Cooperative Games with Partial Observability

Piers R. Williams, Diego Perez-Liebana and Simon M. Lucas

Abstract

This paper describes an ongoing research project that aims to provide the most thorough study to date of cooperative partially observable (CPO) games. In CPO games players must cooperate in order to perform well, and the partial observability provides a strong incentive to cooperate: no single player can obtain optimal information without exchanging messages with its team members. Many existing multi-player games, real-world scenarios, and indeed life itself are well-described by this model. The project will provide novel insights on the scope for using general game playing agents to provide coperative AI, and the effects on player-experience of systematically varying the observability models and the communication protocols.

Introduction

This paper introduces and explains some Cooperative Partial Observability (CPO) domains and evaluates their potential use for Artificial Intelligence (AI) research, and equally importantly, identifies ways in which AI can be used to make better CPO games.

Writing agent AI for a Partial Observability (PO) domain is a challenging task, as is writing AI for cooperative domains. The combination of the two domains presents an even more complex AI challenge, but one that is approachable using recent advances in general game AI including algorithms based on Monte Carlo Tree Search and Rolling Horizon Evolution (Perez-Liebana et al. 2016). Game AI is already being used to aid the design of games using various statistical measures of experience an AI agent has while playing (Isaksen, Gopstein, and Nealen 2015). CPO games would benefit from automatic design or tuning in this manner if there were sufficiently strong agents available for them. Furthermore, these agents provide a ready supply of opponents and collaborators for human players.

Cooperative Partial Observability domains

This section describes an initial set of games used in the study.

Tiny Co-op

Tiny Co-op (Williams et al. 2015) is a simple puzzle domain that required cooperation between players to successfully obtain the maximum score. Two agents in a grid based world with walls, doors, buttons and goals are each tasked with visiting each goal. Score is awarded to both agents for visiting a goal. Doors are only open while an agent is located on top of the linked button - forcing the cooperative element. Tiny Co-op had some experiments with AI performed within a nearly unobservable environment - showing that Monte-Carlo Tree Search (MCTS) was capable of scoring well with a reasonable computational budget provided to it. The agents were only able to observe the score of the game, and simulate the effect of a move or sequence of moves had on the score of the game using the forward model.

Hanabi

Hanabi (Antoine Bauza) is an NP-complete (Baffier et al. 2016) game that features a custom set of cards depicting fireworks.

The game involves playing numbered and coloured cards in order onto the table to increment the number of each colour. Cards are held in hands of the players, but players are not allowed to view their own cards - only the cards of others. Players must choose between a normal turn or a message-passing turn, with a combination of these being necessary for optimal play.

An attempt to develop agents for Hanabi is presented in (Osawa 2015), although the game was restricted to two player versions. Numerous solutions were presented, with varying levels of information available to them from none (random) to full information (cheating).

Further work on agent AI for Hanabi using techniques derived from the Hat Guessing game is presented in (Cox et al. 2015).

Ms Pac-Man Vs Ghost Team

This problem domain is a modification of the original Ms. Pac-Man competition (Lucas 2007), (Rohlfshagen and Lucas 2011) that adds the PO constraints.

The Partially Observable Ms Pac-Man Vs Ghost Team competition (Williams, Perez-Liebana, and Lucas 2016) uses Line-of-Sight (LOS) as the method of PO. LOS is where the agents can see in straight lines up to a limit unless there is an obstacle in the way. Obstacles are considered to be the walls in the maze. Ghosts and pills don't count as obstacles. This applies to both Ms. Pac-Man and the Ghosts and means that they can see both forwards, backwards and

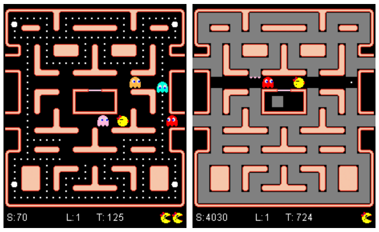

Figure 1: Ms. Pac-Man. Complete Observability on the left with Partial Observability on the right

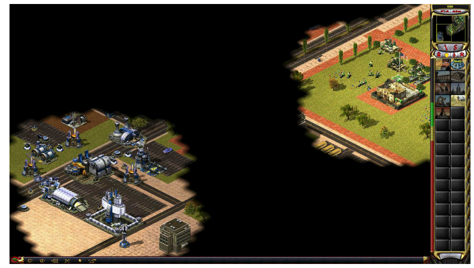

Figure 2: The Fog of War from the Real-Time Strategy game Red Alert 2

sideways. Agents cannot see around corners, just like real people. This is similar to the standard first person view although backwards, left and right sight are also allowed. A representation of the game and the PO constraint is shown in Figure 1.

