

# **Public Perceptions of Popular AI and Society**

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

**Zachary Yahn**

Spring 2023

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

Advisor

Bryn E. Seabrook, Department of Engineering and Society

## **Introduction**

What does humanity do when artificial intelligence (AI) begins to ask: “shall we play a game?” Games are an important part of the human experience, and the people that excel at them are often lauded for their brilliance and skill. It is no surprise, then, that world-champion-level AI would be a noteworthy phenomenon. As early as 1997, Deep Blue rocked the world by defeating then world-champion (and candidate for best player of all time) Garry Kasparov in Chess (Campbell et al., 2002, pp. 57-83). Televised and highly anticipated, the event garnered millions of views. Game-playing systems continued to succeed, reaching another key milestone when AlphaGo defeated top player Lee Sedol at Go in 2016 (Silver et al., 2016, pp. 484-489). Hundreds of millions of viewers tuned in from around the globe to spectate.

Game-playing AI occupies a unique spotlight in that it showcases how the best algorithms directly square up against the best humans. But, as an example of popular AI that has claimed the spotlight, how does it affect the public view of AI in society? Do we overestimate its capabilities, apply it overzealously, gloss over its disproportionate effects, or any other number of potential problems? In other words, how does the publicity of such AI systems affect peoples’ perception of personal and societal applications of AI? This research strives to answer this question using the science and technology studies (STS) frameworks technological fix and technological momentum. The answer may inform how AI systems are publicized in the future and how well people understand the consequences of an increasingly popular technology.

## **Research Question and Methods**

The defining research question for the following work is: How does the publicity of popular AI systems like AlphaGo and DeepBlue affect peoples’ perception of progress towards personal and societal applications of AI? Key aspects of this problem, which inform the

keywords used to search for resources, include artificial intelligence, ethics, public opinion, game-playing, machine learning, and big data. These are the keywords that commonly relate to AI and its interactions with society. The sources used to answer this question fall into three categories: case studies, surveys, and policy analysis. Case studies include examples of popular AI in the past, surveys are limited to those already conducted by other researchers, and policy analysis will include current AI policy and its focus on popular AI. In order to support the components of this research question, crucial sources analyze the ethics of modern AI, such as the work of R. Alberto. These sources support or refute the research thesis by understanding this problem with STS frameworks. Others, like Bunz et al. (2022), discuss how AI is portrayed in the media and how that portrayal affects public opinion (pp. 9-32). These sources instead provide the necessary background information and understanding of the current state of the problem. The remainder of this paper is dedicated to exploring and analyzing these sources in detail, limitations of this work, and finally a conclusion.

### **Background - Contemporary AI in Society**

Artificial intelligence's influence on daily life often flies under the radar, even as institutions continue to deploy it. The same algorithms and techniques that debuted in game-playing AI prove useful in other fields. In education, for example, AI might generate lecture content or guide students looking for informational videos on the internet. One review study shows how AI permeates education through instructional chatbots, personalized curricula, and even embedded systems like robots in the classroom. In theory, these systems are intended to increase engagement with content and tailor learning to the needs of students (Chen et al., 2020, pp. 75264 - 75278). Even popular education platforms like Khan Academy make use of deep learning as well. One study found that Khan Academy's systems displayed several issues,

including “flawed data, partially incomprehensible computational methods, narrow forms of educational knowledge baked into the online environments, and a reductionist discourse of data science with evident economic ramifications” (Perrota et al., 2019, 251-269).

The same systems that continue to be popular in education are also prevalent in another crucial field: healthcare. AI is used for detecting tumors, predicting stroke, developing drugs, diagnosing illnesses, and countless other applications. Morley et al. (2020) note that AI provides “opportunities to cut costs, improve care, and improve the efficiency of health and care systems.” They note that AI-clinician cooperation stands to leverage the analytical prowess of AI by combining it with the widespread knowledge and flexibility of human physicians (pp. 118-126). Babic et al. (2021) take a different approach, arguing that AI tends to disproportionately affect marginalized identities. They also point to the lack of explainability in AI systems, noting the opaque “data in, results out” approach (pp. 284-286). While AI certainly does have great potential to revolutionize patient care, the optimism and enthusiasm surrounding its development might not bode well for equitable and effective implementation.

