

# Design and Modelling of Hanabi-Playing Agents using Artificial Intelligence

Shweta Pramodrao Sontakke<sup>1</sup>, Dr. A. N. Thakare<sup>2</sup>

<sup>1</sup>P. G. Student, Department of Computer Engineering, Bapurao Deshmukh College of Engineering, Sevagram  
Wardha, India

<sup>2</sup>Assistant Professor, Department of Computer Engineering, Department of Computer Engineering  
Bapurao Deshmukh College of Engineering, Sevagram, Wardha, India

## ABSTRACT

In order to affect your own activities, agent modelling entails thinking about how other agents will behave. In this work, we look at how agent modelling is used in the collaborative card game Hanabi, which is based on concealed information. In addition to an Information Set-Monte Carlo Tree Search (IS-MCTS) agent, we implement a variety of rule-based agents from the literature as well as our own inventions. We see that IS-MCTS produces poor results, so we create a new predictor version that incorporates a model of the agents with whom it is matched. This agent outperforms IS-MCTS in terms of game-playing strength, owing to its consideration of what the other agents in the game would do. We also develop a faulty rule-based agent to demonstrate the predictor's capability with such an agent.

**Keywords :** Set-Monte Carlo Tree Search, Hanabi-Playing Agents, Artificial Intelligence

## I. INTRODUCTION

Hanabi is a partially-observable [1] cooperative board game that earned the prestigious Spiel des Jahres award for best board game of the year in 2013. It has been mentioned in a number of recent scholarly articles for the reasons indicated below. The purpose of this article is to see if using agent modelling may increase the strength of the game's agents. Hanabi offers a number of unique characteristics that make it an excellent candidate for agent modelling research. To begin with, the domain is a cooperative one, requiring the agents to collaborate in order to attain a common goal. This disadvantages greedy agents: for example, assisting another player in scoring a point is preferable to playing a hazardous card that may finish

the game. Second, it incorporates well-defined communication acts into its rules. These make use of a communication resource that must be handled by the agents. Finally, the game contains secret information, with no single participant having access to the complete state of the game. Because faulty information must be intelligently reasoned about, this adds to the complexity of agents' tasks. Even when players have perfect knowledge, Hanabi has been demonstrated to be NP-Complete [2]. In the literature, a variety of rule-based ways for developing Hanabi-playing bots have been given, but there have been few attempts to utilise more general tactics. This study contributes to the redressing of this imbalance. Furthermore, because knowledge about other players' methods might assist human participants in

cooperative games, we investigate if such knowledge may assist our general agents in making judgments.

### A. Hanabi

Hanabi is a cooperative card game in which two to five players work together to complete five stacks of consecutively numbered cards (one for each of the game's five suits). A deck of 50 cards, each with a suit and a rank, is used to play the game. White, yellow, green, blue, and red are the colours of the outfits. There are three cards of rank 1, two cards of rank 2, three cards of rank 3, and one card of rank 5 in each suit. There are also two sorts of tokens in the game: an information token and a life token. The game begins with each player having three life tokens and eight information tokens. Every player starts with a five-card hand dealt at random. Players can't see the suit or rank of their own cards since they're held facing away, but they can see the suit and rank of the cards held by the other players. The cards that were not handed out at the start of the game are placed face down in a draw deck that will be accessible throughout play.

Play proceeds with each player taking it in turn to perform an action of their choice. There are three different types of action available:

**Tell** Select a player and point to all their cards of a given number or suit. This costs one information token.

**Play** Choose a card from the player's own hand and play it.

**Discard** Choose a card from the player's own hand and add it to the discard pile

A set of cards may only be recognised by their suit or their rank in a single Tell action; they cannot be recognised by both. Furthermore, cards must be present in the hand to be identified; it is not permissible to claim that another player does not have any cards of a particular suit or rank. When you

play a card, it goes into the stack with the same suit as the one you just played. It is not necessary for the player to know which stack the card belongs to; for example, it is allowed to play a 1 that has been indicated to you blindly at the start of the game. Each card in the stack must be of the same suit as the card below it and have a rank one higher than the one below it (except for 1 cards, which are used to start a stack). The group loses one life token if a card is not played in the correct order. Obtaining an additional information token by completing a stack of cards connected with a specific suit (if the team does not already have the maximum number, eight).

