EXPLORING THE IMPACT OF GENERATIVE AI ON EDUCATION:

Opportunities, Challenges, and Ethical Considerations

Dr. Sidney Shapiro

Faculty, Dhillon School of Business - Policy & Strategy, University of Lethbridge

Abstract

This paper examines the role of generative Artificial Intelligence (AI) in higher education, highlighting its potential uses, obstacles, and ethical considerations. The paper introduces generative AI technology, its functions, and the importance of quality management in AI models for educational purposes. The article discusses AI's roles in facilitating learning, promoting collaboration, and offering personalized tutoring. It also delves into the effects of AI on teaching, emphasizing personalized teaching methods, student feedback, and challenges faced by educators.

Introduction

The evolution of artificial intelligence (AI) has dramatically accelerated, transitioning from a niche element within computer science to a potentially all-encompassing force. Various factors have catalyzed this transformation: the advent of larger, more complex models trained on vast datasets, new technologies, increased user-friendly applications for non-specialists, and the growing range of impressive capabilities AI has to offer—particularly AI-powered chatbots[1], which existed for decades but are now revolutionizing how we interact[2].

Chatbots have significantly advanced from their primitive origins, now exhibiting capabilities far beyond merely accessing and processing data or following predictable pre-trained pathways. Today's chatbots leverage advanced communication methodologies and impersonate an air of authenticity and authority with a unique and engaging personality. We have surpassed the goal of the original Turing Test - creating machines that convincingly mimic human dialogue[3] [4]. Modern AI counterparts reflect human conversation with an intriguing mix of intellect, humour, and charisma, as demonstrated by applications like CupidBot[5], which has automated online dating, or Replika, an AI friendship[6] and companionship bot[7].

Human involvement remains pivotal, underscoring the irreplaceable value of human engagement in learning. Human intervention and monitoring in AI go beyond programming; it involves setting up safeguards[8], curating data[9], and instructing AI systems to discern between beneficial and harmful information[10]. Building advanced AI is a complex, labour-intensive endeavour that could expose workers to exploitation[11] [12]. Human costs include unknowing contributors (users

Arthur C. Clarke's quote that "any sufficiently advanced technology is indistinguishable from magic"[13] resonates in this context, with AI technology seeming nothing short of magical. AI's potential to reshape our lives, alter the work landscape, and transform education are all some seem to talk about. Unlike humans with the capacity to think flexibly for themselves, AI systems can process and sort endless streams of data and create knowledge by shuffling existing data. To generate high-value content, AI systems need to learn about context or sometimes make unpredictable choices that may not seem logical. Sometimes unpredictable choices are wonderful and give us "infinite diversity in infinite combinations"[14].

This paper attempts to clarify and demystify AI and specifically explore its potential impact on teaching and learning. Current AI models include generative text, also known as large language models. The speed at which generative AI has been adopted has caused a profound impact on the academy[15] and on society[16] [17].

