

Indexing Student and Faculty Reactions to GAI

Allan Encarnacion

St. Mary's Ryken High School, USA

ABSTRACT

The following paper analyzes the reactions of academic institutions to the proliferation of Generative AI (GAI) technologies through introducing the GAI Fairness Index (GAIFI), a novel indexing methodology that utilizes survey data to quantitatively assess an institution's alignment with prevailing AI ethics guidelines. In order to demonstrate the validity of the GAIFI, the study applies this methodology to a Catholic high school in Maryland to assess how students and faculty at that institution are adapting to GAI. The GAIFI demonstrates potential to become a practical diagnostic tool for policymakers to develop informed, ethical AI policies, and can be applied to future research across different academic environments.

Introduction

Today, society is moving towards an era of constant and rapid artificial intelligence (AI) proliferation. Generative AI (GAI) models such as ChatGPT and Gemini have completely revolutionized how individuals can collect, process, and relay information by providing users with personalized, human-like outputs according to a vast and comprehensive dataset. While this technology can prove incredibly useful to students and professors alike, many academic institutions are concerned about GAI facilitating academic misconduct (Meckler & Verma, 2022). However, it is imperative for further research to assess how academic institutions are reacting to GAI to ensure that policymakers, specifically within college settings, are navigating and adapting to this new technology in a way that is ethically conscious of GAI. This research project will contribute to the inquiry of how students, teachers, and administrators at academic institutions are currently adapting to the spread of GAI by introducing the GAI Fairness Index, an analysis process designed to effectively assess academic institutions' reactions to the proliferation of AI in terms of their alignment with prevailing AI ethics research.

This project follows a mixed methods approach with an exploratory sequential design; this is because it is necessary to analyze the qualitative traits within the institution's reactions to determine if they meet the necessary criteria for AI integration in schools, but converting these qualities into quantitative representation allows for a refined and efficient evaluation of their reactions. Processed a fairness index, a Catholic high school was assessed through survey data to validate the index. While the observed institution received a relatively neutral score, different aspects of the survey data ultimately revealed divergent perspectives on the inquiry. Nonetheless, this project contributes to the larger spectrum of GAI ethics by providing a novel methodology for assessing institutional reactions to GAI that can be used as a tool for informed decision-making and ethical GAI policies at academic institutions. This paper will explore the methods, results, and implications of the GAI Fairness Index.

Literature Review

Previous literature has primarily defined how GAI can be utilized while adhering to academic principles. This literature review aims to acknowledge some of these works to provide context on the current knowledge of GAI and how these translate to existing standpoints on AI ethics within the educational front.

Compiled from the online databases Google Scholar, EBSCO Host, and arXiv, these sources served as the foundation for establishing the design and content of the index.

GAI Fundamentals

Computer science professor Pedro Domingos (2012) elaborates on the fundamentals of AI in his paper, *A few useful things to know about machine learning*. Researchers often regard his paper as one of the most influential resources in AI due to its conciseness and accessibility, spearheading research, education, and advancement in machine learning and AI technology. In his work, Domingos distills complex machine learning into twelve fundamental concepts. His work provides researchers with essential insights into the field of AI, addressing the intrinsic roles of data, model complexity, overfitting, bias-variance trade-offs, feature engineering, and ensemble methods within the field of machine learning. Domingos's work is crucial to understanding the conceptual traits, methods, and challenges behind machine learning programming.

In the paper, *Efficient Estimation of Word Representations in Vector Space*, Mikolov et al. (2013) introduced the Word2Vec model, a foundational component of language-based GAI. This model aimed to understand representations of words within a vector space. The paper introduced two key architectures within the Word2Vec framework, Continuous Bag of Words (CBOW) and Skip-gram. The CBOW model predicts a target word based on context words whereas Skip-gram inversely predicts context words based on a target word, allowing for processors to be able to capture and understand the intricate connections between words. These models function through algebraic equations to infer semantic relationships (e.g., Paris – France + Italy = Rome). Mikolov et al.'s work significantly improved the understanding of word representations in machine-learning models, providing a novel and practical approach towards word embeddings that is still utilized in GAI models.