Discarding is only allowed if at least one information token can be obtained. This means that you can always do either a Tell or a Discard action.

The player then draws a replacement card from the draw deck after discarding or playing a card. All players may see the cards that have been discarded. When you discard a card, the amount of information tokens you have increases to the maximum. After all of the cards in the draw deck have been pulled, all players are given one more turn before the game is declared over. If a team's life tokens are depleted, the game is also over.

The top card of each stack that has been appropriately played is added together to determine the score. In the normal game, the highest potential score is 25, which is reached by finishing the stacks for all five suits; remaining life or information tokens are not tallied towards scoring in the normal game.

### B. Multi-agent domains

Multi-agent domains can be categorised as either centralized or distributed. A centralized system features a single controller controlling multiple agents; a distributed system has each agent in the world controlled by a separate controller. In this paper, we consider only the distributed approach.

Existing work in this space includes: attempting to reason about what the other agent knows using

answer set program- ming [3]; iterating on a plan that is communicated between agents [4]; and attempting to use plan recognition to allow one agent to assist another in a planning task [5].

Another possibility involves co-operative, multi-agent learn- ing. Within this area, there have been attempts to learn models of teammates in order to make more informed decisions about which action to take. For a review of the literature, see Panait & Luke [6].

The use of embedding agent models into Monte Carlo Tree Search (MCTS) has previously been looked at by Barrett et al in the pursuit domain [7]. They made the assumption that all agents except for their modelling agent would be using the same, fixed strategy, and embedded perfect knowledge of this strategies into their agent. One of their findings was that the system did not perform well with models that didn't represent the behaviour of the agent.

The use of Theory of Mind (ToM) (reasoning about what the other agents know and will do in a given situation) has proven useful in competitive games such as Rock Paper Scissors [8]. In these games, higher-order ToM agents were able to out- perform lower-order ToM players.

### C. Co-ordination in Hanabi

In Hanabi all agents have access to different information; because of this, a centralised approach to multi-agent planning would not make sense in this domain as private information must not be shared between agents.

The fact that the Tell action has an associated cost (an in- formation token) means that information about a player's hand needs to be communicated efficiently. Also, because Hanabi players are limited to a set of well-defined communication actions, communication between them is very limited. This makes using communication between agents to co-ordinate their actions a challenging prospect — which is one reason why Hanabi is increasingly becoming the object of

research. The understanding of other players' strategies forms a core component of a great number of games and has been studied widely [8]. Existing Hanabi research assumes that all agents are playing the same pre-agreed strategy. The ability to reason about the actions that a player would take and their reasons for taking these actions can be used as part of the reasoning process of an agent.

Our approach is to assume that we have access to a model which, given a state, will be able to return a possible action that an agent would perform in that state; if the agent may make multiple moves, then a single action from the set of possible actions will be returned. Given this model, we are able to incorporate the behaviour of the other agent into our model without understanding of that agent's reasoning process. A point to note is that Tell actions can convey more information than just the obvious: because all cards of a given suit or rank must be identified, cards which are not identified therefore must not satisfy the criterion. This negative information can be used to inform the possible values for a given card. Negative information can add up over a few turns, providing enough information to determine what a card is — or at least that it is playable. In the end game, such knowledge becomes very powerful.

Human players of Hanabi often make additional use of Tell actions. In particular, they can restrict their Tell actions by convention only to identify certain cards as playable. For example, suppose that Player 2 had the hand (R, 1), (B, 1), . . . and the current stacks on the table were (R, 1), (B, 0), (G, 0), (W, 0), (Y, 0) . Player 1 may elect to tell Player 2 about the suit rather than the number, to avoid identifying the non-playable red card. Player 2 could then infer that the card being identified was indeed a playable card, as they would know that Player 1 would not have identified a non-playable card. As they were told the suit rather than the number, they could further infer that they have a non-playable 1 in their hand

(although they would not know the location of this card). The use of information in this way requires an understanding of how the player will use the provided information as part of their policy.

