Communication and Shared Mental Models for Teams Performing Interdependent Tasks

Ronal Singh(✉), Liz Sonenberg, and Tim Miller

University of Melbourne, Melbourne, Australia
ronals@student.unimelb.edu.au, {l.sonenberg,tmiller}@unimelb.edu.au

Abstract. Research shows that performance of human teams improves when members have a shared understanding of their task; that is, when teams develop and use a shared mental model (SMM). An SMM can contain different types of information or components and this paper investigates the influence on team performance of sharing different components. We consider two components of an SMM: intentions (e.g. goals) and world knowledge (e.g. beliefs) and investigate which component(s) contribute most to team performance across different forms of interdependent tasks. We performed experiments using a Blocks World for Team (BW4T) testbed for artificial agent teams and our results show that with high levels of interdependence in tasks, communicating intentions contributes most to team performance, while for low levels of interdependence, communicating world knowledge contributes more. Additionally, as is the case with human teams, higher sharedness correlated with improved team performance for the artificial agent teams. These insights can assist in the design of communication protocols that improves team performance when team members are engaged in interdependent tasks and help design artificial agents that can communicate effectively when working with humans as team mates.

Keywords: Task interdependence · Shared mental models · Joint action

1 Introduction

Agents perform tasks that range from independent tasks that does not require interactions with others to highly *interdependent tasks* requiring close and continuous interactions [14]. When faced with interdependent tasks, effective coordination and collaboration of team members become crucial. One of the key foundations of effective coordination and collaboration is having *shared mental models (SMM)*. Shared mental model has been defined as [1]: "knowledge structures held by members of a team that enable them to form accurate explanations

This paper was first published as *Singh R., Sonenberg L., Miller T.: Communication and Shared Mental Models for Teams Performing Interdependent Tasks. In: Osman N., Sierra C. (eds) AAMAS 2016 Workshops. LNCS(LNAI), vol. 10002, pp. 163–179. Springer International Publishing AG (2016).* doi:10.1007/978-3-319-46882-2_10.

S. Cranefield et al. (Eds.): COIN 2016 Workshops, LNAI 10315, pp. 81–97, 2017.
DOI: 10.1007/978-3-319-46882-2_10

and expectations for the task, and, in turn, coordinate their actions and adapt their behaviour to demands of the task and other team members".

More than a decade of research has correlated SMMs with improved team performance in human teams [12]. The basic assumption is that SMMs allow team members to anticipate the needs and actions of other members, thereby increasing team performance. Recent studies in human-agent and artificial agent teams have also found similar correlations [3,5]. SMMs can be broadly classified as either task work model or team work model. Task work concerns the task or job that the team is to perform, while team work concerns what has to be done in order to complete a task as a team [9]. SMMs can also be viewed as having different components [5,9], such as world knowledge and intentions. World knowledge includes knowledge of the current state of the environment and the team while intentions represent what the agents intend to do [4].

Four types of task interdependence have been identified for human activities: pooled, sequential, reciprocal, and team [14,15]. In sequential task interdependence, tasks are performed in a sequential order. For example, in a relay race each runner has to wait for the previous team member to pass on the baton. In reciprocal task interdependence, participants take their turn in completing part of the task. A key property associated with reciprocal task interdependence is interleaved execution: for example, surgical teams often work reciprocally. In team task interdependence, participants execute their individual tasks concurrently and may include *joint actions*. By "action", we mean the atomic actions that make up a task. In joint action, multiple participants execute a particular action concurrently, for example when two people lift a heavy object together. In pooled task interdependence, the participants can successfully execute tasks without any interaction with each other. Due to the simple nature of these tasks, we do not study such tasks in this paper. The four types of task interdependence forms a hierarchy of pooled-sequential-reciprocal-team, with this hierarchy representing increasing levels of dependence between team members as well as increasing needs for coordination [14].

While *sharedness* has been linked with better team performance, central to the notion of SMM is how much and what to share. There has been recent work investigating this question in multi-agent systems research, such as [5,11,17]. However, as far as the authors are aware, with the exception of Li et al. [10], studies in the related work only consider sequentially-interdependent tasks, rather than more tightly linked team and reciprocal tasks. A recent report [16] highlights the need for studies considering other types of interdependence, notably *intensive* task interdependence – a type that we characterise as a *joint action.*

The subject of this paper is the communication content, specifically *what* to share when team members engage in interdependent tasks. We investigate the influence of the two components of the SMM (world knowledge and intentions) on the team performance across different forms of interdependent tasks. We used search and rescue like scenarios for a team of artificial agents for the experiments. The scenarios were generated using a Blocks World for Teams (BW4T) testbed [8]. In BW4T, which is an extension of the classical blocks world domain, the teams' joint task is to find and deliver coloured blocks in a particular order. Using the testbed, we designed and executed two sets of experiments. The first set studies the influence of sharing the two components – world knowledge and intentions – on the team performance for each form of task interdependence. The second set introduces joint actions within sequential and reciprocal tasks and studies the influence of sharing the two components on team performance. Introduction of joint actions allows for a shift from sequential or reciprocal to team task interdependence where members execute individual actions concurrently.