Ethical Considerations and Educational Applications

As exemplified by the previous sources, GAI innovation presents opportunities across a variety of fields. However, there have been notable limitations and ethical considerations to these structures that, when utilized on a large scale, can yield negative social consequences. These ideas have been at the core of AI ethics research and are a point of controversy among researchers, as there has been a diverse array of opinions as to how AI should be employed, if at all. This section of the literature review is dedicated to exploring the ethical limitations of GAI and the current perspectives of GAI in education, primarily within the context of OpenAI's ChatGPT.

As previously established, word embedding is a process in machine learning that uses mathematical expressions to denote or predict the relationships between words in a written language. However, these outputs may not always be correct. As a result, these inaccuracies have occasionally been proven to perpetuate existing societal biases and harmful stereotypes. Because of these limitations, bias has been among some of the most pressing concerns of GAI proliferation. One often cited example of this limitation is the paper, *Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings* by Bolukbasi et al. (2016). This paper demonstrates how vector arithmetic in word embeddings can unintentionally encode harmful gender norms. For example, "Actor – Man = Woman" may yield "Housewife" instead of "Actress." Bolukbasi et al.'s work highlights the potential dangers of implementing language processing content on a larger scale, and may thus pose a risk to many domains, including the educational front. While the authors introduce new debiasing techniques for word embeddings, the prevalence of bias in GAI remains, and not just in gender stereotypes: other research has shown that this issue extends to racial (Manzini et al. 2019) and political (Geng 2022) contexts as well. Furthermore, other research demonstrates GAI's limitations in producing factually accurate responses (Cheng, H.-W. 2023).

Both Rahman & Watanobe (2023) and Hwang & Chen (2023) have provided useful academic perspectives on ChatGPT's and GAI's influence in education and outlined different positive and negative ways GAI

can be utilized in educational settings. As briefly aforementioned, ChatGPT is a large language model developed by the company OpenAI. Their GPT 3 and 4 language models have emerged as particularly innovative AI technology due to their complexity and comprehensiveness. Rahman & Watanobe and Hwang & Chen illustrate ChatGPT's ability to improve and enhance the academic experience for students and teachers alike by encompassing personalized feedback, broadening the accessibility of valid resources, and providing interactive conversations, streamlined lesson preparation, evaluation mechanisms, and innovative instructional approaches to complex concepts. However, alongside its benefits, they underscore certain threats ChatGPT brings to educational research, including enabling academic dishonesty on online tests, the production of human-like text that blurs the line between human-versus-machine-produced content, and the potential to reduce critical thinking skills in over-reliant users.

The majority of mainstream media also seems to be expressing these same sentiments of ChatGPT and other language models, in the sense that they offer new opportunities while emphasizing the need for curriculum adjustment. Among the more optimistic perspectives, Villasenor (2023) advocates for the use of ChatGPT to empower students to achieve more fulfilling academic careers. Svrluga & Natanson (2023) explore how ChatGPT and other AI models have been influencing students both in and outside the classroom. Meckler & Verma (2022) highlight the concern among teachers about students using GAI for cheating. Overall, it seems that the outlook on GAI in schools is multi-faceted and presents both challenges and opportunities. It is crucial that educational policymakers can effectively integrate AI to improve their curriculum without encouraging unethical practices.

The responses among educational institutions have been just as nuanced as the media; some school systems have chosen to ban ChatGPT on all school devices and networks (Everyday AI 2023). Other academic institutions, such as the International Baccalaureate, are so lenient that they even permit its usage on school exams (Schafer 2023). One noteworthy response to ChatGPT has been Cambridge University's article, *ChatGPT: We Need to Talk* (Kirk 2023) which explores and advocates the potential of ChatGPT in a Q&A format while acknowledging and addressing its threats. Its combination of both proponential and skeptical outlooks offers a balanced and clear outline as to how students and professors should ethically utilize GAI in an educational context.

Addressing the Gap: How Are Schools Reacting?

The diverse and nuanced college responses to GAI beg the question: what *is* the best response to GAI proliferation? There cannot be one universally "correct" response, but existing literature suggests useful guidelines — that GAI is not inherently harmful, and it can conform to academic principles and further the educational front, so educational institutions should reap the useful and innovative benefits of AI while being aware of its costs. This points to a major gap in research, which is understanding in a deeper context how well educational institutions are adapting to GAI proliferation according to these guidelines.

