Communication in Human-Agent Teams for Tasks with Joint Action

Sirui Li, Weixing Sun, and Tim Miller(✉)

Department of Computing and Information Systems,
University of Melbourne, Melbourne, Australia
{siruil1,w.sun}@student.unimelb.edu.au, tmiller@unimelb.edu.au

Abstract. In many scenarios, humans must team with agents to achieve joint aims. When working collectively in a team of human and artificial agents, communication is important to establish a shared situation of the task at hand. With no human in the loop and little cost for communication, information about the task can be easily exchanged. However, when communication becomes expensive, or when there are humans in the loop, the strategy for sharing information must be carefully designed: too little information leads to lack of shared situation awareness, while too much overloads the human team members, decreasing performance overall. This paper investigates the effects of sharing beliefs and goals in agent teams and in human-agent teams. We performed a set of experiments using the BlocksWorlds for Teams (BW4T) testbed to assess different strategies for information sharing. In previous experimental studies using BW4T, explanations about agent behaviour were shown to have no effect on team performance. One possible reason for this is because the existing scenarios in BW4T contained joint tasks, but not *joint actions*. That is, atomic actions that required interdependent and simultaneous action between more than one agent. We implemented new scenarios in BW4T in which some actions required two agents to complete. Our results showed an improvement in artificial-agent team performance when communicating goals and sharing beliefs, but with goals contributing more to team performance, and that in human-agent teams, communicating only goals was more effective than communicating both goals and beliefs.

Keywords: Human-agent collaboration · BlocksWorld for Teams · Joint action · Interdependence

1 Introduction

Over the past decade or so, there has been a realisation that "autonomous" intelligent agents will offer more value if they work semi-autonomously as part of a team with humans [4]. Semi-autonomous agents must therefore be designed to explicitly consider the human in the loop to work effectively as part of a team.

In a joint task, a team has a joint aim to achieve a goal, and they must work together to do achieve this goal. While in some simple scenarios, team

V. Dignum et al. (Eds.): COIN 2015, LNAI 9628, pp. 224–241, 2016.
DOI: 10.1007/978-3-319-42691-4_13

members may be able to operate individually to achieve the joint goal, in most scenarios, the individual actions within a task are *interdependent* [17]. However, to successfully operate on an interdependent task, team members must have a shared situation awareness of at least part of the task, and must coordinate the actions that comprise the task. As such, communication between team members is important to efficiently complete a task.

This is the case in human-human teams, but also in human-agent teams. For example, Stubbs et al. [18] observed over 800 h of human-robot interaction and noted that as the level of autonomy in the robot increased, the efficiency of the mission decreased as operators started increasingly questioned the robot's decision making. Stubb et al. concluded that having an agent explain relevant parts of the behaviour to maintain a *common ground* on the task is important for effective collaboration. The process to achieve common ground requires communication, taking into account what is necessary and important, and what the other team members already know.

The aim of our work is to identify the types of and amount of information that are relevant for interdependent tasks. We use the BlocksWorlds for Teams (BW4T) [11] test bed for this. BW4T is a simulation tool that allows experimentation of scenarios involving humans and agents. The joint goal of the human-agent team is to locate and retrieve a sequence of coloured blocks in a given order. Harbers et al. [5,6] have experimented with the same concept in BW4T, however, they found that communication did not have a significant impact on team efficiency in completing the task. They hypothesise that one reason may be the simple nature of the task, and that more complex scenarios show results similar to those seen in field experiments such as the ones by Stubbs discussed above.

In this paper, we develop a new scenario for BW4T that contains *joint actions*, rather than just joint *tasks*. By "action", we mean the atomic actions that make up a task. Our simple extension is to introduce a type of heavy block that requires one agent to hold the door to a room for another agent, meaning that moving the block out of the room is a joint action. That is, no individual agent can move the block out of a room. In terms of the model proposed by Saavedra et al.[15], this extension moves the task from one of a team merely working in parallel towards a common goal, called *pooled interdependence*, to one of *team task interdependence*, where team members must execute actions jointly.

We performed initial experiments to assess different communication strategies teams of artificial agents, demonstrating that sharing goals improves task efficiency better than sharing beliefs. Then, we used this to determine experimental parameters for human-agent experiments on similar scenarios, and showed sharing goals in the scenarios does indeed increase the efficiency of the team in completing the task. Further, we observed that sharing too much information resulted in decreased performance due to information overload.

This paper is outlined as follows. Section 2 presents the most closely related work, and Sect. 3 presents relevant background on the BW4T simulator.

Section 4 outlines the agent communication models used in our experiments, including how agents handle communication for joint tasks. Section 5 presents the experimental evaluation, including results, for both sets of experiments (agent teams and human-agent teams), while Sect. 6 concludes.

2 Related Work

In this section, we discuss the most closely related work to the work in this paper. While there is a large body of work investigating how human teams work together on interdependent tasks [14,15] and how the process of *grounding* a common ground [3,12], and their relation to shared cognition of a team [2,16], this section will focus on related work on interdependence in human-agent teams.

The primary questions of work in this domain are: (1) "how much" autonomy should we grant a semi-autonomous agent, and; (2) given this, what information needs to be communicated between the agent and the human for efficient task completion. In this paper, we look mostly at the second question.

