Modelling the game "The Mind" using Default Logic

Bachelor's Project Thesis

Jasmine O'Callaghan, s3303152, j.m.ocallaghan@student.rug.nl,
Supervisor: Dr. H.A. (Harmen) de Weerd

Abstract: Formal logic can be used to describe real-world problems in order to get an insight into human behavior. "The Mind" (Warsch, 2018) is a cooperative game with simple rules, but from this complex strategies arise. Players must cooperate to play numbered cards in ascending order, without knowledge of the cards the others are holding. The catch is, no communication is allowed during or before gameplay. Default logic allows assumptions to be made based on known rules, which will allow us to model players behavior when they have incomplete information about the situation. This paper aims to define an optimal strategy for players, allowing us to transform complex behavior into its simplest form.

1 Introduction

The game "The Mind" (Warsch, 2018) is a cooperative game with a deceptively simple task, work together to play numbered cards in the correct (ascending) order. In this paper will will consider a simplified game of The Mind played by only two players. They have a deck of cards, each card numbered 1 through 100. Each round the players draw one card and see only their own numbered card. Once the round is called to start, the only objective is for the players to play their cards in the correct order, lowest first. Players may play their card at any point during the round. The catch is, no communication, verbal or physical is allowed during the game and no discussion of strategy is allowed before or during play. This means the player must not only decide on a strategy for themselves, but try to predict the other player's strategy and adapt to it. Finding the optimal strategy is simple, the complexity lies in getting the players to coordinate while respecting the rule of no communication.

It becomes apparent very quickly that the strategies in this game are based on the concept of timekeeping. In simple terms, assuming the goal is to win, lower numbered cards should be played earlier in time than higher cards. However, humans perception of time is not perfect (Eagleman, 2008), so even if two players are playing the same strategy, for example playing card numbered x at t=x(seconds), they may still become mismatched and lose a round. In order to model this game using a logic system we will shun the human error by assuming the two players are perfect time keepers.

The Mind is a game with imperfect information, as the players do not know each others cards. Therefore, we need a logic system that can deal with some level of uncertainty and make assumptions. This model will be based off of Reiter's default logic as described in his 1980 paper (Reiter, 1980). Default logic allows for assumptions, referred to as "default assumptions" (Reiter, 1980), to be made when we have an incomplete set of information about the world. This will then allow us to model the assumption that the human players are making about the other's strategy and allow us to define strategies in a world with incomplete information. Specifically, it will allow players to make a default assumption about the strategy the other player is using, when there are multiple possibilities. We make an extension to Reiter's default logic by adding a time component which will count the rounds. This is because the information available to the player is time dependant, as more information about the other players strategy is revealed as they play further rounds.

This paper describes two models of players decision making when playing "The Mind". The original model uses purely default logic. It did not perform well in the sense of win-rate when playing using the rules of this model. However, it did give rise to some interesting behavior and served as a base for the second model. The second model combines default logic with some simple propositional rules. Players win more often using this model, however it is still not perfect. While neither managed to describe the "perfect" gameplay, they do give us an insight into imperfect gameplay.

2 Original Model Description

The first model we tried is based on Reiter's default logic (Reiter, 1980). We take the very simple case of a game with two players and two possible strategies. We will simply say that the game is won when the players play the same strategy in the same round. With this game it makes sense to add a temporal component, t, where $t \subseteq \mathbb{Z} | t > 1$. Here t does not directly describe a round, but instead when a card is discarded. Since with two players with one card each, two cards are played, the rounds each span two time events. Round one will start with $t = 0$ and will encompass $t = 1$ and $t = 2$. In round two card-playing events triggering the tick over to $t = 3$ and $t = 4$ will happen, and so on. Here each player holds a distinct model of the same structure. Their model initially contains the rules we will define here and in section 3.1, and a rule defining their initial strategy in the form $strategy(p_x, s_y, 0)$. This represents that they, player x, will play strategy y at $t = 0$. As card-playing events occur rules will be added to their models which can be applied to rules 2.1, 2.2, 3.1 and 3.2.