Methodology

This case study follows a mixed-methods participatory social justice design, utilizing thematic coding and index construction methodology. The data from this experiment was qualitatively collected through surveys and public document reviews, then quantitatively processed and assessed through a fairness index based on the observed qualities the colleges possess or lack concerning adapting to GAI proliferation. Students, faculty, and administrators from the observed institution were selected for participation in this study. The selected institution for this study was a Catholic, private high school in Maryland with a student population of approximately 650. This makes collecting an adequate number of responses more efficient due to the smaller student-to-faculty ratio. To comply with ethical research practice, the name of this institution will remain undisclosed.

Data Collection

Data from the observed institution was collected on two different fronts: a survey administered to students and faculty (Front 1), and a survey administered to a high-level administrator at the school (Front 2). Both fronts are designed to assess how GAI has influenced a certain aspect of that college in relation to established GAI ethical guidelines to gain a more in-depth understanding of college adaptation to GAI. The key difference is that Front 1 asks respondents how they use GAI in specific forms, whereas Front 2 asks the administrator for which of those same forms they allow GAI to be utilized. As such, Front 1 is geared towards gaining an understanding of how the college population is utilizing GAI and how much they are conforming to their institution's GAI policy, whereas Front 2 is geared towards gaining an understanding of the ethical alignment of the institution's GAI policy with existing guidelines.

Both Fronts are surveys designed via SurveyMonkey, an online platform for creating and distributing surveys. This platform was chosen due to its organizational tools and user-friendly interface. Surveys were distributed to the entire institution's population via email. To comply with ethical research practices, all participants were required to sign an informed consent form, agreeing to their participation in the study. The surveys, through a checklist, assess how students, faculty, and policymakers have reacted to utilizing GAI in the following positive forms: research assistance, fact-checking, and learning aid. The checklist similarly assesses student, faculty, and policymaker reactions to GAI in the following negative contexts: submitting GAI content with self-credit, utilizing GAI as a sole source, and unmonitored usage on exams and papers. An equal number of positive and negative forms were assessed to ensure that the index does not mathematically tip towards one side or the other. These forms were chosen due to their relevance in identifying the nuances in GAI usage according to how Rahman & Watanobe (2023) and Hwang & Chen (2023) outline ethical GAI usage within education. Respondents were also provided the option in the survey to list other forms in which they use GAI and give additional insights on their responses. These insights will be utilized to validate the accuracy of the scores of their reactions. With this data collection method, assessing reactions in the Fronts and Forms method is logical because surveying how popularly researched forms of GAI are being used allows for a comprehensive positive-to-negative ratio of how students and faculty members are reacting to GAI in various contexts.

Data Analysis and Representation: The GAIFI

Survey data was processed through an index containing two scales of scores: the Form Score and the Front Score. Forms in this context refer to the six aforementioned forms of GAI usage that were analyzed through the checklist in the surveys. The Form Score is derived from the percent, in decimal form, of each respondent who claims to have utilized GAI in that form. This applies to both the Front 1 and Front 2 surveys. The Front Scores are the sum of all the positive Form Scores minus the negative Form Scores from each survey. Each score on the Index was graded on a scale of -10 to 10, with the former indicating poor adaptation to GAI, and the latter representing excellent adaptation. This range allows for simplistic and comprehensible numerical representations while still being able to properly display any nuances in adaptation. The scores from each Front were then mathematically averaged out to provide the college's Mean Index Score (MIS) in GAI adaptation. This data analysis method has been titled the GAI Fairness Index (GAIFI). This indexing method was chosen due to the complex quantitative and qualitative factors involved in survey data collection, making it easier to interpret a large number of responses in an organized and insightful fashion. Each form assessed in Fronts 1 and 2 will receive a numerical rating based on the alignment with established GAI ethical guidelines and then get summed together to form the individual Front Score. The formulas for the index construction process in mathematical representation go as follows:

$$\text{form score} = \frac{10(\% \text{ of respondents answered})}{3}$$

$$\text{front score} = \text{sum of positive form scores} - \text{sum of negative form scores}$$

$$\text{mean index score} = \frac{\text{front 1 score} + \text{front 2 score}}{2}$$

Overall, the goal of this formula process is to amalgamate qualitative insights from surveys and document reviews by categorizing respondents according to their use of GAI and transforming that into a numerical representation that allows for easier digestion of how the observed institution is adapting to GAI proliferation.