In recent years, the realisation that human-agent teams offer more than agent-only teams has lead to many empirical studies of human-agent teams [1,4,13] that address the issue of what and when to communicate to team members. For example, Stubbs et al. [18] discuss their experience observing over 800 h of human-agent teamwork in a scientific setting. Their team remotely deployed a robot in the in Chile's Atacama Desert to investigate microorganisms that grow there, with the view that such a deployment would be similar to deploying a semi-autonomous robot on other planets. The team changed the level of autonomy of the deployed robot, giving it more responsibility on some tasks in certain cases, and observed the scientific teams' response. Stubbs et al. found that as the level of autonomy increased, the effectiveness of the team reduced. This was mostly caused by a lack of transparency in the robot's decision-making, resulting in cases where the scientific team spent more time discussing and trying to understand why the robot had made certain decisions, rather than on the scientific aims related to microorganisms. Stubbs et al. hypothesis that establishing a *common ground* between the relevant parties on tasks is essential.

Bradshaw et al. [1] hypothesise that human-agent teams will become more effective if agents are considered peers and team members, rather than just tools to which to delegate tasks. They later discuss the concept of *coactive design* [9], and argue that the consideration of interdependence between agents in performing joint tasks is key to effective human-agent teams. They define interdependence as the relationships between members of a team, and argue that these relationship determine what information is relevant for the team to complete a task, and in that sense, the interdependent relationships define the common ground that is necessary. In more recent work [10], they present the Coactive Design Method for designing intelligent agents that must interact with humans. In this model, interdependence is the organising principle. Human and artificial agents worked together through an interface that is designed around the concepts of Observability, Predictability and Directability (OPD). The model was applied to the design of a simulated teleoperated robot for the DARPA Virtual Robotics Challenge, and obtained an excellent score due to the advantages the coactive system model. They describe scenarios in which the identification of interdependent tasks improved their agent design, such as the robot having to attach a hose to a spigot. The robot is unable to identify the hose — a task done by the human —, but attaching the hose itself was a joint task, in which the robot positioned the hose and the human directed the arm to the spigot.

Other recent work looks at how to simulate such scenarios in a laboratory setting to allow for more controlled experimentation. In particular, the BlocksWorld for Teams (BW4T) testbed [11], used in our work, was developed to support experimentation of human-agent teaming in joint activities.

Harbers et al. [5,6] use the BW4T testbed to experiment with explanation in human-agent teams. In particular, they looked at the effect of sharing beliefs and intention within teams, providing the humans with the ability to exploit information about intentions to improve their understanding of the situation. Their results showed that, while participants reported increased awareness of what the agents were doing, there was no improvement in team effectiveness measured by completion time. Thus, their explanation model did indeed explain the situation, but this information was not useful for the human players to coordinate their actions. Harbers et al. hypothesise that this may be because the team tasks are so straightforward that the human player can easily predict what behaviour it requires, and thus processing the explanations has a cost that is similar to what the explanation is worth. We agree with this analysis. Our experiments are similar in spirit to these experiments, however, the introduction of joint action helps to provide a more complex scenario without increasing the complexity to a point that confuses the human players or requires extensive training.

In other work, Harbers et al. [7] used BW4T to investigate communication in agent-only teams, and found that sharing intentions and taking advantage of this knowledge increased the team efficiency, while sharing beliefs had minimal impact — a finding consistent with the work in this paper.

Wei et al. [19] study the construction and effectiveness of shared mental models between artificial agents using BW4T. They designed four scenarios with different numbers of artificial agents and environment sizes, and measured completion time as a proxy for the effectiveness of different communication strategies. Their results showed that communicating between team members improved efficiency, especially in the case in which there were sequential interdependencies between tasks; that is, the tasks had an explicit order in which they must have been completed. Further, they also found that communicating more information lead to more interference between agents, indicating that even in agent teams where processing is not a large issue, it is important to communicate only the most relevant and important information. Our work goes further than the experiments by Wei, Hindriks, and Jonker by looking at joint actions and including humans in the loop.

3 BlocksWorld for Teams

BlocksWorld For Team (BW4T) is a simulator that extends the classic blocks world domain, written specifically for experimenting with human-agent teams. The overall goal of the agent team is to search for the required blocks in a given set of rooms. The task can be performed by a single agent or a group of agents. Agents can be either artificial or human. The role of the agents can be distinguished based on how it is programmed.

