Delft University of Technology, Bachelor Seminar of Computer Science and Engineering

The Influence of Information Sharing on the Predictability of the Human to an Agent

Stephen van der Kruk Carolina Jorge Ruben Verhagen
Dr. Myrthe Tielman
TU Delft
S.A.vanderKruk@student.tudelft.nl,
{C.Jorge, R.S.Verhagen, M.L.Tielman}@tudelft.nl

June 27, 2021

Abstract

Mutual predictability shows itself as a contributing factor to mutual trust and is known to improve the effectiveness in a human-agent teamwork setting. As team members communicate to coordinate the team through the task, the question arises as to what information the human should share to be predictable to an agent. To experiment with measuring the predictability, defined as to what extent the agent can anticipate the actions of the human, we used the Blocks World for Teams (BW4T) task. The two information types shared are intentions and world knowledge. The predictability is evaluated in the background by the agent logging, with automatic help from the human agent, a sequence containing the human performed actions. The agent can assign 2 probabilities to a human performed action. The first indicates with what probability the agent could say the human chose this action. The second indicates with what probability the agent could say this action led to this outcome. These probabilities then vary based on whether the information sent beforehand implied this action. The experiment has 4 different cases: sharing no information, sharing intentions, sharing world knowledge, and sharing both types. It is shown that sharing intentions contributes the most to higher predictability, due to these messages being the most effective at implying the future action of the human. World knowledge and sharing both types have less effect on predictability. We speculate that this is because of the larger amount of messages to share as well as when to share them, which overloads the human.

1 Introduction

With AI becoming increasingly more complex, the interaction between humans and AI becomes increasingly more collaborative. As humans and AI perform well on different fronts to each other (Bradshaw et al., 2012), having a beneficial collaboration helps to extract the most out of the specializations of both parties. To achieve this beneficial collaboration, mutual trust is needed. This mutual trust is identified by Salas et al. (2005) as one of three coordinating mechanisms for laying the ground for proper teamwork. In this research, we focus on mutual predictability, which is a contributing factor to mutual trust (Johnson and Bradshaw, 2021). Furthermore, the third presented challenge from Klein et al. (2004)

for effective human-agent teamwork is mutual predictability, which shows the relevance of mutual predictability as a research topic.

The second main topic of this research is the sharing of information in human-agent collaboration. This communication has to do with the two other coordinating mechanisms identified by Salas et al. (2005). First and foremost, there are shared mental models which are used to organize the shared knowledge of the team. This helps team members coordinate better by having shared knowledge about the environment and expectations of further performance of the team. Second, there is closed-loop communication which is to ensure the communication between two or more parties is performed well under all circumstances. This is important for keeping the teams' shared mental model continuously up to date. For this research, is of interest to know how these coordinating mechanisms contribute to predictability in human-agent collaboration.

In this paper, we focus on addressing the question: *What information does a human need to share to be predictable to an agent?* This question focuses on the communication from the human to the agent. To answer this question, it is important to define first what predictability is in our context. Therefore, a sub-question to ask is: *What does it mean to be predictable?* The next sub-question to ask is: *How does predictability correlate to trust?* This will give more definition to where predictability lies in its overarching topic of trust. In the case of Blocks World for Teams (BW4T) (Johnson et al., 2009), the framework for experimenting, it will be important to know what kind of information can be shared. Therefore, the next sub-question to ask is: *What information can a human share with an agent in BW4T?* The final sub-question builds on sub-questions 1 and 3 and asks: *How does information impact the human's predictability?* Experiments are performed in BW4T to answer this sub-question which will then run up to answering the main research question.

The hypothesis for the experiment which we are going to perform in BW4T will be: *Sharing human intentions will give higher predictability of the human compared to no shared information.* The main idea behind this hypothesis is that messaging through intentions in BW4T contributes more knowledge about the human's future actions. Section 2 will cover the explanation of used information types (intentions and world knowledge) and how they contribute to predictability.

The main contributions of this research to the topic of human-agent collaboration and the overarching topic of trust will be a set of information types that the human can share to contribute to its predictability towards an agent. Furthermore, it is shown to what extent an information type affects predictability. Finally, it is discussed how predictability correlates to the overarching topic of trust.

As for the structure of the paper: section 2 will discuss the background by discussing relevant work related to this research. Section 3 will discuss the concepts behind measuring predictability, followed up by section 4 which discusses the implementation of these concepts in BW4T. Then after that, section 5 will showcase the experiment set up, followed up by the results. Section 6 reflects on the ethical responsibility, integrity and reproducibility of the research. This is followed up by section 7 which discusses the results and puts them in the broader context. Finally, section 8 will conclude this research.

2 Background

In related work about predictability in (human-agent) teams, predictability is defined in two parts. To be predictable to someone means that they know about your course of action/behaviour and that they can plan themselves around this. According to Johnson et al. (2014): "predictability means one's actions should be predictable enough that others can reasonably rely on them when considering their own actions" (p.52). Ahrndt et al. (2016) also calls predictability a twofold problem where the first part mentions anticipating one's actions and the second part mentions to integrate them in the planning process. Mentioned here is that the second part has to do with human-aware planning. The team member is modelled using a Markov-Decision Process (MDP) and the goal of predictability is to gain information about the policy of the team member. In terms of the MDP, this is then to find out about what the transition probability is for a certain state given an action. For this research, a similar definition of predictability is used. For the human to be predictable from the perspective of the agent, the agent should be able to anticipate the actions of the human, and the agent should integrate this when considering its next course of action. We consider this to be the answer to the first sub-question.

When connecting predictability to trust, interdependence relationships are the main contributors to mutual trust (Johnson and Bradshaw, 2021). Mutual predictability is one of such an interdependence relationship and therefore contributes to mutual trust. It is according to Johnson et al. self-explanatory that being able to make others anticipate what actions you are performing, results in a higher trust. An interesting observation made by Lewis et al. (2018) is how trust in the early phase of the relationship is based on predictability. Which shows how predictability can have varying contributions to trust.