Limitations in Methodology

While this project is theoretically valid in identifying the intricacies within the observed institution's adaptation to GAI proliferation, it is imperative to acknowledge several limitations in the chosen methodology; first, the study's focus on one specific school may restrict the generalizability of findings beyond this institution. Furthermore, reliance on self-reported data through surveys introduces the possibility of response bias, potentially altering the accuracy of the collected information. This is especially serious because students may be less willing to admit to using AI in a harmful practice. However, anonymity is ensured and enforced to encourage respondents to be as transparent as possible.

Another potential limitation is that the GAIFI can only be valid under the assumption that there are no major differences between the scores of Fronts 1 and 2. This is because, when averaged out, that scenario will present an MIS that might not accurately depict the institution's reaction. For example, it is theoretically possible that Front 1 could score a -10 and Front 2 could score a 10. When averaged out, this would produce an MIS of 0, indicating that the institution had a neutral reaction, which would be inaccurate according to the survey data.

Despite these limitations, this design still maintains its validity; the targeted focus on one specific institution allows for an in-depth exploration of GAI adaptation within a specific context, offering valuable insights that can inform localized strategies, and while the data cannot be generalized, the proposed indexing method presents a novel approach that can be applied to analyze other academic institutions. While self-reported data may introduce biases, the triangulation of information from administrators, students, and faculty enhances the reliability of findings. Furthermore, potential scores that may skew the accuracy of index findings are easily identifiable and can still provide insight into the institution's reaction, but simply have to be judged according to the individual Front scores rather than the MIS.

Findings

The collected data reveals interesting trends regarding the observed institution's reaction to Generative AI. According to the GAIFI, the institution demonstrated an overall neutral response. Further assessment of the individual Front and Form scores suggests that students and teachers tend to utilize GAI in a way that follows existing ethical guidelines, but administrator policies are disallowing potentially useful applications of GAI. This section will highlight the GAIFI findings and additional relevant qualitative information gathered from the survey that further verify the validity of the index, followed by its limitations.

Index Findings

The collected survey data was processed through the GAIFI to quantitatively assess how the observed institution reacted to GAI. Below is the survey data processed through the index:

Table 1. GAI Fairness Index Findings

GAI Forms	Front 1	Front 2	MIS
Research Assistance (+)	0.9	0	0.45
Fact-Checking (+)	0.9	0	0.45
Learning Aid (+)	0.8	3.3	2.05
Self-Credit (-)	0	0	0
Sole Source (-)	0.1	0	0.05
Unmonitored Usage (-)	0	0	0
Total	2.5	3.3	2.9

As denoted by the table, the observed institution received an MIS of 2.9. This indicates that the respondent pool had a relatively neutral reaction to GAI with a slight lean towards ethical habits. This score interpretation is supported by the survey data as an overwhelming amount (39.39%) of respondents stated that they do not utilize GAI at all, contributing to the low scores observed in Front 1. Moreover, most of the respondents in Front 1 (96.97%) who utilized GAI strictly did so within the surveyed positive forms (research assistance, fact-checking, and learning aid). The most significant detractor from the institution's MIS was from Front 2, as the surveyed administrator only allowed for GAI to be utilized within the context of Learning Aid, decreasing the overall MIS. While ultimately undeterminable from the available data, this may suggest that GAI policies at the observed institution may be too strict, preventing students and faculty from using AI in the most ethically optimal manner. With all the aforementioned data taken into consideration, it can be reasonably argued that the statistics from the survey successfully assert that the index is accurate in assessing the reactions among the pool of respondents.

Additional Findings

The additional clarifications provided by survey participants offer a deeper insight into how GAI is interpreted at this institution. Firstly, approximately 38% of respondents who provided additional clarifications mentioned some form of distrust in GAI's capabilities to produce a reliable output, which may imply distrust could be a leading factor behind the 39.39% of respondents who do not use GAI. One particular response that rationalizes this distrust in GAI goes as follows:

I'm a student who has extensive experience in the field of GAI advocacy, education, and deployment. This makes it easier for me to know that GAI hallucinates and doesn't produce factual data, but many students may not know that.