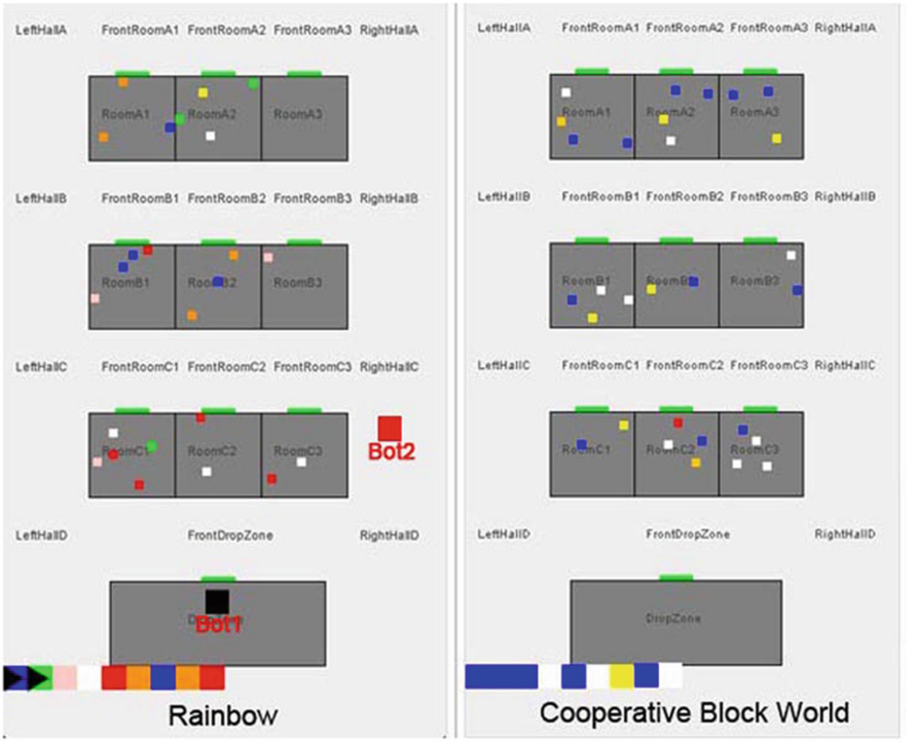

Fig. 1. The BW4T environment (Color figure online)

Figure 1 displays the three different BW4T maps we used in our experiments. The environment of BW4T consists of rooms and coloured blocks scattered in different rooms. Each room has one door, which is represented by the small green bold line. The dark area on the bottom is the drop zone, where blocks are dropped once collected. The small black squares with red labels represent agents. At the bottom, the sequence of colours specifies the blocks that the team is tasked with collecting. The team must put down the block with the right colour into the drop zone, otherwise, the block will disappear. The sequence is represented by the colourful bar on the bottom of the environment. The small triangles on the colourful bar means the completed tasks.

The agents within BW4T are programmed using the GOAL programming language [8], and the BW4T simulator provides specific constructs for interacting with GOAL agents. Agents can perceive the environment using an environmental sensor, including information such as the next target block, or the blocks in the room they are in.

Agents communicate to each other using messaging, and the contents can be arbitrary. On receiving a message, it is stored in a "mailbox" for reading. When an agent representing a human (which we call the *supervisor* agent) receives a message, it translates the message into a human-readable format, and displays this on the GUI that is viewable by the human. The human player can inform and direct the supervisor agent using a drop-down menu of commands; e.g. telling the agent which room a particular-coloured block is in.

4 Agents and Joint Actions in BW4T

In this section, we present the scenario and models of agents that we used to experiment with human-agent teams in joint activity. We model how an artificial agent communicates with artificial team members, and then with humans.

4.1 The Scenario

From the perspective of the rules of the BW4T game, we alter only one aspect: we introduce *types* of block. In the BW4T simulator, blocks have colours, and the sequence of target blocks must be returned according to a specific colour in each slot. In our model, blue blocks are given a special status, in that they are considered *heavier* than other blocks, and they require two agents to get the block from its location to the drop zone. As part of our experiments, we implemented a simple scenario in which, when an agent wanted to take a blue block from a room, a second agent was required to hold the door open for them (because the block is too heavy to hold in one arm, and the carrying agent therefore has no hand to open the door). As soon as the carrying agent exited the room, the second agent holding the door was free to return to another task.

This represents an *interdependent action* [17]: an agent can only take a blue block from a room if another agent opens the door, and the agent opening the door receives no value from this unless the block is taken from the room and back to the drop zone. One can imagine different implementations; e.g. two or more agents must carry blocks together, but this simple variation is enough to test out joint actions in BW4T.

4.2 Agent Models

In this section, we outline our model for dealing with the joint activity of collecting a blue block. We adopt a basic model of searching and retrieving blocks, and extend this with the ability for agents to *request* and *offer* assistance for heavy blocks. In our basic model, all agents know the sequence of blocks to be found. They search rooms in a semi-random fashion, not revisiting rooms known to have been visited by themselves or a team member, and maintain a belief set of the locations of blocks (which colours and in which rooms) that they have perceived. An agent's "default" goal is to find and retrieve the next block in the sequence, until they receive a request for assistance, or until another agent finds the block and broadcasts this fact, in which case they adopt the goal of finding the next block.

4.2.1 Requesting and Offering Assistance

Requesting help and offering assistance are required for the particular joint activity of heavy blocks. As mentioned before, blue blocks represent heavy blocks. This process for request and offering assistance for a blue block is shown in Fig. 2.

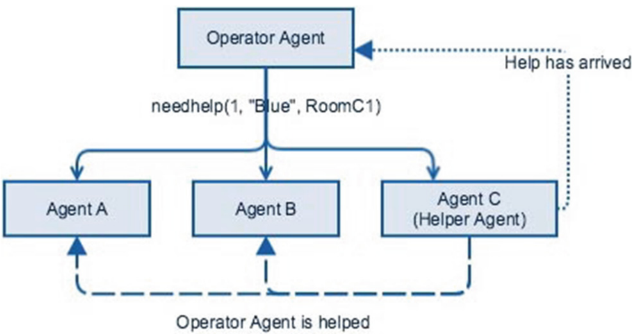

Fig. 2. Request help and offer assistance

All available agents will search for the blue block. The first agent to find one, who we call the *operator agent* will broadcast the `needhelp` message to all other agents. Any artificial agent ready to assist will move towards the room, and send a message ("Help has arrived") indicating they are at the help position (e.g. holding the door at `Room1`). The first agent to arrive will inform all others, who adopt their default goal of searching for the next block in the sequence. All agents attempting to help may not be an efficient use of their time, but we opt for a simple policy here to avoid any possibility of this policy influencing results about communication. As this policy is consistent across all experiments, we do not consider that this meaningfully affects the results.

