REINFORCEMENT LEARNING-BASED IMPLEMENTATION OF NON-PLAYER CHARACTERS IN POKEMON

ALPHA ZERO REINFORCEMENT LEARNING AS A BOUNDARY OBJECT

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
William Bradford

November 7, 2024

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

ADVISORS

Kent Wayland, Department of Engineering and Society

Felix Lin, Department of Computer Science

Introduction

Modern artificial intelligence has been used to automate the playing of popular games since the mid-1900s, and it has shaped both the games and the communities around which it has been designed. Artificial intelligence has affected much more than the ways that games are played and has begun to shape the games themselves, providing novel methods of testing new game rulesets and evaluating their effectiveness. This paper aims to combine my technical research on novel methods for training artificial intelligence to play complex games with a boundary object-based approach to analyzing the development of artificial intelligence and the communities that have been shaped by modern AI frameworks such as AlphaZero.

Technical Project

My technical project aims to use modern machine learning techniques to develop a novel artificial intelligence model for playing Pokémon video games. While prior works focus on providing AI agents that can play against humans in competitive Pokémon battles (*Future Sight AI*, 2021; Kothari, 2024), my work aims to expand upon the current research by providing an agent capable of battling against players by taking on the role of a non-player character (NPC) in one of the main series Pokémon video games for the Nintendo Game Boy, Pokémon Red. This provides a more compelling player experience and allows the user to more closely identify with game characters by making them feel more lifelike while helping players refine their skills more over time relative to classic NPCs by stimulating more difficult decision-making processes.

Pokémon AI research is such an immature field that it cannot provide enough data to assess how AI affects the Pokémon community and therefore could not be used to extrapolate the effects of AI on more generalized game communities. Accordingly, it may be more fruitful to assess the mutual shaping between chess and AI, which would provide a greater understanding of how AI and computer science research affect both modern games and the communities surrounding them. This brings me to my STS topic.

STS Project

Chess is a cultural phenomenon that shares its roots in many cultures and continues to be adapted and changed throughout human history. The culture surrounding chess both invites optimization and encourages redesign of the game itself (Larson, 2018; Tomasev et al., 2022). By virtue of its history in European spheres, chess culture also encourages a comparison between one's ability to play the game and one's ability to think intelligently, making it a logical target for AI engineers (Ensmenger, 2012; Chowdhary et al., 2023). This project aims to answer the question, "How has the perception of the agency of decision-making shifted in chess culture as a result of the Alpha Zero machine learning framework, and how have subsequent derivatives of this framework been shaped as a result of this shift in perception?"

Historically, chess has been considered a "fundamentally human endeavor" (Ensmenger, 2022) and beginning in the late 1900s, scholars began to assess the cultural implications of machines playing chess autonomously. The first recorded instance of "automated" chess was when, in the mid 1800s, a device called a "Mechanical Turk" was used to fool crowds into

believing that a machine could beat human players at chess (Bloomfield & Vurdubakis, 1997). In reality, the machine simply concealed a human who mechanically controlled the chess board. Accordingly, scholars considered the agency of the decision-making process to be entirely human (Bloomfield & Vurdubakis, 1997). In modern times, this dynamic has shifted, and modern automations of chess were considered entirely machine throughout the 1990s, often thought of as 'cold calculation,' incapable of emulating the real and varied patterns of human decision-making (Bloomfield & Vurdubakis, 1997). Yet, as computer science has transformed artificial intelligence, novel AI frameworks driven by teams of human engineers have begun to simulate increasingly human decision-making, thus making decision-making seem more human in the eyes of the modern general public (Wilkenfeld, 2019). However, if chess is such a human endeavor, why has it become so entrenched in the literature surrounding AI?

Since Europe adopted and modified chess from its roots in chaturanga and later shatranj, originating in Indian and Muslim communities respectively, the Euro-centric culture surrounding the game has hinged around optimizing it to be both simple and creatively challenging by European standards (Larson, 2018). Although hotly debated in modern times, one's ability to play chess has historically been considered a strong metric for intelligence, making it an obvious target for artificial intelligence developers due to its easily interpretable metrics of success (Chowdhary et al., 2023; Ensmenger, 2012). If an AI agent wins a game of chess, it can be generally considered a successful agent. The greater the number of pieces remaining after its victory, the better it performed, making it easy to evaluate or 'score' the success of the agent overall, attaching a number to the 'intelligence' of the AI agent. While the relative ease of chess as a scoring metric may seem attractive, it serves to drastically reduce the dimensionality of intelligent expression.