This surveyed student offers a rationale against GAI usage within an academic context which is in alignment with GAI's limitations outlined by Rahman & Watanobe (2023) and Hwang & Chen (2023). Their response exemplifies why people may hesitate or not utilize GAI due to its inability to consistently produce factual information. However, because 23% of students provided a similar line of reasoning for not utilizing GAI, combined with the fact that 96.97% of respondents reported utilizing GAI strictly within a positive context, there may be an overall greater awareness of GAI knowledge, contrary to what the respondent originally

implied. Nonetheless, the respondent's rationale seems to be one of the most popular reasons against GAI usage in an academic context.

Another primary reason against GAI usage according to the additional information was that some respondents (37.7%) felt it was unnecessary in their line of work. This pattern was found strictly among faculty, implying that students at this institution have a greater need for GAI than teachers. The reason behind this is inconclusive with the provided data, but it may be due to an accumulation of various factors such as students' relative inexperience in the subject of their courses as compared to teachers, and differences in workload, among other reasons. However, that is not to say that teachers did not utilize GAI entirely, as roughly 42.3% of teacher respondents reported utilizing GAI in at least one of the aforementioned forms.

It is important to note that because this is a case study, there are also a variety of internal characteristics that may have uniquely influenced GAI usage at this institution. Firstly, the institution's Catholic alignment opens the possibility for the respondents to have been more influenced by Catholic social teachings, which could potentially encourage more ethical practices of GAI among the student and faculty population. Secondly, the institution provides students with an iPad, making GAI resources further accessible. This makes the relatively neutral MIS surprising because more access to GAI technology should reasonably yield a more extreme reaction. However, it is also possible that these results may have been influenced by networking restrictions, as the observed institution's student Wi-Fi network has firewalls and IP address restrictions that block websites with GAI tools. However, these ideas are all speculation and ultimately difficult to substantiate due to lack of evidence and no other institution to compare the studied institution to. It would be interesting to see how these index results compare with a public, secular school, or a school that does not offer technological services to its students.

Limitations in Findings

Although the observed findings present compelling information regarding GAI reactions, it is important to clarify the information that the index and collected data do not represent. Firstly, the MIS should not be generalized to other institutions due to the small respondent pool (39) and the fact that only one institution was observed. As established, reactions to GAI can differ based on cultural and demographic contexts that are not outlined in the survey. However, the index still maintains validity in ranking a group's reaction to GAI within academic contexts as it accurately scores and analyzes the data.

Moreover, the GAIFI only analyzes six ways that GAI can be used. A small number of forms of GAI usage were selected according to research from Rahman & Watanobe (2023) and Hwang & Chen (2023) to encourage a higher participation rate, but this comes at the expense of potentially excluding other ways GAI could have been utilized. In attempts to resolve this, participants were provided with an "Other" box to list additional ways in which they use GAI, but most respondents chose not to utilize this option, serving no broader significance.

Discussion

While the observed institution scored a 2.9, suggesting an overall positive reaction to GAI, it is difficult to arrive at one precise conclusion to the research inquiry, as interpretations of what is "ethical" vary greatly from person to person, especially in the field of AI ethics. Furthermore, the presented findings can be distilled to support numerous interpretations of how students and teachers are ethically adapting to GAI depending on one's personal beliefs and what aspects of the findings are highlighted. To address these nuances, this section will elaborate on various interpretations that can be derived from the data. This section will also discuss what these findings entail for future research on GAI and its ethical implications in the educational field.

Interpretations

Firstly, the presented findings can support the idea that students and teachers at the observed institution are ethically adapting to GAI. The MIS of 2.9 indicates that, while minimal, responses among students and faculty tend to lean towards positive habits that align with the insights provided by Rahman & Watanobe (2023) and Hwang & Chen (2023). Also, 96.97% of students reported using GAI strictly within these positive academic forms and many respondents who provided additional insights demonstrated a heightened awareness of GAI's limitations. Furthermore, the interviewed administrator representative in Front 2 did not permit GAI to be utilized in any negative contexts. This perspective is valuable because it highlights the benefits of GAI for education and how this institution is utilizing them. Moreover, this perspective disproves some potential skepticism regarding GAI, such as the perception that students do not know how to use GAI properly or that it cannot benefit the classroom.