#### D. Monte Carlo Tree Search

MCTS [9] is a widely-used tree-search algorithm that can operate without domain-specific knowledge. This gives MCTS the anytime property: the algorithm can be stopped at any time and can provide an answer for the next move. Given more time, it will typically produce a more accurate answer.

MCTS proceeds using multiple iterations of the four main stages shown in Figure 1. The iterations typically continue until a predetermined end condition is met, such as running out of time. In the selection stage, the current tree is traversed using the tree policy to select the best child of each node. In the expansion stage, a new node is added to the tree. In the main, simulation phase, a simulation (rollout) of future moves is undertaken from the state represented by the new node until an end condition is met. Moves are selected according to the default policy (which is often to select at random from all possible moves). In the backpropagation phase, nodes in the tree that were selected are updated with the result of the rollout.

#### E. Monte Carlo Tree Search and Theory of Mind

Zero-order theory of mind [8] agents are capable of using an agent's history in order to inform future actions. A first-order theory of mind agent is capable of using a model of a zero-order theory of mind agent to inform its own future decisions. Our selection of MCTS for use in this domain came from a particular desire to find an algorithm that could be easily modified to operate with predictions of what other agents would do. This makes it a zero-order agent.

This approach has been tested before in the Tiny Co-op domain [11] by Walton-Rivers, who found that prediction worked best with a deterministic

agent that did as it was instructed [12]. The Tiny Co-op domain is a simple, grid-based world containing a number of agents, goals, doors and buttons. Each agent must visit each goal individually for successful completion. Doors separate different areas in which the agents can move, and each door will only open if an agent is standing on its associated button. This forces the agents to co-operate to succeed overall.

While MCTS was a good performer in Tiny Co-op when paired with itself (and even with random agents), it struggled when trying to co-operate with a particular agent that was designed to follow direction indications. Essentially, this follower agent moved to where it was instructed to move, but MCTS didn't pick up on this. The root cause was that it didn't model such behaviour in its search tree, leading to inaccurate states in the majority of the search space. The author added agent modelling to MCTS and found that the performance of MCTS when paired with the follower agent improved significantly. In this paper we used this approach to create a Hanabi-playing agent to assess the effectiveness of agent modelling in this domain.

#### F. Previous research

1) Imperfect Information AI: Games with imperfect information are a complex challenge for AI. Poker is often chosen as an application, because it is a game that many people are familiar with on some level. Poker contains an unusual dynamic for games, as a strong player doesn't so much play the game as play the opponents. Winning requires a player to understand their opponents and to adopt a strategy that will counter their strengths while exploiting their weaknesses. Rule-based agents feature strongly in this, as do simulation-based agents such as MCTS. Poker has been extensively studied — see the review conducted by Rubin & Watson [13]; one of their notable finds was that a simulation-based approach is inferior to the formula-based approach, despite expectations. Whitehouse et al [14] looked into using MCTS for the card game Dou Di Zhu, which (like

Hanabi) also features imperfect information. Here, they apply determination and IS-MCTS to the problem and conclude while the IS-MCTS is superior in some cases, no overall difference was observed.

2) Hanabi AI: There has been a small amount of research into using artificial intelligence techniques to play Hanabi. Osawa [15] devised a number of rule-based agents for the 2-player version of the game, the mechanisms for which are described in Sections II-A2 and II-A3. Osawa found that the incorporation of consideration of the other agent's strategy and why they did what they did allows an agent to perform better than do the other non-cheating agents.

Cox [16] derives strategies for the game of Hanabi using the hat guessing game as inspiration. The agents all use an agreed encoding strategy to indicate what any particular Tell action specifically means, enabling them to co-operate so as to work around the limited view of their own hands. The encoding strategy does require the 5 player version of the game, however, as it won't work unless the hand size matches the number of other players in the game. We considered using this agent in the tests, as its unique strategy could have been the perfect test for agent modelling. However, there is an issue with the encoding strategy: every agent must know what every other agent has in their hands. This is cannot be used in agent modelling. If the Predictor IS-MCTS is agent 1, then it has access to the hands of agents 2, 3, 4 and 5 — which in Hanabi it does indeed have. Unfortunately, its internal copy of agent 2 needs access to the hands of agents 1, 3, 4 and 5. Agent 1 cannot give this information without breaking the rules of Hanabi. For this reason, we did not run tests with this agent. Van den Bergh et al [17] analyse Hanabi and define a number of rules for the game. The amount of time it would take to test every possible combination of these rules was too large, however, so they used an iterative approach to explore the search space intelligently. They note that some rules are far more effective than others, as

well as observing that a risk-taking rule does have some value. They found that the use of a Discard action when there is a possible hint is not optimal. In a follow-up paper [18], the authors present their best rule-based agent along with one using a Monte Carlo search.