We can describe the reasoning of players as follows:

$$\frac{\neg strategy(other, 1, t) : strategy(other, 2, t))}{strategy(other(2, t))} \tag{2.1}$$

and the reverse

$$\frac{\neg strategy(other, 2, t) : strategy(other, 1, t)}{strategy(other(1, t))} \tag{2.2}$$

Or in natural language, if we know the other player is not playing strategy 1 at time t, then they must be playing strategy 2 and the reverse. For p_1, other would represent p_2 and vice versa. As we will see later players represent themselves as *me*. $strategy(other, 1, t)$ represents that the strategy of the other is 1 at time t. So for p_1, this would represent that p_2 is playing s_1 at time t.

3 Original Model Results

To evaluate this model there are two cases we must consider, depending on whether the round at time t was won or lost. Rounds are won when the players successfully play their cards in ascending order. If a card is played out of order, i.e. too early, the player with the lower card will immediately reveal that the round is lost and the round will end. At the end of the round, regardless of win or loss, the cards are reshuffled back into the deck and then players are dealt new cards to begin the next round.

3.1 Case 1: Win

In the previous section we described the reasoning of the players, but we must also describe their behavior or strategy. We use the information gained in the previous round, to predict what they should do in the next round. To do this we define the players as *sender* and *receiver*. The act of putting down a card reveals information about that players strategy to the other player. Thus, the one to play a card is the sender, as they are essentially sending a piece of information (their strategy) by playing a card. The receiver is receiving this information from the sender. For example, if p_1 plays a card at time t, then both players receive the information $sender(p_1, t)$, $receiver(p_2, t)$ and $strategy(p1, s_n, t)$. The strategy of the sender, here p_1, is already known to them though, so this is only new information to the receiver. Once the card is played, the receiver then knows the senders strategy and can adjust their strategy to coordinate with it. This can be described by the following rules:

$$\frac{sender(me, t) : strategy(me, x, t)}{strategy(me, x, t + 1)} \tag{3.1}$$

$$\frac{receiver(me, t) : strategy(sender, x, t)}{strategy(me, x, t + 1)} \tag{3.2}$$

This model is based on the important assumption that complete information is given to the receiver when a card is played. In other words, the act of the sender playing their card reveals their strategy to the receiver.

In this case, where the round is won, there is a simple sender/receiver dynamic, with the player putting down the card first being the sender and the other the receiver. Two card-playing events happen during the round. When the initial player puts down their card information is revealed and we have our sender/receiver dynamic. However, in the second event, no new information is revealed to either player, so we can simply ignore it. Then we can easily see with our rules that the receiver will stick with their strategy (which matches the strategy of the sender already) and all subsequent rounds will be won.

3.2 Case 2: Loss

In this case the model does not so easily describe the behavior. When one player plays their card too early, the other player reveals their own card to be lower and both players receive new information about the other's strategy. By revealing a loss, both players are now aware that they are playing a different strategy to one another. The first player to put down a card has revealed full information about their strategy, and the second player to discard the card reveals incomplete information, specifically that they are playing a strategy that is slower. However, in the case of only two strategies the second player is actually revealing full information. Here, we have two card-playing events that reveal information. In this case both players will at some point act as the receiver. For example, take the case where p_1 plays first at $time = t$, acting according to s_1 and p_2 is playing according to s_2. Initially p_1 will act as sender when they discard their card, and p_2 gains the information $strategy(p_1, t, s_1)$, so will chose to play s_1 at $t + 1$. As receiver, p_1 will not change strategy. Then, p_2 will discard their card revealing the loss. Even though this is not consistent with the strategy they are now playing, they have no choice but to discard their card. Now p_1 becomes the the sender and p_2 is the receiver. p_1 will receive the incorrect information that $strategy(p_2, t+1, s_2)$. So, p_1 will switch their strategy to s_2 for $t + 2$. p_2 is the sender so will continue to use s_1 in $t + 2$ Thus they will end up exchanging strategies and continue to play opposing strategies in the next round. Switching will continue to happen every round that there is a loss. Since we have described a win as only when the players have matching strategies, the players will get stuck in a loop where their strategies are always mismatched and thus they will lose every subsequent round after a loss.