While the findings support the notion that students and teachers are ethically adapting to GAI at this institution, some of the data can be used to suggest the opposite; the primary argumentation behind this perspective lies in Front 2, as the interviewed administrator representative only allowed for GAI to be utilized for learning aid, disallowing all other listed positive forms. Those against the idea that students and teachers are ethically adapting to GAI at this institution would likely emphasize the restrictive policies outlined in the findings. This restrictiveness poses an ethical dilemma for students and teachers because their ethical GAI usage is not in alignment with their institution's policies, despite it being justified according to Rahman & Watanobe (2023) and Hwang & Chen (2023). Contrary to the former perspective, this viewpoint highlights the flaws of the institution's reaction to GAI, demonstrating the need for lighter policies. Overall, this perspective shows that the observed institution is not ethically adapting to GAI.

Another interesting interpretation of the data is that students and teachers at this institution may not be reacting to GAI as much as originally anticipated to effectively determine whether their reaction was either ethical or unethical; this is supported by the MIS of 2.9, which, while positive, indicates a more neutral reaction. Furthermore, 39.39% of respondents reported not utilizing GAI at all, and some teachers and faculty reported GAI not being necessary in their field of work. Additionally, the small respondent pool, despite the survey being sent out to the entire student and faculty population, could imply that individuals at this institution simply did not engage enough with GAI to provide any insights. However, this is just speculation and is ultimately not supported by the data, as the choice for individuals to not participate could be due to various reasons. Nonetheless, the provided data may also imply that the institution's reaction to GAI was too minuscule to be considered ethical or unethical. This is interesting as compared to the previous two interpretations because it contradicts the previous speculation for GAI to have an extreme impact on educational institutions, as outlined by the Literature Review.

In short, the multifaceted nature of ethical considerations surrounding GAI adoption in educational settings presents a complex variety of interpretations depending on what aspects of the data are emphasized. Although the MIS of 2.9 suggests a generally neutral outlook with generally positive tendencies, ethical interpretations can still diverge significantly, reflecting the inherent subjectivity in defining ethical practices within AI. These interpretations emphasize the need for nuanced analysis to accurately address ethical dilemmas and optimize GAI's potential in education.

Conclusion and Implications

The key findings collected from the study show that the observed institution scored a 2.9, suggesting an outwardly neutral reaction with an inclination towards ethical habits. These findings could cultivate a significant impact on the research community for AI ethics by presenting a novel form of analyzing GAI reactions amongst educational institutions, facilitating the development of ethical GAI policies. These provide crucial implications

for executives at the observed institution, as they serve as a justification for enacting a lighter policy on GAI, like permitting access to GAI through the student network.

However, while the GAI findings are only applicable to the observed institution, the implications of this study extend far beyond it; as established, one of the major limitations of this study is that it is non-generalizable to other academic institutions due to environmental factors that may influence one's adaptation to GAI. However, the GAI Fairness Index has applicable insights and guidelines for future research and policymaking in AI ethics; it contributes to a novel understanding of GAI ethics by presenting a promising methodology for assessing GAI adaptation across different educational settings, facilitating informed decision-making and the development of ethically aligned policies and practices. Future researchers could also include more Fronts and Forms to analyze for improved accuracy of future scores. Moreover, this study opens the potential for future researchers and policymakers to observe how reactions to GAI differ between academic institutions with different characteristics (e.g. prestige, location, access to technology, etc.) using the presented GAIFI. As such, applying this index on a broader spectrum can provide academic researchers and policymakers with the insight required to optimally calibrate their institutions to established guidelines on GAI ethics.