Generative AI

AI technologies include but are not limited to areas like machine learning, machine vision, and natural language processing. New techniques including transformers, deep learning, and neural networks have brought forth a wave of innovation, transforming how we engage with data and information. Recently, the AI landscape has seen a significant shift with the emergence of generative AI, spurred on by advancements in specialized hardware and software technologies which support the construction of large language models (LLMs).
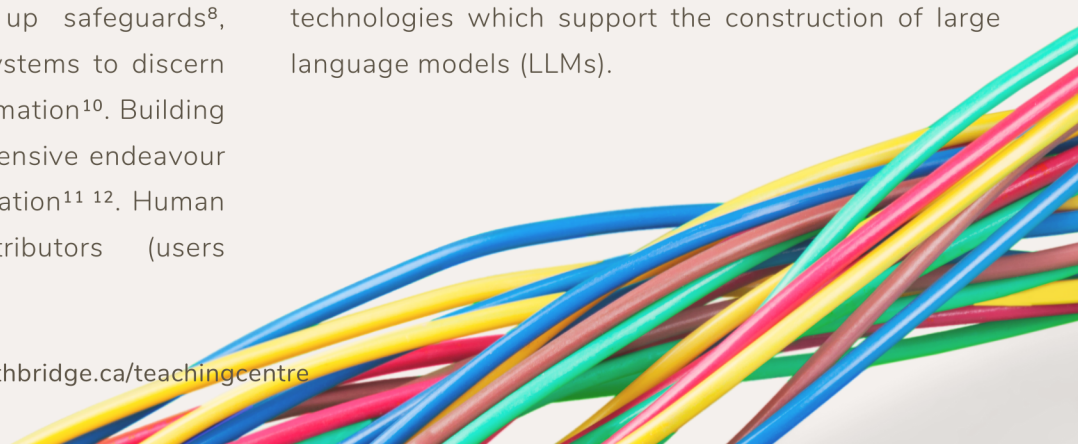
LLMs function as an extensive web of mathematical relationships among many words. Their core functionality is rooted in the idea of prediction for both text and image data. Similar to how we might anticipate a friend's behaviour based on repeated patterns, LLMs are designed to anticipate what comes next in a data sequence based on previous patterns. With large amounts of data, models can become more "knowledgeable" and complex. These systems use the idea of predictive text that we currently use on smartphones and email but at a much larger scale, generating more complex sequences like sentences, paragraphs, pages, and entire conversations.

It's pivotal to emphasize that generative AI doesn't create new knowledge from scratch. Rather, it draws predictions from existing data patterns. This is much like understanding a musical genre and then composing a fresh tune that resonates with that genre's traits. Both the processes of discovery and teaching are fluid, often veering away from established patterns to make space for unexpected turns and creative shifts. Generative AI excels in sifting through massive data sets in unique ways, like detecting trends in a collection of resumes and predicting how new data might mesh with that trend. Yet, its scope is limited to the boundaries of existing information and doesn't explore the unknown realms of discovery. The true strength of generative AI emerges when we use it to format and showcase information in a new light, setting the stage for novel insights and deeper comprehension.

Generative AI is making waves across numerous domains, spanning from generating text and code to crafting images and videos, converting speech to text, creating 3D models, game design, and music composition. While its implementation in many of these areas remains in its early stages, we haven't yet reached a juncture where a wholly original and captivating TV show script can be whipped up with a simple click. Meanwhile, with TV writers currently on strike[18], there's growing apprehension about the potential transformation of their craft. This shift might see AI systems recycling old scripts into formulaic patterns, potentially sidelining human ingenuity.

A future perspective on employment suggests a notable shift, where workers and students harness AI to bolster their workflow and enhance their innate capabilities. This trajectory has both technological and societal ramifications, with many expressing concerns that swift AI advancements could culminate in societal challenges[19]. Although AI innovations are reshaping, rather than replacing, human tasks[20], they do provide avenues for automating specific functions. Nevertheless, human intervention remains pivotal, especially in sectors that demand analytical prowess, inventiveness, and social interactions[21 22]. Despite AI's impressive leaps, human workers retain a crucial role in offering the intricate understanding and compassionate communication that AI is yet to achieve.
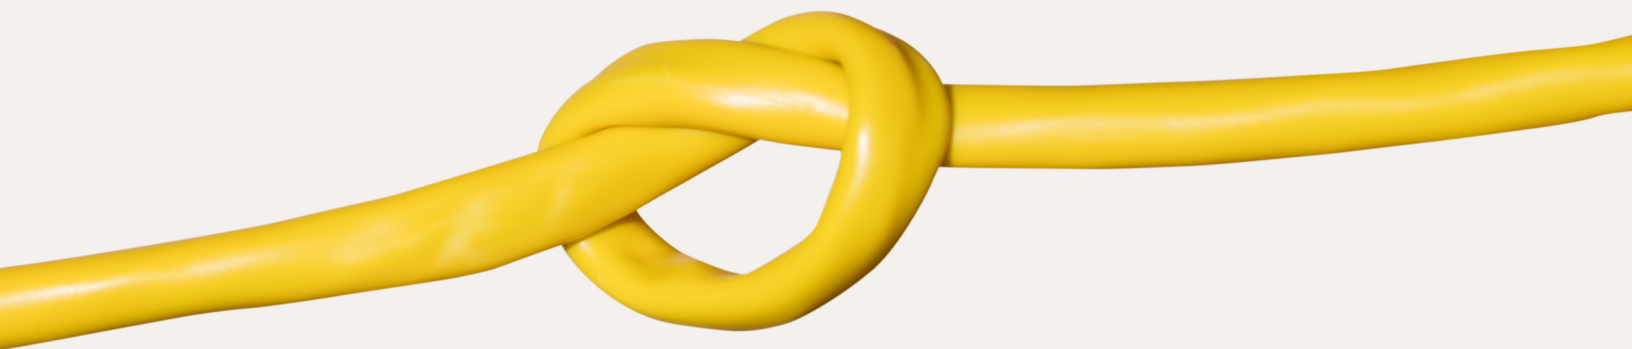
Quality Control in AI Models