The outline of the paper is as follows. Section 2 introduces SMM, along with related work. Section 3 describes the task and the testbed and provides the details of the artificial agents that we implemented. Section 4 details the experimental setup while Sect. 5 discusses the results. Sections 6 and 7 conclude the paper with a discussion.

2 Background and Related Work

Mental models are simplified representations used by individuals to explain and predict their surroundings [13]. These models comprise content and structure or relationships between the content. In addition, individuals can simultaneously hold multiple mental models. In a team setting, when team members interact, their mental models converge resulting in shared mental models.

To extend the concepts of SMM that has been well studied for human teams [12] to human-agent teams, Jonker et al. [9] proposed mental model ontologies. They view a team as a system. A team performs team activities and has physical components, e.g. team members. A team member is an agent with a mind comprising many mental models: all but one of which represent the mental models of others in the team. Based on this conceptualisation, they proposed a measure that could be used to assess the similarity or the overlap of agents' mental models. We discuss this measure in the next section.

2.1 Measuring SMM

While several methods exist for measuring SMMs for human teams [2], one for teams comprising artificial agents is Jonker et al. [9]. Harbers et al. [5] extended Jonker et al's similarity measure so that it could be applied to teams of agents and performed experiments to show that their similarity measure can be used to predict team performance. We discuss the extended version of the measure next. In the following discussions, similarity refers to the overlap of the mental model contents of the agents. We consider the SMM to be made of two components – world knowledge and intentions.

Figure 1 shows an example of SMM. Assume Bot 1 and Bot 2 are two agents engaged in a joint task. Each has its mental model. While engaged in their task, the agents may communicate their beliefs and goals, making their own beliefs and goals known to others. For example, notice that each agent has it's own

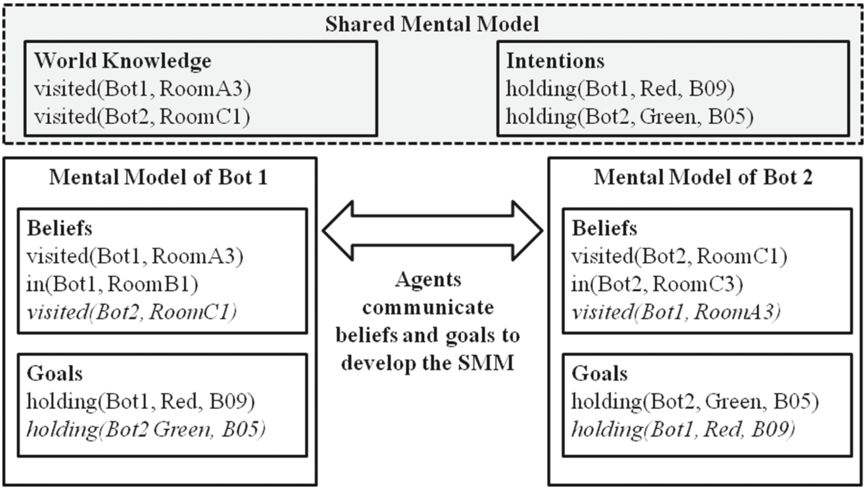

Fig. 1. Example SMM. The beliefs and goals of other agents are shown in italics. An agent has certain beliefs and goals that it is not required to communicate, e.g. *in (agent, room)*, and these may not part of the SMM.

as well as others' beliefs and goals, which are shown in italics. The SMM is a theoretical construct that can be used to represent the overlapping content of the mental models of the two agents. In the example, the SMM is composed of the components - world knowledge (beliefs) and intentions (goals).