One way around this is to always consider the first player to put down a card to be the sender and to ignore the second card-playing event in a round. In this case the receiver would switch strategies after the round, and the subsequent round would be won. However, this does seem like "communication", which is explicitly banned in the rules of the game.

3.3 Overall conclusions

While this model seems sensible at first, we can see that the behavior it models is not very complex. Once one round is won, all subsequent rounds will be won. If the first round is a loss, then similarly all subsequent rounds will be lost, unless we relax the rule of "communication". If we relax the communication rule, then still the behavior is quite boring, as the players play too perfectly, always winning the game on at most the second round.

While this model did not achieve the original goal of describing an optimal gameplay strategy, it gives us some good insights into making a better model. The players only succeed in aligning their strategies when they either win a round, or coordinate through communication. If we want to avoid communication, we need a better way of dealing with the case of both parties getting information during a loss.

While the switching behavior is far from optimal, it is interesting in the sense that human players may get stuck in these "switching" loops. However, human players would probably eventually adapt their behavior after recognising this pattern. Something our model is unable to do.

4 New Model Description

We will keep the sender and receiver dynamic in the new model, however we will consider the information the receiver gets to be incomplete. This means that the receiver will not know the sender's strategy just from that round, but will instead be able to reject strategies based on the information learned. For this model we will continue to consider the game played with two players only, each perfect time keepers. However we will now consider the possibility of more than two strategies.

Each player p_x has a list of n possible strategies S_{p_x}, with $S_{p_x} \subset (s_1, ...s_n)$. This list of strategies is the list of both the possible strategies they are using and that they think the other player may be using. In a simple case, both players have the same list of strategies, however in more complex games we may have the case that one player is playing a strategy unknown to the other. This would give rise to cases where perfect, always win, gameplay is unattainable.

When a card is played during a round information is given. This information can be considered in terms of formal logic rules. This rule can then be applied by the player to reject possible strategies from their list of possibilities. To give an example, when the *sender* puts down a card in the round, the other player, the *receiver*, gains information about their strategy. For example that their strategy s is faster than n (in seconds). They can then apply that rule to the list of their possible strategies and reject some.

Unlike the last model, we do not consider a win to only happen in the case that both players have matching strategies. In some cases a round may be won even if players have different strategies. We will refer to these as *complementary strategies*. Strategies that are complementary for certain combinations of cards, and win those rounds, will not necessarily win for every possible combination. In the case of matching strategies the players will always win every round, as we are still assuming they are perfect time keepers.

The players strategy is dependent on whether they were the sender or receiver. The sender will never change their strategy. We can describe the behavior of the sender as follows:

$$sender(me) \wedge strategy(me, s_x, t) \implies strategy(me, s_x, t+1) \quad (4.1)$$

The receiver's strategy is described by this rule:

$$reciever(me) \wedge strategy(other, s_x, t) \implies strategy(me, s_x, t+1) \quad (4.2)$$

The receiver will not necessarily change their strategy the next round, only in the case of a win. Here the senders strategy, s_x, is not actually known by the receiver. It is instead assumed. This assumption can be handled by the default logic:

$$\frac{\neg strategy(other, s_1, t), ..., \neg strategy(other, s_n, t) : strategy(other, s_{n+1}, t)}{strategy(other, s_{n+1}, t)} \quad (4.3)$$

Essentially, if the receiver knows the other is not playing strategies 1 to n, i.e. they have already been rejected, and strategy n+1 has not yet been rejected, then they assume that strategy n+1 is being played. They then apply this information to the previous rule. Thus, they will decide to play the strategy they assume the other is playing in the next round. The order of a players list of strategies is arbitrary, but does not change.

