect, its “intelligence” was about as basic as a flatworm's: it did the absolute minimum to get a task done. Those programs made no pretence of generality. You could program one to play checkers or diagnose a specific condition, but it could not suddenly start recognising photos or driving a car. Each was specialised, and weak outside its narrow domain.

For decades, the field of AI research fluctuated between periods of optimism and periods of cold disappointment, commonly referred to as “AI winters”. People would publicise AI, invest a lot, and when progress stalled, funding would dry up. But underlying all those developments was a quiet revolution in technology. Between 1950 and 2000, computers became a billion times faster, smaller, and much cheaper. Programming languages became more powerful, databases grew. With each advancement in hardware and data, AI experienced a renewed sense of potential. Milestones appeared one by one. By 1972, a robot named “Shakey” could navigate a room on its own. By 1989, neural networks had improved enough that a computer could read handwritten numerals better than most people. Then in 1997 came a shocker: IBM's Deep Blue beat the reigning chess world

champion, Garry Kasparov, in a six-game match. For years we had thought chess required human genius. For humans it still does; for a computer, it was just a calculation. Many shrugged it off as “just chess,” the work of silicon speed, not real intelligence — but they were underestimating what was coming.

Modern machine learning and neural networks brought about a significant breakthrough. Around the 2000s, AIs went from “random gibberish” to masterful performance by training on more data. These AIs now often function as mysterious entities. We give them audio, images, or text, and after training, they suddenly recognise faces, translate languages, diagnose tumours, or predict the stock market. The astonishing part is how quickly it happens. A machine learning model might start out terrible at a task, but with enough training it can self-improve to superhuman levels. By contrast, a person might take years to learn something; an AI does it in days or hours. The 2010s saw rapid, almost dizzying advances. In 2014, Facebook's AI researchers unveiled DeepFace, a neural network that could match people's faces with about 97% accuracy, rivalling the best human identification. In 2016, Google DeepMind's AlphaGo defeated the world champion at the ancient game of Go, a game so complex that mere brute force is not enough. In 2018, AlphaZero went even further: it taught itself chess from scratch in just four hours and then defeated the reigning champion computer program Stockfish. Astonishingly, it started from knowing only the rules, and no human intuition. Within hours, it was better than years of human grandmaster play.

Then came late 2022, when everything changed focus: ChatGPT. By training on essentially the entire written internet, this AI language model could produce essays, poems, computer code, and more, with fluency never seen before. You could give it a topic, anything at all, and it would summarise it, explain it, or argue a point. It learned natural language patterns at an astonishing scale. Suddenly, anyone could have a conversation with an AI that felt knowledgeable and even creative. ChatGPT itself does not “understand” meaning; it simply predicts word after word based on statistical patterns. It generates plausible prose by mimicking its training data, a process often referred to as “stochastic parroting.” Yet it can mimic understanding convincingly. You can ask it to translate from Chinese to Swahili, to write a sonnet about the moon, to draft a business plan, or to help debug code. Many seasoned writers and coders found that ChatGPT did these tasks almost as well as they did it sometimes even better, if given the right prompt. Why is this episode a turning point? This is not due to ChatGPT's perfection; rather, it often makes mistakes or fabricates details, but it proves the power of multipurpose AI. It is not limited to a single function. It is one system doing thousands of tasks, none needing human-like insight, yet all working astonishingly well. This versatility has ignited a frenzy of investment. Billionaires, tech giants, and governments around the world are now pouring billions of dollars into chasing the next breakthrough.

The Current Scenario

AI is already changing our world today; it is not some far-off future. AI decides what posts and news we see on social

media, influencing elections and public opinion. AI is used by marketing departments to target advertisements and customise campaigns for each individual. Algorithms in the entertainment industry suggest what we should watch or listen to. Lawyers use AI tools to sort through documents. Financial firms use AI for trading and risk assessment. From education to transportation, from manufacturing to science, the effects are everywhere. The result is a mix of promise and peril. The future holds many new opportunities, including more rapid medical research, individualised education, safer transportation, and unbridled creativity. However, there are also risks, such as heightened disinformation, surveillance-eroded privacy, biases in algorithms, disruption of employment, and more. Numerous seemingly stable industries are now experiencing disruption or obsolescence. This is still narrow AI! Each of the systems above is brilliant at its specific task but utterly clueless outside it. They are not fully intelligent; they are specialised geniuses.

