Pandemic as a Challenge for Human-AI Cooperation

Pablo Sauma Chacón and **Markus Eger**
Universidad de Costa Rica
San José, Costa Rica
pablo.saumachacon@ucr.ac.cr, markus.eger@ucr.ac.cr

Abstract

Cooperation between human players and AI agents in games is a subject of great interest in the research community. In this paper we propose a new domain as a challenge for future AI research: the cooperative game Pandemic. This game represents a challenge because it requires cooperation between players in an environment with incomplete information. Additionally, we propose a first approach for a cooperative agent for the game, that uses planning in conjunction with plan recognition. We also propose an experiment to test the mentioned approach. In this paper we argue why Pandemic makes a compelling domain for AI research, the status of our project, as well as which challenges remain to be addressed.

Introduction

Cooperation in games is a subject of increasing interest for AI research. However, the problem of engaging in cooperative games, instead of adversarial ones like for example Go (Silver et al. 2016) or StarCraft II (Vinyals et al. 2019), and specially in scenarios where communication may be restricted by the game rules has proven to be challenging.

Most recent advances in the research of cooperative games have been research using the game Hanabi (Bauza 2010; Bard et al. 2019), where a group of players must cooperate with each other to win the game by using hint-actions defined by the game rules as the only means of communication between them to convey information. This game has been used in various studies that range from a team of bots interacting with each other to get a perfect score (Cox et al. 2015; Bouzy 2017), agents that try to collaborate with other agents of a different type (Canaan et al. 2019), to using belief modeling and communication theory to determine the actions of an AI agent playing with a human player (Eger, Martens, and Cordoba 2017; Liang et al. 2019).

In this paper, we present a domain that we believe constitutes a new challenge for future AI research: The cooperative game Pandemic. We will discuss what makes this game an interesting application for AI, and which problems have to be addressed by AI agents that are to play the game with human players. We will then discuss a proposed solution to these challenges, and how we plan to validate it.

Related Work

Planning (Fikes and Nilsson 1971) has been used in games before, an example of this is the case of the AAA game F.E.A.R developed by Monolith Productions, which uses planning as a tool to create smarter and more versatile intelligence for the computer controlled players giving the players a more challenging experience (Orkin 2006). Other examples are applications of decompositional planning in the form of hierarchical-task networks in the creation of bots for the game Unreal Tournament (Hoang, Lee-Urban, and Muñoz-Avila 2005) and non-playable characters for the game The Elder Scroll IV: Oblivion (Kelly, Botea, and Koenig 2008).

Plan recognition is the process where, given a sequence of observed actions, the agent has to identify a goal which explains the observed actions. Typically, it is assumed that the observed actions are part of an optimal plan. One of the ways to perform this is by treating the plan recognition problem as a planning problem in which the observations are included as a new set of actions which must be completely performed as part of a plan before reaching a goal state (Ramirez and Geffner 2009; Cardona-Rivera and Young 2015). This process can be useful for cooperative agents, since cooperation requires understanding of other agents' goals, but it has not been studied extensively yet in the context of cooperative games.

Pandemic as a Challenge for AI Agents

In this section, we will present why Pandemic constitutes an interesting challenge for future AI research. We will start by explaining rules of the games, including its objectives and mechanics, and then the reasons behind its consideration as a challenge for the research in cooperation between human and AI players.

Game Rules

Pandemic (Leacock 2008) is a cooperative board game designed to be played by 2 to 4 players. The group of players are considered members of a disease control team tasked with discovering the cure for 4 diseases that are ravaging the world (an example of the game board can be seen in figure 1). To do this they use their characters (represented as pawns on the board) to travel the world treating the diseases, building research stations and gathering enough knowledge to discover a cure and, at the same time, prevent outbreaks (players lose if 8 or more outbreaks occur) and the propagation of the diseases (players lose if one of the diseases propagates too much), whilst being in a race against time since the numbers of turns in each game is limited by cards drawn from a player card deck.

At the beginning of the game each of the players is assigned a role which may impact their play style, changing the way in which they can perform certain actions and their general purpose in the game. There are two decks: the player deck and the infection deck. Both consist of one card corresponding to each city on the board, and are shuffled at the beginning of the game. The player deck also contains special event cards and so-called "epidemic". The players are dealt an initial amount of player cards and the game board begins with an initial infection across the world determined by cards taken from the infection deck.

