Reflections from UK and US Classrooms on Building Responsible AI Literacy

Justina J. Wang
State College Area High School, State College, Pennsylvania, USA.

This position paper highlights the importance of fostering responsible AI literacy in secondary education, drawing on personal observations from contrasting technology-use environments in UK and US schools. The author critiques simplistic regulatory approaches, emphasizing that unchecked AI use can amplify misinformation, limit student creativity, and impair critical thinking. The paper advocates for comprehensive AI education through curriculum enhancements and targeted training. Such measures aim to help students understand AI's capabilities, limitations, and ethical implications, and to encourage informed, balanced engagement with the technology.

In my first year of high school, I attended a school in the UK, about an hour and a half by train north of London. It was the first-ever all-girls grammar school in the country, established 175 years ago, an unsurprising claim well matched by its historical architecture: striking buildings with crimson brick and sharp spires. Like their architecture, their teaching methods also felt historical, with strict regulations on phones, paper assignments, and the use of whiteboards. Having come from a school district where even kindergarteners had access to Chromebooks, I was astonished at their lack of technology. I wondered how they could keep track of assignments or look things up when they were in doubt; surely they didn't write everything on paper in this day and age. However, this was the reality for students attending this school: handwriting essays, using textbooks, and paper notes. I was soon thrown into this environment, although at first I was appalled by the fact that I could no longer pull out a computer to Google simple questions. However, as I acclimated to the new atmosphere, I realized I didn't need technology as much as I thought, or even at all. When I had questions, I could simply ask the teacher, and when I had paper assignments, it became an opportunity to improve my handwriting. I was quickly peer-pressured by everyone else's graceful handwriting, developed over years of training. A year later, when I moved back to the US, I was once again overwhelmed by how much our school system relies on technology.

These two vastly different approaches to technology use have driven me to reflect on AI's role in the education system. While AI can be a powerful tool, it is often used irresponsibly, with users blindly following its suggestions and remaining oblivious to its complex implications. High schools should move beyond simply banning or permitting certain software; instead, they should implement comprehensive programs that educate students about AI and its responsible use. The education system must take responsibility for shaping future generations to become AI literate. It is vital for students to understand, at an early stage, the potential of AI, its limitations, and its long-term effects on intellectual growth.

According to my current school's AI policy, our approved tools are limited to four applications: Khanmigo, Brisk, Eduaide, and Magic School. When I first heard of this, I didn't have any opinion because I didn't use AI very often, and I remained unopinionated even after the policy had been in place for a few months. Perhaps this was because I had never become reliant on applications like ChatGPT. Without the recent AI tools, I found it adequate to rely on Google search as I have done for years. However, this simplicity soon changed when Google implemented the AI overview; I suddenly felt I had lost the autonomy of using information

online. The AI overview provides an AI summary of all of the primary sources. Although meant to offer more specific answers to queries, I found it to be a nuisance, as the AI overview frequently lacked clarity, and worse, I sometimes noticed it was incorrect. Unfortunately, I could not work around it because it is not intuitive to see the source from which the response was pulled. The AI overview was especially infuriating when I had a medical inquiry, and I later discovered the answer that the AI provided was from Reddit, an online forum. For someone not carefully checking, the information provided could be misleading and even potentially dangerous. The only solution to avoid the risk of misinformation would be to click on each of the potential sources of the summary and read through the articles, a time-consuming task usually overlooked. In contrast, with the search engine before the AI overview, Google would highlight important information and make it immediately clear exactly where the source came from. Although I was aware of the risks, the AI overview continued nagging at me to take the easier route and hope that it was correct.

Things only became worse when more people around me grew dependent on AI. Students began using AI to circumvent learning, rather than using it to learn. I became increasingly frustrated when group members would lean on ChatGPT for lab reports, as I would often notice their parts were incorrect because ChatGPT had spread misinformation regarding facts such as melting points, stoichiometry, and the color of precipitates. When I confronted them about it, they would say, 'the AI told me,' and that the AI is always correct. Suddenly, it seems as if the people around me have become much more stubborn and harder to reason with. As a result, I now spend immense amounts of time in group projects fact-checking everything due to paranoia that something incorrect from ChatGPT might sneak in. Furthermore, this AI-trusting stubbornness is not only an issue among teenagers, as it has affected my own parents. My father will often consult ChatGPT when trying to back up his failing arguments and then refuse to admit he is wrong because ChatGPT told him so. Interestingly, Chong et al. (*1*) found that human self-confidence outweighs confidence in AI when deciding to accept AI suggestions.