Jonker et al. [9] and Harbers et al. [5] proposed a compositional measure of sharedness. We reproduce their definitions here with some simplifications. They view SMMs as having components, which can include sub-components. For example, Fig. 1 shows an SMM with two components. Examples of sub-components can be found in Sect. 3.3. The (sub)components can be queried by posing questions that all team members should be able to answer. The answers are used to compute the model agreements, which is a measure of the similarity of the answers provided by each agent for each question. Formally, let M be the set of all mental models, Q be the set of all questions, and $ans(m, q)$ be the answer of model $m \in M$ with respect to question $q \in Q$. The agreement between models M for questions Q is:

$$Ag(M, Q) = \frac{1}{|Q|} \sum_{q \in Q} \frac{| \cap_{m \in M} ans(m, q)|}{| \cup_{m \in M} ans(m, q)|} \tag{1}$$

If $| \cup_{m \in M} ans(m, q)| = 0$ then the agreement for question q is 0. Given a set of agents A, a set of mental models M_A (a model for each agent), and questions Q, we say that the model m is shared *to the extent* θ, denoted by $Sh(M, A, Q, \theta)$, with respect Q, iff $Ag(M_A \cup \{m\}, Q) >= \theta$. The compositional measure CS is:

$$\begin{aligned} CS(M, A, Q) &= max\{\theta \mid Sh(M, A, Q, \theta)\}, \text{ if M is not composed} \\ CS(M, A, Q) &= c(\{CS(m, A, Q) \mid m \in M\}), \text{ if M is composed} \end{aligned} \tag{2}$$

Where m is a component of M and c is composition function, for example: $\sum_{m \in M} w_m CS(m, A, Q)$. Each component and sub-component can be weighted to model the relevance of each (sub)component. The weight of each (sub)component is $w_m \in [0, 1]$ and CS can be normalised to $[0, 1]$ by setting $\sum_{m \in M} w_m = 1$.

2.2 SMM and Task Interdependence

Interdependence is the central organising principle of *Coactive Design Method*, from Johnson et al. [7], which is a method aimed at designing systems in which humans and agents collaborate as teammates. They define interdependence as relationships between members of a team, and argue that these relationships determine what information is relevant for the team to complete (interdependent) tasks, and in that sense, the interdependent relationships define the *common ground* that is necessary. A number of studies have considered some of the different forms of task interdependence [5, 10, 17], and some have also measured sharedness [5, 9]. Generally, higher sharedness of mental models produces better team performance. For example, Harbers et al. [5] found higher sharedness correlated with better team performance. In their work, SMM were composed of world knowledge and intentions, which is how we view SMM in this work. Similarly, task interdependence has naturally been part of these studies. However, almost all involve sequentially interdependent tasks. The exception is Li et al. [10], who introduced joint action in sequentially interdependent tasks and Wei et al. [17], who studied tasks that were not very strongly sequential. They did this by creating subtasks that multiple agents could complete simultaneously. None of these have explicitly employed reciprocally interdependent tasks.

Mixed results have been reported for studies involving sequentially interdependent tasks in terms of which type of information or component contributes more to team performance, that is task completion times. Harbers et al. [5] reported that when agents communicated their intentions with others, the team performance improved more than if they shared world knowledge. However, Wei et al. [17] reported that beliefs contributed more to team performance than goals. While [17] did not measure sharedness, they view the agents mental models to comprise of two components, goals (intentions) and beliefs (world knowledge). We perform further experiments involving sequentially interdependent tasks and may help explain the difference between the two studies.

In a separate study, Li et al. [10] introduced joint action in sequentially interdependent tasks. They studied search and retrieval tasks using the BW4T testbed. In one setup, agents collaborated on a task in which some blocks were heavier, and required two agents to collect. The agents exchanged goals, beliefs, and both. Their experiments revealed that with joint actions, exchanging goals improved team performance, measured as completion time, more than sharing beliefs only. When agents shared their goals that fulfil the current team sub-goal with others, the other team members could start on a new task. This allowed the team to finish the team task more quickly.

These works show that sequentially interdependent tasks have been investigated, but other forms of task interdependence have not. This work aims to fill that gap.

3 Scenario: Blocks World for Teams

We used a BW4T testbed [8] for our experiments. As explained next, we modified the testbed to be able to setup tasks with joint actions.

Basic BW4T: In BW4T, teams find and deliver coloured blocks in a particular order. The environment has a set of rooms, each containing coloured blocks, and a drop zone. The agents search the rooms, find the required blocks and drop these in the drop zone. Agents have a map of area but do not know the location of the required blocks. Agents have to go to each room to perceive the blocks that are present in it. Agents cannot see each other but can communicate with others. A simplified map is shown in Fig. 2. Each room has one door. The teams' joint task, i.e. the sequence of colours, is displayed at the bottom left. A black triangle appears on top of a colour if the colour is dropped off. The room above the joint task is the drop zone. The agents are represented by either black squares or the colour the agent is holding, and their names are displayed in red. The basic version is well suited to perform experiments for sequential and reciprocal tasks. However, it does not explicitly support joint actions.

