

On the Effects of Artificial Intelligence in Classroom Settings

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Spencer Hernandez

Spring 2024

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisors

Richard D. Jacques
Department of Engineering and Society

Technical Advisor

STS Research Paper

Introduction

Since late 2022, anyone with an email and internet connection can ask ChatGPT whatever questions they like, no matter how obscure, and generally expect a rational response. While this is an important milestone in the development of AI, it has not come without concerns. Schools across the U.S. are not only struggling to prevent their students from using AI software to cheat on assignments but are overwhelmed when faced with the question of how AI should be integrated into the classroom. In this paper, I will analyze the social and economic implications of AI in the classroom setting.

Since the dawn of the Personal Computer (PC), schools have needed to find a solution to either integrate or eliminate a variety of new technologies from classrooms. After the PC, a couple of these include internet access, smart boards, and Chromebook programs. Perhaps the biggest example is the adoption of Zoom, the online video conferencing software that was used ubiquitously after the rise of COVID-19 in 2020. These technologies are not without their issues (such as being distracting to the classroom and expensive), but they do provide major pedagogical benefits, such as ease of information sharing. To introduce another possible technology to benefit classrooms, a gesture-controlled light-emitting diode (LED) Matrix was developed through the technical capstone project. This LED Matrix was developed specifically to increase K-12 student engagement and get them interested in Electrical and Computer Engineering.

With the widespread adoption of laptops and Chromebooks in U.S. classrooms, this free-to-use software has become an effortless way for students to cheat on assignments. Not all use

cases of this Generative Pre-training Transformer (GPT) chatbot are of ill intent: AI could help subsidize teacher's lessons by tailoring to the needs of individual students. This could certainly bolster classroom learning and engagement and assist with keeping all students up-to-date on content. This idea may sound promising, but it as well is not without its own issues. Schools with Chromebooks could choose to encourage the use of AI in their classrooms, but schools without in-classroom Chromebooks will not be able to. Perhaps this is due to the school's rural location, or maybe the school does not have the funds to implement a costly personal computer program. Either way, schools without the ability to leverage AI could rapidly fall behind their counterparts. The goal of the STS paper is to perform a deep-dive into the drastic effects, both positive and negative, that AI could have in US K-12 schools.

Methodology

In order to properly analyze the effects of Artificial Intelligence on education, it is first necessary to analyze and then determine which research methodologies are appropriate to use. To draw a meaningful conclusion, which is to say that the implications of Artificial Intelligence in the classroom have been properly analyzed, an applied breakdown of current findings will be presented alongside a literature review and an ethical analysis of AI in the classroom.

The literature review is a crucial part of analyzing these effects especially since a novel experiment was unable to be conducted. Artificial Intelligence is a relatively new development that has been the subject of many research papers, with the total number of AI publications doubling from 2010 to 2021 (Artificial Intelligence Index Report 2022, 2022). From the lens of Computer Science, studying AI in the classroom would be classified as Human-Computer

Interaction (HCI), the study of how computers can influence the way humans utilize them. From the 350,000 AI publications in 2021, only 10,000 of them were related to HCI, or about 3%. Even if all ten thousand of these publications were about the effects of AI in the classroom, it is still important to see how trends have changed over time, especially with the release of ChatGPT and Google Gemini. With the release of newer AI technologies, the need for an up-to-date literature review on AI in the classroom becomes increasingly prevalent.

People's outlook on AI is shifting. A study conducted by the Pew Research Center reveals that the percentage of U.S. adults who are concerned with the increased use of AI in day-to-day life has increased significantly from 2022 to 2023, with the percentage jumping from 38% to 52% (Tyson & Kikuchi, 2023). Since there is a growing concern about AI, it seems like an appropriate time to conduct an analysis on the ethical concerns of AI using an ethical analysis methodology. This means that the literature review will not only analyze current findings, but also analyze them from an ethical perspective to propose the optimal solution.

The sources utilized for this paper come from highly reputable organizations. These include The National Center for Biotechnology Information, Pew Research Center, and the Institute of Electrical and Electronics Engineers. As such, the analysis of these findings stands to draw major conclusions that could be of assistance when developing policy around the practical use, and the ethical implications of AI in education.