2. AI and Education

Related Work

To ensure that AI reflects human values, Bostrom and Yudkowsky (2014) stress the need for ethical reasoning. Taddeo and Floridi (2018) argue that AI should act as 'a force of good' by using an ethical framework. Baker and Smith (2019) categorise the roles of AI technologies in education, noting that these tools should support teachers rather than replace them. Gianni et. al. (2022) warns about growing AI risks that require key ethical values like equity and privacy. Silva and Janes (2023) mention that AI's development may 'devalue' educators and create a highly automated education system. Chan and Hu (2023) and Wang et al. (2025) highlight that while students' value AI for information, they still need human qualities, especially empathy and emotional support, which AI cannot offer. Derakhshan and Ghiasvand (2024) describe ChatGPT as a 'double-edged sword' with the potential to threaten integrity and innovation. Pikhart and Al-Obaydi (2025) share teachers' concerns that AI might hinder critical thinking and cannot replicate the 'emotional and behavioural engagement' that educators provide. Additionally, Banihashem et al. (2024) argue that AI should improve human instruction rather than replace it, for example, by providing feedback. All these sources support the idea that AI lacks the intelligence, compassion, and trust needed to replace the essential human elements of education, such as moral judgment, emotional support, and thorough mentoring.

Education is the process of developing intelligence, character, and judgement. We risk destroying the entire base of learning if we give these things to computers that do not have a conscience, context, or creativity. While we think about these big futures, AI is quietly changing something that affects us every day: our schools. There are changes happening in the classroom, the lecture hall, and the relationship between students and teachers. This is not the next big stride forward; it could be our last mistake. It swaps wisdom for algorithms and converts learners into passive consumers in a system that is not meant to teach, but to improve. The reasons for the same are as follows:

1) AI cannot act with Ethics or Integrity:

AI systems lack the moral awareness necessary for honesty

and ethics, particularly in the classroom. AI does not have a conscience, so it does not evaluate information. It takes in both good and bad, including scientific findings and harmful propaganda, true facts, and false conspiracy theories. When students use AI to author essays, solve problems, or complete homework, they give away their work and their accountability. Teachers may receive material that looks decent, but it could be copied, fabricated, or filled with false citations. This violates academic honesty. AI does not recognise bias or unfairness; it simply replicates patterns from the data it was trained on. AI does not feel remorse, does not correct itself unless prompted, and does not consider context or nuance like a human teacher would. This creates a hazardous hole in responsibility. No one really takes the blame when something goes wrong. It is not easy to see, but trust is slowly fading in the classroom. Integrity is lost if neither students nor teachers completely understand or control the outcome. Schooling becomes less about learning and more about figuring out how to use things that do not make sense. In the end, teaching ethics relies on people's ability to think, feel, and judge. AI does not have any of these things at its core. Putting AI in jobs that require certain traits goes against the basic moral basis of teaching and learning. This is not just a technical mistake; it is a big moral mistake.

2) AI hinders the ability to be creative:

Education is not just about filling the mind with facts; it is also about encouraging creativity. But AI is changing that balance. Creative thinking, original thought, and problem-solving are important parts of education. When students face the unknown, deal with frustration, and produce their own ideas, these skills grow. That problem is going away now that AI is here. Why spend all night trying to figure out a hard problem when a chatbot can do it in seconds? Why think about a philosophical question when AI can just give you the answers? Why bother authoring your own poem or story when you can get a lot of them with just one click? This seems like an extreme version of the "calculator effect." Do you remember when calculators became popular and teachers saw that students' math skills were getting worse? Now picture that happening on a much bigger scale in all areas of creativity. Here are some harmful effects: A lot of students now use AI to author essays or figure out maths problems. They get the results, but they do not get to think about them. People do not do the useful things of debating an argument or producing a narrative voice. Instead of learning about perspective, shading, or brushwork, art students might use AI image generators to make pictures. They push a button and get an image, but they do not know how it was made. Computer science students might use AI coding assistants that do most of the work for them. Do they learn algorithms over time, or do they just learn how to ask the tool for help? We risk losing our creative abilities if we let machines do all the work. Creativity needs to be hard. The blank page, the first draft that is not sure what to say, and the revision after a mistake all make it harder. AI takes these experiences away. It does not get writer's block or painter's doubt; it just makes things. Even worse, AI cannot really make anything new; it just rearranges ideas that are already there. A machine only processes information: it does not wonder, dream, or feel awe. But curiosity, dreams, and emotions are what make people creative. If students start to think of creation as just something to do, they lose something important. They stop finding their

own voice because it is easier to follow AI's "voice." This is a real problem: the highs and lows of being creative; the thrill of a hard-won insight or the letdown of a project that did not work out, are what really that makes one learn.

3) It costs more to use AI in Education:

It is easy to see AI in schools as a way to save money by automating grading, making tutoring easier, and cutting down on the number of teaching assistants needed. But this point of view leaves out some essential information. There are hidden costs in terms of money, education, and social life that need to be looked at. The following paragraphs provide a detailed explanation on these points.