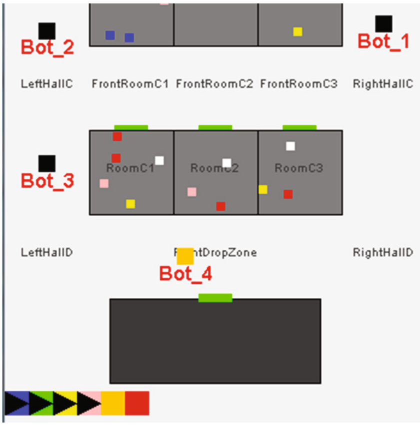

Fig. 2. Sample BW4T environment.

Modified BW4T: To design joint tasks that would be a fair representation of the different forms of task interdependence, we modified the testbed. In the original version, only one agent could be in a room at any one time. To implement joint actions, we follow Li et al. [10] and introduce "heavy blocks", which required two agents to carry to the drop zone. This means in our version, two agents can carry the same block simultaneous, and therefore can be in the same room at the same time. Secondly, for team task interdependence, the blocks could be delivered in any order, that is, we removed the sequential delivery requirement.

3.1 Task Design

We designed tasks to be able to test the effects of communication content on the team performance for each type of the task interdependence as well as later include joint action within other forms of task interdependence and test the effect of communication content on the team performance for each combination. Variations of two basic joint tasks (Fig. 3) has been used to realise the different forms of task interdependence.

(a) Task 1 (b) Task 2

Fig. 3. Basic joint tasks used to simulate different types of task interdependence (Color figure online).

Team Task: In team tasks, agents execute their actions concurrently. The joint task had some *heavy* blocks. The heavy blocks required one agent to help the other lift it, and afterwards the first agent delivers it to the drop zone. The act of lifting the heavy block *together* is the joint action. Additionally, the agents could lift any colour. Consider the task shown in Fig. 3a. In this task, agents can lift both colours. The red blocks are heavy blocks. In order to remove the underlining sequential interdependence from this task, the agents could deliver the blocks in any order, for example, the second (red) block can be delivered before the first (yellow) block. Green, pink and red are heavy blocks in Task 2.

Reciprocal Task: In a reciprocal task, each agent takes it's turn in completing part of the task. In this task, the agents deliver a sequence of alternating colour sets in the order the colours appear in the task. Furthermore, each agent can lift colours from only one of the two distinct colour sets. Consider the task shown in Fig. 3a. For this task, one agent would be delivering yellow blocks while the other red ones. The blocks must be delivered in the order they appear. This means that agent delivering the red block now depends on the agent delivering the yellow blocks and vice-versa, making them reciprocally interdependent.

Sequential Task: In sequential task, the first three colours are delivered by one agent while the remaining three by another agent. The blocks must be delivered in the specified order, but the second agent is free to search for its coloured blocks while the first agent is delivering.

3.2 Agent Teams and Agent Behaviours

We had two team compositions; (1) 2-agent team and (2) 4-agent team. The 4-agent team was a 2×2-agent team, i.e. 2 sub-teams of 2 agents each.

This composition was required for certain tasks, such as reciprocal tasks in which we needed to have at least one agent for each of the two colour sets.

Agents were programmed in GOAL [6]. The BW4T testbed provides interfaces that enable GOAL agents to interact with it. Using these interfaces, the agents can perceive specific details of the environment, such as the blocks present in rooms, and can perform actions, such as picking up a block. The abstract decision cycle of an agent is shown in Fig. 4. The basic steps each agent takes are: (*1*) decide the colour to search for; (*2*) choose a room; (*3*) go to and search room; (*4*) if required block is found and is not heavy, pick it up; (*5*) if required block is found and is heavy, ask for help and wait. When help arrives, pick up the block; (*6*) deliver the block to the drop zone; (*7*) if help is requested, go to the particular room and help lift the heavy block.