The idea of "faded copies" offers a vivid analogy for understanding AI's operational mechanics and the paramount importance of the data it ingests. Imagine taking "a piece of paper and repeatedly copying it on a photocopier; with each subsequent copy, the quality degrades"[23]. This simile mirrors how AI language models work. These models' integrity, dependability, and precision hinge heavily on their training data. Garbled input predictably results in garbled output.

The repercussions of this mechanism become especially alarming when AI propagates erroneous information. Consider a website with multiple solutions to a query but only a single correct response. Without discerning which data is accurate and prioritizing it, this data can be recycled into the system, furthering the spread of false information. Beyond such benign instances, the deliberate input of malicious data amplifies the risk of disseminating incorrect or detrimental content. AI can inadvertently magnify falsehoods and misleading information without vigilant oversight and curating. Another looming hurdle in the AI domain is the phenomenon termed "hallucinations"[24], where the AI fabricates plausible but inaccurate data[25]. This becomes gravely important when dealing with factual content like medical conditions, flight schedules, or product prices. AI can swiftly disseminate false data without a verified truth reference, potentially leading to harm and miscomprehension.

AI chatbots are being developed to use factual data during discussions, bridging the gap between generative conversation and fact-checking[26]. There is a risk that people might perceive AI-generated content as truth, potentially leading to the spread of further misinformation[27]. This underlines the importance of robust and accurate training data as the foundation of AI systems. Through proper training, safe practices, and reliable databases, it is possible to harness the true potential of AI while safeguarding against risks.

AI in Academic Research

Drawing a comparison between the nuanced structures of language and the fundamental principles underpinning scientific exploration might seem apt on the surface. Yet, this comparison plunges us into a deeper philosophical mire: is our daily reality merely a reflection of our linguistic descriptions, or is there a more complex mechanism at work? This mechanism could demand a more profound grasp of reality, subsequently shaping our perceptions of the world. Beyond the realms of philosophical discourse, one might posit that by recognizing consistent linguistic patterns, we could be better equipped to forge the right words, possibly leading to revolutionary insights, or revealing previously unnoticed real-world connections. However, real-world data implies that this task is more intricate than it might initially seem.

A recent project by Meta using the Galactica model aimed to compile knowledge from various libraries of published articles to identify gaps[28]. Though the AI model successfully gathered a wide range of

knowledge, it needed help differentiating between important insights and irrelevant content. This limitation highlights that many modern AIs see information primarily as word and phrase patterns, sometimes missing more profound or critical concepts. Such limitations can result in misinterpretations. Using AI for routine tasks like evaluating grant proposals or writing academic papers can be beneficial. If used correctly, these systems can enhance human capabilities, offering consistent quality in outcomes. Just as we rely on spellcheck for writing, a sophisticated language AI could serve a broader, refined function.