### I. Artificial Intelligence

A number of the controllers used in this experiment were implemented as production rule agents. Many of these share individual rules, so each rule will be described here independently. All rules have additional pre-conditions that ensure they can only fire if it is legal to do so within the game rules (for example, a Discard action would necessitate a check that an information token was available). To avoid verbosity, we assume that the rules of Hanabi will be properly followed (so, for example, if the rule says to inform a player about a card, then the player will also be informed about other cards that satisfy the Tell's stated criterion).

- **PlaySafeCard:** Plays a card only if it is guaranteed that it is playable




Fig. 1: The four steps of Monte Carlo Tree Search [10]

• **OsawaDiscard:** Discards a card if it cannot be played at the end of the turn. This will discard cards that we know enough about to disqualify them from being playable. For example, a card with an unknown suit but a rank of 1 will not be playable if all the stacks have been started. This rule also considers cards that can not be played because their pre-requisite cards have already been discarded.



- **TellPlayableCard:** Tells the next player a random fact about any playable card in their hand.
  - **TellRandomly:** Tells the next player a random fact about any card in their hand.
  - **DiscardRandomly:** Randomly discards a card from the hand.
  - **TellPlayableCardOuter:** Tells the next player an un-known (to that player) fact about any playable card in their hand.
  - **TellUnknown:** Tells the next player an unknown fact about any card in their hand.
  - **PlayIfCertain:** Plays a card if we are certain about which card it is and that it is playable.
  - **DiscardOldestFirst:** Discards the card that has been held in the hand the longest amount of time.
  - **IfRule( $\lambda$ ) Then (Rule) Else (Rule):** Takes a Boolean  $\lambda$  expression and either one or two rules. The first rule will be used if the  $\lambda$  evaluates to true. If it is false, and a second rule was provided, then that will be used instead.
  - **PlayProbablySafeCard(Threshold  $\in [0, 1]$ ):** Plays the card that is the most likely to be playable if it is at least as probable as Threshold.
  - **DiscardProbablyUselessCard(Threshold  $\in [0, 1]$ ):** Discards the card that is most likely to be useless if it is at least as probable as Threshold.
  - **TellMostInformation(New?  $\in \{\text{True}, \text{False}\}$ ):** Tells whatever reveals the most information, whether this is the most information in total or the most new information.
  - **TellDispensable:** Tells the next player with an unknown dispensable card the information needed to correctly identify that the card is dispensable. This rule will only target cards that can be identified to the holder as dispensable with the addition of a single piece of information.
  - **TellAnyoneAboutUsefulCard:** Tells the next player with a useful card either the remaining unknown suit of the card or the rank of the card.
  - **TellAnyoneAboutUselessCard:** Tells the next player with a useless card either the remaining unknown suit of the card or the rank of the card.
- ^ -
- A. Agents
- 1) Legal Random: This agent makes a move at random from the set of legal actions available to it at any given time step.
    - PlaySafeCard
    - OsawaDiscard
    - TellPlayableCard
    - TellRandomly
    - DiscardRandomly
  - 2) Internal: This is a clone of the agent presented by Osawa that shares the same name. It features memory of the information it has been told about its own hand but does not remember information about what other players have been told. The rules used in order are:
    - PlaySafeCard
    - OsawaDiscard
    - TellPlayableCard
    - TellRandomly
    - DiscardRandomly
  - 3) Outer: This is a clone of the agent presented by Osawa with the same name. It features knowledge of what the other agents have been told already, to avoid repeating Tell actions. The rules used in order are:
    - PlaySafeCard
    - OsawaDiscard
    - TellPlayableCardOuter
    - TellUnknown
    - DiscardRandomly
  - 4) Cautious: This is an agent derived from human game-play. The agent plays cautiously, never losing a life. The rules used in order are:
    - PlayIfCertain
    - PlaySafeCard
    - TellAnyoneAboutUsefulCard
    - OsawaDiscard
    - DiscardRandomly
  - 5) IGGI: This agent is a modification of Cautious. The alteration to a deterministic Discard function greatly aids the predictability of this player. The rules used in order are:
    - PlayIfCertain
    - PlaySafeCard
    - TellAnyoneAboutUsefulCard
    - OsawaDiscard
    - DiscardOldestFirst
  - 6) Piers: This is an agent designed to use IfRules to improve the overall score. Otherwise, it is similar to IGGI. The rules used in order are:
    - IfRule (lives > 1  $\wedge$  deck.hasCardsLeft) Then (PlayProbablySafeCard(0.0))
    - PlaySafeCard
    - IfRule (lives > 1) Then (PlayProbablySafeCard(0.6))
    - TellAnyoneAboutUsefulCard
    - IfRule (information < 4) Then (TellDispensable)