4.2.2 Supervisor Agent

Recall that humans are represented by a *supervisor* agent, who can direct other agents to perform tasks. This agent acts as an interface between the human and artificial agents, but is also a player capable of finding and retrieving blocks.

Human players direct their representative agents using high-level commands, such as which block to search for; in sense, simulating a basic remote teleoperation of a robot. Human players can request and offer assistance like artificial agents, however, the decision making about whether to offer assistance is left up to the human, rather than coded in GOAL.

Figure 3 shows the models used for a supervisor agent. The first model is used when taking on a new task (Fig. 3a), and the second is used when the player intends to provide assistance to another agent trying to retrieve a blue block (Fig. 3b).

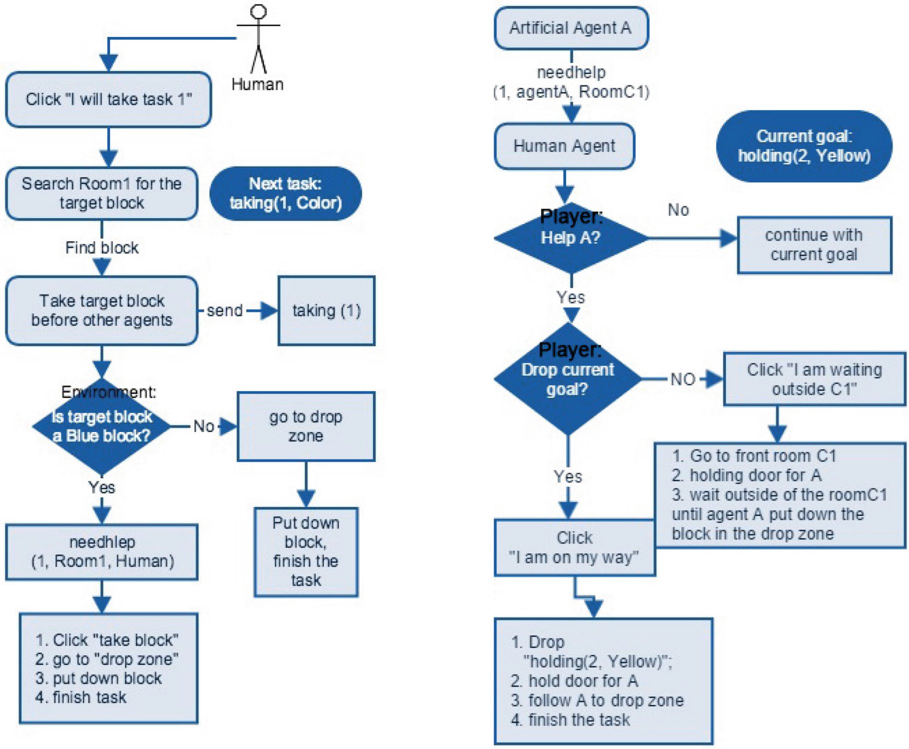

(a) Completing a New Task (b) Helping the Artificial Agent

Fig. 3. Supervisor agent (Color figure online)

From Fig. 3a, one can see that a supervisor agent is idle unless directed by the human player to take a task; that is, to starting searching for a particular colour block. The supervisor agent then searches autonomously for the block. If it finds the block and the block is non-blue, it will update the other agents to inform them that the block has been located and is being taken back to the drop zone, allowing other agents to drop this task[1]. If the block is blue, it will request help and wait. After getting help from other agent, the "*take block*" option is made available on the human player's GUI, and clicking this directs the supervisor agent to take the block to the drop zone autonomously.

The other artificial agents adapt the player's changing actions. For example, if the supervisor agent drops its goal while carrying a block (e.g. yellow) to the drop zone, the other agents will drop their current task of searching for the next block in the sequence, and will adopt the goal of finding a yellow block.

The model outlined in this section is put in place to provide human decision support into the system. In a team with only artificial agents, if all agents have a current goal and one agent finds a blue block, it will be required to wait for one of its team members to complete its tasks.

However, in our model, we offer the human player the possibility to drop its own goal to help complete the tasks. We opted not to have the human player directing other artificial agents to drop goals when other agents find blue blocks, as we believed that the extra decision of *which* agent to direct could increase the cognitive load of the human player to the point where decisions became arbitrary. By allowing the human player to direct only their own agent, this model provides a complex-enough scenario to introduce an interdependent action into BW4T, without the complexity of the scenario overwhelming participants.

Ultimately, we believe that the results from our experiments (see Sect. 5) demonstrate that our decision is justified.

4.3 Information Exchange Between Agents

It is clear that sharing information can improve team efficiency. However, the information shared, and how much of it, is crucial, especially in human-agent teams, where the humans' capacities to process information is reduced compared to its artificial team members.

4.3.1 Information Messaging

In this section, we present the communication protocols between agents, which consist of individual messages. Several types of message can be sent, enabling agents to inform others about part of the environment, or its own goals.

Beliefs are the information about the environment, which are perceived via agents, such as the location of different coloured blocks. *Goals* are mental states that motivate action. To complete a single task, an agent must complete a sequence of goals. We enabled agents to share their beliefs and goals.

Five messages can be transferred amongst agents:

1. `block(BlockID, ColourID, PlaceID)`: block `BlockID` with colour `ColourID` has been found by the message sender at room `PlaceID`. When in a room, an agent broadcasts information about any block colours that are in the goal sequence.