AI's rise has led to academic concerns about its impact on human roles, especially in preserving knowledge and upholding educational values. Students can now produce essays, and researchers can write papers with minimal effort using AI. However, the unique human ability to connect diverse information in creative ways is something AI has not yet mastered. Using unsourced data or content that isn't fact-based raises ethical issues[29]. As AI changes traditional methods like student testing, we must adapt and assess its benefits and limits. For example, using models to generate new science shows AI's challenges in scientific discovery[30]. While skilled at combining knowledge, models are not always distinguish valid hypotheses from baseless claims[31]. Believing such misleading information can have harmful consequences[32]. This highlights the importance of human judgment in evaluating and curating information.

AI has the potential to aid in grant applications[33], peer reviews[34], and academic article writing[29]. Generative AI can act as advanced spellcheck, correcting errors, logical errors and suggesting improvements without creating entirely new content. Ideally, human creativity will combine with AI's pattern recognition for the best results. Achieving this balance requires changes in public policies and societal views on AI. While there are worries about AI taking over academic roles, it can't match humans in combining diverse data into new insights. Ethical issues around AI, particularly when using unsourced data, need attention. As AI impacts traditional teaching practices and student learning, we must reassess and adjust our methods, recognizing AI's benefits and limitations.

AI in Education

AI's potential influence in education extends far beyond mere automation. It has the capability to revolutionize the entire educational landscape, creating personalized learning experiences and optimized curricula[35]. As modern jobs become more intertwined with technology, it's imperative for educational systems to stay in step. This means introducing students to the tools and technologies they'll encounter in their careers and preparing them to think critically about these tools and their implications.

Even as the pace of technological change accelerates in the corporate world, educational institutions often face challenges in keeping up. Bureaucratic hurdles, limited resources, and a natural inclination towards time-tested methods can slow down the integration of AI technologies. Institutions recognize that they must embrace these changes to remain relevant, even if it means overhauling long-standing systems and methodologies.

This transition to an AI-integrated educational system isn't just about the technology itself but also about how it's used. For instance, with the advent of AI tools capable of generating essays, the very nature of assignments may need to shift. It becomes less about the product and more about the process. Educators may need to emphasize the student's thought processes, ability to critically evaluate AI suggestions, and ability to synthesize information from multiple sources, including AI-generated content.

Successfully ushering in this new era of AI-driven education requires vision, collaboration, and innovation from all stakeholders, including educators, technologists, policymakers, and students. In a fully realized AI-enhanced educational environment, AI can serve various roles to enhance the learning experience. Imagine an AI that can brainstorm various potential essay topics tailored to a student's interests, challenge students with thought-provoking questions to deepen their comprehension, or adjust course materials in real-time to match a student's unique learning curve. Such a system wouldn't just be about technology and harnessing technology to foster genuine understanding and lifelong learning.

AI has the capability to function as a tailored tutor, which is particularly beneficial for students seeking assistance in language-related tasks. AI can provide targeted support to students in challenging areas by analyzing individual needs and learning patterns. Furthermore, both students and educators can leverage AI to define clear course objectives, establish transparent evaluation standards, and create content that aligns with each student's unique learning style and pace.

Considering the diverse ways students absorb information, AI's adaptability proves invaluable. For instance, those who favour auditory learning or have educational requirements can benefit from AI's ability to transform traditional written materials. Typically, the output is presented in text and using various tools can be converted into audio presentations, making the content more accessible and engaging for these learners.

AI can offer students a safe and supportive environment where students can experiment with different study techniques to discover what works best for them. It can also serve as an analytical tool, keeping track of student's progress over time, highlighting areas of improvement and those needing more attention. By providing consistent feedback and recommendations, AI can encourage students to remain committed to their learning journey, ensuring they meet their educational goals.

While the potential benefits of AI in education are evident, its integration is still a work in progress. However, given the ongoing advancements in the field, these tools are anticipated to become commonplace in educational settings in the coming years. Transitioning to a more technologically driven education system necessitates preparedness from all involved in the academic sector. It's essential that as we move forward, there's a collective effort from educators, students, administrative staff, and technology professionals. Open dialogues and collaborative planning sessions will be crucial. Such

collective efforts will not only help in determining the most effective applications of AI in education but also in foreseeing potential challenges. By working together, the academic community can ensure that AI's role in higher education is effective and beneficial for all stakeholders.