- OsawaDiscard
- DiscardOldestFirst
- TellRandomly
- DiscardRandomly

The first IfRule is designed as a hail Mary in the end game: if there is nothing left to lose, try to gain a point. This derives from human play, when typically during the end game we make random plays if we know there is a playable card somewhere in our hand. This rule is more accurate, as it uses all the information it has gathered to calculate probabilities.

The second IfRule simply risks playing a card if there is a reasonable chance of its being safe.

The third IfRule is designed to try to provide more intelligent Tell conditions. If there is nothing useful to Tell and we are low on information, we set another agent up to be able to discard cards that are not needed. This means that the agents can burn through cards that are not helpful so as to try to obtain useful cards from the deck.

7) Flawed: This is an agent designed to be intelligent but with some flaws: it does not possess intelligent Tell rules, and has a risky Play rule as well. Understanding this agent is the key to playing well with it, because other agents can give it the information it needs to prevent it from playing poorly. The rules used in order are:

- PlaySafeCard
- PlayProbablySafeCard(0.25)
- TellRandomly
- OsawaDiscard
- DiscardOldestFirst
- DiscardRandomly

Giving information is the key to getting this agent to work intelligently. Without information, the intelligent rules can't fire, thereby leaving this agent to Tell randomly and Discard randomly — not a great strategy.

8) Van den Bergh Rule: This is the best rule-based agent from [18]. It was created by observing from human play that there are four main tasks:

- If I'm certain enough that a card is playable, Play it.
- If I'm certain enough that a card is useless, Discard it.

- Give a hint if possible.
- Discard a card.

Van den Bergh et al used a Genetic Algorithm (GA) to evolve the best options for each section, resulting in the following rules as an implementation:

- IfRule (lives > 1) Then (PlayProbablySafeCard(.6)) Else(PlaySafeCard)
- DiscardProbablyUselessCard(1.0)
- TellAnyoneAboutUsefulCard
- TellAnyoneAboutUselessCard
- TellMostInformation
- DiscardProbablyUselessCard(0.0)

9) MCS: This agent is a simple Monte Carlo Search (MCS) that uses a provided agent for the rollout phase. MCS is a technique that uses the Upper Confidence Bound (UCB) equation to select actions in a single step lookahead, with policy informed rollouts to evaluate those positions. It is essentially MCTS with a tree depth limit of one turn. In this paper, we name the agent MCS-[agent] to indicate which agent provided the rollout policy. For example, a MCS agent using IGGI as a policy would be named MCS-IGGI. The agent has a one-second time limit to return a move.

10) IS-MCTS: This agent uses a MCTS technique for handling games with partial observability as described in the paper by Cowling et al [19].

11) IS-MCTS is a modification to MCTS in which, on each iteration through the tree, the partially-observable game state is determined into a possible fully-observable state. This state remains consistent for the selection, expansion, rollout and BackPropagation phases before being replaced by a new determination. The implementation uses a time limit for returning moves of one second per move.

12) Predictor IS-MCTS: This agent was provided with a copy of each of the agents that it was paired with to use in its prediction. The predicted agents were initialised with random seeds: this corresponds to the predictor's having knowledge of each agent's overall strategy but no knowledge of its internal workings.