[1] Artificial agents are also programmed with this capability in our model.

2. `visited(PlaceID)`: room `PlaceID` has been visited by the message sender. While this can be inferred when a `block` message (above) is sent, the `visited` message is sent when the room does not contain any blocks that are in the target sequence.
3. `hold(BlockID, ColourID)`: block `BlockID` with `ColourID` is held by the message sender.
4. `will_deliver(Index, ColourID)`: a block of colour `ColourID`, which is also the `Index`th block in the sequence of the main goal, is being delivered to the drop zone.
5. `dropped(Index, ColourID)`: the `Index`th block in the task sequence, with colour `ColourID`, previously held by the message sender, has been dropped.

While all messages are sharing information about the task, the intention of the first three is to share belief about the environment, while the intention of the last two is to share goals; e.g. when an agent delivered the task, it will drop this goal.

Agents use the information about where they have visited and what colour blocks are in the rooms to inform their search strategy. We model the artificial agents to used the shared information about block locations and room searching to improve the completion of the task. For example, when the agents share their belief about the location of blocks, others can update their own beliefs with this information, preventing unnecessary searching of rooms.

Our models use shared information about blocks being retrieved and dropped to further improve this. That is, when an agent broadcasts that they have located the next block, others will stop searching for that colour, and when a block in the main sequence is dropped, others will starting searching for this again.

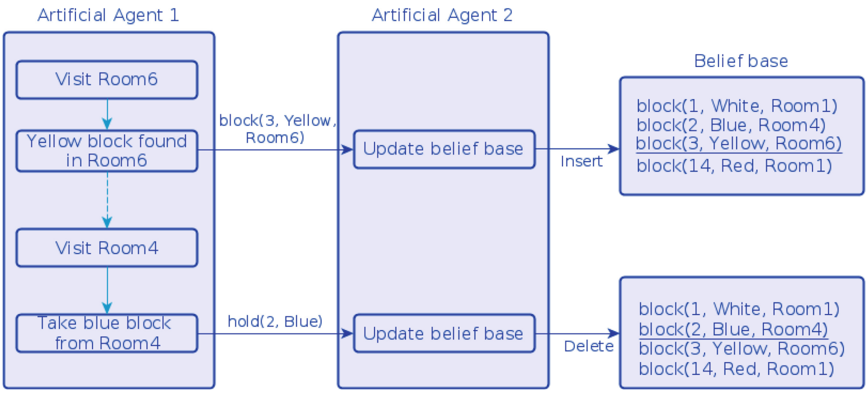

Fig. 4. The effect of sharing belief

Figure 4 shows how information about block location and holding blocks. Suppose artificial agent 2 is stationary, while artificial agent 1 is exploring for a blue target. Agent 1 visits room 6, there is no blue block, but a yellow block is inside of the room. If the agents share beliefs, agent 1 will send a message `block(3, Yellow, Room6)` to all the other agents, and all other agents will update their belief base with this information.

Next, suppose agent 1 goes to room 4, which contains a blue block. The agent will pick up the block and broadcast the message `hold(2, Blue)`. The others agents will then update their belief base to remove the belief `block(2, Blue, Room4)`, if they had this belief in their belief base.

4.4 Filtering for Human Players

The hypothesis in human-agent collaboration research is that explanation from the later can improve team performance in the joint activities. However, it is clear that humans do not have sufficient processing capabilities to use all information shared in the previous section. Despite this, the human player also needs to know some of the critical information such as the environment states and other artificial agents' message.

In our model, the supervisor agent takes on the role of an information broker who is responsible to deliver and translate information for the human player, and to filter the "explanation" from artificial agents. From the artificial agents' perspective, a supervisor agent is another artificial agent that receives and sends messages, and supervisor agents are the bridge between the environment, artificial agents, and human players.

Figure 5 shows two examples of translating information between the artificial agents and human agents — one for each direction. In Fig. 5a, the supervisor agent accepts two parts of input: (1) from the environment, including information such as the room occupancy and current team target; and (2) from another agent, including the requesting assistance and sharing the goals. The supervisor agent selects some of the incoming information and "explains" this to the human agent.

The key part of any design is what information should be filtered out, and what should be filtered in and explained. In the next section, we describe an experiment design that looks at three levels of filtering, and their effect on the performance of the overall system.

5 Experimental Evaluation

In this section, we outline two sets of experiments to provide evidence towards our hypothesis that communication can improve the team performance in joint activities, and report the results. The first set of experiments runs three BW4T scenarios using a team made entirely of artificial agents, while the second set includes a human player in the loop, along with its supervisor agent. Within each experiment, the information that is shared between team members is changed to measure the effect of information exchange.