After the game begins, players take turns in order, and perform 4 actions on each of their turns. Actions players may take include: movement actions (there are 4 different types of movement actions, some requiring the use of a player card), building a research station in the city the player is in, treating a disease in the current city, sharing knowledge (giving or receiving a card from other players in the same city), discovering a cure or a special action depending on the player role. After a player has performed the 4 actions, the player will draw 2 cards from the player deck, which will allow them to perform actions that require discarding a player card, and a specific amount of infection cards (2 to 4, depending on the game state). Each city drawn in the infection cards gets one disease counter of the specified color.

The main challenge the players are presented with by the game is that they are in a race against the previously mentioned loss conditions while trying to discover the cure for the diseases. To be able to discover a cure a single player is required to be in a city with a previously built research station and discard 5 cards which share the same color (of the disease for which the cure is going to be found). However, the players only draw 2 cards at the end of their turns and have a hand size limit of 7 cards, which forces the players to coordinate and help each other so that all 4 cures can be discovered in time. This requires the players to strike a balance between the need of staying together to exchange cards and spreading out to fight the diseases.

Also, as previously mentioned, the game starts with an initial infection and at the end of each players' turn more infection cards are drawn. This has the effect of adding disease counters of a specific color to one of the cities. If a city which would be infected would end up with more than 3 counters of a same color then an "outbreak" occurs, meaning that instead of infecting the city itself, the disease spreads to all surrounding cities (and the outbreaks themselves are a lose condition). This forces the players to stay on watch and have to spend much of their effort trying to prevent the propagation of diseases. There are also special epidemic cards in the player deck which, when drawn, add 3 disease counters to a random city and then reshuffle the discarded infection cards, which are put on top of the infection deck, causing the same cities to be infected multiple times over the course of the game.

In conclusion: Pandemic is a game in which the players must balance between the long term goals (finding the cures) and their immediate needs (preventing propagation of diseases), while also trying to cooperate with each other either by helping each other or dividing tasks efficiently. And lastly, the players have to make all of these decisions while having incomplete information regarding the game state, because of randomly shuffled player card and infection decks.

Cooperating with a Human Player

The problems an AI agent faces when playing Pandemic can be presented as the problem of figuring out which is the best plan or set of actions to achieve victory in the game, but, forced by the game's mechanics, figure out how to cooperate with the other players specifically since only their actions are observable and failure to work together properly would most likely result in defeat. As communication is restricted the AI agent requires to have some sense of what the human players may be trying to do, which presents an interesting challenge for developing systems that collaborate with humans with unknown (intermediate) goals.

Plan recognition is one approach to solving how to cooperate properly with players in this game that does not require a means of explicit communication between players. One of the main issues regarding this approach arises from the fact that human players are not completely predictable and are likely to make sub-optimal decisions. This provides an additional challenge for AI agents trying to predict what the human will do, because simply using its own reasoning method (which is likely desired to be tuned for optimal play) may not be sufficient to account for sub-optimal play on part of the human player. Another issue is the unknown information caused by the randomness of the card decks, and the fact that the space of possible actions is rather large. To successfully play the game with a human player, all of these challenges need to be taken into account.

Proposed Approach

We will now present one possible approach to Pandemic, in the form of the creation of a planning agent which incorporates plan recognition as the tool to interpret and cooperate with the human player. We also propose an experiment which could indicate the effectiveness of such an agent and measure if the cooperation is improved by the incorporation of plan recognition into the system.

Planning Agent

For the development of the planning agent, an encoding that allows to describe the game state of Pandemic needs to be developed and the actions for the game with their preconditions and effects for the planning agent to use have to be defined.