Education and healthcare highlight how AI directly impacts daily lives, but unforeseen consequences may also arise in how it is developed and powered. In *Atlas of AI*, Kate Crawford notes that the AI we so effortlessly interact with requires significant natural resources in the form of electricity, data centers, and computer hardware (p. 75). She notes that mining rare earth metals and lithium decimates portions of the planet and poses a significant health risk to nearby communities (p. 76). Crawford further cites that the complex supply chains necessary for assembling AI-enabling technology results in emissions and waste pollution. In the same way that the flaws in the applications of AI often go unnoticed by the general population, the drawbacks of its development rarely reach the average consumer. Given the subtle ways by

which AI is produced and disseminated, the research discussed herein explores whether popular, innocent-seeming systems influence public perception of the aforementioned drawbacks and also attempts to understand the specifics of the public view.

### **Technological Fix and Momentum**

This work makes use of two frameworks to approach the question of game-playing AI and public opinion: technological fix and technological momentum. Technological fix, coined by Alvin Weinberg in 1966, describes the practice of approaching problems with a technological solution (Weinberg, p. 7). Technological momentum, on the other hand, was proposed as a metaphor for how the interaction of technology and society changed over time when describing hydrogenation in Germany (Hughes, 1969, pp. 106-132). According to the theory, technology and society influence each other as technology becomes more or less accepted as time passes, which in turn changes the dynamic between the two. It is important to note that both frameworks were proposed in the 1960s. While still relevant, multiple criticisms of the theories have emerged and are worth exploring.

Perhaps the strongest critique of technological fix is its narrow-minded focus on technology as the solution to all problems without considering societal factors. In a paper evaluating the intersection of technological fix and climate change, Peter Brewer notes that the solution cannot be purely technological. He describes how many of the causes of climate change are societal (rampant consumerism, misinformation, etc.) that cannot be directly addressed by technology (Brewer, 2007, pp. 9915-9916). A similar view is expressed by Sean Johnston, who in a criticism of technological fix writes that engineers “should recognize that complex modern societies incorporate multiple values and forms of expertise. Modern problems cannot be reduced to mere engineering solutions over the long term; human goals are diverse and

constantly changing” (Johnston, 2007, pp. 47-54). In recognition of this conclusion, this project turns traditional technological fix on its head by analyzing how technological solutions often fail to account for non-technological factors.

As a framework that considers the complex and interwoven roles of technology and society, technological momentum endures less criticism than technological fix. One paper uses it to assess the differences between American and Danish wind energy systems. The authors describe how through society-centric practices like transparency, information sharing, and publicized experimentation created public interest in the technology that led to superlinear growth (momentum) in its implementation (Sovacool et al., 2010, p. 18). This study illustrates how technological momentum can be a force for good by reducing friction and increasing public interest for helpful technologies like renewable energy. In the following sections technological momentum will be used to see whether a similar approach is possible for AI, and how society’s perception of AI changes as it becomes more integrated into daily life.

## **Results and Discussion**

### *Overview*

Based on policy review, case studies, and surveys, results show that widely-publicized artificial intelligence systems affect public perception by causing people to overestimate their knowledge and understanding of how AI affects their daily lives, influencing policy that addresses abstract ideas in AI without acknowledging specific situations or problems, and encouraging the implementation of the technology in fields like education and medicine where its effectiveness may be misunderstood. Popular terms like “deep learning” and systems like AlphaGo inspire oft-misplaced enthusiasm in AI as a whole, lead people to overestimate the capabilities of other systems, and can cause potentially harmful anthropomorphization. The

following sections explore in depth the effect publicized AI has on policy, case studies in medicine and education, and finally survey data reflecting the awareness, understanding, and opinions of the general public.