Challenges and Ethical Concerns in Implementing AI

Introducing AI in academic settings brings several challenges and ethical considerations, especially concerning data origins and potential biases. It's essential to recognize that data used in many academic AI systems might reflect social biases, carrying forward inherent discriminatory or exclusionary traits. To address these challenges, institutions should establish policies and guidelines that emphasize using AI data that aligns with fairness principles and represents diverse demographics. Ensuring the data's inclusivity and understanding its broader context helps prevent the misuse of AI for spreading misleading information. Instead, the goal should be to present a clear, fact-driven model suitable for educational purposes.

The emergence of AI technologies in academia has heightened concerns about academic integrity. AI tools can produce high-standard content, presenting advantages and pitfalls. While they can benefit learning and research, there's potential for misuse, such as students presenting AI-generated work as their own. Institutions must create clear guidelines on using AI in academics, educating students on harnessing these tools without compromising intellectual honesty. This would involve defining the scope of AI's role in academic projects and emphasizing the significance of students' original input.

With the increasing prevalence of AI, there are concerns about individual privacy. AI models often need substantial data to perform efficiently, and there's a risk that this could encroach on personal privacy. The information fed into these models might not only be stored but could also be used in subsequent model training. Given this, the possibility of sensitive data being accessible in different contexts to various users exists. It's vital to have stringent privacy policies and secure mechanisms to guard against the inappropriate use of confidential data. AI tool development should prioritize data anonymization and secure data storage to uphold user trust and privacy.

From my experience, delegating specific responsibilities to students within the class yields tangible benefits. For example, in certain courses I have taught, I assign a different student each week to oversee the interactions of the online group and promptly relay any questions or concerns to me. This approach grants me the flexibility to move around the physical classroom, deliver the lecture, and ensure that important points are addressed promptly, even when my focus is not directly on the online screen.

Challenges related to bias, gender, and diversity in AI are paramount. Data biases can cause AI systems to make skewed judgments, potentially favouring certain groups. It's essential for educational institutions to ensure their AI tools don't perpetuate biases and are inclusive in terms of gender and diversity, encapsulating a broad range of human experiences. Incorporating and respecting diverse knowledge systems, like indigenous wisdom, into AI models is also crucial. These technologies should not just adapt but actively promote such knowledge, creating a more inclusive and varied academic environment. Such a practice celebrates cultural diversity, recognizes the

richness of varied knowledge bases, and contributes to a comprehensive and well-rounded educational experience.

AI and the Classroom

As AI becomes more prevalent in educational settings, it's vital to strike a balance between technology and the human touch. Despite being in their infancy, AI tools have shown tremendous promise in revolutionizing how we teach and learn. They come with a set of challenges, especially regarding ensuring academic integrity. For example, with AI's ability to craft and proofread essays, there's a heightened concern about academic misconduct and plagiarism. As a result, educators might need to rethink assessment methods. Instead of solely relying on essays, they could incorporate oral presentations or supervised written tests, allowing students to showcase their understanding.

AI can also aid in designing teaching materials. Consider the creation of rubrics: instead of teachers spending hours determining the grading criteria, AI can suggest comprehensive rubrics based on the learning objectives, ensuring clarity and fairness. Similarly, AI can analyze course content in designing quizzes and generate questions that appropriately gauge students' understanding. AI can assist in marking papers by comparing student submissions to set benchmarks, ensuring consistent and unbiased grading. Additionally, AI can create visually appealing lecture slides, ensuring that key points are highlighted. It can even build student guides, summarizing key materials and offering practice questions tailored to individual learning paths.

A particularly promising feature of AI is its ability to give personalized feedback. For instance, after a student submits an assignment, AI can analyze the content about the assignment's goals. This analysis allows for feedback that's not just about right or wrong answers but delves deep into the nuances of a student's work, offering insights on areas of improvement. Imagine a student receiving feedback that doesn't just state an error but explains why it's wrong and how to correct it. Integrating AI into education is not about replacing teachers but equipping them with powerful tools to enhance teaching and learning. Over time, as AI continues to mature and refine, it will provide both students and educators with innovative and effective means to facilitate the educational journey.