The Predictor IS-MCTS agent modifies the selection, expansion and rollout phases of MCTS

when considering nodes for other agents turns. The modifications remove Upper Confidence bound for Trees (UCT) for other agents' turns and replaces it with a query to the agent model to discover what that agent would do in that situation. The rollout phase is similarly modified. When making moves for its own turn, the predictor agent defaults to the legal random selection method used by IS-MCTS. The implementation maintains the one-second-per-move limit of IS-MCTS.

## II. METHOD

### A. Validation

In order first to validate our framework and AI implementations, we performed experiments using re-implementations of the Osawa and Van den Bergh agents. This involved recreating the experiments that they described in their papers and checking that we obtained similar results.

### B. Full Test

The set of agents under test contained a mix of current research on Hanabi as well as some rule-based agents of our own. There is also a mix of strong and poor agents for balance. We tested all the agents from this list:

- Legal Random
- Outer
- IGGI
- Piers
- Flawed
- Van den Bergh Rule
- MCS-Legal Random
- MCS-IGGI
- MCS-Flawed
- IS-MCTS
- Predictor IS-MCTS

In each experiment, one of the agents was selected from the list above and the remaining agents were selected as a group from the list below. For example, in the first experiment the Legal Random agent would be alone among four IGGI agents a concept we call pairing. The agents above were all paired in turn with:

- IGGI
- Internal
- Outer
- Legal Random

- Van den Bergh Rule
- Flawed
- Piers

200 random seeds were chosen, and for each seed every agent under test played two games with every agent with which it was paired. It did this for standard Hanabi rules with 2, 3, 4 and 5 players. Each agent under test played from a randomised position (first, second, third, fourth or fifth) determined by the seed. This ensured that each agent under test was in the same position for the same seed. Every agent therefore played  $200(nSeeds) \cdot 4(2, 3, 4 \text{ or } 5 \text{ Players}) \cdot 7(nAgentPaired) \cdot 2(reruns) = 11200$  games.

The configuration, final score and other basic state information were logged to a file upon completion of the game. The results were collated per agent and the mean score and number of turns taken were calculated. We also stored additional information about the final state of each game including the number of lives remaining and the information tokens remaining. When there are no lives remaining at the end of the game, this indicates that the players ran out of life tokens. The full (human readable) game traces for each game are also stored, for evaluating agent behaviour and the effectiveness of strategies.

Finally, the configuration and results of each game are processed to obtain the mean score, mean number of moves per game and the mean remaining life and information tokens.

## III. CONCLUSION

In conclusion, we found that agent modelling improves playing strength for tree search algorithms such as MCTS in the game of Hanabi. These results are consistent with the findings of [7].

## IV. FUTURE WORK

There is a lot of scope for future work in this area. Hanabi has some additional variants in its rules that focus on the addition of a multi-coloured suit of cards. This suit also contains 3 1's, 2 2's, 3's and 4's as well as a single 5. The different variants are:

- Variant 1** Add the multi-coloured suit as a sixth suit to the game. Maximum score is boosted to 30.



**Variation 2** Same as Variation 1, but only a single tile of each number from the multi-coloured suit is added to the game.

**Variation 3** The multi-coloured suit now functions as a wild card in Tell actions, and cannot be directly called out. For example, if Player 1 tells Player 2 (M, 2), (Y, 2), (B, 5), (B, 3) about

all the blues, then cards 1, 3 and 4 will be indicated. With this setup, the multi-coloured cards can only be identified by contradicting information given, requiring 3 pieces of information to fully identify one.

Variation 1 would be simple to implement and test, but was omitted from this paper as being too off-topic. Variation 2 adds a little extra strategy, but is very similar to Variation 1. Variation 3 would require some additional work to implement, as well as appropriate modifications to the AI agents.

The Predictor IS-MCTS has a number of limitations that we aim to address. The agent requires access to an accurate model of the co-operators in advance. It would be better if the agent could instead attempt to learn agent strategies based on observations in the game state. This would lead naturally to a more complicated agent that started with a more generic capability but was able to build models of its team members and update those models as games go on. Testing how much information is needed to learn enough to significantly improve the scores that a team achieve would then need to be done.

