Improvements in Imitation Learning for Overcooked

AUTHOR: Duuk Niemantsverdriet SUPERVISOR: Robert Loftin RESPONSIBLE PROFESSOR: Frans Oliehoek

1 Research Question

How well can a collaborative AI agent that was trained by directly imitating human-generated data perform in *Overcooked*?

2 Background

- Imitation learning
  - An approach to machine learning where you train an AI agent to replicate the behavior of a human expert. [1]
- Behavioral cloning
  - Uses supervised machine learning algorithms (commonly neural networks) on state-action pairs.
- *Overcooked* [2]
  - Collaborative game focused around running a kitchen and serving dishes quickly.
  - Simplified environment with discrete movement options and timesteps and only one recipe.

[1] Y. Yue and H. M. Le, "Icml2018: Imitation learning tutorial." https://sites.google.com/view/icml2018-imitation-learning/, 2018.
[2] Ghost Town Games, "Overcooked." https://store.steampowered.com/app/448510/Overcooked/, 2018.
[3] M. Carroll, R. Shah, M. K. Ho, T. L. Griffiths, S. A. Seshia, P. Abbeel, and A. Dragan, "On the utility of learning about humans for human-ai coordination," NeurIPS, 2019.

3 Environment

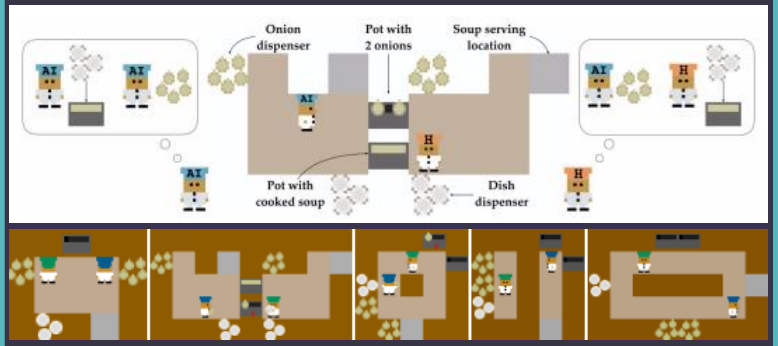

Figure 1: The *Overcooked* environment [3]

4 Experimental Setup

- Experimental agents
  - Baseline agent, trained on 96-dimensional featurization of the state
  - Two agents with smaller feature vectors
  - Three agents with larger feature vectors, which include data from previous states
- Evaluated against three agents
  - "Human proxy" agent, also trained with behavioral cloning
  - Self-play agent
  - Scripted agent

5 Results

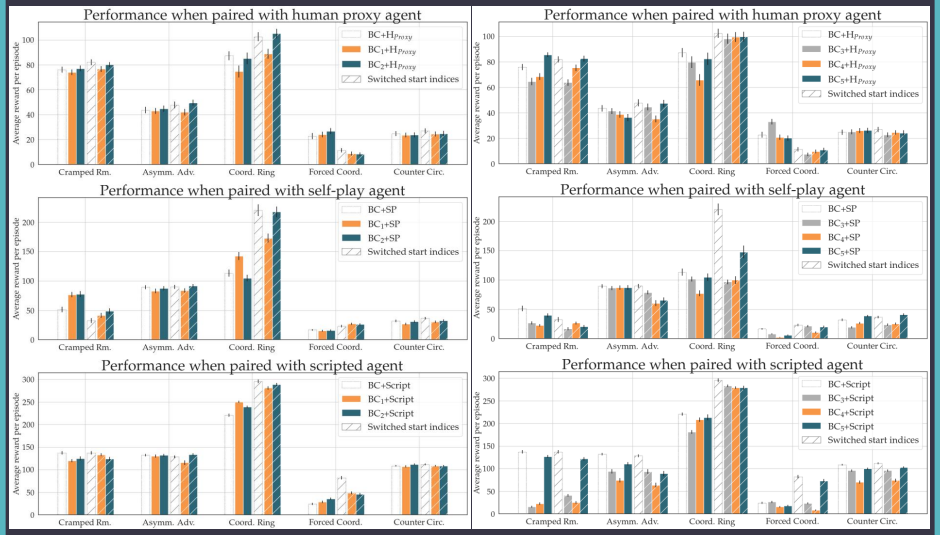

Figure 2: Results of each pairing of agents, average score in 100 simulated games of 400 timesteps.

6 Conclusions

- Taking historical actions into account hurts performance, especially with unfamiliar agents
- The dataset is the biggest limiting factor for behavioral cloning, both its quality and its size to support larger feature vectors.