- a) Financial costs: A school may save on paper and pencils, but they often pay steep subscription fees to big tech companies for AI platforms. To run AI, you need computers or cloud servers, and many schools will have to buy new hardware that costs a lot of money. To support AI use, bandwidth needs to go up and security needs to get stricter. Teachers and staff also need training to use these new systems well. A lot of schools hire outside experts or pay for workshops. Eventually, the savings from grading could be wiped out by a lot of new costs for training and infrastructure.
- b) Hidden expenses for students: Families and students also share some costs. Many "learning apps" powered by AI are free for schools but come with in-app purchases and premium features. Parents may have to pay for advanced tools or go into debt for school to get the newest AI tutors. Essay mills and tutoring services have sadly turned into AI writers for hire. Students can pay a fee to get an AI-generated assignment that is "perfect." What seems easy has a moral cost and will hurt a generation of teens eventually if they buy parts of their education. AI wants metrics and will ask for more of them. Schools may give more standardised tests as AI becomes more common because "the machines can grade it quickly." This makes the curriculum smaller because teachers only teach what an algorithm can test. The variety of open-ended projects and field trips may be replaced by endless drills with multiple-choice questions and data reports. This means that education could turn into a game of points that looks good on paper but does not teach kids how to think critically.
- c) Social costs: The idea that "AI will level the playing field" is already wrong. Schools with more money are using AI tools faster because they have the money and the technology. Schools that are not as rich are behind, but it is not because the students are not smart; it is because they do not have enough access. There is a growing digital divide. Ironically, AI in education is making learning more of a business. Instead of substance, the focus is on scale and profit. When big companies run the schools and decide what to teach and how to teach it, learning has become just another product. In the short term, it may seem like "easy learning" but overall, we will have on generations of "vibing vermin" or "foolish philosophers" who never learnt how to work hard or produce innovative ideas. That is too much for any society to pay.

The Psychological Toll: Dependency Over Discipline

"Education is not just about acquiring knowledge, it also tries

to build character, discipline, fortitude, and a sense of self." In a classroom with AI, students start learning that answers are commodities to be bought or accessed instantly, rather than puzzles to be solved through perseverance. Imagine a pupil who never learns to deal with frustration because the AI always knows the solution. They may get a well-written paragraph instead of spending hours thinking about a concept in their head. This convenience may be appealing, but it alters how young people handle challenges. When they have an issue for the first time, their initial thought will be to Google or AI it away instead of looking for other solutions. The psychological impacts are small yet significant. As students get used to getting things done right away, they lose their ability to deal with frustration. Because producing something new is now the same as taking a risk, innovative thinking is becoming less common. Why take the chance when a machine can give you a perfectly convenient answer? Eventually, students tend to see themselves as people who buy ready-made knowledge instead than people who produce current ideas. The most harmful change may be going from 'learner' to 'user of answers'. "I'll just check the solution this time," is how it starts out innocently. But if the AI is better at reasoning, they might start to wonder whether they should even bother. With each question they give to a machine, their confidence in their own minds fades a little. Young adults who do well on examinations but cannot handle uncertainty will come from a generation that thinks of education as something that can be gotten rather than something that is earned through demanding work. Taking risks is not usual because trying something new could make you fall behind the AI. Over time, passivity takes the place of innovation and curiosity. The temperament essential for innovation is progressively diminishing.

3. Conclusion

Intelligence can be dangerous without wisdom, cold without compassion, and pointless without integrity. No matter how smart AI develops, it will never be able to entirely replace teachers, mentors, and the rich tapestry of human relationships that shape learning. Letting AI run our classrooms might make students "more efficient" on paper, but they might not get as much done in real life. They may score higher on standardised tests, even though they are less resilient, creative, and curious. We need to remember that learning does not happen automatically and that education is not something you can download. Knowing the answer is different from understanding what it means. A youngster who knows all the data that an AI can provide them may not understand the most important parts of learning, such how to establish a moral compass, produce new questions, or take responsibility for their own thoughts. Education instills morality, the bravery to question authority, the honesty to question sources, and the joy of finding beauty in maths or poetry. You cannot 'download' the feelings of empathy, purpose, and amazement. An AI can compose an essay about compassion, even though it does not have any. AI is one of the most powerful tools that people have ever made. But if we do not keep an eye on how we use it, it could become our sole tool. That would be a pyrrhic victory in a culture where people are good at taking tests but not so good at thinking. It means the end of human innovation as we know it. This should not be our legacy. Let us use AI with a purpose and

enthusiasm, and make sure that the product we make helps people instead of replacing them.