### Public Policy Analysis

Public policy is a crucial aspect of understanding the public reaction to popularized artificial intelligence. The statements, principles, and policies outlined by government bodies dictate how officials perceive the dangers and benefits of AI and how they plan to address them. The White House has outlined five principles and dangers for handling the future of AI: safe and effective systems, algorithmic discrimination, data privacy, notice and explanation, and human alternatives (The White House, 2022). They also acknowledge potentially positive impacts of the technology, writing that, “Fueled by the power of American innovation, these tools hold the potential to redefine every part of our society and make life better for everyone” (The White House, 2022). While statements such as these are important for informing a broad approach to AI, they do little to establish concrete strategies or regulations

Another approach taken by the U.S. government is the formation of the National Artificial Intelligence Research Task Force and the National AI Initiative, whose mission is “to ensure continued U.S. leadership in AI research and development, lead the world in the development and use of trustworthy AI in the public and private sectors, and prepare the present and future U.S. workforce for the integration of AI systems across all sectors of the economy and society” (National AI Initiative, 2022). The National AI Research Task Force reported on the development of public AI projects, showing that the percentage of commits (a contribution to a code base) to projects run by the government from 2011 to 2022 has steeply declined (2023).

These data show that AI is becoming increasingly privatized. Simultaneously, the U.S. government has increased the quantity of policies involving AI. In 2018, five policies mentioned artificial intelligence, two of which were specific to military applications. This increased to eight with two military in 2019, and ten with two military in 2020 (Center for Data Innovation, 2020). Globally, 30 nations have published policies relating to various aspects of AI governance (Stanford University, 2021).

While it is clear that governments are aware of AI technology, whether they publish effective strategies or simply cater to public opinion is unclear. According to Dwivedi et al. (2019), “The probabilistic analysis of the economic impact of AI and automation has been assessed by the World Economic Forum (WEF), where they predict that 20% of existing UK jobs could be impacted by AI technologies” (p. 2). It is clear that AI will affect the everyday lives of common people, yet many policies (especially in the United States) tend to focus on abstract concepts rather than concrete implementations. Krafft et al. (2020) identify two key issues with governments’ approach to AI policy. First, while AI researchers tend to “emphasize empirical functionality, policy-makers instead use definitions that compare systems to human thinking and behavior” (p. 1). This anthropomorphization leads to incorrect conclusions about the capabilities of AI systems and how quickly they develop and change. Krafft et al. (2020) theorize that this may be a result of the popularization of specific systems like AlphaGo, which may be the only example of AI with which policy-makers are familiar. Second, a major barrier in policy is a lack of consensus on definitions. (p. 7). For example, international discussions have been largely marred by disagreement on the definition of autonomous weapons.

Through the lens of technological momentum, it appears that the momentum of AI makes it difficult to anticipate and react to the slow-moving mechanisms of policy. The dynamic

between AI and society changes rapidly as new models are released and popularized, in many cases forcing the government to work as fast as the industry to keep up. While it is clear that AI has a major effect on society that the government intends to address, it is less clear how society in turn affects AI. The technology appears to be shaped more by financial pressures than policy, leading to development in lucrative but under-regulated sectors like social media. In this way, technological momentum highlights the mismatch between the influences of AI and society on each other: while AI grows ever dominant in people's daily lives, a major mechanism for the same people to direct the course of this development (i.e. public policy) is weakly available at best. While policy is not the only way to shape AI, the people who are most affected by the technology are often powerless to use other methods such as steering a company, boycotting the technology, or guiding research as an AI scientist. The situation is further complicated by the technical complexity and knowledge required to understand AI, as well as the rapid pace of development. The tug-of-war of technological momentum in part relies on time for societal pressures to actually react to changes in technology, yet AI often advances so quickly that responses lag behind. In other words, AI innovation is largely left unchecked by society. While the increased rate of AI-related policy shows that the government is aware of the emerging relevance of AI, the lack of specific action suggests that they are unsure how to respond to individual cases or where to allocate resources. Instead, policy focuses on broad strokes and abstract guidelines: meeting public AI with public messages without targeting the underlying issues.