In Lea's article from 2020, he warns against the use of chess as a metric for AI development, claiming that social constructivism can lead to significant dangers in the field of AI and using machine learning's fixation on chess as the poster child for these dangers. As Lea claims, modern chess is dominated by males, which is in part due to the use of intimidation tactics to increase the player's odds of winning, thus implying that, while helpful, intelligence is not a sufficient condition for winning in chess (Lea, 2020). Therefore, because chess victories cannot be solely attributed to intelligence, using chess to assess the progress of artificial intelligence is an overly limiting, mechanical interpretation of intelligence that may not provide a sufficient indication of true, overall intelligence.

As AI has become more tightly associated with chess culture, it has affected both the game itself and the way that the game is played in modern circles. Novel techniques such as the Alpha-Zero framework have emerged in which AI agents play games against themselves to learn what moves may or may not be effective (Tomasev et al., 2022). These non-human, chess-playing agents then influence top level chess players' move selection, reducing the variety of moves selected in favor of choosing more 'optimal' moves, thus making chess less interesting to watch and less creatively expressive to play (Wilkenfeld, 2019). Anecdotally, this makes professional chess gameplay significantly less interesting because players use the same few

openers and overarching strategies in every single game, which makes the viewing experience feel monotonous.

To solve this problem, scientists turned to these same machine learning agents to help assess which modifications of chess might make the game more interesting to both play and watch at the highest level of gameplay (Tomasev et al., 2022). To accomplish this, scientists change the ruleset of chess, then allow the machine learning agent to learn how to play this new version of the game and use qualitative and quantitative models to assess how many distinct and varied gameplay strategies may be effective in high level play, thus indicating how interesting it may be to watch (Tomasev et al., 2022). This use of AI has far-reaching implications, and it may provide insight into how artificial intelligence might be used to alter the principles of other fields going forward. Contemporary scholars are now considering what it means to make society's humanities less 'human' by allowing mechanized reasoning to replace human thought in the arts, such as chess play. Significant thought has been given to how different cultures might conceive of the sociotechnical imaginaries associated with artificially intelligent creative works, such as game play, and how they might affect the implementations of AI or the use of AI in the design of modern games.

AlphaZero as a Boundary Object

The purpose of my STS paper is to use actor-network theory to analyze the mutual shaping between the Alpha-Zero framework and chess culture to assess how chess theory and culture have unified with computer science to produce novel actants in the machine learning and chess networks that function as boundary objects between the two fields. According to Stewart and Watson (2020), a boundary object is a "heuristic," or estimation model, that allows scholars to analyze the development of technological artifacts as a product of the independent pursuit of sociotechnical imaginaries by two or more groups. As modern computing has become more prevalent in society, both casual computer scientists and computer science researchers have become involved in the chess network, and the involvement of these actants has produced new actants including machine learning agents, which fulfill a unique and theoretical role in the shaping of chess culture. Machine learning specialists approach chess AI as a method of improving machine learning techniques and emphasize the importance of the technology whereas chess enthusiasts approach chess AI as a means of learning about and further developing the game itself (Tomasev et al., 2022). Accordingly, AI has shaped chess gameplay to increase move selection precision at the cost of creativity as a direct product of these two mostly independent groups' pursuits of the development of chess AI (Wilkenfeld, 2019).

These two differing sociotechnical imaginaries from machine learning specialists and chess experts produced several mutually beneficial arrangements for chess specialists and machine learning researchers, but also produced several conflicting interests within the field, as previously discussed. The general public also plays a role in the chess network that is almost entirely separate from that of both chess experts and AI researchers. When artificial intelligence first rose to prominence in the field of chess, an IBM machine known as "Deep Blue" beat the reigning chess world champion, Gary Kasparov (Ensmenger, 2022). News articles and the public both referred to the machine as some sort of monster, but nonetheless the machine served as a

comfortable entry point for chess enthusiasts to begin discussing machine learning (Bloomfield & Vurdubakis, 1997). Deep Blue was dismantled only a short time after its match against Kasparov, but the cultural impact of its matches has endured into the current day, which has irreversibly shifted the dynamic between machines and chess experts (Ensmenger, 2022; Wilkenfeld, 2019).