In summary, the GAIFI proved to be an insightful method for examining the adaptation of educational institutions to GAI proliferation, as it proved to be valuable and accurate in evaluating the ethical alignment of GAI policies and practices. Findings from the index and additional qualitative insights from respondents revealed diverse perspectives and interpretations regarding GAI adoption. Consequently, while the observed institution demonstrated a relatively neutral response through its MIS of 2.9, other aspects of the collected data can reveal different interpretations, varying from assertions of ethical adaptation to concerns over institutional policies and the necessity of GAI usage. These divergent perspectives highlight the need for a multi-faceted analysis of the GAIFI and the consideration of contextual factors while evaluating GAI adoption and its ethical implications. Overall, this study contributes to a more refined understanding of the complexities inherent in GAI adoption within educational institutions and underscores the importance of incorporating a comprehensive approach toward understanding how educational institutions are adapting to GAI while harnessing the potential of AI technology for learning and innovation.

Acknowledgments

I would like to thank my advisor for the valuable insight provided to me on this topic.

References

- Bolukbasi, T., Chang, K. W., Zou, J. Y., Saligrama, V., Kalai, A. T. (2016). *Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings*. In Advances in Neural Information Processing Systems (NeurIPS), 29. <https://doi.org/10.48550/arXiv.1607.06520>.
- Cheng, H.-W. (2023). Challenges and limitations of CHATGPT and Artificial Intelligence for scientific research: A perspective from organic materials. *AI*, 4(2), 401–405. <https://doi.org/10.3390/ai4020021>.
- Domingos, P. (2012). *A Few Useful Things to Know About Machine Learning*. Communications of the ACM, 55(10), 78-87. <https://doi.org/10.1145/2347736.2347755>.
- Everyday AI. (2023, August 25). Back to School: The Students vs. Teachers and ChatGPT Dilemma. *Everyday AI Podcast - An AI and ChatGPT Podcast*. Episode 88. <https://www.youeverydayai.com/back-to-school-the-students-vs-teachers-and-chatgpt-dilemma/>.
- Geng, Y. (2022). Media bias detecting based on word embedding. *Highlights in Science, Engineering and Technology*, 12, 61–67. <https://doi.org/10.54097/hset.v12i.1367>.

- Hwang, G.-J., & Chen, N.-S. (2023). *Editorial Position Paper: Exploring the Potential of Generative Artificial Intelligence in Education: Applications, Challenges, and Future Research Directions.*, 26(2), 1–18. [https://doi.org/10.30191/ETS.202304_26\(2\).0014](https://doi.org/10.30191/ETS.202304_26(2).0014).
- Kirk, T. (2023, April 5). *ChatGPT (We need to talk)*. Retrieved from <https://www.cam.ac.uk/stories/ChatGPT-and-education>.
- Manzini, T., Yao Chong, L., Black, A. W., & Tsvetkov, Y. (2019). Black is to criminal as Caucasian is to police: Detecting and removing multiclass bias in word embeddings. *Proceedings of the 2019 Conference of the North*. <https://doi.org/10.18653/v1/n19-1062>.
- Meckler, L., & Verma, P. (2022, December 28). Teachers are on alert for inevitable cheating after release of ChatGPT. *Washington Post*. <https://www.washingtonpost.com/education/2022/12/28/chatbot-cheating-ai-chatbotgpt-teachers/>.
- Mikolov, T., Chen, K., Corrado, G., Dean, J. (2013, January 16). *Efficient Estimation of Word Representations in Vector Space*. In Proceedings of the 1st International Conference on Learning Representations (ICLR).
- Rahman, Md. M., & Watanobe, Y. (2023). Chatgpt for Education and research: Opportunities, threats, and strategies. *Applied Sciences*, 13(9), 5783. <https://doi.org/10.3390/app13095783>.
- Schafer, J. (2023, March 3). ChatGPT Unbanned in School Exams. *AI Applied: Covering AI News, Interviews and Tools*. Episode 14. <https://podcasts.apple.com/us/podcast/ai-applied-covering-ai-news-interviews-and-tools/id1669799110?i=1000602679206>.
- Svrluga, S., & Natanson, H. (2023, June 1). *All the unexpected ways ChatGPT is infiltrating student lives*. The Washington Post. Retrieved from <https://www.washingtonpost.com/education/2023/06/01/students-chatgpt-ai-tools/>.
- Villasenor, J. (2023, February 10). *How ChatGPT can improve education, not threaten it*. Scientific American. <https://www.scientificamerican.com/article/how-chatgpt-can-improve-education-not-threaten-it/>.