Related experimentation in multi-agent collaboration in BW4T has shown experimentation with two types of information (Jonker et al., 2012; Li et al., 2016; Singh et al., 2017). These two types are intentions (or goals) and world knowledge (or beliefs). The experimentation focused on testing this information on the team performance or on testing this in joint actions (Li et al., 2016; Singh et al., 2017). Sharing intentions has overall shown to be the most effective in increasing team performance. In terms of how these types of information relate to anticipating actions, Li et al. (2016) mentions that intentions are indicative of action, therefore sharing intentions gives knowledge about future actions. For answering the sub-question about what information an agent can share in BW4T, intentions and world knowledge seem the most appropriate categorization of information types to use for experimentation in BW4T.

Previous research in BW4T has done experimentation with the aforementioned information types and has shown the impact on the team effectiveness in either agent-agent or human-agent teams. However, there has not been a closer look at how these information types exactly influence the knowledge of the agent about the future actions of the human. Therefore, this research will focus on how intentions and world knowledge contribute to, as our predictability definition states, the agents' ability to anticipate the actions of the human. This is a more concrete definition of the fourth sub-question, with which we will construct the experiment.

3 Measuring Predictability during Information Sharing

3.1 Measuring predictability with an MDP-like graph

The following idea for measuring predictability was created whilst working on the implementation of BW4T in MATRX (TNO, 2021a). This idea is therefore influenced by the constraints faced during implementation. However, the idea itself is discussed in a general manner such that it can be applied to different use cases.

Ahrndt et al. (2016) mentions that an agent being predictable means that other agents can build up experience from that agent and predict its future actions. Therefore, they labelled it as a reinforcement learning problem, using Markov-Decision Processes (MDPs) to formalize the states and actions of the agent. Considering the duration of this research, making a reinforcement learning agent was not feasible as MATRX does not support this out of the box (TNO, 2021b). Furthermore, the BW4T task was down-scaled in size and duration because of the short duration of the research. This did not allow much room for a reinforcement learning problem. Performing reinforcement learning over multiple games was also not an option since the experiment setup had multiple participants, therefore a learning agent could not learn one type of behaviour.

The default use case for an MDP problem is to find the most cost-effective or profitable policy given a graph where actions have a probabilistic outcome of ending up in a certain state (this probability we call the 'action outcome probability'). This policy is used to decide which action to choose when the decision-maker finds itself in a certain state. Using an MDP to model the agent you are teaming with bodes a way to express the predictability of your teammate, as it can model its policy. Having this policy be more accurate of its actual actions helps to plan your actions around that. This accuracy can be achieved by knowing how probable the action is to result in a particular state.

Taking the above into account, we introduce an MDP-like graph structure to measure the predictability of a teammate who is sharing information. This MDP-like graph is extended from a default MDP in two ways. First, an extra probability is attached to the choice of the teammate performing this specific action in this specific state. This likeliness of the teammate performing this action we call the 'action choice probability'. Second, the 'action choice probability' or 'action outcome probability' are allowed to change. Both probabilities can be put to one when the teammate shares information that implies the choice of action and the outcome of that action. Therefore, knowing for sure that this action is being performed allows for a better knowledge of the policy of the teammate. When the task is finished, the result is a sequence of states and actions through this MDP-like graph. The transitions logged in the sequence contain the 'action choice probability' and 'action outcome probability' of every performed action by the teammate. With this action sequence, measures of predictability can be generated by aggregating the probabilities.

3.2 Team effectiveness

The definition of predictability used for this research also includes an agent integrating the knowledge about one's future actions into its own course of action. This we can connect to showcasing 'backup behaviour' and 'adaptability' which Salas et al. (2005) defined as the two factors from the 'big five' of teamwork that directly influence the team effectiveness. Therefore, measuring team effectiveness can give more insight into this aspect of the predictability of one agent to the rest of the team. Note that this is a rough measure as more factors come into play that affects the team effectiveness. For our experiment, such a factor is the backup behaviour and adaptability of the human. To try to make this factor more constant, the autonomous agent shares the same types of information for every task run to level out the differences in backup behaviour and adaptability of the humans. This makes the team effectiveness more dependent on the agent that plans around the knowledge it has of the human's future actions, while it is by default dependent on both. However, it is still expected that it will be partly dependent on both, making it likely that these findings from this measure will be speculative.

4 BW4T Implementation

4.1 Blocks World for Teams

Blocks World for Teams (BW4T) (Johnson et al., 2009) is a testbed for joint activity in which human or autonomous agents, collaborate to deliver coloured blocks in the right order (bottom to top) to a drop zone at the bottom of the map. The blocks are hidden in rooms which agents can search. An agent has a limited vision radius of one tile and can perform different kinds of actions such as walking around the environment and picking up and dropping off blocks. The collaboration comes from that a team of at least two agents, need to communicate and coordinate via messages. This communication allows for experimentation between agents. The BW4T implementation used for this experiment is within the MATRX framework (TNO, 2021a), which allows the use of a custom front end. We use this to teleoperate an agent via buttons for the experiment. This is to allow the human to focus more on completing the task and sharing information rather than typing messages or walking with the keyboard.

As for the experiment BW4T configuration, two agents will be created. First, there is the autonomous PredictorAgent. Since we are interested in the predictability of the human from the perspective of the agent, the PredictorAgent keeps track of the human sequence of actions explained in section 3.1. Furthermore, it plans its actions around the information the human shares through heuristics. Second, there is the TeleOpHumanAgent (i.e. the human) which is an agent teleoperated by the human to make it perform actions or send messages. The world layout can be seen in figure 1. There are a total of six rooms that can be visited and four blocks that need to be dropped off at the drop zone. For the experiment, the agents are moving slowly to give the participant time to look at the map and the chat and to think about which action it is going to take. Therefore, this number of rooms keeps the duration of searching rooms acceptable. Having four target blocks gives enough opportunities for teamwork between the human and agent but keeps the amount of data to be generated with the human action sequence to an acceptable length per task.