However the planning problem in this case is non-trivial because the game contains hidden, and random, information in the form of the infection and player decks. While some information can be deduced (by keeping track of the game

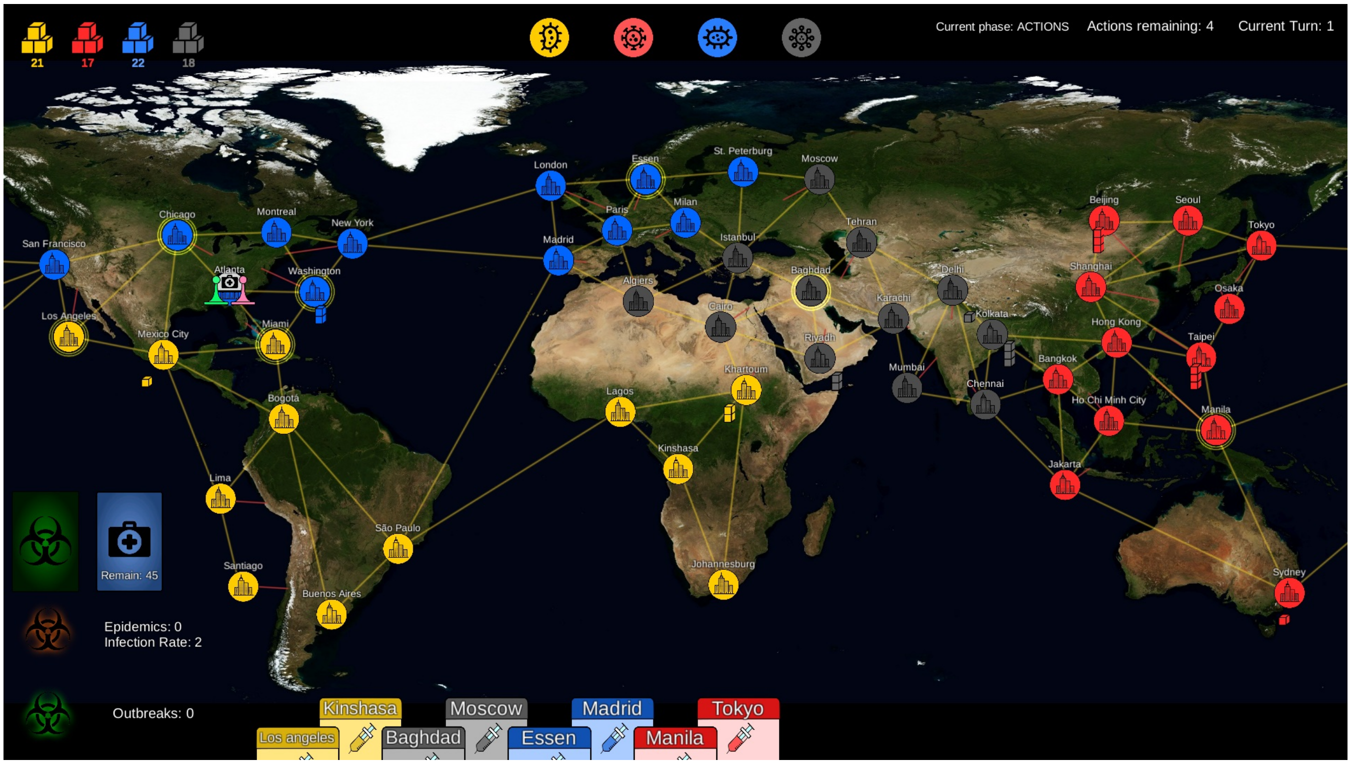

Figure 1: Screenshot of the web browser client developed for the game Pandemic

state), most of the information remains unknown. To solve this issue, we plan to use Monte Carlo Tree Search (Chaslot et al. 2008) which can be applied at the times when cards need to be drawn (random events) to simulate possible outcomes for the game. This approach has the advantage that, by using an anytime algorithm, the execution can be stopped arbitrarily after a set amount of time in order to keep the game reactive for play with a human player.

Plan-Recognition

To implement plan recognition for the AI agent, it is first important to create a way of extracting, storing and accessing the previously performed actions done by the player. This will be necessary since that information will be passed to the plan recognition system which allows the planning agent to determine which of the possible goals the human player may be trying to achieve.

Nonetheless, identifying the current goals of the player might not be enough, since the planning agent needs to take those predictions into account for its own planning, to avoid a duplication of efforts and instead assist the player if needed or capable. For example, if the player is about to discover a cure, and is moving towards the agent's pawn because they require a card in the agent's hand, the agent should infer that goal and cooperate with it, instead of "running away" from the human player.

