

I see what you see: Integrating eye tracking into Hanabi playing agents

Eva Tallula Gottwald

egottwald@mills.edu

Mills College
Oakland, CA, USA

Markus Eger and Chris Martens

meger@ncsu.edu, crmarten@ncsu.edu

Principles of Expressive Machines Lab
NC State University
Raleigh, NC, USA

Abstract

Humans' eye movements convey a lot of information about their intentions, often unconsciously. Intelligent agents that cooperate with humans in various domains can benefit from interpreting this information. This paper contains a preliminary look at how eye tracking could be useful for agents that play the cooperative card game Hanabi with human players. We outline several situations in which an AI agent can utilize gaze information, and present an outlook on how we plan to integrate this with reimplementations of contemporary Hanabi agents.

Introduction

Humans often give non-verbal cues to indicate their intentions (Land and Hayhoe 2001) or augment their verbal communication, often subconsciously. It would therefore be beneficial for the usability of computational systems to be able to interpret such signals. However, the subtle, subconscious use of signaling and lack of simple test domains make interpreting these signals very challenging. For many other AI techniques, games have served as a test environment, because they provide a low-risk, high-fidelity environment and often have a clear performance metric that can be used to measure success. We propose that games involving communication can be used as test environments for the interpretation of non-verbal cues given by humans.

One example for a game that relies heavily on inter-player communication is Hanabi (Bauza 2010), a cooperative card game in which players collaborate to build fireworks represented by cards with ranks from 1 to 5 in five colors. Unlike in traditional card games, players hold their cards facing away from them, i.e. every player sees every other players' card, but not their own. On a player's turn, they may give a hint to another player about the contents of that player's hand. These hints are limited to either telling the other player which of their cards have a particular color, or a particular rank. For example, player *A* may tell player *B* which of their cards are red and which are not, but not a subset thereof. Giving a hint expends a hint token, of which the players initially collectively have eight. Instead of giving a hint on their turn, players may also opt to play a card by choosing any card

from their hand and putting it on the table. If the played card is the next card in numerical order of its corresponding color, e.g. if a blue 4 was played and the highest blue card currently on the table is a 3, the card is added to the board, otherwise it is discarded and a mistake is noted. When there is no card of a particular color on the board, a 1 is considered to be the next card in numerical order. Players may also opt to outright discard a card instead of playing it; this recovers one hint token. After players play or discard a card they draw a new card from the deck to restock their hand. The game ends once the players collectively have made 3 mistakes, or when the deck has been exhausted, plus one extra round. The score of the players equals the number of cards on the board, for a maximum of 25 points if all five cards in each of the five colors have been played successfully.

Even though the game provides the players with very limited communication, when human players play the game, they typically follow the same strategy as in normal conversation, using Grice's maxims of communication (Grice 1975):

- The maxim of quantity by giving necessary hints, but not more
- The maxim of quality is enforced by the rules (players may not lie)
- The maxim of relation by not giving hints that are not relevant to the current state of the game
- The maxim of manner by trying to avoid hints that could be misinterpreted

However, when games are closely observed, players also often provide clues about their behavior in ways that are not strictly part of game play, such as hesitation, visibly deciding between two players to give hints to, etc. While there has been significant research into Hanabi game play, including how to build agents that play the game well with human players, the interpretation of non-verbal communication during game play has been understudied.

In this paper we present preliminary work that aims to integrate eye-tracking into agents that play Hanabi with human players. We have implemented a 2-player version of Hanabi in Unity that integrates with a Tobii eye tracker. We will present our hypotheses of how eye tracking information could be utilized by AI agents, the eye tracking information

we have available, and some initial observations about players' gaze behavior.

Related work

Hanabi has been of interest for several AI researchers because of its cooperative nature, the hidden information and limited communication channels. One approach to the game is to purely optimize the score the agents obtain. Cox et al. (2015) have devised a logical/mathematical protocol to convey a large amount of information using the limited communication Hanabi allows, scoring close to a perfect score in most games. While this approach only works in games with 5 players, Bouzy (2017) presents an improved version that also works with fewer players. Walton-Rivers et al. (2017) present a comparison of several different approaches, including several based on Monte Carlo Tree Search (Browne et al. 2012), focusing on how they perform in simulated games. However, while agents using the techniques discussed by these authors obtain very high scores when playing with each other, the protocols they use are very hard to follow for humans, and certainly not what a human player would intuitively expect.

Another approach for building Hanabi playing agents is more in line with how human players approach the game. Van den Bergh et al. (2015) present several agents using simple if-then rules defined by experts. Osawa (2015) describes agents that follow an expert-informed protocol, while also deducing how information obtained from the other players should be interpreted by having a model of possible interpretations. Note that Walton-Rivers et al. also included several rule-based approaches, including Osawa's and van der Bergh's, in their comparison, some of which did not perform much worse than their Monte Carlo Tree Search variants. Eger et al. (2017) specifically investigated how AI agents interact with human players, noting that agents that exhibit intentional behavior score higher when playing with a human player than those simply following their own protocol.

