Undergraduate Thesis Prospectus

Improving Human-AI Interaction: An Assistive LED Chessboard

(technical research project in Computer Engineering)

The Struggle over AI in Healthcare

(sociotechnical research project)

by

Srikar Chittari

December 4, 2022

technical project collaborators:

Ramie Katan
Iain Ramsey
James Weeden

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.

Srikar Chittari

Technical advisor: Harry Powell, Department of Electrical and Computer Engineering

STS advisor: Peter Norton, Department of Engineering and Society

General Research Problem

How can artificial intelligence systems better serve human needs?

Artificial intelligence (AI) allows computers to think and act rationally to mimic the decision-making capabilities of the human mind. AI can automate tasks that humans do easily and solve complex problems that surpass human capabilities. While AI is most commonly known for its applications in intellectually challenging games like chess, it is also implemented in systems where stakes are higher such as banking, healthcare, agriculture, and the military (Keskinbora, 2019). AI can contribute significantly to society and help humans solve new problems while also running the risk of endangering humanity (West, 2018).

Improving Human-AI Interaction: An Assistive LED Chessboard

How can AI help a person learn chess?

For my computer engineering capstone, I am working in a team with Ramie Katan, Iain Ramsey, and James Weeden, and our advisor is Prof. Harry Powell from the Electrical and Computer Engineering department.

AI in chess is often portrayed as a machine playing chess against a human. While the purpose of AI is to emulate a human chess player, "AI-augmentation isn't designed to replace humans" (Huang, 2022). Rather, the goal of AI in chess is "to take advantage of the best capabilities of humans and technology" (Huang, 2022). AI not only has tremendous potential in defeating humans in chess but also in helping humans improve their skills. Viewing chess move recommendations generated by AI is analogous to receiving advice from highly skilled chess players. Both cases can build the intuition of players and accelerate the growth of chess skills. However, recommendations must be tailored to the ability of the chess player to maximize

learning. AI can also be employed to recognize the skill level of the human and produce personalized recommendations (McIlroy-Young et al., 2020). Historically, AI has aimed to defeat humans in chess, but new AI applications are being developed to assist chess players.

The goal of the computer engineering capstone project is to produce an embedded system, which integrates hardware and software, to accomplish a real-world task. My team and I are designing and programming an assistive LED chessboard that provides a user with recommendations while playing against an AI engine. The recommended chess move will be illustrated by LEDs on the chessboard. Assistive chess boards are widely available from a multitude of different companies. The most prominent chess board, ChessUp by Bryght Labs, interfaces with an AI engine and illuminates chess tiles using capacitive touch. It indicates a good move or a blunder based on the configured strength of the AI. Existing projects differ in how the game state is detected. Most products have used capacitive sensors or relied on the strengths of magnets to differentiate between chess pieces, but these methods lack reliability. For example, one product, which used magnetic strengths for piece detection, achieved 80% reliability (Muji et al., 2016). My capstone project differs from the previous projects since it uses a configuration of multiple magnets on each chess piece to improve reliability.

The capstone project will involve designing printed circuit boards to detect the chess pieces and illuminate the chess board, programming a microcontroller to process the game state, and interfacing with the AI chess engine. A prototype will be built and tested for accurate piece detection and AI recommendations. By the end of the semester, a successful capstone project would be able to produce accurate and dynamic recommendations in real-time to one or more chess players. This project will demonstrate how AI can be integrated into a popular game to assist the user and help hone their skills.

The Struggle over AI in Healthcare

In the U.S., how have advocates of healthcare equity resisted discriminatory biases encoded in medical AI?

AI is portrayed as a self-correcting system, but it is only as accurate and precise as its training data and model. Unrepresentative training data and a biased model may produce inaccurate AI applications that propagate in industry. Engineers who develop algorithms and datasets may neglect biases in large-scale AI applications (Ahmed, Athyaab, & Muqtadeer, 2021). In industries where AI applications decide the lives of their clients, algorithmic bias can be detrimental. Specifically, discriminatory biases in training sets for medical AI have resulted in healthcare inequity in the United States (Thomasian, 2021).