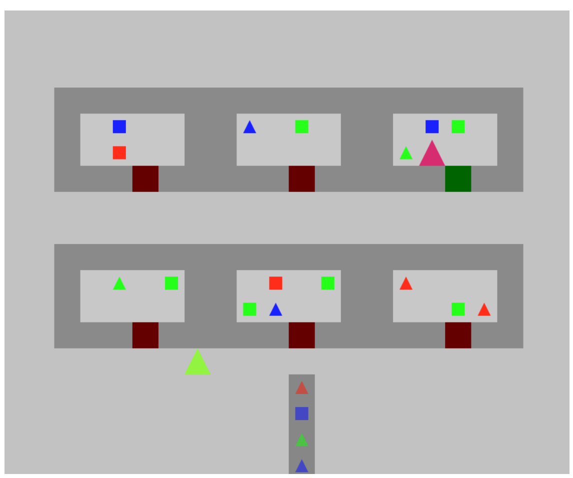

Figure 1: The 'God' view of the BW4T world used for the experiment. The human (pink triangle) is in upper-right room. The PredictorAgent (green triangle) is below the rooms. The agents only see the blocks in a one tile radius around them.

4.2 Modelling BW4T into states and actions

To use the human action sequence for measuring predictability as explained in section 3.1, the BW4T task is modelled into states and actions. First, there are four actions, each consisting of multiple smaller MATRX actions (e.g. walking a step into a direction, dropping a held block) which are performed in one go. These actions are:

- **a0 - Walking to room X opening the door (open door)**: This action makes the agent walk to a specified room and open the door. This action is only possible if the room is not traversed already by the human or the agent.
- **a1 - Traversing the room and checking for blocks (traverse room)**: This action makes the agent walk through the room, checking for every tile whether any target block from the drop zone is located there, if that is the case, it is added to its memory. This action is only possible if the agent performed an 'open door' action beforehand.
- **a2 - Walking to a found block and grabbing block (grab block)**: This action makes the agent walk to the location of a found target block and grab that block. (The location of the target block is gotten from either traversing rooms or information shared by the other agent) This action is only possible if the target block which the agent is currently targeting is found.
- **a3 - Walking to the drop zone and delivering the held block (drop block)**: This action makes the agent walk to the appropriate location on the drop zone and drops the held block. This action is only possible when it performed the 'grab block' action beforehand. Since the blocks need to be delivered in order the agent will wait with dropping if the block below is not delivered yet. The PredictorAgent will always focus on the most downward undelivered block unless it plans around the actions of the human. This is to ensure that the blocks are placed in order.

These actions are grouped in this way for multiple reasons. First, it is to simplify the task more for the human, as these actions can be performed by pressing a button. (See the red buttons in figure 3) It would most likely lead to a task overload for the human if it had to focus on walking to exact locations, opening doors etc. on top of interacting with the PredictorAgent and thinking of a strategy. Second, the points where these actions end allow for the most part for a point of decision to be made for the next action. For example, if the 'open door' action and 'traverse room' action were merged, there would be fewer transitions where decisions could have been made, leading to less data being generated. Four actions lead to the right amount of data generation whilst keeping the human task load manageable. Third, this number of actions keeps measuring the predictability of the human to a manageable level. Splitting up these actions creates more states where these actions should lead to, which the system that measures the human predictability should account for. This was rather infeasible to implement considering the duration of this research.

Second, the states the agents can find themselves in are based on the outcome of the actions. They are as follows:

- **s0 - Initial**: The starting state. In this state the agent is idling at the left top corner of the map.
- **s1 to s6 - Opened door room 1-6**: In these 6 states the agent is standing in front of the corresponding room with the door opened. From here it can perform the 'traverse room' action, or the 'grab block' action if the targeted block was found.
- **s7 to s12 - Traversed room 1-6**: In these 6 states the agent has traversed the corresponding room. From here it can perform the 'open door' action to another room, or the

'grab block' action if the targeted block was found.

- **s13 to s16 - Grabbed block 1-4**: In these 4 states the agent has grabbed their corresponding target block. From here they can perform the 'drop block' action.
- **s17 to s20 - Dropped block 1-4**: In these 4 states the agent has dropped their corresponding target block. From here it can perform the 'traverse room' action, or the 'grab block' action if the targeted block was found. Dropped block 4 (s20) is the final state of the game. If one agent has reached it the game ends since the blocks are delivered in order.

Figure 2 showcases the state transitions in the 'building blocks' of the MDP-like graph which the PredictorAgent uses to construct the human action sequence. These transitions are shown with the 'action choice probabilities' and 'action outcome probabilities' when the human has not shared a message. The probabilities with a bracket are only valid when working with a world layout of 6 rooms and 4 target blocks.

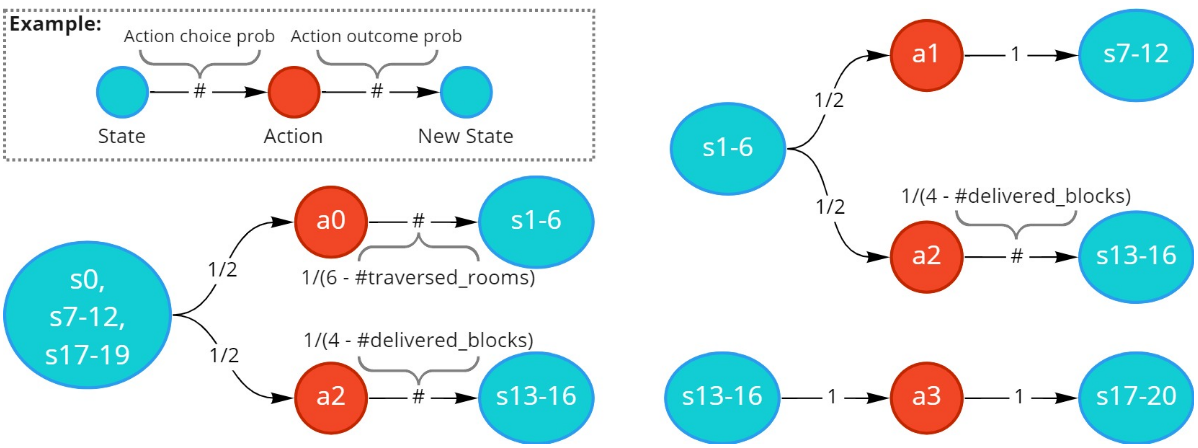

Figure 2: The 'building blocks' of the MDP-like graph. Note: multiple states are showcased in one such building block, which is a shortcut for writing all possible transitions.