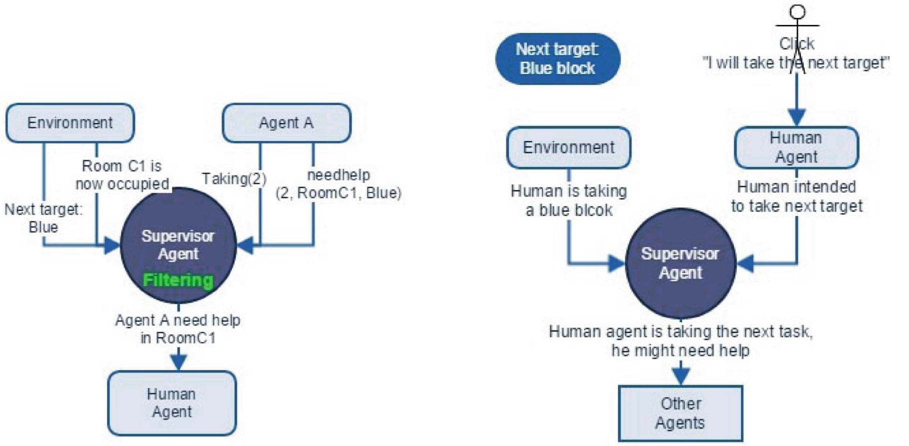

(a) Supervisor transfers artificial agent intentions to human player

(b) Supervisor transfers human intentions to artificial agents

Fig. 5. Filtering by the supervisor agent

5.1 Artificial Agent Team Experiment

5.1.1 Experiment Design

The aim of this experiment is to study which type of information sharing between artificial agents effects the team performance: sharing beliefs, goals, or both.

Independent Variable. We modify two independent variables: (1) communication strategy; and (2) the environment type.

For the communication strategy, we use four values: (a) minimal information shared: the only communication is to ask for help moving a blue block; (b) belief only: minimal plus belief about the environment (items 1–3 in Sect. 4.3.1); (c) goals only: minimal plus agent goals (items 4–5 in Sect. 4.3.1); and (d) belief and goals.

For the environment, we use three different maps: (a) cooperative block world; (b) rainbow; and (c) simple. The first two are shown in Fig. 1 (page 5). Cooperative block world contains seven block colours, but only three occur in the main goal, and these are randomly allocated to the main goal in a uniform manner. Rainbow contains seven coloured blocks, and all seven colours can appear in the main goal. Simple contains randomly allocated blocks, but with no blue blocks; and therefore, no joint action.

Measures. We measure completion time of the entire scenario as a proxy for the effectiveness of each communication strategy.

Setup. For each map, we run all four communication strategies giving us 12 combinations. Each combination is run 30 times, with different random seeds to generate different block locations, resulting in 360 experiments run in total.

All experiments were run with two artificial agents, nine rooms, and nine blocks in the main goal.

5.1.2 Results

Figure 6 shows the average completion time for all combinations of scenarios and communication strategies. This figure demonstrates several interesting findings from our experiments.

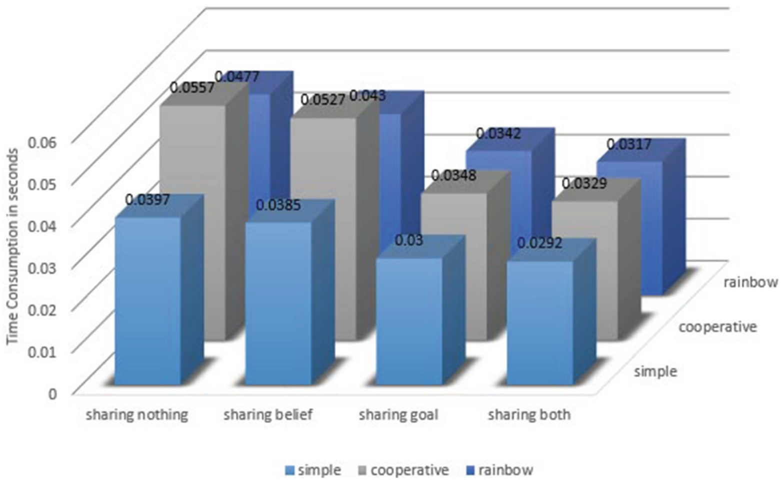

Fig. 6. Average task completion time for the artificial agents team (Color figure online)

With regards to the three scenarios maps, cooperative blocks world consumes more time than other two, and the simple map, with no strictly joint action, took the least time to finish on average. This supports our hypothesis that having joint actions in a scenario increasing the complexity more than simply joint tasks. The largest gap (40 %) between the cooperative block world and simple world results is in the scenario where "nothing" is shared (recall that agents still request help once they pick up a heavy blue block), indicating that sharing beliefs and goals is useful in this environment. Further, for the cooperative blocks world scenario, there is a large step between sharing belief and sharing goals, indicating that sharing goals is far more valuable that sharing just belief. This is further backed up by the small decrement from sharing goals to sharing both belief and goals. In all three maps, sharing belief had only a small impact. This finding is interesting, because while agents share their knowledge of the environment, meaning that searching for the right coloured block can be reduced, it is in fact coordinating the joint action early that increases efficiency the most in this scenario.

Table 1 shows the outcomes of a two-way factorial ANOVA to examine the influence of the two different independent variables. The p-values for the rows (maps), columns (communication strategy), and the interaction, are all < 0.001,

Table 1. The two-way factorial ANOVA results for the artificial agents team

Source of variance	SS	F	P-value
Communication strategy	6669.31	110.39	9.20E-33
Map	1948.40	48.38	9.97E-16
Interaction	585.17	4.84	2.04E-04

indicating that the results are statistically significant to this level. Comparing the sum of square errors (SS), we see that communication has more impact than the scenarios, but both factors have a significant influence on the results.

The results show that communication is beneficial for improving cooperative team work, and sharing goals has the largest impact. We drilled down into the experiment data and found that the primary reason for this was labour redundancy. An agent will update its team members once a block is placed in the drop zone, limiting the team members' knowledge of task progress. By sharing the goal that they have collected a block suitable to fulfil the current team sub-goal, the other team members can start on a new task.