Current research indicates advocates of healthcare equity have defined their principles and have formed advocacy groups to encourage the government to regulate AI practices. For example, AI4PEOPLE has created a framework for holding corporations accountable for AI for biases concerning human rights and ethics (Rossi, 2018). The principles of AI4PEOPLE are mapped to "concrete recommendations for potential action that EU policymakers can follow" such as developing a new EU agency to evaluate and supervise AI (Rossi, 2018). Federal regulations in the U.S., however, remain inadequate, urging more advocacy groups and researchers to compile instances of AI bias and influence legislation (Thomasian, 2021). Therefore, more research can help analyze the methods employed by advocacies exposing medical AI bias and healthcare providers that want to maintain public trust in their AI services.

Participants addressing biases in medical AI include advocacies for marginalized patients such as the American Civil Liberties Union, advocacies for transparency in AI development such as the Algorithmic Justice League, AI reviewers such as researchers at the Harvard School of

Public Health, federal AI regulators such as the U.S. Food and Drug Administration, and healthcare providers such as Optum.

Several instances of AI discrimination in the health sector have surfaced recently. For example, Arkansas residents with cerebral palsy experienced extreme budget cuts to in-home care, resulting in hospitalizations and disruptions to their lives. With the aid of a local American Civil Liberties Union (ACLU) branch and a lawsuit, errors in the algorithm's characterization of patients with disabilities were found to be responsible (Grant, 2022). As demonstrated by ACLU's intervention to aid the residents of Arkansas, the advocacy group engages in strategic lawsuits to expose unethical AI and protect underrepresented patients in the health sector. ACLU also aims to initiate "policy changes and collaboration among key stakeholders, including state and federal regulators, medical, public health, and clinical advocacy groups and organizations… to address these gaps and inefficiencies" in medical AI (Grant, 2022). In addition to lawsuits against AI developers, ACLU acknowledges the ubiquity of AI bias and looks to engage in coalition building across the healthcare industry.

While some advocacy groups improve the ethics of patient treatment, other groups, such as the Algorithmic Justice League (AJL), advocate for increased transparency in the AI development process. Joy Boulamwini, the founder of the AJL, has "launched the Safe Face Pledge to prevent the lethal use and mitigate abuse of facial analysis and recognition technology" (Boulamwini, 2019). Multiple companies have agreed to sign the pledge, which facilitates transparency during AI development to prevent the weaponization of AI and harmful discrimination. Although the pledge minimizes bias for facial recognition AI in law enforcement, the AJL remains active in the healthcare industry by encouraging similar transparency among healthcare AI developers. Like the ACLU, the AJL calls on AI companies to improve their

development practices, but the AJL uses rules enforcement as their primary method to achieve transparency.

Medical AI has drawn attention from algorithmic reviewers in academia who study healthcare inequity. Researchers from the Harvard School of Public Health (HSPH) also point to a lack of transparency in AI development as the root cause of bias but propose a course of action differing from the ACLU and the AJL. Trishan Panch, the president-elect of the HSPM Alumni Association, asserts that "there will probably always be some amount of bias," so incentives can be implemented to expose biases (Igoe, 2021). For example, "if researchers or other professionals can prove that data analysis is biased, they can utilize legislation via class action lawsuits," and "this will incentivize private companies to change or to preemptively look at bias before this occurs" (Igoe, 2021). The approach HSPH proposes relies on installing a reward system for third parties to discover AI biases, contrasting the methods outlined by the ACLU and the AJL which pursue legal action and voluntary corporate pledges.

With growing calls to reform AI regulation, the U.S. Food and Drug Administration (FDA) has expanded its oversight of medical AI to verify training sets and combat algorithmic bias. The agency employed a new framework that would require corporations to send "periodic updates to the FDA on what changes were implemented as part of the approved pre-specifications and the algorithm change protocol" (CDRH, 2021). The FDA publishes action plans to regulate AI to demonstrate its receptiveness to advocacy. However, the FDA approves AI applications under the 'Software as Medical Devices' classification, and as a result, "a large proportion of artificial intelligence algorithms are exempted by this definition and are already in widespread use throughout the health sector" (Thomasian, 2021). Advocacy groups have been

critical of this FDA policy for only regulating a portion of AI applications, and the FDA continues to iterate its policy to verify more AI applications.