### Case Studies in Education and Medicine

In order to understand the effects of AI on daily lives, this work goes in-depth on two case studies: AI in education and AI in healthcare. Both are essential parts of everyday

livelihood, and both are popular areas for artificial intelligence innovation. In these cases, experts note that the popularization of terms like “deep learning” and “expert systems” make algorithms sound intelligent and competent (Bunz, 2021, p. 1). Chen et al. (2020) note how “the mention of artificial intelligence brings to mind a supercomputer, a computer with immense processing capabilities” (p. 3). Game-playing systems like AlphaGo and Deep Blue contribute to the same perception by appearing superhuman at very specific tasks and by turning the technologies behind them (deep learning included, in the case of AlphaGo) into buzzwords.

This perception of AI technologies is especially relevant in medicine, where an analysis of 365 articles about AI showed that there is a tendency to portray medical AI as surpassing human expertise, regardless of whether this perception was rigorously demonstrated (Bunz, p. 1, 2021). Popular media outlets also contribute to the enthusiasm, such as an article from the Wall Street Journal titled “The AI doctor will see you now.” Bunz notes that such titles anthropomorphize AI, potentially placing it above criticism. (p. 4, 2021). Anthropomorphization also misdirects attention away from problems in AI like dataset bias and resource consumption, instead casting the limelight on specific errors that systems make or quirks in their outputs. Though similar to other vague comments discussed in the policy analysis section above, a White House statement emphasizes the risks of AI in healthcare: “In America and around the world, systems supposed to help with patient care have proven unsafe, ineffective, or biased” (2022). While some researchers, including Chen et al., (2020) define artificial intelligence as “the ability of computers and machines to mimic human cognition and actions,” the tendency to overestimate the abilities of AI because of its popularization indicates the importance of considering AI as distinctly non-human entities with unique and subtle flaws (p. 3).

These trends continue in the educational domain, where AI is rapidly integrating into classrooms, tailoring individual study plans, and even substituting for teachers. As Guilherme (2019) writes, “A defining aspect of our modern age is our tenacious belief in technology in all walks of life, not least in education” (p. 1). Education has often been a proving ground for new (and sometimes poorly understood) technologies, especially when they might relieve the strain of limited public school budgets. As with education, people tend to overestimate the capabilities of AI in educational scenarios. School administrators and other education professions who determine how to implement AI in the classroom hope that algorithms that perform well for certain tasks may also generalize well, even if the data does not support that (Perrota, 2019, p. 4).

In many cases, using new technology in schools is labeled “progress” because it conforms to the popular idea that AI is capable and applicable in all scenarios. Though there are many cases where AI does perform well (object recognition, game-playing, machine translation), that does not indicate its effectiveness in entirely different and often more complex domains like education. As one paper elucidates, “these changes in schools and in education systems are associated with a notion of progress and this might lead us to ask questions, such as: ‘Whose idea of progress? Progress for what?’” (Guilherme, 2019, p. 4).

For both medical and educational situations, the overestimation of AI’s accuracy and overenthusiasm for its widespread implementation appears to encourage technological fix. Rather than addressing the nuances of a particular situation, AI is deployed as a catch-all tool for replacing dynamic and complex roles like doctors and teachers. The evidence from AI-relevant discussions in the medical context shows how labels like “expert systems” and anthropomorphizing titles make AI seem like an attractive fix. Likewise, general public interest in AI has led to its implementation in education, where its effectiveness is largely unproven.

Attempting to fix complicated problems with broad strokes sometimes makes things even more difficult. One author writes, “Focusing on the importance of technology for education has had a direct impact on teachers and teacher education, because they are expected to combine students’ and their own development and intellectual excellence within a globally technological and economically demanding society” (Guilherme, 2019, p. 2). While AI can be effective in well-understood context with thoroughly-demonstrated results, it is also a prevalent technological fix.

