

Injecting Machine Learning into the Apprentice Learner Architecture, Project Milestone Report 4 15-400, Spring 2020

Cayden Codel (ccodel)
Project website: www.contrib.andrew.cmu.edu/~ccodel

February 26, 2020

1 Major Changes

No major changes have occurred in the project.

2 What I've accomplished since last meeting

As the last two weeks have been filled with midterms and an out-of-town tournament for my extracurricular activity, I haven't spent too much time on my research. However, in the last two weeks, I have finished updating RumbleBlocks to a new version of Unity, I read Christopher MacLellan's Ph.D. thesis on apprentice learning and the AL framework (with TRESTLE), and I received access to player actions and data from previous user studies on RumbleBlocks. My goal is by next AL research meeting to get RumbleBlocks to interface with AL with example tracing (by using the player actions). Since a lot of that infrastructure is already in place, the engineering effort will be focused on how AL has changed in the last two years.

In addition to the thesis, I've read plenty of related papers to AL and apprentice learning. I've also identified a few papers that will be relevant to future work with reinforcement learning agents, such as "Apprenticeship Learning via Inverse Reinforcement Learning," which can be found here.

3 Milestone progress

I have completed the milestone **Update RumbleBlocks to a new version of Unity**. I hope to accomplish **Update RumbleBlocks to interface with the new version of AL** soon, now that my midterms are over. Since I am now comfortable with both the RumbleBlocks and AL code bases, I don't anticipate this taking too long.

4 Surprises

No surprises to report.

5 Looking ahead

In the next two weeks, I hope to get an interactive trainer with AL up and running. There are several files in the AL_Core repository with examples of interactive training (including one with Tic-Tac-Toe) that I can use as a skeleton for the one with RumbleBlocks (although I may run into some issues if a GUI is required, as I am not as familiar with Unity. However, an interactive trainer could be made using the log output in a pinch).

6 Revisions to future milestones

Milestones have been pushed back approximately one week, since the last two weeks have been so busy for me.

7 Resources needed

No additional resources are required at this time.