Background

Before discussing the effects of Artificial Intelligence, it is necessary to gain a better understanding of how AI actually works. Artificial Intelligence is defined as the development of

computer systems that can perform tasks that typically require human intelligence, such as image recognition, speech recognition, and translating languages. OpenAI's ChatGPT, xAI's Grok, Google's Gemini, and Meta's Llama 2 are some of the most well-known AI chat bots currently on the internet. All of these AI are what is known as Large-Language Models (LLM). A LLM is a type of AI that can read, process, and produce text. They can do this by being trained, or having a large amount of information, such as books, websites, and blogs input into them to allow the models to learn from the data and find similarities. The LLM then inputs this data into a neural network, a network consisting of neurons, or inputs, and edges, which connect neurons to other neurons. This architecture allows the model to find the similarities and differences between the data that was input, all by itself. This is both a good and bad thing. On one hand, it allows the model to learn by itself without anyone telling it what the output should be given an input. On the other hand, it means that it is exceedingly difficult to determine why a neural network came to a certain conclusion. By repeatedly running tests from the input data and utilizing various techniques to minimize the error of the output, an entire model is trained.

These LLMs are capable of extremely complex calculations that allow them to read and process text, a process known as Natural Language Processing (NLP). With NLP, LLMs are able to do several various processes, one of which is sentiment analysis. Sentiment analysis consists of processing textual input to conclude if the input had a positive, neutral, or negative tone to it. This is used to gain even more input from its user, which can then be fed back into the model to train from.

With these cutting-edge AI models, virtually anyone with an internet connection and a computer can conversate with any model of their choice. The purpose of this paper is to analyze some of the consequences of these models, specifically in the context of educational settings.

Literature Review

With the increased use of AI, the various effects it could have on society must be analyzed. Since there are so many different possible side effects of AI, the theories that will be focused on are latent and manifest functions and dysfunctions, and Cultural lag.

The idea of Latent and manifest functions and dysfunctions was created by sociologist Robert C. Merton in 1968. This concept says that technologies have manifest (explicit) and latent (unintended) consequences (Merton 1968). These effects are either good (in which case they are functions) or bad (dysfunctions). This framework of analyzing social phenomena shines light on the fact that new ideas and technologies can have consequences that are difficult to predict. With the rise of a technology that is proposed to be extremely disruptive, using such a framework to analyze the effects of AI is crucial.

Artificial Intelligence, even in its preliminary stages, has already been shown to have extraordinarily strong manifest and latent consequences. An example of a manifest function of AI is to be able to solve large, complex problems in a short timeframe. Perhaps the most notable of these is Google's DeepMind project, where in 2016, an AI beat the world's greatest Go player (Lee et al., 2016). Google spent hundreds of millions of dollars developing this AI, with its only purpose being to play Go. While AI is to be used to solve difficult problems and automate tasks, it has shown to have unanticipated (but welcome) applications. Researchers have utilized AI to analyze biological and medical data, which has led to breakthroughs that include the ability to treat diseases, detect tumors, and assist with surgical procedures (Cheng-Tek Tai, 2020). Even though this was not the original purpose of AI, it is promising to see that it does not just have

negative consequences. An example of a manifest dysfunction in the context of AI has to do with the effects of automation. One of AI's main selling points is that it can automate jobs to reduce labor costs. It is estimated that roughly "30 percent of the activities in 60 percent of all professions could be automated" (Badet, 2021). While this may allow businesses to cut overhead and increase profits, experts argue that the effects of automation will disproportionately affect workers without college degrees (Badet, 2021). As such, if nothing is done to help curve this consequence then there may be a sharp increase in unemployment. Another dysfunction of AI is the major bias that comes from training models. AI models are trained on human-generated content on the internet, which is plagued with human biases. This bias is then transferred into the model and can be exacerbated by AI if unchecked. In a case study that analyzed leadership narratives using an AI tool, male leaders were depicted as "strong, charismatic, and sometimes intimidating" whereas female leaders were portrayed as "emotional, ineffective, overly focused on pleasing others, and fearful" (Newstead et al., 2023). As previously mentioned, AIs follow a black box model, making it difficult to figure out why they made a certain choice. With the prevalence of bias in AI, people are becoming increasingly worried that the black box model will be a big roadblock in overcoming this disparity. This latent dysfunction could echo harmful and outdated ideologies if not addressed, which is why it is crucial to think about the unintended consequences, as well as the intended consequences, of AI.

There is no arguing that the discussed consequences are not of concern: it will take lots of policy, awareness, and forethought to mitigate these negative consequences of AI. To this extent, these consequences must be thoroughly studied before AI can appropriately be applied to K-12 classrooms. Studies have shown that children have significantly more credulity than adults, such as being quick to believe in machines that can shrink rooms or clone an object (Jaswal et al.,

2010). Philosopher Anette Baier (1986) provides one possible explanation: children are dependent on others around them to survive, so it is in their best interest to trust them (Baier, 1986). This in turn, can lead to bias if children are being taught harmful ideologies by a trusting adult, or perhaps, an AI chatbot. Since AI can propagate harmful ideologies that children need to be protected from, AI must be studied much more extensively before pedagogical integration to reduce biases.