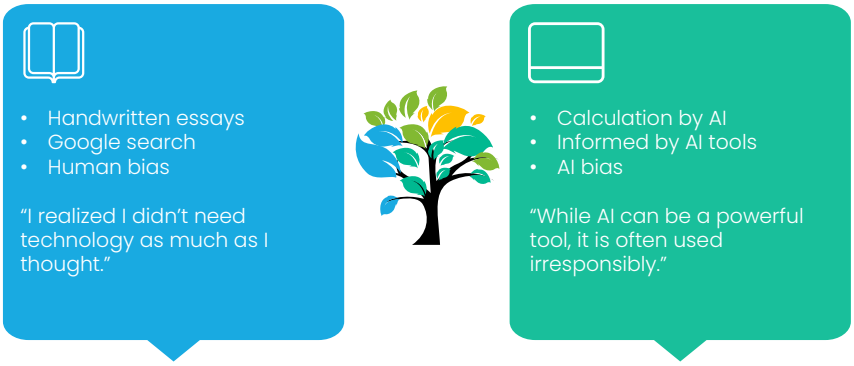

Figure 1: Comparison of traditional education and AI-enhanced learning.

AI can also harm the creativity of students. For example, when students consult ChatGPT for essay and speech ideas, or even ask it to write the content, I often find the responses are similar, leaving no room for surprise. When I read these essays written by ChatGPT, I wasn't able to conjure a feeling of connection to the author and their opinions; instead, it was as if I were watching a performer making all the right moves, but blankly staring at the audience. A report by the Hopelab found that 51% of teenage users of generative AI use it to brainstorm ideas (*2*). This demonstrates how human creativity may be declining as they struggle with the ability to create something new on their own. Cremer, Bianzino, and Falk (*3*) thoroughly discussed how AI could impact human creativity. Although the risk of AI monopolizing creativity is apparent, these authors also presented scenarios in which AI assists creative work and fosters greater appreciation for human ingenuity—possibilities not mutually exclusive with the more negative scenario. They noted that a key issue is determining how to prepare for generative AI. In a broader context, Touretzky et al. (*4*) explored both positive and negative societal impacts of AI.

The narrow regulations set by my school are not effective in promoting informed and optimized use of AI. First, the school cannot control students' practices at home. Despite that most AI tools are not permitted at school, students will exploit AI tools at home for essay ideas. As Gerald Letendre said regarding AI tools,

"the horse is out of the barn" (*5*). Although the horse may have escaped the barn, we can still attempt to tame the horse without stuffing it back into the barn. I am fairly certain that we will never be able to go back to the days of an AI-free learning environment, but through education initiatives, we can help students see the bigger picture regarding AI's pros and cons. Gaining a fundamental understanding of major AI systems helps students see that AI is neither something to fear nor something to rely on excessively. Second, due to the human tendency, especially among teenagers, to do the opposite of what they're told, imposing rigid rules can often backfire. As a study showed, teenagers tend to see rebellious behaviors as a form of fun (*6*).

Developing best practices for K–12 AI education has attracted increasing attention from both educators and AI experts. One promising example of such educational efforts is the AI4K12 initiative, a collaboration between the global AI research community (AAAI) and the K-12 education community (CSTA) (*4*). It provides national guidelines for teaching AI in K–12 education, organized around five core 'big ideas' in AI. The initiative offers educators concrete tools and frameworks to help students understand how AI systems perceive, reason, learn, and interact with society. Incorporating this model into school programs could support the kind of structured, age-appropriate learning needed to foster responsible and informed engagement with AI technologies. Integrating such programs early may ultimately shape a generation of students who not only use AI properly but also bring new perspectives on its development.

I would propose a few initiatives to improve the AI literacy of future generations in my school district, which could also be applicable in other regions. First, it is vital for schools to teach students how to use AI effectively. Although, as implied by their nature, AI tools strive for ease of use, not every student automatically understands the broad range of AI tools available or, importantly, which tools are best for what purposes. Instead of letting students rely on fragmented pieces of information regarding AI that they come across when scrolling through Instagram or talking with friends, schools should provide in-depth courses on how to use AI with maximum efficiency. Numerous online resources and books (*7*) explain how to use ChatGPT and other AI tools; schools mainly need to synthesize these materials, keeping preparation minimal.