Practical Considerations for AI Implementation

Integrating AI into educational settings requires careful preparation, including setting up the necessary technological infrastructure. AI models need specialized tools and considerable computing resources. However, education may only sometimes need broad, all-purpose AI models. Instead, more specific applications, like chatbots designed using course content, might be more suitable. These focused AI tools can work well on regular computer systems, making them more user-friendly and responsive.

For AI to be effective in classrooms, it's essential for both teachers and students to have a basic understanding of how these systems function. The output from AI largely depends on the quality of the input data. Sometimes, AI might process or connect data differently than expected, needing corrections or additional clarifications. AI can produce results more relevant to the educational topic by providing more detailed context or metadata.

It's also crucial to have checks in place to review and manage the content produced by AI, ensuring it meets educational standards. As advancements in AI technology continue, the expenses related to creating and using these models are expected to decrease. This trend points to AI tools' increasing feasibility and utility in education, especially those designed for specific teaching purposes.

Conclusion

AI is becoming increasingly common in various sectors, including education and business. Students use AI to assist in their assignments and broaden their understanding. Companies are integrating AI into their services, reflecting the rapid growth in this domain. However, it's important to understand that AI, despite its capabilities, doesn't have general intelligence or emotions and can sometimes be incorrect. The rise of AI prompts discussions about job automation and its impact on the workforce. AI has become a notable aspect of technology and is here to stay.

Incorporating AI into educational settings requires careful thought. Using AI effectively is essential while preserving the human aspects of teaching and learning. By adopting a balanced approach to AI, we can ensure it complements human abilities, such as creativity, empathy, and critical reasoning, rather than replacing them.