5.2 Human-Agent Teams

The results from the artificial agent teams helped to inform the design of the human-agent team experiments. In this section, we outline the experimental design and results for the human-agent team scenarios.

5.2.1 Experiment Design

The aim of this experiment is to study how the type of information shared between the human player and other agents effects the team performance. Due to the introduction of a human into the loop, the experiment is much simplified compared to the experiments in the previous section, as we aimed to keep total completion time to under 30 min for each participant.

Independent Variable. The independent variable in the experiments is filtering strategy used by the supervisor agent to exchange information with the human player: (1) *full info*: everything is shared as in the artificial team; (2) *partial info*: only information that will change the goals of the human player are shared; and (3) *silence*: only information that a block has been delivered to the drop zone. Table 2 outlines what information is shared in each of the three cases.

Measures. As in the artificial team experiment, we use completion time of the entire scenario as a proxy for the effectiveness of each communication strategy.

Setup. We recruited 12 participants to perform three runs of the experiment — one with each communication strategy. No participant had used or heard of the BW4T simulator previously. To avoid bias, the order in which the participants used the various communication strategies were systematically varied.

Table 2. The information shared in the three scenarios

Information	Full info	Partial info	Silence
Next target	✔	✔	✔
Other agent's current task	✔	✔	
Request assistance	✔	✔	
Offer assistance	✔	✔	
Task completion	✔	✔	
Block location	✔		
Room occupancy	✔		
Other agents' state	✔		

Due to the relative difficult of recruiting participants and running the experiments, we used only one map in all three scenarios: the cooperative block world map (Fig. 1). We chose this map because the results of the agent-team experiments demonstrate that this best simulates a reasonably complex scenario with joint action. The speed of the BW4T simulation is adjusted to be slow to provide the human player with sufficient time to make decisions. Each experiment consisted of two artificial agents, one supervisor agent, and one human players. There was no time out for completion of the experiment, and none of the participants failed to complete the scenario.

5.2.2 Results

Figure 7 shows the results for the human-agent team experiments. Due to the relatively smaller number of data points, results for each participant is shown. The x-axis is the communication strategy, the y-axis are the individual participants, and the z-axis is completion time. Results are sorted roughly by completion time. The overall average completion time for the three scenarios are: full info = 4.92 min, partial info = 4.36 min, and silence = 4.72 min.

From the figure, it is clear that results differ among people, but that the difference between the strategies per person establishes a trend. From the average scores, having full information took the longest time, followed by silence, and finally, the partial information. These results support the hypothesis that explanation can improve the team performance in scenarios with joint action, and further, that too much explanation can hinder a human players ability for decision making.

To test the effect of different communication scenarios, we performed a repeated measures two-way ANOVA between groups, and a pairwise Tukey HSD comparison between all pairs of groups. Relevant values for the ANOVA are shown in Table 3. For the pairwise Tukey HSD test, the full information vs. partial information results are significant at the 0.05 level, while the other two pairs are not. These tests demonstrate that the results between groups is significant (Table 3).

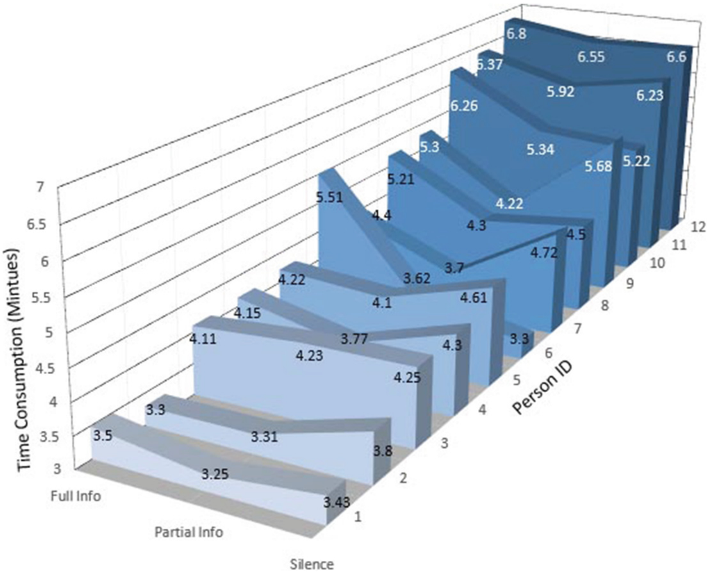

Fig. 7. The results of human-artificial agents' team

Table 3. ANOVA analysis of human-artificial agents' team results

Source	SS	df	MS	F	P-value
Treatment (between groups)	1.985	2	0.9925	5.07	0.0154
Error (within groups)	4.305	22	0.1957		

6 Conclusions and Future Work

In this paper, we studied the effectiveness of communication in artificial agent teams and human-agent teams using the BW4T testbed. Extending previous studies using BW4T, we added the concept of a *joint action* — a single atomic action that requires more than one agent to complete.

For the artificial agents team, we performed extensive simulation experiments to assess the value of sharing beliefs, sharing goals, and sharing both belief and goals. The results showed that sharing goals, namely, agents exchanging their immediate goals, increase team efficiency significantly more than sharing beliefs.