While technological fix does elucidate motivations behind rapidly applying AI in critical contexts such as medicine and education, it is also important to use technological momentum to consider the timing and contrary pressures of the technology and its rapid development. Such eagerness to apply AI methods to medicine may result from how quickly AI changes relative to societal response times. While pushback from regulators or public outcry might normally slow a new medical technology, AI benefits from all of the rapid innovation with none of the red tape. Organizations like the Food and Drug Administration (FDA) take years to approve of new medicines, yet AI can be implemented as soon as a new model is developed. In some cases, society even administers an encouraging pressure on AI through overconfidence in its performance. As technological momentum predicts, society does indeed influence technology; in the case of AI in medicine and in education, however, it may do more to reduce AI friction than increase it.

#### *Trust in and Knowledge of AI Through Surveys*

The final section of this methodology concerns how the general public perceives artificial intelligence, including awareness, understanding, and opinions about its use. One Pew Research poll showed that overall Americans are more concerned than excited about AI, with 37% more

concerned than excited, 45% equally concerned and excited, and 18% more excited than concerned (Rainie et al., 2022). The same survey also asked respondents about whether three issues (facial recognition by police, automated fact-checking on social media, and automated driving) were a good idea, bad idea, or not sure. For all three issues the responses were roughly split evenly, showing that Americans are broadly unsure about how to interpret these popular implementations. (Rainie et al., 2022). It is important to note that distrust in artificial intelligence varies for different marginalized groups, especially those who are often directly and disproportionately impacted by AI technologies such as criminal justice algorithms for facial recognition for surveillance.

Other studies show the extent to which American and English citizens think they know about AI and how much they actually understand. In a Bristows study of over 2,000 people, 75% of respondents claimed to have limited knowledge of AI or better, with 8% claiming to have significant knowledge or be an expert (Bristows, 2018). The same study shows that, for the same group of respondents, 62% thought they had never used an AI-powered technology, and 23% were not sure. Only 15% of respondents were confident that they had interacted with AI (Bristows, 2018). Given that in both countries almost everyone has interacted with artificial intelligence in some way, whether through smartphones, social media, music, email, or something else, the disparity between claimed knowledge and actual understanding of AI in their lives shows that respondents think they are more AI-literate than they actually are. Along these lines, only 39% of respondents to the same study thought AI would use their personal data, with 39% being unsure, showing that the privacy implications of AI are particularly misunderstood (Bristow, 2018).

A similar study by Pew Research investigated whether or not Americans thought six technologies used AI. The six technologies were: wearable fitness trackers, chatbots, product recommendation systems, facial recognition for security cameras, music recommendations, and email spam detectors. Only 30% of Americans correctly identified that all six use AI, while 31% thought two or fewer made use of the technology (Kennedy et al., 2023). Meanwhile, another study showed that Americans were generally aware of major AI breakthroughs including weather prediction and protein folding, with 40-50% of respondents saying that they were at least somewhat familiar with the technology (Funk et al., 2023). The same study reports that “Among Americans with some awareness of AI uses, majorities see AI-driven surgical robots and AI-based skin cancer screening as major advances for medical care” (Funk et al., 2023). Though according to Tyson et al. “60% of Americans are uncomfortable with their provider relying on AI in their own health care” (Tyson et al., 2023).

These results indicate that Americans are unaware of the many ways AI is in their daily lives despite reading about popular breakthroughs in artificial intelligence. People have also formed opinions about specific implementations of the technology, but are not well-versed in the technical details of AI or the problem space. As this research establishes, AI is often used as a technological fix, yet many respondents are in favor of using AI to solve varied complicated problems. There appears to be disparities between how much Americans claim to know about AI and how much they actually know, and between how much Americans would like to see AI-based solutions and how much they understand the true details of AI or the problems it would solve. As discussed in the previous section, technological momentum also shows that this ignorance changes how society affects AI technology. Even as it continues to have a substantial

presence in their lives, people are ill-equipped to properly regulate AI development, identify and mitigate its effect on their lives, and generally steer it in a beneficial direction.