Note

This paper's outline, Exploring Generative AI in Education, was presented at the University of Lethbridge Spark Teaching Symposium on May 3rd, 2023.
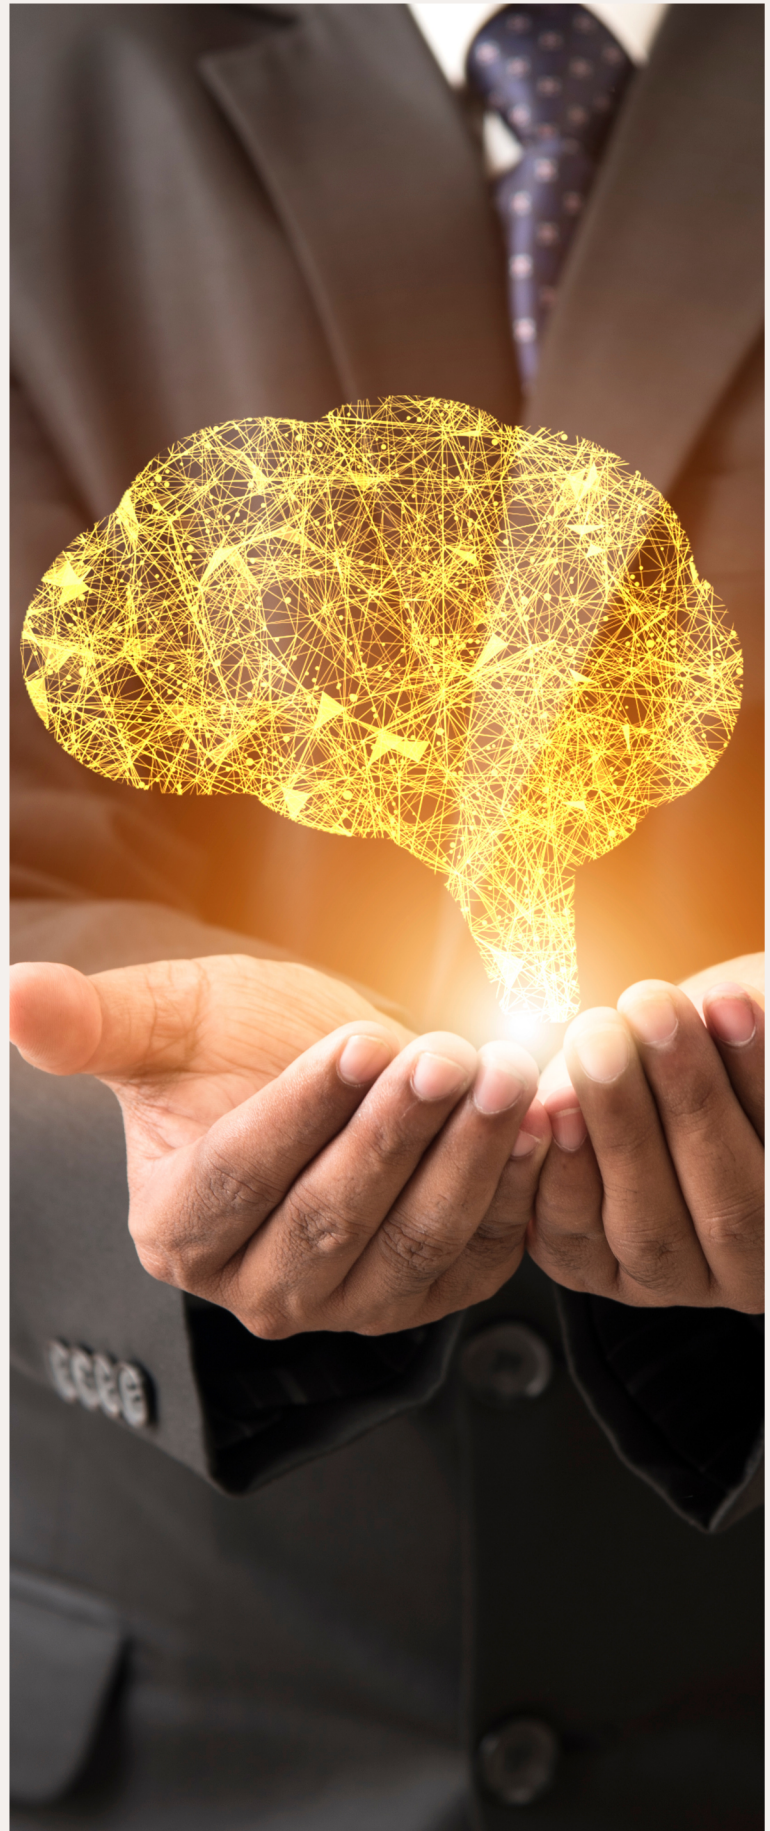
Works Cited

1. George, A. S. & George, A. H. A review of ChatGPT AI's impact on several business sectors. *Partners Universal International Innovation Journal* 1, 9–23 (2023).

2. Bryant, A. AI Chatbots: Threat or Opportunity? in vol. 10 49 (MDPI, 2023).

3. French, R. M. The Turing Test: the first 50 years. *Trends in cognitive sciences* 4, 115–122 (2000).

4. Natale, S. *Deceitful media: Artificial intelligence and social life after the Turing test*. (Oxford University Press, USA, 2021).

5. Scucci, R. AI Chatbots Are Now Being Used To Flirt On Dating Apps. *Giant Freakin Robot* https://www.giantfreakinrobot.com/tech/ai-chatbots-flirt-dating-apps.html (2023).

6. Brandtzaeg, P. B., Skjuve, M. & Følstad, A. My AI friend: How users of a social chatbot understand their human–AI friendship. *Human Communication Research* 48, 404–429 (2022).

7. Xie, T. & Pentina, I. Attachment theory as a framework to understand relationships with social chatbots: a case study of Replika. (2022).

8. Bærøe, K., Miyata-Sturm, A. & Henden, E. How to achieve trustworthy artificial intelligence for health. *Bulletin of the World Health Organization* 98, 257 (2020).

9. Muller, M. *et al*. How data science workers work with data: Discovery, capture, curation, design, creation. in 1–15 (2019).

10. Fügener, A., Grahl, J., Gupta, A. & Ketter, W. Will humans-in-the-loop become borgs? Merits and pitfalls of working with AI. *Management Information Systems Quarterly (MISQ)-Vol 45*, (2021).

