Learning from Learners: Adapting Reinforcement Learning Agents to be Competitive in a Card Game

Pablo Barros, Ana Tanevska and Alessandra Sciutti
Italian Institute of Technology

Chef´s Hat Multiplayer Competitive Card Game and Simulation Environment

Reinforcement Learning Agents for Learning a Competitive Strategy

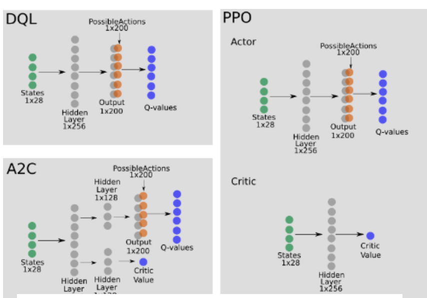

Three types of learning agents

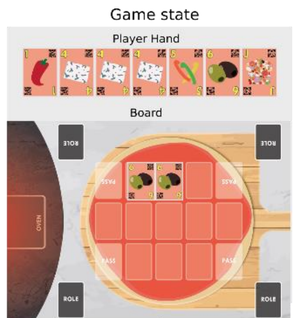

Gamestate based on cards at hand and current cards at the board

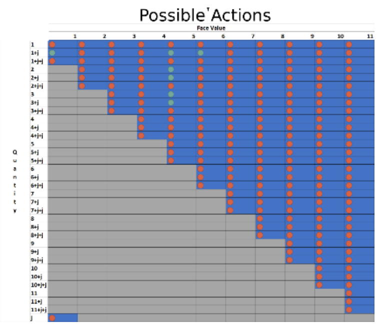

Action space: 199 different combinations of a discard action and one pass action

Three evaluation scenarios and performance measures over 100 games

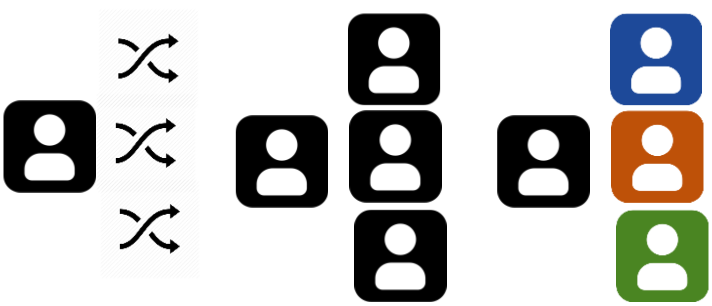

Vs Random

Vs Myself

Vs Others

vs. Random				
Model	Victories	Random1	Random2	Random3
DQL	**66.8** ±5.69	9.7 ±3.13	12.9 ±4.66	10.6 ±1.8
A2C	**65.1** ±5.19	9.3 ±3.1	12.1 ±4.35	13.5 ±3.58
PPO	**83.1** ±4.18	4.7 ±2.19	6.0 ±2.28	6.2 ±1.83
vs. Myself				
Model	Gen-1	Gen-25	Gen-50	Random
DQL	19.4 ±4.78	24.8 ±4.98	**42.9** ±7.06	12.9 ±6.64
A2C	25.4 ±4.39	29.1 ±6.14	**34.5** ±7.12	11 ±2.86
PPO	16.9 ±3.36	32.5 ±3.75	**40.3** ±3.52	10.3 ±4.1
vs. Others				
Model	Before training		After training	
DQL	35.9 ±3.11		35.9 ±3.11	
A2C	18.9 ±3.51		4.9 ±2.84	
PPO	**42.8** ±5.06		**48.5** ±40.6	
Random	2.4 ±0.8		3.3 ±1.85	

Performance results after training for a thousand games.

Strategy Emergence Observation

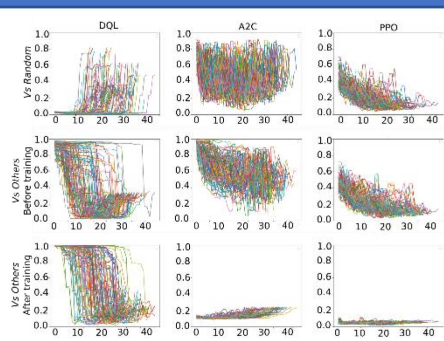

Each agent adopted a specific strategy, and adapted it when playing against each other.

Contributions

Simulation Environment

Trained Agents

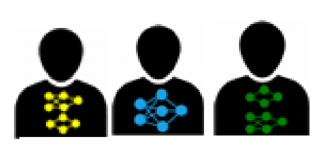

DQL A2C PPO

https://github.com/pablovin/ChefsHatGYM