Another reason to be wary of the use of Artificial Intelligence in education, is the fact that the United States already has one problem to address when it comes to the use of technology in the classroom. The Digital Divide in the context of education is the idea that there is a growing disparity in learning opportunities due to certain socioeconomic groups having easier access to technology than others. Pew Research findings show that nearly 20% of US teens are not always able to finish schoolwork at home since they do not have access to at-home internet (Anderson & Perrin, 2018). Unfortunately, communities with lower median incomes will have less access to the Internet, which disproportionately affects African American people. African Americans have significantly higher levels of poverty (18.8%) when compared to Caucasians (7.3%), despite making up less of the U.S. population. It therefore comes as no surprise that schools with a majority of African American students are less likely to have classroom-level internet access compared to predominantly Caucasian schools (46% and 31%, respectively) (Creamer, 2020). The Digital Divide is already a big issue, and AI only stands to worsen it. While AI could certainly benefit those who do have access to at-home internet, it stands to also punish unfairly those who do not. This is yet another latent dysfunction of AI that could further widen educational outcomes in the US.

The application of the latent and manifest functions and dysfunctions socio-technical framework highlights the benefits and pitfalls of AI, specifically in an educational context. While AI has had positive consequences, like the development of DeepMind and medical breakthroughs, it also has negative consequences, such as the removal of jobs due to automation and the propagation of bias'. It is crucial, for children's development, to thoroughly understand what AI is capable of and what threats it may pose before use in a classroom.

Cultural lag is the idea that societal norms, policies, etc. develop at a slower pace than technologies. This term was invented by Philosopher and Statistician William F. Ogburn, and specifically, refers to the period of "maladjustment" that comes from the "non-material" culture (policy, values, norms, etc.) trying to adapt to "material" (physical objects, resources, etc.) culture (Ogburn & University of California Libraries, 1922). Ogburn believed that understanding the causes and effects of cultural lag was crucial for anticipating societal changes that came from the development of new technology. He concluded that society needs to make proactive efforts to adjust to cultural lag, otherwise societal problems and conflicts may develop. As with many developments, the U.S. government struggles to develop policy quickly enough to effectively combat new physical changes.

Perhaps the most robust example of this would be the COVID-19 outbreak in 2020, where the U.S. government did not react quickly enough to the outbreak. This ultimately led to 28.5% of all COVID-19 deaths coming from the U.S., despite the U.S. population only accounting for 4.25% of the world's population (Nowroozpoor et al., 2020). Some of the explanations for this include the lack of epidemic outbreaks in the U.S. (the effects of Zika and Ebola, for example, were very minimal), and fear of economic recession due to forcing businesses to close down. This demonstrates that even though the U.S. may have had some

reason to not close down, it ultimately led to nearly 1.2 million deaths, emphasizing the importance of minimizing cultural lag. While the immediate effects of AI will almost certainly not compare to those of Covid-19, there could still be significant issues that arise, which makes it truly relevant.

Currently, there is almost no literature on the effects of cultural lag of AI, making it difficult to find reputable sources to cite in order to discuss how these effects may have implications for classrooms. As such, it was necessary to slightly widen the scope of this research paper to analyze how the cultural lag of AI may affect students. On January 31, 2024, CEOs of several social media platforms, including Mark Zuckerberg, testified in front of the U.S. Senate about the effects social media is having on children and teenagers. Platforms such as X, Instagram, Meta, and TikTok are currently under fire due to the mental health consequences they are having on their younger audience. Studies have shown that young people who are on social media for more than three hours a day face twice the risk of worsened mental health (Foundation, 2023). Since 8th to 10th graders have reported being on these social media platforms for an average of 3.5 hours a day, there is grave concern about such mental health effects (Foundation, 2023). These effects of social media on mental health are drastic, as mental health issues at young ages have been linked to failing grades and not being college-eligible (Agnafors et al., 2020). With the highly sophisticated recommendation algorithms that social media platforms are using, especially those that utilize AI, children and adults must be made aware of these effects (Zhang et al., 2020). Given that AI recommendation algorithms can lead to worsening mental health outcomes for young people due to perpetuating a cycle of harmful content, schools must be made aware of the effects these platforms have on students. This is an instance in which cultural lag poses a serious threat. If a policy is not made that assists public K-

12 schools in helping and providing aid to their students, then not only will students' performances fall, but the mental health consequences may be severe.

By utilizing the frameworks of latent and manifest functions and dysfunctions, as well as cultural lag, the effects of AI are better understood. These frameworks revealed that there are several unintended consequences of AI that pose serious threats to the well-being of students' mental health, and quality of education.