4.3 Message Sharing

To test the hypothesis of whether sharing intentions gives higher predictability of the human compared to sharing world knowledge, the experiment requires participants to be divided into four cases. The following listing explains per case what messages are shared there.

- **Case One**: No information sharing. The data measured serves as a baseline comparison.
- **Case Two**: Sharing intentions. These messages need to be shared before the actions are performed. These are the purple buttons in figure 3. The messages shared are as follows:
  - **Intent of Targeting block x**: This message implies the outcome of the 'grab action' (a2) since the PredictorAgent knows which block is being grabbed. Therefore, this outcome probability is set to 1 and the PredictorAgent will as a heuristic target another block if it was targeting block x or skip it when selecting a new block. An important note here is with the way the PredictorAgent plans around the actions of the human, when the human has set a block to target, (done by the blue buttons in figure 3) it should fulfil the goal of delivering that block before another block can be targeted.

– **Intent of Traversing room x**: This message implies that the 'open door' action is performed followed up by the 'traverse room' action. Then also the outcome of the 'open door' action is known since it knows which door is being opened. Therefore, these probabilities are all set to 1 and the PredictorAgent will as a heuristic skip searching room x.

– **Intent of Delivering found block x**: This message implies that the 'grab block' action is performed followed up by the 'drop block' action. Then also the outcome of the 'grab block' action is known since it knows which block is being grabbed. Therefore, these probabilities are all set to 1 and if the PredictorAgent did not know yet that this block is targeted, it will as a heuristic target another block if it was targeting block x or skip it when selecting a new block.

- **Case Three**: Sharing world knowledge. These messages need to be shared when the actions are being performed. When shared, it puts the probabilities of the implied action to 1. These are the yellow buttons in figure 3. The messages shared are as follows:

  – **Door x opened**: This message implies that the 'open door' action is performed. The PredictorAgent will as a heuristic put room x last in line for searching with the assumption that this action will be followed up by a 'traverse room' action.

  – **Room x traversed**: This message implies that the 'traverse room' action is performed. The PredictorAgent will as a heuristic skip searching room x.

  – **Found block(s) [*list*]**: This message implies nothing in terms of which action is being performed currently since it can be shared at any moment. If this list contains the targeted block of the PredictorAgent which it has not found, the PredictorAgent will as a heuristic deliver that block.

  – **Block x grabbed**: This message implies that the 'grab block' action is performed. The PredictorAgent will as a heuristic skip searching for block x.

  – **Block x dropped**: This message implies that the 'drop block' action is performed. The PredictorAgent will as a heuristic skip searching for block x.

- **Case Four**: Sharing both intentions and world knowledge.

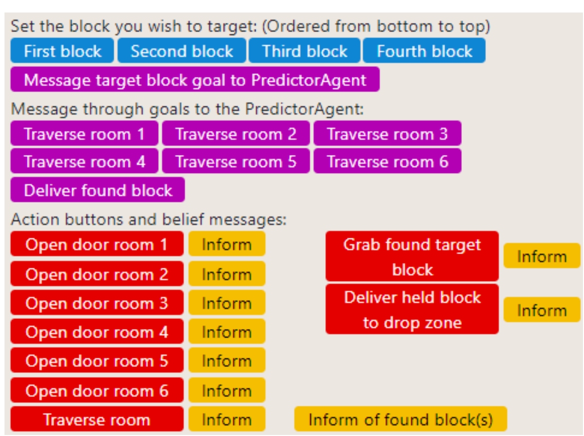

Figure 3: The interface used by the human. Blue buttons for targeting blocks, red buttons for actions, purple buttons for messaging intentions, and yellow buttons for messaging world knowledge.

4.4 Generating the human action sequence

The human action sequence is generated by the PredictorAgent. However, since the PredictorAgent might not know what action the human performed, the actions are evaluated with help from the human agent. The human agent automatically shares an evaluation message containing the previous state, the performed action, and the current state of itself. This is done after the human agent has finished that action and is not visible to the human performing the task. The PredictorAgent uses this information solely for measuring human predictability and does not use it for planning its actions around that. Through chaining the states and actions in these messages, the human action sequence is generated. Depending on whether it received a message which implied something about the action it received in the evaluation message, it either logs the 'action choice probability' and 'action outcome probability' of this message to 1 or applies the probabilities as shown in Figure 2. Multiple measures for the human predictability can then be taken from this entire sequence. These are shown in section 5.2.

5 Experiment

5.1 Method

Design. The experiment has a between-subjects design. Participants are allotted to one of four cases. These are in order: no information sharing, sharing intentions, sharing world knowledge, and sharing both intentions and world knowledge. The case a participant is allotted to is rotated through this order to prevent the participant from being allotted a specific case.

Variables. The independent variable for this experiment is the type of messages being shared in the task run. The dependent variables are the probabilities in the state transitions of the human action sequence and the achieved team effectiveness in the task run. There are confounding variables. First, there is human error which can be messages of the wrong type being shared, clicking buttons multiple times, etc. This is mitigated by having constraints when pressing buttons, blocking the action or message when it cannot be performed. However, it might be that not every constraint was found and implemented. It is also a point of discussion whether putting in too many constraints can also influence the decision making of the human, which we do not want to compromise. Second, there is familiarity with BW4T and MATRX, as the course Collaborative Artificial Intelligence (CAI) from the TU Delft Computer Science and Engineering bachelor included a practical assignment in BW4T. It is expected that several participants will have done this course before participating in this experiment. However, the implementation of this task and its purpose are different. The participants are given an explanation, which for the most part includes material that has not been covered in the practical assignment. Everyone performs the same number of tasks, which levels the playing field.