It has been noted that humans use the gaze of other people to determine their intentions, feelings, etc. starting from as young as 3-4 years (Baron-Cohen 1997). In order to make the interaction with computers more natural, it is therefore of great interest to research the integration of gaze into human-computer interaction (Poole and Ball 2006). Bader and Beyrer (2011) report how user's mental models change their gaze behavior to be more forward-looking to indicate their intentions as they become more familiar with a task. Hristova and Grinberg (2005) showed that players that are more likely to cooperate in an Iterated Prisoners' Dilemma scenario are also more likely to look at the payouts, while players less likely to cooperate looked more at the computer's moves. While this indicates that player behavior can be predicted from their gaze, the game under consideration was very simple. In the next section we will explore how eye-tracking could be used in a more complex domain.

Eye Tracking for Hanabi

In Hanabi, when receiving a hint from another player, it is essential to determine the intention behind that hint, in order



to interpret it correctly. The work cited above does that by either assuming that the player follows a fixed protocol, or by explicitly or implicitly enumerating all possible current game states given what is known about the hidden information, and determining in which game state a player would give the hint they gave. However, as mentioned, humans often indicate their intentions with their gaze. We therefore postulate that an AI agent with access to eye tracking information can perform better than one without.

Consider, for example, the case where a hint can be interpreted to indicate either a playable card, or a card that should be discarded. By using the player's gaze, the AI agent might be able to disambiguate between the two options. If the card should be played, it is more likely that the player looked at the board where that card should go, whereas a card that should be discarded might prompt the player to look at the other discarded cards more.

What we are interested in is going beyond this very basic example, and look at more complex cases. Hypothetically interesting scenarios include:

- The player's gaze going back and forth between two cards in the AI agent's hand before giving a hint including one of them. Depending on what the AI agent already knows about the two cards, they may be able to infer additional information. For example, since there is only one copy of each 5, players often give hints to prevent them from being discarded, especially when they think that the person holding the 5 is likely to discard it. However, this is at tension with giving hints that have a more immediate effect on game play. An AI agent could deduce this tension by observing which options a player is considering.
- Because players know how many copies of each card are in the deck, counting cards that were discarded, played or are in the other players' hands can be used to narrow down which cards are in a player's own hand. By tracking which cards a player looks at before performing a play or discard action, it is possible to determine which possibilities they are considering. For example, consider that a player knows that they have a 4, but not which color it is. When they look at all discarded 4s and a particular card in the AI agent's hand before playing their own 4, it is possible to deduce that the card in the AI agent's hand might also be a 4. In particular, if the color of the player's 4 was ambiguous, the AI agent might infer that their card is a 4 of a color that would help disambiguate the color of the player's 4.
- When the AI agent draws a new card, the duration of the gaze of the human player can be used to determine how immediately useful a card is likely to be. This is particularly true if the players are waiting to draw a specific card, such as a missing 1, or if a card that can only be played later in the game, such as a 4 is drawn early. We believe that players' gaze will linger shorter on cards that are not immediately useful. However, if a card is useful, the player has to scan the other cards in the AI agent's hand to determine which hint to give to unambiguously indicate the usefulness of the new card.

To be able to integrate these scenarios into an agent that



(a) A screenshot of our Hanabi implementation during game play (b) A heatmap of player gaze behavior overlaid over the game screen

Figure 1: Unity implementation of Hanabi with eye tracking

plays Hanabi, we need to be able to track the player’s gaze (to determine what they are focusing on), including which of multiple options it changes between (to determine decision making), and how long it rests in a particular spot (to determine interest/ disinterest in a particular option). Additionally, because of the inherent uncertainty of the information obtained the agent must not take this information as a fact, but rather only use it as guidance to help determine player intentions.

Existing Hanabi agents interpret hints that they are given by determining in which situation, or in service of which goal, the other player would give such a hint. If the agent determines that there are multiple applicable situations, it needs to break this tie in some way. The conservative approach would be to refrain from choosing any particular situation and continuing game play with the information obtained, as is done by Osawa’ (2015) *Outer State Player*. Alternatively, in the approach used by Eger et al. (2017), ambiguities are resolved by assuming that players prefer actionable hints that advance the game. Eye tracking data can be used in addition to these options to provide additional weight to each possible situation, without being the sole deciding factor. This is particularly appealing because it would allow our approach to be integrated in multiple existing and new agent designs.