Employers of medical AI such as Optum, contest research on algorithmic bias, deflecting the inequity to other causes. In a statement, Optum claimed the research is "misleading" and that the model is just one of many data elements used to determine treatments (Ledford, 2019). Optum places importance on its public relations to respond to its critics and maintain public trust in its AI services. Optum also conducts annual surveys to promote confidence in its AI services. Their survey found that "99% agree that AI can be trusted for use in health care" (Optum, 2022). Optum also found that "94% [of health executives] agree that they have a duty to ensure responsible use of AI," which deflects responsibility for AI bias from the healthcare providers to other individuals (Optum, 2022). Optum has publicly downplayed research on algorithmic biases and has bolstered its clients' confidence in AI through the use of surveys.

References

Ahmed, S.; Athyaab, S.A.; and Muqtadeer, S.A. (2021, Jan.). Attenuation of Human Bias in Artificial Intelligence: An Exploratory Approach. *2021 6th International Conference on Inventive Computation Technologies (ICICT)*, 557-563. Web of Science

Buolamwini, J. (2019, Feb. 7). Artificial Intelligence Has a Problem With Gender and Racial Bias. Here's How to Solve It. *Time*. time.com/5520558/artificial-intelligence-racial-gender-bias

CDRH (2021, Sep. 22). Center for Devices and Radiological Health. Artificial Intelligence and machine learning in software. U.S. Food and Drug Administration. www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-software-medical-device

Grant, C. (2022, Oct. 3). Algorithms Are Making Decisions About Health Care, Which May Only Worsen Medical Racism. American Civil Liberties Union. www.aclu.org/news/privacy-technology/algorithms-in-health-care-may-worsen-medical-racism

Huang, T. (2022, Jan. 10). Why Computer-Assisted Humans Are The Best Chess Players And What That Means For Technology Operations. *Forbes*. www.forbes.com/sites/forbestechcouncil/2022/01/07/why-computer-assisted-humans-are-the-best-chess-players-and-what-that-means-for-technology-operations/?sh=7ed334fe14e0

Igoe, K. (2021, March 12). Algorithmic Bias in Health Care Exacerbates Social Inequities — How to Prevent It. Harvard School of Public Health. www.hsph.harvard.edu/ecpe/how-to-prevent-algorithmic-bias-in-health-care

Keskinbora, K. H. (2019, March 14). Medical Ethics Considerations on Artificial Intelligence. *Journal of Clinical Neuroscience 64*, 277–282.

Ledford, H. (2019, Oct. 24). Millions of black people affected by racial bias in health-care algorithms. *Nature*. www.nature.com/articles/d41586-019-03228-6

McIlroy-Young, Anderson, Kleinberg, and Sen (2020, Nov. 30). The human side of AI for chess. Microsoft Research. www.microsoft.com/en-us/research/blog/the-human-side-of-ai-for-chess

Muji, Wahab, Ambar, and Loo (2016, Nov.). Design and implementation of electronic chess set. *2016 International Conference on Advances in Electrical, Electronic and Systems Engineering (ICAEES)*, 451-456.

Optum (2022). Responsible use of artificial intelligence. www.optum.com/business/insights/health-care-delivery/page.hub.responsible-artificial-intelligence.html

Rossi, F. (2018). Building Trust in Artificial Intelligence. *Journal of International Affairs 72(1)*, 127–134. JSTOR

Thomasian, N.M.; Eickhoff, C.; and Adashi, E.Y. (2021, Nov. 22). Advancing health equity with artificial intelligence. *Journal of Public Health Policy 42*, 602–611. Web of Science

West, D. M. (2018, Oct. 4). What is Artificial Intelligence? Brookings. www.brookings.edu/research/what-is-artificial-intelligence