The rejection of strategies happens at the point of a card being played. This reveals information to the receiver and they can update their list of strategies accordingly.

The model is unforgiving, meaning that once a strategy has been rejected it will never be used again, as described by this rule:

$$\neg strategy(me, s_x, t) \implies \neg strategy(me, s_x, t+1) \quad (4.4)$$

5 New Model Results

At every round we will describe the behavior of the players in the case of a win or loss. Unlike the last model, we can see that it is now possible players will win a round after losing the previous round, and vice versa.

5.1 Case 1: Win with matching strategy

In the case of a win, the first player to put down the card is the sender and the other the receiver. The sender gains information, namely that their strategies are either the same, or similar enough to win together. This information is received by the sender at the point the other player puts down a card. From this information they may still be able to narrow down their list. However, as we can see from the rules described in the previous section, both players should keep the same strategy for the next round in the case of a win. This is because receiver will not gain the information $\neg strategy(s_x, t)$ where s_x is the strategy they are currently using. Therefore, the receiver will not reject their current strategy and will continue playing it. When a round is won, the players do not know if they have matching strategies or not, as they may have won with complementary strategies. In fact, the players will not know that they have matching strategies rather than complementary strategies unless every possible combination of cards is played and won. This is a limitation of this model. If they are playing the same strategy, all further rounds will be won.

5.2 Case 2: Win with complementary strategies

This case is the same as the previous, the players will stick to the same strategy, but the receiver will continue to update their list of strategies if they receive information from the sender.

Complementary strategies do not necessarily win for every possible combination of cards, but it is possible that they will. Two strategies may be similar enough that they indeed win for every possible combination of cards, but still distinct. For the purpose of this game I will refer to these strategies as *functionally the same*. The opposite we will call *totally mismatched*.

In the case of strategies that are functionally the same, we have the same outcome as with matching strategies. Both players will stick to their strategy and thus all following rounds will be won.

The case of complementary, but not functionally the same strategies, will eventually result in losing a subsequent round (assuming the players continue playing the game infinitely). After this subsequent round is lost they may eventually reach (functionally) the same strategies or they will continue to lose until all strategies are rejected. What happens after a lost round is described more in depth in section 5.3.

5.3 Case 3: Loss

In the case of a loss there is a switch in dynamic. Now the second player to put down a card is the sender and it is the first player who receives information. Specifically, the first to put down the card receives information that their strategy is faster than the other player when the second player discards their card and announces the round has been lost. The first player, acting as the receiver, can update their list of strategies. They gain at least the information that the strategy they are playing is not matching or complementary to the other players strategy and will reject that strategy from their list. They may also gain enough information to reject further strategies as well.

In this model we assume players rejection of strategies to be *unforgiving*, that is the players never re-add strategies that they have rejected back into their list of possible strategies. In this case one of three things will happen. In one case, by the receiver updating their list and switching strategy, the players will reach complementary, or the same, strategies and win a subsequent round. Another possibility is the players will not reach complementary, or matching, strategies in the next round. Thus, they will lose the subsequent round and we will loop back to what happens in the case of a loss. In the other case, we have the possibility that the players lose a round with only one strategy left on their list, so we know these strategies are totally mismatched. In this case all possible strategies will be rejected and there is no way for the players win any subsequent rounds. Further, once all strategies are rejected they cannot continue to play the game as with the unforgiving rejection of strategies there is no way for them to re-add a strategy to their list to continue play.

This may happen in the case that the players did have matching or functionally the same strategies on their two lists, but they were not playing them in the same round. However, it will not happen if at some point they were playing matching or functionally the same strategies. This is an unfortunate limitation of unforgiving rejection of strategies.