### Limitations

Two key limitations of this research must be acknowledged: lack of scope in case studies, and America-centric perspective. While medicine and education are essential fields with a significant AI presence, other domains including policing, social media, and criminal justice would provide an additional dimension to the analysis. Each of these requires a thorough analysis which went beyond the scope of this work. Criminal justice and policing in particular disproportionately affect members of marginalized groups who are already vulnerable to increased violence from institutionalized prejudice. Future work should take into account the perspectives of members of these groups in particular rather than simply aggregating the American population. Furthermore, while the surveys used in this work accurately portrayed the American population in general, they did not address AI in a global context. Although the United States is home to many of the companies responsible for the rise of AI, it is unmistakably a global technology that should also be viewed from the perspective of other nations and cultures. The same can be said for policy analysis, which solely focused on the American government like the White House and national agencies.

### Future Work

In addition to including other domains, future work might also explore public AI through other STS frameworks. Social construction of technology might be particularly useful for assessing how the public affects AI instead of vice versa. This study also leaves out very recent breakthroughs in natural language processing, including OpenAI's ChatGPT and other large language models. Reaching 100 million users in two months, ChatGPT is undoubtedly the most

well-known AI to date. ChatGPT is so new that it is too soon to comment on its consequences, though in time it will be essential to understand its popularization.

## **Conclusion**

AI is increasingly popularized for breakthroughs in specific domains, with noteworthy success in game-playing environments. This success broadly leads to public perception of AI as widely competent and applicable to much more complicated tasks like medicine and education. Policy related to AI reflects interest in addressing and controlling it, but typically through abstract statements rather than specific solutions. While the public is aware of artificial intelligence and major breakthroughs, many people do not fully comprehend how thoroughly AI is integrated into their daily lives. It is essential for people to grasp how AI affects their decisions, government, health, education, and to understand the momentum that it has in society. It is up to AI experts to effectively communicate and disseminate their knowledge to make an understanding of AI more accessible for the everyday person. AI is also broadly used as a technological fix for many of the world's most complicated problems, and it is imperative that people understand the nuances of these problems and when to use (and not use) AI to address them. Artificial intelligence is becoming increasingly important and powerful, and people must be aware of how it will change their lives.

## Works Cited

- Babic, B., Gerke, S., Evgeniou, T., Cohen, I. (2021). Beware explanations from AI in health care. *Science* 273(6552), 284-286. doi: 10.1126/science.abg183.
- Brewer, P. G. (2007). Evaluating a technological fix for climate. *Proceedings of the National Academy of Sciences* 104(24), 9915-9916. <https://doi.org/10.1073/pnas.0703892104>.
- Bristows (2018). *Artificial intelligence: public perception, attitude, and trust*  
<https://www.bristows.com/app/uploads/2019/06/Artificial-Intelligence-Public-Perception-Attitude-and-Trust.pdf>
- Bunz, M., Braghieri, M. (2022). The AI doctor will see you now: assessing the framing of AI in news coverage. *AI & Society* 37(1), 9-32. doi:10.1007/s00146-021-01145-9
- Campbell, M., Hoane, J., Hsu, F. (2002). Deep blue. *Artificial Intelligence* 134(1), 57-83.  
[https://doi.org/10.1016/S0004-3702\(01\)00129-1](https://doi.org/10.1016/S0004-3702(01)00129-1).
- Center for Data Innovation. (2020). *AI legislation in 2020*. <https://datainnovation.org/ai-policy-leadership/ai-legislation-tracker/ai-legislation-in-2020/>
- Chen, L., Chen, P., Lin, Z. (2020). Artificial intelligence in education: A review. *IEEE Access* 8, 75264-75278. doi: 10.1109/ACCESS.2020.2988510
- Crawford, K. (2021). *The atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., et al. (2019). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, 57. <https://doi.org/10.1016/j.ijinfomgt.2019.08.002>.