In order to allow the ghosts to maintain a co-operative aspect, communication is allowed through a controlled medium. The game allows the transmission of messages through a basic postal system, that get delivered to the recipient later on in the game. The competition is running at the IEEE Computational Intelligence and Games Conference (CIG) 2016 conference.

Initial experiments in this domain comparing the same basic rule based strategy between a ghost team with communication and a ghost team without communication have been run. The addition of communication to the ghosts caused Ms. Pac-Man to drop down to a mere 33% of her previous score. Initial results including experiments with human players indicate that varying observaility and communication mechanisms provides a valuable extension to the normal design space of Ms Pac-Man-like games. Indeed, some players reported surges in excitement on turning a corner and suddenly finding a ghost a few pixels ahead!

Full Real-Time Strategy Games

Most Real-Time Strategy (RTS) games released for the computer feature multiple teams of players acting co-operatively as well as featuring the Fog of War (FOW) as shown in 2.

These games are incredibly complex for AI to handle, with hundreds of units at once over a large and diverse map. This makes them unsuitable for research into CPO with only mild success in some older RTS games currently recorded. Full size RTS games have been studied in AI research, though typically only controlling a single player acting greedily.

RTS games feature many levels of abstraction for agent AI. Units are typically controlled by a single AI that issues orders to groups of units. This is simpler to write and also bears some resemblance to most military control structures. Writing cooperative AI in this genre is typically restricted to the large tactical AI systems cooperating with each other as opposed to individual units cooperating with each other.

Conclusion

This paper described the motivation for research in CPO games and discussed a number of examples. Further study of the area will provide insight into the relative strengths of various AI algorithms for this type of game, and also how AI can be used to provide novel game experiences by automatically tuning the details of the observability and the communication models.

References

Baffier, J.-F.; Chiu, M.-K.; Diez, Y.; Korman, M.; Mitsou, V.; van Renssen, A.; Roeloffzen, M.; and Uno, Y. 2016. Hanabi is NP-complete, Even for Cheaters who Look at Their Cards. *arXiv preprint arXiv:1603.01911*.

Cox, C.; De Silva, J.; Deorsey, P.; Kenter, F. H.; Retter, T.; and Tobin, J. 2015. How to make the perfect fireworks display: Two strategies for Hanabi. *Mathematics Magazine* 88(5):323–336.

Isaksen, A.; Gopstein, D.; and Nealen, A. 2015. Exploring game space using survival analysis. *Foundations of Digital Games*.

Lucas, S. M. 2007. Ms pac-man competition. *ACM SIGEVOlution* 2(4):37–38.

Osawa, H. 2015. Solving Hanabi: Estimating Hands by Opponent's Actions in Cooperative Game with Incomplete Information. In *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*.

Perez-Liebana, D.; Samothrakis, S.; Togelius, J.; Schaul, T.; and Lucas, S. 2016. Analyzing the Robustness of General Video Game Playing Agents. In *Proceedings of the IEEE Conference on Computational Intelligence and Games*.

Rohlfshagen, P., and Lucas, S. M. 2011. Ms Pac-Man versus Ghost Team CEC 2011 Competition. In *Evolutionary Computation (CEC), 2011 IEEE Congress on*, 70–77. IEEE.

Williams, P. R.; Walton-Rivers, J.; Perez-Liebana, D.; and Lucas, S. M. 2015. Monte Carlo Tree Search Applied to Co-operative Problems. In *CEEC2015 - IEEE Conference on Computer Science and Electronic Engineering*, IEEE CEEC, 219–224. IEEE Computer Society.

Williams, P. R.; Perez-Liebana, D.; and Lucas, S. M. 2016. Ms Pac-Man Versus Ghost Team CIG 2016 Competition. In *Proceedings of the IEEE Conference on Computational Intelligence and Games*.