Procedure. The experiment is performed in the BW4T environment explained in section 4. First, the participant is shown the consent form online for their consent for participating in this experiment. Second, the participant is read a script (Appendix A) that explains the BW4T task to the extent of selecting target blocks with the blue buttons and performing actions with the red buttons. Then the participant performs a 'dummy' task in a fixed practice world to get accustomed to the BW4T task environment. Third, when the participant has finished the practice task, the world is then switched to another fixed layout. See Figure 1 for the 'God' view of this layout. The participant is not allowed to see the 'God' view for the actual experiment task run, instead only the agent view is shown. Now, the message buttons corresponding to the case the participant is allotted to are explained. Fourth, the participant performs the task run which generates the data used for analysis. Finally, when the participant has finished the task, any remaining questions about the experiment or research are answered.

Participants. The experiment had 20 people participating with an average age of 21.1 with a standard deviation of 0.718. Every case got 5 participants assigned.

5.2 Results

First, we have the 'mean action choice probability' and 'mean action outcome probability'. The former indicates on average how well the PredictorAgent knew the choice of the next action of the human. The latter indicated on average how well the PredictorAgent knew about the outcome of the performed action of the human. Per task, the mean of these probabilities was calculated. The results in figure 4 are the average of 5 mean values per case with the standard deviation.

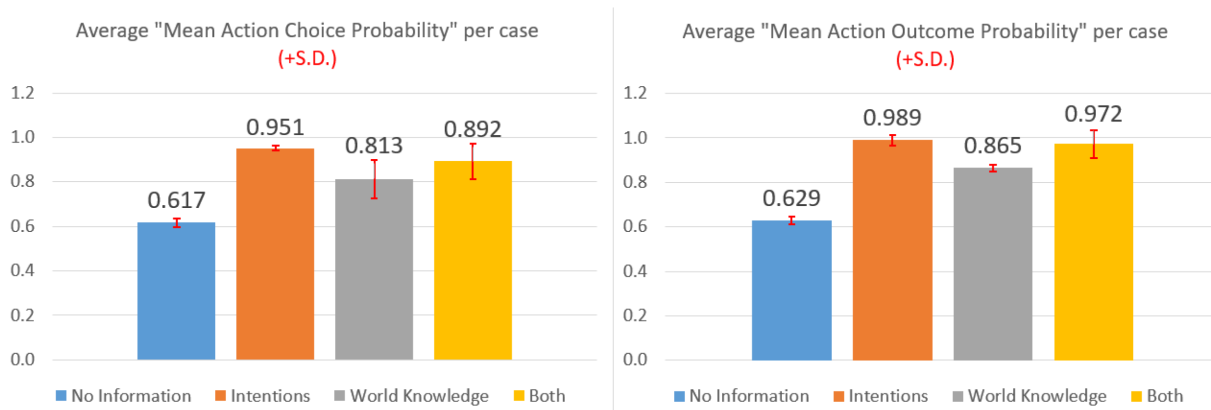

Figure 4: The average 'mean action choice probability' and 'mean action outcome probability' and their standard deviations per case.

All performed ANOVAs and t-tests (unpaired, assuming unequal variance) are taken at the $p<0.05$ level. The average 'mean action choice probability' gave significant differences from the one-way ANOVA [$F(3,16)=30.5$, $p=7.41E-7$]. The average 'mean action outcome probability' gave significant differences as well [$F(3,16)=109.6$, $p=7.01E-11$]. From the baseline case with no information sharing, a lower average mean probability is displayed. Intentions show the highest average, with world knowledge showing the lowest average of the 3 cases where information was shared. In both measures, the difference was significant between intentions and world knowledge. (unpaired t-test: $p=0.023$ for action choice and $p=2.95E-5$ for action outcome) In both graphs, the case where both information types were shared had a lower average than the case where only intentions were shared. However, this had no statistical significance. (unpaired t-test: $p=0.17$ for action choice and $p=0.60$ for action outcome)

Second, we have the probability of the entire human action sequence from the perspective of the PredictorAgent. This number is the product of all the 'action choice probabilities' and 'action outcome probabilities' of the logged human action sequence. This indicates the probability from the perspective of the PredictorAgent that this entire human action sequence was taken. Looking at this from the opposite direction, we have the number of possible sequences the human could have taken from the perspective of the PredictorAgent. This is the product of all the possible action choices and possible action outcomes in the human sequence. These measures shrink or grow exponentially and therefore the differences between the cases will be large. But simply put for both measures, the closer the number is to 1, the more likely it is that the actual human sequence was taken in the eyes of the PredictorAgent. Since this measure is dependent on the length of the sequence, the points of data can lay orders of magnitude away from each other, resulting in a high standard deviation. For these reasons, the results are shown in a table. Table 1 shows the average of the 5 values per case with the standard deviation.

	Avg sequence probability	SD	Avg possible sequences	SD
No information	8.35E-08	1.13E-07	209741414.4	312316587.4
Intentions	0.433	0.149	2.8	1.789
World Knowledge	0.00612	0.00841	547.2	618.064
Both	0.283	0.211	8	9.274

Table 1: The average probability of the human action sequence per case and the average number of possible sequences the human could have taken per case.