5.4 Overall conclusions

This model is a marked improvement on the previous one. By making sure that only one player changes strategy per round, we have eliminated the cyclic behavior. So, now players can transition from a win to a loss, or the reverse, between rounds.

However the model certainly isn't perfect and there are two large drawbacks to take into account. The first of these is the point at which players are playing the same or functionally the same strategies. While they have reached the ultimate goal of perfect gameplay and will win all subsequent rounds, they are not aware of this fact. This does not pose a problem for the validity of the model, however it does limit the transferability to humans. There is a finite combination of cards that can be drawn by the two players, so if they continue to play infinitely they can eventually be certain their strategies are (functionally) the same. However, this takes them playing every possible combination of cards, which is very inefficient as there are ${}_{100}C_2 = 4950$ combinations. While we used the assumption of infinite play in our model, when we are talking about real people, we can assume they are not going to continue playing for that long.

The second, and larger, drawback of this model comes from it's unforgiving rejection. As mentioned in section 5.2, it is possible that players may have matching or functionally the same strategies on their list, but still reach a point where they cannot continue the game as they have not played these strategies in the same round. So, one or both of the players has rejected the strategy needed to achieve perfect gameplay. To fix this a mechanism would need to be added to allow players to reintroduce strategies back into their list, making the rejection forgiving.

6 Conclusion

While both models put forward showed interesting behavior, they both had their limitations. As we saw, the first model gets stuck in cyclic behavior, with the players always winning or losing rounds depending on the outcome of the first round. The second model improves on this, but does not always reach perfect gameplay, even when that is achievable from the constraints of the players list of strategies. An improved model would always reach perfect gameplay when players have matching or functionally the same strategies on their lists. Further, it would do this in minimal rounds. Both models also had limitations arising from the assumptions made in order to limit the model. For example, assuming perfect time keeping makes the model less applicable to humans. However, assumptions always need to be made when creating models such as these and that is unavoidable.

The most obvious improvement to our second model is to add forgiving rejection of strategies. However, this is not a simple task. How you would choose to re-add strategies, so to minimise the chance of infinitely looping without ever managing to match up, would require experimentation to discover. This could be simple, like a first in, first out approach, simply re-adding the first strategy to be rejected. There are also more complicated possibilities, such as giving each strategy that is already rejected a "strike" every further round where it would have been rejected again. Then, re-adding the strategy with the least strikes.

However, even with forgiving rejection added to the second model, it still would not achieve perfect gameplay. To demonstrate this let us take two players, p_1 and p_2, with the lists, (s_1, s_2) and (s_2, s_3). In the case that s_1, s_2 and s_3 are totally mismatched and s_2 and s_3 are both "completely faster" strategies than s_1. By completely faster we mean that s_2 and s_3 would play the card 100 before s_1 would play 1. we have a situation where p_2 is always playing first and every round will be lost. However, p_2 would also always be the receiver, so would be the one to switch strategy every round. In this case, regardless of whether we have forgiving rejection and how the forgiving rejection is handled we have a case where p_1 will never switch strategy to s_2 and the players can never match strategies to win.

Overall, while the models were not perfect they did both give interesting insight into players behavior. It would be interesting to see how the addition of forgiving rejection could improve the second model. I do believe that, even if it was not reached in the constrains of this project, a model that achieves perfect gameplay if players have matching or functionally the same strategies within their lists is possible. Even though the models did not model perfect gameplay, this results are still interesting. Humans are not perfect and they do not always make optimal choices. So, while it was not the initial aim, from these models we can see some ways in which players fail. From this, conclusions can be drawn about the "faulty" logic humans are employing that causes these losses. We can see how some "meta-strategies" that humans could chose to use can cause them to be stuck in no-win situations.

References

D. M. Eagleman. Human time perception and its illusions. *Current Opinion in Neurobiology*, 18: 131–136, 2008.

R Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

Wolfgang Warsch. The Mind. Nürnberger-Spielkarten-Verlag, Germany, 2018. [Card game].