11. Altenried, M. The platform as factory: Crowdwork and the hidden labour behind artificial intelligence. *Capital & Class* 44, 145–158 (2020).

12. Chan, A., Okolo, C. T., Terner, Z. & Wang, A. The limits of global inclusion in AI development. *arXiv preprint arXiv:2102.01265* (2021).

13. Clarke, A. C. *Profiles of the Future: An Inquiry Into the Limits of the Possible.* (Holt, Rinehart, and Winston, 1984).

14. IMDB. 'Star Trek: The Animated Series' The Infinite Vulcan (TV Episode 1973) - IMDb. http://www.imdb.com/title/tt0832421/characters/nm0000559 (2023).

15. Dwivedi, Y. K. *et al*. "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management* 71, 102642 (2023).

16. Hacker, P., Engel, A. & Mauer, M. Regulating ChatGPT and other large generative AI models. *arXiv preprint arXiv:2302.02337* (2023).

17. Stokel-Walker, C. & Van Noorden, R. What ChatGPT and Generative AI Mean for Science. *Nature* 614, 214–216 (2023).

18. Chmielewski, D. Striking Hollywood writers expected to respond to studios' proposal | Reuters. https://www.reuters.com/world/us/striking-hollywood-writers-expected-respond-studios-proposal-2023-08-15/ (2023).

19. Weidinger, L. *et al*. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359* (2021).

20. Yarlagadda, R. T. Future of Robots, AI and Automation in the United States. *IEJRD-International Multidisciplinary Journal* 1, 6 (2015).

21. Daugherty, P. R. & Wilson, H. J. *Human+ machine: Reimagining work in the age of AI*. (Harvard Business Press, 2018).

22. Dwivedi, Y. K. *et al*. Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management* 57, 101994 (2021).

23. CBC News. 'I am terrified,' How 2 Sudbury, Ont. college instructors area dealing with AI in the classroom. CBC https://www.cbc.ca/news/canada/sudbury/ai-chat-gpt-college-classrooms-1.6818524 (2023).

24. Alkaissi, H. & McFarlane, S. I. Artificial hallucinations in ChatGPT: implications in scientific writing. *Cureus* 15, (2023).

25. De Angelis, L. *et al*. ChatGPT and the rise of large language models: the new AI-driven infodemic threat in public health. *Frontiers in Public Health* 11, 1166120 (2023).

26. Dumpala, S. Check-It-Chatbot. (2021).

27. Chen, C., Fu, J. & Lyu, L. A pathway towards responsible ai generated content. *arXiv preprint arXiv:2303.01325* (2023).

28. Snoswell, A. & Burgess, J. The Galactica AI model was trained on scientific knowledge–but it spat out alarmingly plausible nonsense. *The Conversation* (2022).

29. Hosseini, M., Rasmussen, L. M. & Resnik, D. B. Using AI to write scholarly publications. *Accountability in research* 1–9 (2023).

30. Jo, A. The promise and peril of generative AI. *Nature* 614, 214–216 (2023).

31. Heaven, W. D. Why Meta's latest large language model survived only three days online. *MIT Technology Review* https://www.technologyreview.com/2022/11/18/1063487/meta-large-language-model-ai-only-survived-three-days-gpt-3-science/ (2022).

32. Wodecki, B. Meta's Galactica AI Criticized as 'Dangerous' for Science by Renowned Experts. https://aibusiness.com/nlp/meta-s-galactica-ai-criticized-as-dangerous-for-science (2022).

33. M Alshater, M. Exploring the role of artificial intelligence in enhancing academic performance: A case study of ChatGPT. *Available at SSRN* (2022).

34. Liang, W. *et al.* Advances, challenges and opportunities in creating data for trustworthy AI. *Nature Machine Intelligence* 4, 669–677 (2022).

35. Bhutoria, A. Personalized education and artificial intelligence in the United States, China, and India: A systematic review using a human-in-the-loop model. *Computers and Education: Artificial Intelligence* 3, 100068 (2022).