- Funk, C., Tyson, A., and Kennedy, B. (2023) How Americans view emerging uses of artificial intelligence, including programs to generate text or art. *Pew Research Center*.  
<https://www.pewresearch.org/fact-tank/2023/02/22/how-americans-view-emerging-uses-of-artificial-intelligence-including-programs-to-generate-text-or-art/>
- Guilherme, A., (2019). AI and education: The importance of teacher and student relations. *AI & Society* 34, 47-54. <https://doi.org/10.1007/s00146-017-0693-8>
- Hughes, T. (1969). Technological momentum in history: Hydrogenation in Germany 1898-1933. *Past & Present* 44, 106-132, <http://www.jstor.org/stable/649734>.
- Johnston, S. F. (2007). The technological fix as social cure-all: Origins and implications. *IEEE Technology and Society Magazine* 37(1), 47-54. doi: 10.1109/MTS.2018.2795118.
- Kennedy, B., Tyson, A., and Saks, E. (2023). Public awareness of artificial intelligence in everyday activities. *Pew Research Center*. <https://www.pewresearch.org/science/2023/02/15/public-awareness-of-artificial-intelligence-in-everyday-activities/>
- Krafft, P. M., Young, M., Katell, M., Huang, K., and Bugingo, G. (2020). Defining AI in policy versus practice. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 72–78. <https://doi.org/10.1145/3375627.3375835>
- Morley, J., Machado, C., Burr, C., Cows, J., Joshi, I., Taddeo, M., Floridi, L. (2020). The ethics of AI in health care: A mapping review. *Social Science & Medicine* 260, 118-126.  
<https://doi.org/10.1016/j.socscimed.2020.113172>.
- National Artificial Intelligence Initiative (2022). *National artificial intelligence initiative: Overseeing and implementing the United States national AI strategy*. <https://www.ai.gov/>
- National Artificial Intelligence Research Resource Task Force. (2023). *Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem*. <https://www.ai.gov>

/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf

Perrota, C., Selwyn, N. (2019). Deep learning goes to school: toward a relational understanding of AI in education. *Learning, Media and Technology* 45(3), 251-269. <https://doi.org/10.1080/17439884.2020.1686017>.

Rainie, L., Funk, C., Anderson, M., and Tyson, A. (2022). How Americans think about artificial intelligence. *Pew Research Center*. <https://www.pewresearch.org/internet/2022/03/17/how-americans-think-about-artificial-intelligence/>

Silver, D., Huang, A., Maddison, C. et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature* 529, 484–489. <https://doi.org/10.1038/nature16961>.

Sovacool, B. K., & Sawin, J. L. (2010). Creating technological momentum: lessons from American and Danish wind energy research. *Whitehead Journal of Diplomacy and International Relations*, 11(2), 43-58. <https://heinonline.org/HOL/P?h=hein.journals/whith11&i=204>

Stanford University (2021). *Artificial intelligence index report 2021*. [https://aiindex.stanford.edu/wp-content/uploads/2021/03/2021-AI-Index-Report-\\_Chapter-7.pdf](https://aiindex.stanford.edu/wp-content/uploads/2021/03/2021-AI-Index-Report-_Chapter-7.pdf)

The White House. (2022). *Blueprint for an AI bill of rights*. <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>

Tyson, A., Pasquini, G., Spencer, A., and Funk, C. (2023). 60% of Americans would be uncomfortable with provider relying on AI in their own health care. *Pew Research Center*. <https://www.pewresearch.org/science/2023/02/22/60-of-americans-would-be-uncomfortable-with-provider-relying-on-ai-in-their-own-health-care/>

Weinberg, A. M. (1967). Can technology replace social engineering? *American Behavioral Scientist*, 10(9), 7–7. <https://doi.org/10.1177/0002764201000903>.