Portfolio Deep Dive
for **Teach – It**

[Will Update as Needed]

June 13, 2022

Introduction

My interest in machine learning speckles my website, but I've become especially interested in computer agents capable of mimicking human-like decision making in constrained environments. The previous statement may sound a lot like artificial intelligence, but that's not necessarily my end goal. Though, admittedly, my interest and the goal of artificial intelligence do seem to run parallel to one another.

For clarity's sake, I generally interpret the aim of most machine learning approaches as approaches seeking to optimize some procedure or outcome. Organizations want the most profitable stock trading agent, the undefeatable gaming agent, or the agent that accurately predicts the most effective and efficient business decisions. However, what I seek is human mimicry, which unfortunately allows the possibility for an agent to make "mistakes". That said, I think the range of questions we can ask of an agent that has been trained to decide in human-like ways in constrained environments / situations become a tad more inspired. We can move beyond asking for tomorrow's lottery numbers (i.e., tomorrow's winning stock ticker) and ask instead, "By the standards of society X, was the verdict of legal trial Y a fair one?" That question is markedly different from asking what the verdict and – if any – the associated penalties should be (because models making those decisions are already at play).

The more ephemeral questions – at least for a machine learning model and the surrounding ethics – that generally require some level of human consensus gain tangibility. From an idealistic viewpoint, a time consuming and potentially expensive survey of thousands of locals regarding the next restaurant to build and where to put it is replaced by a model that has been trained on the predilections of the people from its area of deployment. My thoughts about such a model and what society stands to gain from it might be in poor taste and even possibly butt-heads with the popular ethics of the field, but I can't help but think that value can be extracted from such an approach, which finally brings me to my "Teach – It" collection of projects.

Portfolio Deep Dive
for **Teach – It**

[Will Update as Needed]

June 13, 2022

What are Teach – It projects?

Aside from acting as contributions to my portfolio, the Teach – It projects are also mini-feasibility studies aimed at my goal of producing models in certain environments that sufficiently mimic human decision-making. By using common games and gathering move or sequence data from human players, my hope is to create a handful of computer agents that go from making uninformed / random moves to making actions that 'feel' more human. (Unfortunately, I don't have a litmus test for computer agents 'feeling' like human players in my game environments. So, I'm flying by the seat of my pants on this one.)

For those readers a bit more steeped in the machine learning field, my choice of words may steer you towards reinforcement learning – essentially using strategically constructed rewards and penalties to incentivize certain behaviors for a model. I do plan on using those techniques to some extent, but I'm also interested in evaluating the sequence data (e.g., more in line with recurrent neural networks) while attempting to avoid declaring value to certain moves.

The logic on my part is that in some scenarios – such as when playing a game – people naturally tend to choose self-maximizing options as they understand them. Eight-year-olds and thirty-eight-year-olds tend to play chess differently and think about the game differently, but when they play, they make moves that will inch them closer to victory. I am banking on the value human players naturally assign to moves in certain states to shine through in the sequence analysis of the game play, which may have broader implications for the types of situations this logic may be applied to. And, if there is a less 'gamey' situation that still seems to meet the self-maximization framework, then the very models I'm feasibility testing now might be able to one day shed light on those areas. A potentially controversial self-maximization arena would be the realm of politics. The game of chess and the political arena are leagues apart, but there are shared pinpoints in the frameworks of both – at least in my opinion.

Portfolio Deep Dive
for **Teach – It**

[Will Update as Needed]

June 13, 2022

Technicals

The Teach – It projects are still works in progress, but the bulleted list below should act as a decent summary of how they're currently stitched together.

- Link from main website to website(s) hosted on Google's Cloud Platform.
- Human player moves and/or game board states are received server-side and fed to a python-based machine learning model.
    - Currently, the TicTacToe model is already trained using development data I generated. However, once I improve the response speed of the process, I'll generate a fresh model and train it daily ~ weekly using data from the last 1,000 ~ 10,000 games. (In case you're wondering, a computer agent that makes random moves can learn how to play TicTacToe provided its playing against an opponent using a better strategy. I'll write more on that later.)
    - The chess agent will require more work because I'm not as familiar with the approach I want to use for it, but that's what the learning process is all about! Currently, it's in a playable state where the computer opponent makes random -yet chess legal – moves. However, that's largely due to the great programming skills of minds greater than mine (and acknowledged on the chess project site).
- The model returns the computer opponent's move, and the website's display is updated showing which move the opponent made.
- (Eventually) During the completion of a game – and only during the full completion of a game – the game data will be stored in a game database that will be used to train the next version of the model.
- After enough games have passed and the model appears to reach a level of convergence from its data and the human-player data, I'll generate a report detailing my findings from this series project collection.