Another issue to take into account is the fact that the game state is changing every turn, and as such a player's plan may change as a result of a card drawn at the end of their turn. The planning agent may need to readjust its interpretation of the human player's plan when given more information to provide a teammate-like behaviour in an environment in which there is no communication, while also facing the fact that the player's actions might not be part of an optimal plan. To account for this, the AI agent needs to strike a balance between playing optimally itself, and cooperating with the human player's goal, rather than always choosing one over the other.

Experiment Design

For the purpose of performing an experiment with the developed agent, and obtaining measures of the impact of plan recognition on the cooperation between an AI and a human player, we propose the comparison of the agent with and without using plan recognition in a controlled environment.

To this end, a platform that allows participants to play with the agent will be used to gather information regarding player satisfaction, via a questionnaire at the end of each game, as well as information regarding win rate and the final state for each of the games. The questions participants will be asked will focus on their impression about the agent and its reaction to changes in the game state, as well as the participant's own actions. This can provide insight about the overall quality of the agent, as well as the perceived effectiveness of the plan recognition module.

In the current state of our project the platform for the experiment has already been implemented, consisting of a client built in Unity and a Python server, based on similar setups that were used to evaluate Hanabi agents (Eger and Martens 2017) and agents for One Night Ultimate Werewolf (Eger and Martens 2019) in the past. The client runs in a web browser to allow easier access to a plethora of players to interact with the game and participate in the evaluation process of the agent. A screenshot of the client can be seen in figure 1. The server, meanwhile, manages the game logic by handling all game requests made by the clients and was built with the objective of providing a platform in which a variety of agents can be easily developed and deployed making use of the existing library, in order to encourage future research.

For the implemented game logic, it is important to mention that the game was implemented with some minor restrictions to reduce the game to its core mechanics. In the original game, the player cards contain so-called event cards which may be used by the players at any moment to perform a special action (each event card has its own distinctive action). These cards were removed in our implementation of the game as they would greatly increase the complexity of the problem without showing up often and consistently enough to significantly impact the core game play experience. Also, two out of the seven original game roles were omitted in our implementation: the "contingency planner" and the "dispatcher". The contingency planner has a special ability that uses event cards, which were removed, and could therefore never be used. The dispatcher, on the other hand, has the ability to move other players pawns as if it were their own, which could potentially make the search space explode so the role was removed as part of the implementation, since there are enough roles available to keep the game interesting. However, these restrictions do not significantly change the nature of the problem and can be addressed in the future.

Additionally, the data from the game, such as actions performed and the game state at every action (both for the human player and the computer) are recorded by the server (with the participant's consent) in order to produce data sets for future research.

Conclusion and Future Work

In this paper we present the domain Pandemic as a new challenge for future AI research. We describe the characteristics that make this problem specifically difficult and different from previously studied domains. We also describe our proposed approach for AI agents that can play the game with human players without requiring explicit communication.

As stated, the implementation of the planning agent and the application of the plan recognition to its planning is currently work in progress, as is the application of the proposed experiment. The game client and server, however, are finished, and are available on GitHub for use by the research community [1].

As part of future work, there are several possibilities that might be considered. The restrictions applied to the game in the implementation could be removed, allowing the use of the complete game. Another option for this modification would be to include the event cards but only allow players to use them on their turn which would provide a more controlled search space for the agent.

Explicit or direct communication could also be considered as a way of approaching the problem and would be an interesting subject for future research. This would involve a way to make the player convey its plans to the agent and vice-versa, possibly by predefined messages or cues. The gathered data from the games played by participants in the planned experiment could also be used as part of a study regarding actions and decision making in games.

Finally, the proposed methodology could also be applied to other cooperative games such as Legends of Andor (Menzel 2012) or the Big Book of Madness (Rambourg 2015), which are turn based games in which players must cooperate with each other to win the game though planning and teamwork.