## V. REFERENCES

- [1]. P. R. Williams, D. Perez-Liebana, and S. M. Lucas, "Cooperative games with partial observability."
- [2]. J.-F. Baffier, M.-K. Chiu, Y. Diez, M. Korman, V. Mitsou, A. van Renssen, M. Roeloffzen, and Y. Uno, "Hanabi is np-complete, even for cheaters who look at their cards," arXiv preprint arXiv:1603.01911, 2016.
- [3]. C. Baral, G. Gelfond, T. C. Son, and E. Pontelli, "Using answer set programming to model multi-agent scenarios involving agents' knowledge about other's knowledge," in Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1, ser. AAMAS '10. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2010, pp. 259–266. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1838206.1838243>
- [4]. Torreno, E. Onaindia, and O. Sapena, "Fmap: a heuristic approach to cooperative multi-agent planning," in Proc. of DMAP Workshop of ICAPS, vol. 13, 2013, pp. 84–92.
- [5]. Geib, B. Craenen, and R. P. A. Petrick, "Generating collaborative behaviour through plan recognition and planning." [Online]. Available: <http://icaps16.icaps-conference.org/proceedings/dmap16.pdf#page=101>
- [6]. L. Panait and S. Luke, "Cooperative multi-agent learning: The state of the art," *Autonomous Agents and Multi-Agent Systems*, vol. 11, no. 3, pp. 387–434, 2005. [Online]. Available: <http://dx.doi.org/10.1007/s10458-005-2631-2>
- [7]. S. Barrett, P. Stone, and S. Kraus, "Empirical evaluation of ad hoc teamwork in the pursuit domain," in *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 567–574.
- [8]. H. De Weerd, R. Verbrugge, and B. Verheij, "How much does it help to know what she knows you know? an agent-based simulation study," *Artificial Intelligence*, vol. 199, pp. 67–92, 2013.
- [9]. C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. Cowling,
- [10]. P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, S. Colton et al., "A Survey of Monte Carlo Tree Search Methods,"

- Computational Intelligence and AI in Games, IEEE Transactions on, vol. 4, no. 1, pp. 1–43, 2012.
- [11]. G. M. J. Chaslot, M. H. Winands, H. J. V. D. HERIK, J. W. Uiterwijk, and B. Bouzy, “Progressive strategies for monte-carlo tree search,” *New Mathematics and Natural Computation*, vol. 4, no. 03, pp. 343–357, 2008.
- [12]. P. R. Williams, J. Walton-Rivers, D. Perez-Liebana, and S. M. Lucas, “Monte Carlo Tree Search Applied to Co-operative Problems,” in *CEEC2015 - IEEE Conference on Computer Science and Electronic Engineering*, ser. IEEE CEEC. IEEE Computer Society, September 2015, pp. 219–224.
- [13]. J. Walton-Rivers, “Controlling co-incidental non-player characters.”
- [14]. J. Rubin and I. Watson, “Computer poker: A review,” *Artificial Intelligence*, vol. 175, no. 5, pp. 958–987, 2011.
- [15]. D. Whitehouse, E. J. Powley, and P. I. Cowling, “Determinization and information set monte carlo tree search for the card game dou di zhu,” in *2011 IEEE Conference on Computational Intelligence and Games (CIG’11)*. IEEE, 2011, pp. 87–94.
- [16]. H. Osawa, “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*, 2015.
- [17]. C. Cox, J. De Silva, P. Deorsey, F. H. Kenter, T. Retter, and J. Tobin, “How to make the perfect fireworks display: Two strategies for hanabi,” *Mathematics Magazine*, vol. 88, no. 5, pp. 323–336, 2015.
- [18]. M. Van Den Bergh, F. S. MI, and W. Kusters, “Hanabi, a co-operative game of fireworks,” 2015.
- [19]. M. J. van den Bergh, W. A. Kusters, and F. M. Spieksma, “Aspects of the cooperative card game hanabi.”
- [20]. P. I. Cowling, E. J. Powley, and D. Whitehouse, “Information set monte carlo tree search,” *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 4, no. 2, pp. 120–143, 2012.