This STS project aims to consider the relationship between humans and the culture of optimization in chess to address the ways in which artificial intelligence has subverted the dominance of human-developed strategies in favor of those generated by computers and how this change in culture has already begun to shape the way that chess is played as well as the rules surrounding the game. As discussed in the work by Tomasev et al. (2022), artificial intelligence is currently being used almost as a preview for how new rulesets for chess may be played at a high level, and contemporary chess theorists are proposing revisions to the existing ruleset based on the data generated from such experiments.

Methods and Evidence

This STS project will employ an in-depth review of the scholarly literature surrounding modern integration of AI in chess and review several pop-culture sources to evaluate the opinions and effects of the general public as actants on AI chess development. This study will reference publicly available interviews with both machine learning specialists and chess experts, which relate to the use of Alpha Zero and chess development and will aim to glean how the Alpha Zero framework and other machine learning models function in the fields of chess and machine learning. This will establish a clear set of goals for the two largely independent groups of chess specialists and machine learning researchers and allow the paper to analyze how these goals may influence both the development of future technologies within machine learning and future iterations of the chess ruleset.

Conclusion

This paper proposes an analysis of AlphaZero as a boundary object between the fields of chess and machine learning. Because chess as a game, chess strategy, and machine learning are in a constant state of development, they are all very prone to induce change between disciplines and AlphaZero is a prime example. This paper will first provide historical context as motivation for the STS paper, then analyze AlphaZero as a boundary object between chess and machine learning. The technical project associated with this paper aims to create a machine learning agent that can play Pokémon. The conclusions reached by this paper might then provide a more general understanding of how the communities surrounding other games could be affected by machine learning technologies and may deepen the understanding of the technical project's role in society.

References

Bloomfield, B. P., & Vurdubakis, T. (1997). THE REVENGE OF THE OBJECT? On Artificial Intelligence as a Cultural Enterprise. *Social Analysis: The International Journal of Social and Cultural Practice*, *41*(1), 29–45. http://www.jstor.org/stable/23171730

Chowdhary, S., Iacopini, I., & Battiston, F. (2023, March 23). Quantifying human performance in chess. *Scientific Reports*, 13(1), 1 - 8.

Ensmenger, N. (2012, February 1). Is chess the drosophila of artificial intelligence? A social history of an algorithm. *Social Studies of Science (Sage Publications, Ltd.)*, 42(1), 5 - 30.

Januário, J. d. A., Giglio, S. S., & Amaral, S. C. F. (2024, August 1). Grandmasters, distinct elite: Taste submitted to discussion from the social conditioning factors of the predilection for chess. *International Review for the Sociology of Sport*, 59(5), 787 - 802.

Kothari, T. (2024, August 29). *Using reinforcement learning to play competitive Pokémon VGC*. Medium. https://medium.com/@tanujkothari/using-ai-to-play-pokemon-vgc-4d087f03e465

Larson, M. (2018, January 1). Optimizing Chess: Philology and Algorithmic Culture. *Diacritics*, 46(1), 30 - 29.

Lea, G. R. (2020). Constructivism and its risks in artificial intelligence. *Prometheus*, *36*(4), 322–346. https://www.jstor.org/stable/10.13169/prometheus.36.4.0322

Luchitskaya, S. (2015). Review of *Chess in the Middle Ages and Early Modern Age. A Fundamental Thought Paradigm of the Premodern World. Fundamentals of Medieval and Early Modern Culture, 10*, by D. O'Sullivan. *Mediaevistik*, *28*, 387–390. http://www.jstor.org/stable/44163605

Stewart, H., & Watson, N. (2020). A Sociotechnical History of the Ultralightweight Wheelchair: A Vehicle of Social Change. *Science, Technology, & Human Values*, *45*(6), 1195–1219. https://www.jstor.org/stable/26933350

The Third Build. (2021, July 19). *Future Sight AI*. Pokemon Battle Predictor. https://www.pokemonbattlepredictor.com/FSAI

Tomašev, N., Paquet, U., Hassabis, D., & Kramnik, V. (2022, February 1). Reimagining Chess with AlphaZero: AI is driving the next evolution of chess, giving players a glimpse into the game's future. *Communications of the ACM*, 65(2), 60 - 66.

Wilkenfeld, Y. (2019). Can Chess Survive Artificial Intelligence? *The New Atlantis*, *58*, 37–45. https://www.jstor.org/stable/26609113