One-way ANOVA shows a significant difference for the average sequence probability across the cases[F(3,16)=13.6, p=0.0001]. No significant difference was given for the average possible sequences due to high variances [F(3,16)=2.26, p=0.12]. As with the average mean probabilities, the predictability per case was similar. Sharing no information has the worst performance, sharing intentions scores the closest to 1, and world knowledge had a significantly worse performance than intentions in the average sequence probability (unpaired t-test: p=0.0030) but this comparison had no significance due to the high variance in the average possible sequences. (unpaired t-test: p=0.12) Sharing both information types performed on average worse than only sharing intentions, but had no statistical significance in both measures (unpaired t-test: p=0.23 for average sequence probability and p=0.28 for average possible sequences)

Third, figure 5 shows the team effectiveness with the average number of ticks and average total moves made. (PredictorAgent + human) The data from the total moves is pre-processed as the slowdowns from the agents caused a single action to be logged multiple times. Finally, in the same figure, the average number of human messages shared per case is displayed. One-way ANOVA gives no significant difference to the average ticks [F(3,16)=0.16, p=0.92] and the average total moves [F(3,16)=3.14, p=0.054]. For the total moves, there is a significant difference considering the no information case with the intentions case and sharing both types case. (unpaired t-test: p=0.02 for no information and intentions, p=0.03 for no information and both types) As for the average human message sharing, there was a significant difference shown by the one-way ANOVA [F(3,16)=11.15, p=0.0003]. However, there is no significant difference when comparing the information sharing cases. (unpaired t-test: p=0.13 for intentions and world knowledge, p=0.091 for intentions and both types, p=0.37 for world knowledge and both types)

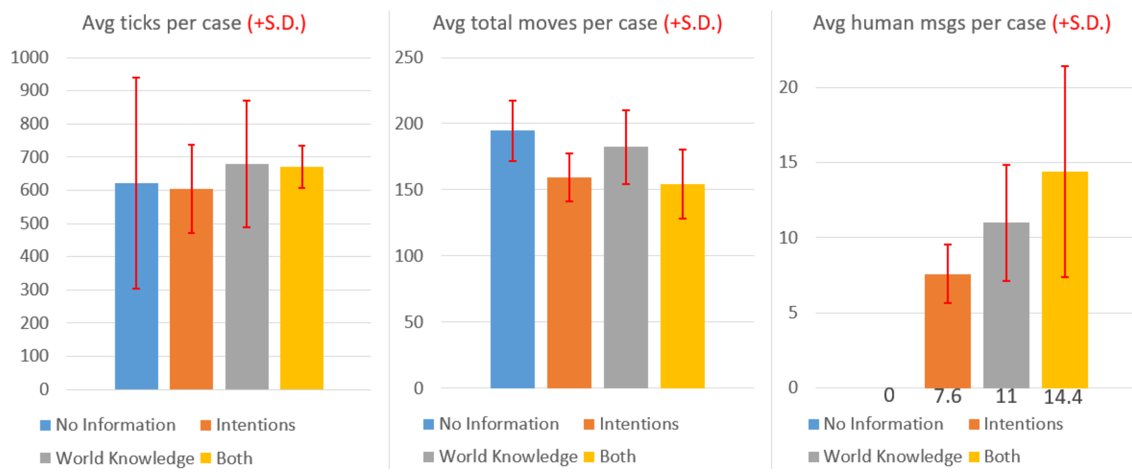

Figure 5: The average total ticks, total moves, and number of human messages

6 Responsible Research

As a large part of this research was performing experiments with humans, multiple things had to be considered to ensure that the experiment was performed ethically responsible. This research got TU Delft HREC confirmation which is a screening to confirm that this research is ethically sound. During the experimentation phase of the research, the participant is shown beforehand an informed consent form which informed the participant of the aim of this research and what will be done with their age and anonymous generated output as data. Furthermore, the consent form confirms whether the aim of the research was understood, the possibility of withdrawal from the study at any time, and whether they understood the use of their data for the research. During the experiment, the participant is allowed to ask further questions about the aim of the experiment as well as the future analysis of the data. This is to fully inform the participant about the research process to allow them to have a well informed consent.

To ensure the integrity of the research, attention is given to how data was collected as well as how it was analysed. As experimentation within the BW4T task has a learning factor to it, it is of importance to create a level field across the participants. All participants are given the same explanation about the task and perform a practice run regardless of their previous experience with BW4T. It was not considered influential that some participants had experience in BW4T from CAI as this BW4T task has a different purpose and interfacing for which those participants also needed explanation. However, performing multiple runs for data gathering after the practice run is not allowed. If the second run fails to generate the data. Then the participant's data cannot be used since they have learned too much.

For analysis and interpretation of the results, statements are only made about comparisons between cases where there was a statistically significant difference. Other discussed findings which showed an insignificant difference were left as speculation. Therefore, a clear indication is made to what is concretely found in this research and what could be further explored to confirm the speculation.

For reproducibility of the research, it is of importance that the information and tools from this research can reproduce the same findings when performed by others. For that, the research provides an explanation of the concepts used in the experiment as well as the implementation details of the BW4T task in MATRX used for the experiment. Furthermore, details about the experimental setup are discussed. However, because of the brevity of this paper; we cannot go over all the details in the BW4T implementation and experiment. Therefore, the tools used in the experiment (the explanation script and the source code with used world layouts) can be made available upon request. With this, the research provides both conceptually and in practice the components needed to help others to reproduce these findings.

7 Discussion

7.1 Interpretation of results

The results show indeed as our hypothesis states that sharing intentions gives higher predictability compared to sharing no information. During the implementation it was clear that sharing no information did not result in any changes in the PredictorAgent's knowledge about the probabilities in the human action sequence, leading to a clear difference in the measures. An interesting observation is comparing sharing intentions and sharing world knowledge. Here we can see that sharing world knowledge contributes to higher predictability but has a lower performance when compared to sharing intentions. In most measures this difference was found to be significant, with the exception of the 'average possible sequences' measure. However this was more due to the measure itself as all t-tests showed insignificance. This was caused by the huge variance as the measure grows exponentially with the length of the human action sequence. The number stays small if there is enough knowledge about future actions but grows large and varies greatly when no knowledge is shared about future actions.

To explain the difference in result for intentions and world knowledge, there are two factors. First, to prevent the experiment from being a self-fulfilling prophecy, full predictability for the human was achievable in both the sharing intentions case and the sharing world knowledge case. This resulted in more messages having to be sent in the world knowledge case as every action needed a message which implied the manipulation of the world by that action. This is reflected in the average human messages sent in figure 5. Second, when considering the definition of predictability in our context, it is about anticipating the actions of the human. Therefore, a world knowledge message had to be sent before the action was completed but be sent after we are sure that the world is going to be manipulated in this way by the action. This made shared world knowledge only effective for predictability when shared while the action was being performed. This is a short time frame compared to sharing intentions, which are shared before the human performs these actions. Also, on average more actions could be implied by sharing an intent, which meant that fewer messages needed to be shared to achieve the same effect. This is reflected in the average intent messages shared in figure 5.