Using these results, we designed an experiment using the same joint action scenario, but with a human player in the loop. We recruited 12 people to each play in three scenarios using three different communication strategies: (1) update only when a block sub-task has been completed; (2) share goals; and (3) share goals and beliefs. We observed that sharing goals and beliefs lead to information overload of the human, resulting in a less efficient team than just sharing goals, and that sharing almost nothing is more efficient than sharing all goals and beliefs, most likely because the scenario is still straightforward enough to guess the optimal next movement.

We identify two areas of future work: (1) a more fine-grained study on the types of goals that are shared; and (2) study of tasks in which communication is necessary to complete a task.

Acknowledgements. The authors thank Christian Muise for his valuable feedback on drafts of this paper. This research is partially funded by Australian Research Council Discovery Grant DP130102825: *Foundations of Human-Agent Collaboration: Situation-Relevant Information Sharing*.

References

1. Bradshaw, J.M., Feltovich, P., Johnson, M., Breedy, M., Bunch, L., Eskridge, T., Jung, H., Lott, J., Uszok, A., van Diggelen, J.: From tools to teammates: joint activity in human-agent-robot teams. In: Kurosu, M. (ed.) HCD 2009. LNCS, vol. 5619, pp. 935–944. Springer, Heidelberg (2009)
2. Cannon-Bowers, J.A., Salas, E.: Reflections on shared cognition. J. Organ. Behav. **22**(2), 195–202 (2001)
3. Clark, H.: Using Language. Cambridge University Press, Cambridge (1996)
4. Cummings, M.M.: Man versus machine or man + machine? IEEE Intell. Syst. **29**(5), 62–69 (2014)
5. Harbers, M., Bradshaw, J.M., Johnson, M., Feltovich, P., van den Bosch, K., Meyer, J.-J.: Explanation and coordination in human-agent teams: a study in the BW4T testbed. In: Proceedings of the 2011 IEEE International Conferences on Web Intelligence and Intelligent Agent Technology, pp. 17–20. IEEE Computer Society (2011)
6. Harbers, M., Bradshaw, J.M., Johnson, M., Feltovich, P., van den Bosch, K., Meyer, J.-J.: Explanation in human-agent teamwork. In: Cranefield, S., van Riemsdijk, M.B., Vázquez-Salceda, J., Noriega, P. (eds.) COIN 2011. LNCS, vol. 7254, pp. 21–37. Springer, Heidelberg (2012)
7. Harbers, M., Jonker, C., Van Riemsdijk, B.: Enhancing team performance through effective communication. In: Proceedings of the 4th Annual Human-Agent-Robot Teamwork Workshop, pp. 1–2 (2012)
8. Hindriks, K.V.: Programming rational agents in GOAL. In: El Fallah Seghrouchni, A., Dix, J., Dastani, M., Bordini, R.H. (eds.) Multi-Agent Programming, pp. 119–157. Springer, New York (2009)
9. Johnson, M., Bradshaw, J.M., Feltovich, P.J., Jonker, C.M., van Riemsdijk, B., Sierhuis, M.: The fundamental principle of coactive design: interdependence must shape autonomy. In: De Vos, M., Fornara, N., Pitt, J.V., Vouros, G. (eds.) COIN 2010. LNCS, vol. 6541, pp. 172–191. Springer, Heidelberg (2011)
10. Johnson, M., Bradshaw, J.M., Feltovich, P.J., Jonker, C.M., Van Riemsdijk, M.B.: Coactive design: designing support for interdependence in joint activity. J. Hum.-Robot Interact. **3**, 43–69 (2014)
11. Johnson, M., Jonker, C., van Riemsdijk, B., Feltovich, P.J., Bradshaw, J.M.: Joint activity testbed: blocks world for teams (BW4T). In: Aldewereld, H., Dignum, V., Picard, G. (eds.) ESAW 2009. LNCS, vol. 5881, pp. 254–256. Springer, Heidelberg (2009)

12. Kashima, Y., Klein, O., Clark, A.E.: Grounding: sharing information in social interaction. In: Fiedler, K. (ed.) Social Communication, pp. 27–77. Psychology Press, New York (2007)
13. Langan-Fox, J., Canty, J.M., Sankey, M.J.: Human-automation teams and adaptable control for future air traffic management. Int. J. Ind. Ergon. **39**(5), 894–903 (2009)
14. Puranam, P., Raveendran, M., Knudsen, T.: Organization design: the epistemic interdependence perspective. Acad. Manag. Rev. **37**(3), 419–440 (2012)
15. Saavedra, R., Earley, P.C., Van Dyne, L.: Complex interdependence in task-performing groups. J. Appl. Psychol. **78**(1), 61–72 (1993)
16. Salas, E., Prince, C., Baker, D.P., Shrestha, L.: Situation awareness in team performance: implications for measurement and training. Hum. Factors: J. Hum. Factors Ergon. Soc. **37**(1), 123–136 (1995)
17. Singh, R., Miller, T., Sonenberg, L.: A preliminary analysis of interdependence in multiagent systems. In: Dam, H.K., Pitt, J., Xu, Y., Governatori, G., Ito, T. (eds.) PRIMA 2014. LNCS, vol. 8861, pp. 381–389. Springer, Heidelberg (2014)
18. Stubbs, K., Wettergreen, D., Hinds, P.J.: Autonomy and common ground in human-robot interaction: a field study. IEEE Intell. Syst. **22**(2), 42–50 (2007)
19. Wei, C., Hindriks, K., Jonker, C.M.: The role of communication in coordination protocols for cooperative robot teams. In: International Conference on Agents and Artificial Intelligence, pp. 28–39 (2014)