Implementation

In order to test our hypotheses about how to integrate gaze into Hanabi agents, we implemented the 2 player version of Hanabi in Unity with support for a Tobii eye tracker¹. Using the eye tracker, we are able to determine where a player’s gaze lingers with reasonable precision to determine which card they are focusing on. Figure 1 shows the user interface of our implementation, as well as an example for where a player’s gaze lands on the screen. Note that this data comes from a rough development version which does not currently filter out any noise. We are still in the process of tweaking gaze duration thresholds, and considering additional tech-

niques to reduce the noise, starting with a simple low-pass filter. However, even with this noisy data, one can already see that players focus on particular cards more than others. Another, not entirely unexpected, observation we have made is that a player’s gaze is drawn towards UI elements that move or pop up, such as when they are given a hint, or when they click on a pop-up menu.

In our current version of the game, the AI agent performs its moves randomly, with our main focus being obtaining and interpreting eye tracking data. In the next section we will discuss how we plan on incorporating this information into the agent’s decisions.

Future Work

So far, our efforts have been focused on creating an implementation of Hanabi that incorporates eye tracking. For future work, we want to explore how players’ gaze behavior lines up with the situations outlined above, and test the hypothesis that player’s gaze can be used by an agent to improve its behavior. While we have already identified that players tend to look at certain UI elements when they become interactable, we have yet to determine how closely a player’s gaze correlates with their intentions, and in what way.

However, the main advantage of having eye tracking information available is not that it is necessarily an accurate indicator of a player’s intention, but rather that it can be used in addition to other techniques. We therefore plan to reimplement Osawa’s (2015) and Eger et al.’s (2017) agents, and use eye-tracking for cases in which their approaches have to decide between two or more possibilities. To validate this approach, we plan on performing a user study to compare the different approaches. Each participant will play several games with the same agent type, where the agents will ignore the eye tracking information in some games, while for others it is taken into account. We believe that taking gaze into account will allow the AI agents to perform better when playing with human player, but the games could also provide insights into how gaze differs between different players, if at all. Additionally, the score in the game is not the only rele-

¹<https://developer.tobii.com/tobii-unity-sdk/>

vant variable. We will therefore also perform a survey to ask participants if they perceived the AI to understand them better, or play more rationally.

One limitation of this approach, and an ethical experiment setup in general, is that the participants are necessarily aware of the eye tracker, and may adapt their behavior. Our experimental design therefore only compares games in which eye tracking information is present, but ignored by the AI agents, with games in which it is utilized. By comparing the scores from games in which gaze information is ignored with prior work, we can determine whether players also changed their in-game behavior, though.

Finally, we are also considering applications beyond games, such as assisting users of software tools. By determining user's intentions, help can be given in a more contextual way.

References

- Bader, T., and Beyerer, J. 2011. Influence of users mental model on natural gaze behavior during human-computer interaction. In *2nd Workshop on Eye Gaze in Intelligent Human Machine Interaction*, 25–32.
- Baron-Cohen, S. 1997. *Mindblindness: An essay on autism and theory of mind*. MIT press.
- Bauza, A. 2010. Hanabi.
- Bouzy, B. 2017. Playing hanabi near-optimally. In *Advances in Computer Games*, 51–62. Springer.
- Browne, C. B.; Powley, E.; Whitehouse, D.; Lucas, S. M.; Cowling, P. I.; Rohlfshagen, P.; Tavener, S.; Perez, D.; Samothrakis, S.; and Colton, S. 2012. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games* 4(1):1–43.
- 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.
- Eger, M.; Martens, C.; and Alfaro Córdoba, M. 2017. An intentional ai for hanabi. In *Computational Intelligence and Games (CIG), 2017 IEEE Conference on*, 68–75. IEEE.
- Grice, H. P. 1975. Logic and conversation. *Syntax and semantics* 41–58.
- Hristova, E., and Grinberg, M. 2005. Information acquisition in the iterated prisoners dilemma game: An eye-tracking study. In *Proceedings of the 27th annual conference of the cognitive science society*, 983–988. Lawrence Erlbaum Hillsdale, NJ.
- Land, M. F., and Hayhoe, M. 2001. In what ways do eye movements contribute to everyday activities? *Vision research* 41(25-26):3559–3565.
- 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*.
- Poole, A., and Ball, L. J. 2006. Eye tracking in hci and usability research. In *Encyclopedia of human computer interaction*. IGI Global. 211–219.
- van den Bergh, M.; Spijksma, F.; and Kusters, W. 2015. Hanabi, a co-operative game of fireworks. Bachelor's thesis, Universiteit Leiden.
- Walton-Rivers, J.; Williams, P. R.; Bartle, R.; Perez-Liebana, D.; and Lucas, S. M. 2017. Evaluating and modelling hanabi-playing agents. In *Evolutionary Computation (CEC), 2017 IEEE Congress on*, 1382–1389. IEEE.