The elephant in the room is generative AI. Schools can design activities to inform students about its potential biases and long-term influence. A study showed that societal stereotypes are often magnified by AI systems (*8*). For example, when generating images of different occupations, the AI would lean towards stereotypes, rather than produce random or balanced representations. If students are unaware of these biases, gradual exposure may lead them to unintentionally adopt prejudiced views. As a society, we may become stuck in traditional ideas, eliminating the possibility of progress. Personally, I often notice that ChatGPT's responses depend heavily on how the prompt was worded. Even when the question is a factual inquiry, if the prompt hints slightly for confirmation, ChatGPT will often confirm that whatever is stated is true, even if it is blatantly incorrect. Users, including myself, are often unaware of this bias. Instead of using AI as a tool, they use it to confirm their incorrect beliefs and convince themselves that they are correct. In order to mitigate the effects of this issue, schools have to raise awareness of its existence. Being more informed would allow students to make more thoughtful decisions. Discussions around AI bias will allow students to let their individual biases balance out, and thus, learn from each other.

Schools can create programs such as short courses or interest clubs to teach students about the computational principles behind AI software, specifically, what forms the data sources for training AI models, how AI generates its responses, and what can limit the scope and accuracy of AI's responses. Insights gained from the effort to understand these questions will help students realize that AI tools are computational models that rely on data sources, make decisions in fundamentally different manners from humans, and may generate incorrect information. Thus, students should be strongly encouraged to fact-check the information they obtain and, as a byproduct, learn more about the topic they are researching. In addition, exposing students to knowledge of how AI works can open young minds to potential improvements in AI systems, inspiring them to become the next generation of AI researchers.

Improving the AI literacy of future generations begins with education in schools. Some people may argue that it is not necessary; however, AI will only continue to develop, and understanding how to incorporate

it into daily life is essential to keep up with a functioning society. Currently, many students lack a clear understanding of what AI actually does, believing that it may someday be the downfall of human beings or that it is simply 'stupid.' Addressing these misconceptions through stronger awareness of AI technology will help students become thoughtful, responsible users—and perhaps even creators—of the tools that will shape their future.

References

1. L. Chong, G. Zhang, K. Goucher-Lambert, K. Kotovsky, J. Cagan, Human confidence in artificial intelligence and in themselves: The evolution and impact of confidence on adoption of AI advice. *Computers in Human Behavior* **127**, 107018 (2022).
2. Hopelab, Common Sense Media, Center for Digital Thriving, Teen and young adult perspectives on generative AI (2024), `https://hopelab.org/stories/teen-and-young-adult-perspectives-on-generative-ai/`, accessed: 2025-05-04.
3. D. De Cremer, N. M. Bianzino, B. Falk, How generative AI could disrupt creative work. *Harvard Business Review* **13** (2023).
4. D. Touretzky, C. Gardner-McCune, C. Breazeal, F. Martin, D. Seehorn, A year in K–12 AI education. *AI Magazine* **40** (4), 88–90 (2019).
5. T. Durkin, SCASD takes first step to addressing AI in the classroom (2024), `https://www.wtaj.com/news/local-news/scasd-takes-first-step-to-addressing-ai-in-the-classroom/`, WTAJ News.
6. N. E. Blankenstein, E. H. Telzer, K. T. Do, A. C. Van Duijvenvoorde, E. A. Crone, Behavioral and neural pathways supporting the development of prosocial and risk-taking behavior across adolescence. *Child development* **91** (3), e665–e681 (2020).
7. K. Kahn, *The Learner's Apprentice: AI and the Amplification of Human Creativity* (Constructing Modern Knowledge Press, Torrance, California), 1st ed. (2025).
8. L. Nicoletti, D. Bass, Humans are biased. Generative AI is even worse (2023), `https://www.bloomberg.com/graphics/2023-generative-ai-bias/`, accessed: 2025-05-03.

Acknowledgments J.J. Wang gratefully acknowledges the guidance and inspiration of her teachers and friends at State College Area High School, Loughborough High School (UK), the Delta Program of the State College Area School District, and State College Friends School. She thanks Professor Jia Li in the Department of Statistics at Penn State University for her mentorship and critical feedback throughout the writing of the article. She also thanks her family for their constant support and encouragement.