Cooperative AI and Planning Algorithms in OVERCOOKED!

Jonte Herben

1. INTRODUCTION

- Cooperative AI is AI meant to cooperate well with a human.
- Overcooked is a multiplayer cooking game that tests the players' cooperation, making it prime real estate to evaluate cooperative AI.
- Planning algorithms create a plan from state A to state B.
- One paper already researched this subject[1], creating two planning-based AI, detailed in figure 1. But, they made multiple heuristic optimisations, leading to suboptimal results.
- This research will remove and analyse an important one: the omission of counters.

2. RESEARCH QUESTIONS

- How can cooperative planning within Overcooked be improved by adding counters?
  - How can planning be used to achieve cooperative AI?
  - What is the impact of more focused A* heuristics?
  - What is the impact of incorporating history?

Figure 1: overcooked_ai environment [2]

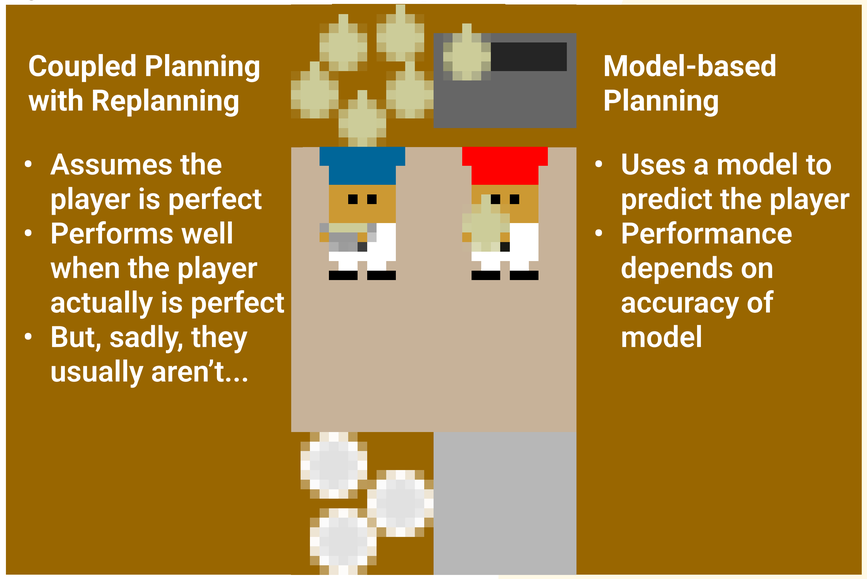

3. METHOD

- Just adding counters would increase the runtime too much, so substitute optimisations had to be made:
  - The A* prediction heuristics were made more specific.
  - The model-based planner was given a history to increase the accuracy of the model.

4. FOCUSED HEURISTIC

- The planning agents use A* pathfinding, which makes use of a heuristic to decide which nodes to explore.
- Changing the heuristic might improve the runtime and/or the performance, as visible in Figure 2.
- The heuristic was changed by including the distance from items to their destination.
- The final results show little impact on performance, but a significant decrease in runtime.

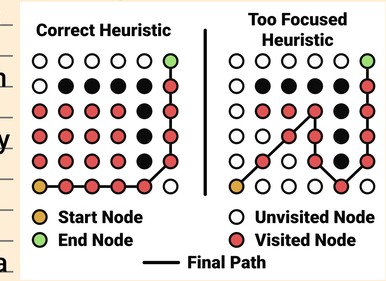

Figure 2: Possible Impact of Heuristic

5. HISTORY

- The evaluation models have a history to detect and break out of loops. However, the model inside the model-based planner does not have one, which leads to inaccuracies.
- This change was only made to increase performance and not the runtime.
- The final results drastically increase the runtime, which limited the amount of evaluations possible.
- The final results do show an increase in performance.

6. COUNTERS

- Even with multiple optimisations, many evaluations with counters were still too slow.
- The results that were obtained, shown in figure 3, showed the AI performed very poorly due to the agents getting stuck.

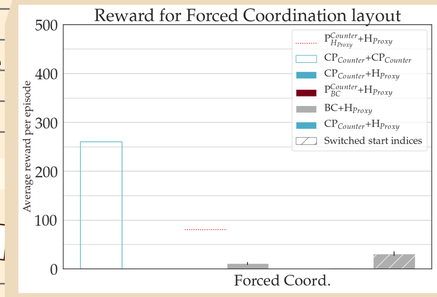

Figure 3: Final Results with Counters

7. CONCLUSION

- The final implementation is still slow and performs poorly.
- The runtime could be improved by adding more heuristic optimisations and the performance by adding more ways to prevent loops, either on the planner side, or the human side.
- The evaluation agent could also be refined, as they perform and adapt much worse than a normal human would.

[1] Micah Carroll et al. "On the utility of learning about humans for human-ai coordination". In:Advances in neural information processing systems 32 (2019).
[2] HumanCompatibleAI/overcooked_ai: A benchmark environment for fullycooperative human-AI performance. Nov. 2022. url: https://github.com/HumanCompatibleAI/overcooked%5C_ai.

TU Delft

Supervisor: Robert Loftin
Responsible Professor: Frans Oliehoek

J.H.J.Herben@student.tudelft.nl