[1] https://github.com/BlopaSc/PAIndemic

References

Bard, N.; Foerster, J. N.; Chandar, S.; Burch, N.; Lanctot, M.; Song, F.; Parisotto, E.; Dumoulin, V.; Moitra, S.; Hughes, E.; Dunning, I.; Mourad, S.; Larochelle, H.; Bellemare, M. G.; and Bowling, M. 2019. The Hanabi challenge: A new frontier for AI research. *CoRR* abs/1902.00506.

Bauza, A. 2010. Hanabi. https://boardgamegeek.com/boardgame/98778/hanabi. Online; accessed 25 February 2019.

Bouzy, B. 2017. *Playing Hanabi Near-Optimally*. 51–62.

Canaan, R.; Togelius, J.; Nealen, A.; and Menzel, S. 2019. Diverse agents for ad-hoc cooperation in hanabi. *arXiv preprint arXiv:1907.03840*.

Cardona-Rivera, R., and Young, R. 2015. Symbolic plan recognition in interactive narrative environments.

Chaslot, G.; Bakkes, S.; Szita, I.; and Spronck, P. 2008. Monte-carlo tree search: a new framework for game ai. In *Proceedings of the Fourth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 216–217. AAAI Press.

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

Eger, M., and Martens, C. 2017. A browser-based interface for the exploration and evaluation of hanabi AIs.

Eger, M., and Martens, C. 2019. A study of AI agent commitment in One Night Ultimate Werewolf with human players. Proceedings of the 15th AAAI International Conference on Artificial Intelligence and Interactive Digital Entertainment.

Eger, M.; Martens, C.; and Cordoba, M. A. 2017. An intentional AI for hanabi. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*, 68–75.

Fikes, R. E., and Nilsson, N. J. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial intelligence* 2(3-4):189–208.

Hoang, H.; Lee-Urban, S.; and Muñoz-Avila, H. 2005. Hierarchical plan representations for encoding strategic game ai. 63–68.

Kelly, J.-P.; Botea, A.; and Koenig, S. 2008. Offline planning with hierarchical task networks in video games. In *Proceedings of the Fourth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, AIIDE'08, 60–65. AAAI Press.

Leacock, M. 2008. Pandemic. https://boardgamegeek.com/boardgame/30549/pandemic. Online; accessed 25 February 2019.

Liang, C.; Proft, J.; Andersen, E.; and Knepper, R. A. 2019. Implicit communication of actionable information in human-ai teams. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 95. ACM.

Menzel, M. 2012. Legends of andor. https://boardgamegeek.com/boardgame/127398/legends-andor. Online, accessed 26 July 2019.

Orkin, J. 2006. Three states and a plan: The A.I. of F.E.A.R.

Rambourg, M. 2015. The big book of madness. https://boardgamegeek.com/boardgame/171233/big-book-madness. Online, accessed 26 July 2019.

Ramirez, M., and Geffner, H. 2009. Plan recognition as planning. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI-09)*.

Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; van den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; Dieleman, S.; Grewe, D.; Nham, J.; Kalchbrenner, N.; Sutskever, I.; Lillicrap, T.; Leach, M.; Kavukcuoglu, K.; Graepel, T.; and Hassabis, D. 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529(7587):484–489.

Vinyals, O.; Babuschkin, I.; Chung, J.; Mathieu, M.; Jaderberg, M.; Czarnecki, W.; Dudzik, A.; Huang, A.; Georgiev, P.; Powell, R.; Ewalds, T.; Horgan, D.; Kroiss, M.; Danihelka, I.; Agapiou, J.; Oh, J.; Dalibard, V.; Choi, D.; Sifre, L.; Sulsky, Y.; Vezhnevets, S.; Molloy, J.; Cai, T.; Budden, D.; Paine, T.; Gulcehre, C.; Wang, Z.; Pfaff, T.; Pohlen, T.; Wu, Y.; Yogatama, D.; Cohen, J.; McKinney, K.; Smith, O.; Schaul, T.; Lillicrap, T.; Apps, C.; Kavukcuoglu, K.; Hassabis, D.; and Silver, D. 2019. AlphaStar: Mastering the real-time strategy game StarCraft II. https://deepmind.com/blog/alphastar-mastering-real-time-strategy-game-starcraft-ii/. Online; accessed 5 September 2019.