A Better Path Forward: Human-Centred Education

- 1) Cultivate Intellectual Independence and Ethical Reasoning: Education should focus on developing creativity, critical thinking, and moral judgment. This approach enables learners to think beyond mere reproduction of information.
- 2) Affirm the Primacy of the Learning Process: Acknowledge effort, experimentation, and revision instead of solely the outcome. True learning comes from struggle and growth, not perfect results.
- 3) Make Curricula That Focus on Creativity, Ethics, and Conversation: Educational frameworks should emphasise creativity, ethics, dialogue, and the humanities. This focus protects areas where human abilities go beyond what technology can replicate.
- 4) Entrust Educators with the Responsible Stewardship of AI: Teachers need to be ready to see AI as a tool, not as a replacement. Human supervision ensures that learning stays relevant, thoughtful, and personal.
- 5) Ensure the Deliberate and Measured Use of AI in Learning: AI should be used with caution for clear purposes. This ensures that education stays rooted in human values, personal connections, and the quest for knowledge.

References

- [1] Bostrom, N., & Yudkowsky, E. (2014). The ethics of artificial intelligence. In K. Frankish & W. M. Ramsey (Eds.), *The Cambridge Handbook of Artificial Intelligence* (pp. 316–334). chapter, Cambridge: Cambridge University Press. <https://doi.org/10.1017/cbo9781139046855.020>
- [2] Sohn, B. K., Thomas, S. P., Greenberg, K. H., & Pollio, H. R. (2017). Hearing the voices of students and teachers: A phenomenological approach to educational research. *Qualitative Research in Education*, 6(2), 121–148. <https://doi.org/10.17583/qre.2017.2374>
- [3] Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science*, 361(6404), 751–752. <https://doi.org/10.1126/science.aat5991>
- [4] Baker, T., & Smith, L. (2019). Educ-AI-tion rebooted? Exploring the future of artificial intelligence in schools and colleges. Nesta Foundation. Retrieved from: https://media.nesta.org.uk/documents/Future_of_AI_and_education_v5_WEB.pdf
- [5] Gianni, R., Lehtinen, S., & Nieminen, M. (2022). Governance of responsible AI: From ethical guidelines to cooperative policies. *Frontiers in Computer Science*, 4. <https://doi.org/10.3389/fcomp.2022.873437>
- [6] Silva, A. de O., & Janes, D. dos S. (2023). Challenges and opportunities of artificial intelligence in education in a global context. *Review of Artificial Intelligence in Education*, 4(00), e1. <https://doi.org/10.37497/rev.artif.intell.education.v4i00.1>
- [7] Al-Obaydi, L. H., Pikhart, M., & Tawafak, R. (2023). Online assessment in language teaching environment through essays, oral discussion, and multiple-choice

questions. *Computer-Assisted Language Learning Electronic Journal (CALL-EJ)*, 24(2), 175–197.

- [8] Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20, 43. <https://doi.org/10.1186/s41239-023-00411-8>
- [9] Banihashem, S. K., Kerman, N. T., Noroozi, O., Moon, J., & Drachsler, H. (2024). Feedback sources in essay writing: Peer-generated or AI-generated feedback? *International Journal of Educational Technology in Higher Education*, 21(1). <https://doi.org/10.1186/s41239-024-00455-4>
- [10] Derakhshan, A., & Ghiasvand, F. (2024). Is ChatGPT an evil or an angel for second language education and research? A phenomenographic study of research-active EFL teachers' perceptions. *International Journal of Applied Linguistics*, 1–19. <https://doi.org/10.1111/ijal.12561>
- [11] Wang, X., Gao, Y., Wang, Q., & Zhang, P. (2025). Fostering engagement in AI-assisted Chinese EFL classrooms: The role of classroom climate, AI literacy, and resilience. *European Journal of Education*, 60(1), Article e12874. <https://doi.org/10.1111/ejed.12874>
- [12] Pikhart, Marcel & Al-Obaydi, Liqaa. (2025). Reporting the potential risk of using AI in higher Education: Subjective perspectives of educators. *Computers in Human Behavior Reports*

Author Profile



Pranav Srinivas is a Final-year Undergraduate student in the Department of Computer Science and Engineering at the B.M.S. College of Engineering. Guided by a profound command of mathematics, he unites his academic rigour with humanitarian purpose to drive transformative progress in the society. He pursues interdisciplinary projects and sustainable initiatives that meaningfully contribute to national development.



Dr. Asha G R is a seasoned educator, currently serving as Professor in the Department of Computer Science and Engineering at the B.M.S. College of Engineering. With over 26 years of distinguished teaching experience, she combines scholarly depth with mentorship, fostering professional excellence. Her unwavering commitment to education, research, and student development continues to inspire students and institutional growth.