Discussion & Conclusion

Artificial Intelligence's implications for society as a whole are already starting to reveal themselves. With increased companies developing their own LLMs and each model having drastic updates on a month-to-month basis, it comes as no surprise that AI keeps expanding and adapting. Whether or not these effects will change the world has already been subject to debate, but that does not mean that it will not still have drastic effects. Healthcare, finance, technology, transportation, and other sectors may be subject to these AI-related changes, but education is a topic that is highly relevant and critical to understand. With the effects of AI being realized by K-12 students, children, and teenagers, it is especially important to understand how this disruptive technology may influence the next generation.

With the culmination of this research paper, hopefully more people will become educated on the effects of AI in classroom settings. Researching viable solutions to help curb these effects was out of the scope, but the in-depth analysis of research articles, studies, and well-documented literature will hopefully be enough to start the conversation of mitigating harmful consequences.

References

- Agnafors, S., Barmark, M., & Sydsjö, G. (2020). Mental health and academic performance: A study on selection and causation effects from childhood to early adulthood. *Social Psychiatry and Psychiatric Epidemiology*, 56(5).
<https://doi.org/10.1007/s00127-020-01934-5>
- Anderson, M., & Perrin, A. (2018, October 26). *Nearly one-in-five teens can't always finish their homework because of the digital divide*. Pew Research Center.
<https://www.pewresearch.org/short-reads/2018/10/26/nearly-one-in-five-teens-cant-always-finish-their-homework-because-of-the-digital-divide/>
- Artificial Intelligence Index Report 2022*. (2022). https://aiindex.stanford.edu/wp-content/uploads/2022/03/2022-AI-Index-Report_Chapter-1.pdf
- Badet, J. (2021). AI, Automation and New Jobs. *Open Journal of Business and Management*, 09(05), 2452–2463. <https://doi.org/10.4236/ojbm.2021.95132>
- Baier, A. (1986). Trust and Antitrust. *Ethics*, 96(2), 231–260.
<https://www.jstor.org/stable/2381376>
- Cheng-Tek Tai, M. (2020). The impact of artificial intelligence on human society and bioethics. *Tzu Chi Medical Journal*, 32(4), 339–343. National Library of Medicine.
https://doi.org/10.4103/tcmj.tcmj_71_20
- Creamer, J. (2020, September 15). *Poverty Rates for Blacks and Hispanics Reached Historic Lows in 2019*. The United States Census Bureau.
<https://www.census.gov/library/stories/2020/09/poverty-rates-for-blacks-and-hispanics-reached-historic-lows-in-2019.html>

Foundation, T. A. E. C. (2023, August 10). *Social Media and Teen Mental Health*. The Annie E. Casey Foundation. https://www.aecf.org/blog/social-medias-concerning-effect-on-teen-mental-health?gad_source=1

Jaswal, V. K., Croft, A. C., Setia, A. R., & Cole, C. A. (2010). Young Children Have a Specific, Highly Robust Bias to Trust Testimony. *Psychological Science*, 21(10), 1541–1547. <https://doi.org/10.1177/0956797610383438>

Lee, C.-S., Wang, M.-H., Yen, S.-J., Wei, T.-H., Wu, I-Chen., Chou, P.-C., Chou, C.-H., Wang, M.-W., & Yan, T.-H. (2016). Human vs. Computer Go: Review and Prospect [Discussion Forum]. *IEEE Computational Intelligence Magazine*, 11(3), 67–72. <https://doi.org/10.1109/mci.2016.2572559>

Merton, R. (1968). *Social Theory and Social structure; toward the Codification of Theory and research*. Glencoe, Ill., Free Press.

Newstead, T., Eager, B., & Wilson, S. (2023). How AI can perpetuate – Or help mitigate – Gender bias in leadership. *Organizational Dynamics*, 52(4), 100998. <https://doi.org/10.1016/j.orgdyn.2023.100998>

Nowroozpoor, A., Choo, E. K., & Faust, J. S. (2020). Why the United States failed to contain COVID-19. *Journal of the American College of Emergency Physicians Open*, 1(4). <https://doi.org/10.1002/emp2.12155>

Ogburn, W. F., & University of California Libraries. (1922). Social change with respect to culture and original nature. In *Internet Archive*. New York : B.W. Huebsch, Inc. <https://archive.org/details/socialchangewith00ogburich/page/n7/mode/2up>

Tyson, A., & Kikuchi, E. (2023, August 28). *Growing Public Concern about the Role of Artificial Intelligence in Daily Life*. Pew Research Center.

<https://www.pewresearch.org/short-reads/2023/08/28/growing-public-concern-about-the-role-of-artificial-intelligence-in-daily-life/>

Zhang, Q., Lu, J., & Jin, Y. (2020). Artificial intelligence in recommender systems. *Complex & Intelligent Systems*, 7. <https://doi.org/10.1007/s40747-020-00212-w>