Although there was feedback from the participants that sharing world knowledge was a demanding task, there was no measurement done of the task load of the human. However, we can speculate that the two factors played a role in overloading the participant, giving a higher human error rate, (i.e. missing out on timing) which lead to a lower predictability for the world knowledge case compared to the intentions case. In the same way it can be speculated why sharing both information types might have had a worse performance compared to only sharing intentions. Feedback from this case was that it was also demanding, which can be traced back to the high standard deviation in the shared human messages. Some participants shared few messages since they missed out on sharing the right information because they were overloaded with thinking about what to share. Other participants could keep up resulting in a higher number of messages shared. These speculations reflect human capacity in a teamwork environment. Therefore, being as efficient as possible with sharing information that implies future action (in this case: sharing intentions), plays a bigger role in the predictability of the human.

Figure 5 shows the average ticks, the average total movement actions (PredictorAgent + human) and the average number of messages shared. Due to the agent slowdowns and humans needing more time for decision making, more could be found about the team effectiveness with the average total moves than with the average ticks. On average, sharing intentions and sharing both information types resulted in significantly less movement needed to finish the task compared to no information sharing. As explained in section 3.2, we could speculate that there was higher predictability of the human as more 'backup behaviour' and 'adaptability' was shown by the PredictorAgent. Showcasing a similar direction with the findings of team effectiveness and information sharing in the works of Jonker et al. (2012) and Li et al. (2016).

Although not significant, the shared world knowledge had on average lower team effectiveness compared to other information-sharing cases. This is speculated to be because of the policy of the PredictorAgent. If participants shared the world knowledge of having found the PredictorAgent's target block in one of the upper rooms, the PredictorAgent would walk very slowly to the upper rooms to grab that block while it would normally be searching the lower rooms first. As participants were observed to search the upper rooms first, this often resulted in decreased team effectiveness through more movement of the PredictorAgent.

To come back to the main contributions and the research question of this research, it is shown that both intentions and world knowledge contribute to higher predictability of the human from the perspective of the agent. However, sharing intentions performs better when compared to sharing world knowledge or sharing both information types. From the experiment, it is indicated that this is because of intentions being the most effective in implying future action. Therefore, the human experiences less task load since it shares fewer messages. As discussed in section 2, predictability is important for building the interdependence relationship, which contributes to a higher mutual trust.

7.2 Limitations and Future work

The main limitation was the duration of this research. The usage of the human action sequence was only to evaluate the predictability. The PredictorAgent made no use of the sequence in terms of planning its own actions around it and it did not create a human action sequence on its own based on the human messages shared. Therefore, it is still a question as to how such a sequence can be created by messages that imply the action taken and how that can positively influence the planning of the agent. The factor of trust was also excluded as the experimental setup was kept small. Due to the experiment having to be set up and performed partially online, it had a small sample size with a relatively similar population.

These limitations are left as directions for future work. An interesting direction to delve into is to integrate the notion of the action sequence into existing planning algorithms used in multi-agent systems. Such as using it to reduce the non-determinism when working with the FOND planner (Muise et al., 2015). Furthermore, a direction to follow is to see whether the findings hold up when the experiment has a bigger and more diverse sample size as well as when the task complexity increases.

8 Conclusion

This research has experimented in BW4T about the effect of sharing intentions, world knowledge or both types on the effect of the human's predictability to an autonomous agent. The main findings of this experiment were that sharing intentions contributes the most to higher predictability relative to no shared information. This is because intentions imply more actions in one go, which added the least complexity to the task for the human. Sharing world knowledge contributes to higher predictability but performs worse compared to sharing intentions. This is because requires more messages to be sent and is only effective for predictability when shared in the appropriate time interval. Therefore, it is speculated that the human is overloaded by this task. Sharing both information types performed better than sharing world knowledge but requires the most messages to be shared. It is speculated that this case failed to perform equally or better than the intentions case due to its complexity.

The main limitation was found to be the duration of this research, which limited the implementation of the PredictorAgent to only the system which evaluates the Predictability of the human. Future work directions are therefore using a planning algorithm which integrates the notion of an action sequence and to experiment on a larger, more diverse participant group with a more complex task.

These findings contribute to the active research topic of mutual predictability, which is an important part between humans and agents in teamwork as it is an interdependence relationship that contributes to mutual trust. With this mutual trust being an important factor in improving beneficial collaboration, it shows how mutual predictability is of importance as human-agent teamwork. Especially as human-agent teamwork becomes more collaborative due to the increasing complexity of AI.

References

Ahrndt, S., Fähndrich, J., and Albayrak, S. (2016). Human-agent teamwork. *Proceedings of the 31st Annual ACM Symposium on Applied Computing*, 04-08-April-2016:284–286.

Bradshaw, J. M., Dignum, V., Jonker, C., and Sierhuis, M. (2012). Human-agent-robot teamwork. *IEEE Intelligent Systems*, 27:8–13.

Johnson, M. and Bradshaw, J. M. (2021). The role of interdependence in trust. *Trust in Human-Robot Interaction*, pages 379–403.

Johnson, M., Bradshaw, J. M., Feltovich, P. J., Jonker, C. M., Riemsdijk, M. B. V., and Sierhuis, M. (2014). Coactive design: Designing support for interdependence in joint activity. *Journal of Human-Robot Interaction*, 3.

Johnson, M., Jonker, C., Riemsdijk, B. V., Feltovich, P. J., and Bradshaw, J. M. (2009). Joint activity testbed: Blocks world for teams (bw4t). *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 5881 LNAI:254–256.

Jonker, C. M., Riemsdijk, M. B. V., Kieft, I. C. V. D., and Gini, M. (2012). Compositionality of team mental models in relation to sharedness and team performance. volume 7345 LNAI, pages 242–251.

Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., and Feltovich, P. J. (2004). Ten challenges for making automation a "team player" in joint human-agent activity. *IEEE Intelligent Systems*, 19:91–95.

Lewis, M., Sycara, K., and Walker, P. (2018). The role of trust in human-robot interaction. *Studies in Systems, Decision and Control*, 117:135–159.

Li, S., Sun, W., and Miller, T. (2016). Communication in human-agent teams for tasks with joint action. volume 9628, pages 224–241. Springer Verlag.

Muise, C., Felli, P., Miller, T., Pearce, A. R., and Sonenberg, L. (2015). Leveraging fond planning technology to solve multi-agent planning problems.

Salas, E., Sims, D. E., and Burke, C. S. (2005). Is there a "big five" in teamwork? *Small Group Research*, 36:555–599.

Singh, R., Sonenberg, L., and Miller, T. (2017). Communication and shared mental models for teams performing interdependent tasks. volume 10315 LNAI, pages 81–97. Springer Verlag.

TNO (2021a). Matrx - accelerating human-agent teaming research together.

TNO (2021b). When to use - matrx.

A Experiment Script

Anything between :: indicates something to be done

:open BW4T practice world in god mode:

Thank you for participating in this experiment about human-agent collaboration. You are going to perform the blocks world for teams task together with an autonomous agent which we call the PredictorAgent. The blocks world for teams task works as follows: You are an agent, working together with another agent in a 2d grid world. The goal is to deliver the 4 coloured blocks to the drop zone at the bottom of the map. These blocks have to be delivered in order from the bottom to the top. So the first block is the one at the bottom, the second block is the one above that, etcetera..

So where are these blocks found? They are hidden within rooms located on the map, of which there are six. With room 1 at the left top, :name rooms: , and room 6 at the bottom right. In these rooms, other blocks can also be found. To access the target room, you need to search the room. For that you need to walk to the room, open the door and traverse the inside of the room and check for blocks. If you found the selected target block, you can grab the block and deliver it to the drop zone.

:switch from god mode to TeleOpHumanAgent view:

So, the main constraint of this game is that you and the PredictorAgent have a limited vision range of 1 tile. Which you see in this view. A small detail however is that you can see which doors are opened by you or the PredictorAgent across the whole map. You can finish this game alone, but collaborating with the PredictorAgent means that you two are like an extra pair of eyes to each other, which allows for faster goal completion. Communication works via a chat system, but more on that later.

:open the chat in the TeleOpHumanAgent view:

So you are controlling the dark pink agent with buttons which when pressed, make the agent perform an action. First, you select which target block you are going to search and deliver to the drop zone. After that, there are 4 actions. Going through these one by one:

- *Walking to a room and opening the door:*
  This action makes the agent walk to the selected room and opens the door, a closed door is brown and an open door is green.
- *Traversing the room and checking for target blocks:*
  After you have done the 'open door' action, pressing 'traverse room' makes the agent walk inside the room and check per step whether the blocks it can see are target blocks that need to be collected. If the target block which you are currently searching for is found, then the agent will notify you via the chat.
- *Grabbing a found target block:*
  When you are notified that your currently selected target block is found, you press the 'grab found target block' action. This action makes the agent select the closest found target block and grabs it. The agent knows the location of this target block, therefore

there is no need to walk a selected room first, the agent does this for you.

- *Walking to the drop zone and dropping the block:*
  When you have grabbed a target block, pressing this button makes you walk to the appropriate location on the drop zone. When at the drop zone, if the blocks below have not been delivered yet, the agent will wait until the PredictorAgent has found and delivered the blocks below you on the drop zone. The PredictorAgent will therefore always deliver the next target block in line to its own knowledge. Therefore focussing on delivering the blocks in order from bottom to top will make your agent wait for less.

When the agent has dropped the block, then you will need to select the next block you want to target after which you can start searching for that block.

Final piece of information is the chat. The chat works as follows, messages in blue in the chat are from your agent giving you feedback on the world and the actions you assigned and the messages you have sent. (give examples) The messages in white are from the PredictorAgent giving you information about its intentions and knowledge about the world. When the task is running, you are expected to look at the map, read from the chat and base your actions around that. You can take your time with this as PredictorAgent is moving at a slower pace to give you time to take in the information from the screen. Do not worry about making mistakes such as clicking too much on single buttons and mis-clicking buttons, but try to avoid them of course.

Before we start the practice run with only actions (red buttons), do you have any questions?

:answer questions if there are any:

You can press the start button in the left top corner to start the task.

:play practice round:

Good job on finishing the practice run! We will now switch to the actual experiment world.

:change world settings to final world, refresh the page, open the chat again in the TeleOpHumanAgent view (Note, do not show the god view!):

As the name of the PredictorAgent implies it will try and assess how predictable you are. This predictability is dependent on the communication you have with the PredictorAgent via the chat. There are two types of message buttons: purple buttons for communicating intentions to the PredictorAgent; and yellow buttons for communicating through world knowledge.

You will be assigned to the case with... :assign participant to case: (if only one information type sharing, then mention to disregard other coloured buttons)

:take time explaining the following section, allow for asking questions:
(Explain this only for intentions and both case) For the purple buttons you have the following: for when you set the target block with the blue buttons then you can directly after that press the "Message target block goal to PredictorAgent" button to communicate your current target block. Then for the other purple buttons, you have to press them before you

perform the actions which the purple button implied. So for example when you press the "traverse room 1" button, means that you must then perform the action sequence of "Open door room 1" and "Traverse room". Then for the purple "Deliver found block" button when you press it you must then perform the action sequence of "Grab found target block" followed up by "Deliver held block to drop zone".

(Explain this only for world knowledge and both case) For the yellow buttons you need to do the following: when you have pressed a red action button, then directly after that you need to press the yellow button to the right of it. Furthermore, there is the "Inform of found block(s)" button which you need to press after you have traversed a room. This is different to the 'inform' button of the traverse room action as this only informs that a room has been visited, whereas the other one informs of the blocks found.

Any questions before we start the second run?

:answer questions if there are any:

:play real round: