Thesis Project Portfolio

Reinforcement Learning for Player Experience Enhancement in Pokemon Red

(Technical Report)

AlphaZero: the Downfall and Salvation of Chess

(STS Research Paper)

An Undergraduate Thesis

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Fulfillment of the Requirements for the Degree

Bachelor of Science

William Bradford

Spring, 2025

Department of Computer Science

Table of Contents

Exectuive Summary

Historically, computers have always been worse than humans when it comes to artistic expression. Now, with the rise of artificial intelligence, computers have become increasingly proficient at mimicking human artistic expression, and even perform better than humans in some areas. Specifically, machines are far superior to humans at playing relatively simple games like Connect Four. In the last few decades, computers have even far eclipsed humans at playing more complex games, such as western chess and the Chinese game, Go using novel machine learning techniques. Even despite this, machine learning has yet to crack the most complex of competitive games, like the Pokémon video games. This body of works aims to both advance machine learning in complex games and address the effects of machine learning on the value of human artistic expression.

The technical project of this portfolio aimed to develop an advancement towards a machine learning model that can effectively play Pokémon. This work aimed to implement the principles of Google Deepmind's AlphaZero, allowing the resulting machine learning model to play games of Pokémon against itself in order to find a model for playing the game that is unlike human gameplay and could eventually surpass human gameplay. To accomplish this, a program was created that allows the computer to interact with Pokémon Red, one of the first video games in the Pokémon series. The computer was then allowed to play games against itself and was scored on its ability to choose moves that reliably brought it closer to victory. This project had many unforeseen difficulties, which have resulted in several restructurings of the project. This project was intended to use a simulated Pokémon game, which could then be generalized to any version of Pokémon, however due to certain unforeseen limitations of the simulator, this was not possible. In its final form, the machine learning model is able to play Pokémon matches in which both parties have a singular Pokémon in Pokémon Red. The final performance of the model had a 13.6% win rate against random move selection. In other words, this model is significantly worse than randomly choosing a move. Further analysis is presented in the technical report, as these results were non-trivial.

The STS Research portion of this portfolio addresses the effects of machine learning on the value of human artistic expression in chess. In recent years, there has been a general sense that machine learning has decreased the variety of moves selected in chess and resulted in a dramatic increase in the number of ties present in high-level chess tournaments. Although this causes games to become more challenging, it also leads the players to feel less satisfied because they feel less able to creatively express themselves, since the most creative move is not always the most optimal move, and machine learning models have a strong preference for optimality. Now that machine learning and artificial intelligence has become a part of the chess ecosystem, its effects are irreversible, but many believe that it is more beneficial than it is harmful. First and foremost, it allows less privileged players to gain a better understanding of their mistakes when playing, thus increasing access to the game by providing feedback to those who might otherwise not have received it. Through recent advances machine learning research, it can be used to break ties in tournaments by analyzing who made generally better moves. More importantly, it has allowed researchers to experiment with different variations of chess that might be more creatively challenging at their highest levels.

Although the results of the technical project were somewhat disappointing, the training framework that was developed as a result of the project is very satisfactory. Such a training framework was previously not available, and further iterations of this framework might lead to a significant advancement of artificial intelligence in Pokémon. If advancements were to be made, the training framework would need to be adapted. While the concepts remain useful, the framework should be rewritten in C so that it can be compiled into the same executable as the game files. Initial results The STS Research project was very fruitful, and its results provide some hope in the sea of negative effects of artificial intelligence. Beginning this project, the author of this work did not expect to find himself in support of machine learning in chess, or by extension other games, but the creation of this work has successfully changed his opinion.