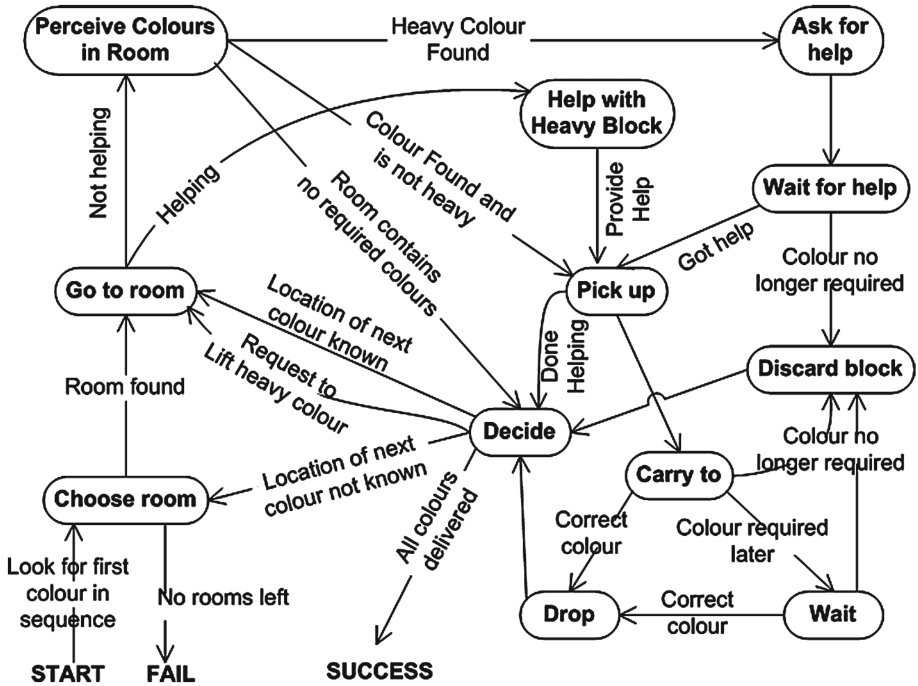

Fig. 4. Abstract decision cycle of an agent.

Initially, agents start searching for the first undelivered colour. However, agents use a two-block look-ahead protocol to determine which colour to deliver. If an agent knows the location of the first undelivered colour and has the intention of collecting it, remaining agents search for the second undelivered colour. If the one or more of the remaining agents know the location of the next required colour, they go to that room. However, only one will be able to collect the block. When the first colour is picked up, one agent collects the second colour while others start searching for the third colour. The aim of this is to ensure that sufficient time is dedicated to search. When required to lift a heavy block, an agent only asks for help when it is physically present at the heavy block. Other (helper) agents could potentially infer that help will be required soon and go to the location of the heavy block before the agent actually asks for help because the agent may tell others that it has the goal of going to the (heavy) block. However, our agents do not perform this level of reasoning and only go to help when asked. Furthermore, if one agent asks for help, all agents that are waiting to drop a block at the drop zone or those that are currently searching for their block will go to help. If the agent knows that the colour that it is searching for, has the intention of holding or is holding is no longer required, then it will discard the colour and go on to deciding what it will do next. Rooms are chosen randomly and the agents avoid visiting a room more than once unless the room contains multiple required blocks.

While the basic behaviours of agents are almost the same across the different forms of task interdependence, there are differences in the way agents reason about which colour to search for:

(1) Sequential and reciprocal tasks: Agents choose the first undelivered colour. If another agent has the goal of holding this colour, the agent chooses the next undelivered colour.
(2) Team Task: Blocks can be delivered in any order. Therefore, agents do not reason about when the block has to be delivered. Instead agents have to determine whether the block is heavy and ask for help.

While certain aspects of agent behaviours are different because of task interdependence, there are differences because of what the agents share with each other. Therefore, while the basic decision cycle shown in Fig. 4 is used by all agents, there are some variations in their implementation. The implementation has been guided by what the agents actually do with the information they receive and has been described later in Sect. 3.4. Therefore, if only one component is exchanged, the agent performs reasoning described for that component only.

3.3 Communication and SMM

Agents exchange messages that are indicative of the world knowledge and the intentions. To develop the shared mental model, agents communicate as soon as they have the required information. Agents exchange six sub-components, three each of goals and world knowledge. These sub-components were selected based on prior research work [5,10] and preliminary experiments revealed that each sub-component had the potential to improve team performance. The sub-components are communicated as messages, which are discussed next. The keyword *imp* stands for imperative and indicates what the agent intends to do.

The messages indicative of intentions are:

(1) imp(in(Sender, Room)): Sender intends to visit Room.
(2) imp(holding(Sender, Colour, Block)): Sender intends to collect Block of Colour.
(3) delivered(Sender, Colour, Block): Sender has delivered Block of Colour - implies agent has dropped current goal and may have a new goal.

The messages indicative of world knowledge are:

(1) blockLoc(Sender, Block, Colour, Room): Sender has perceived Block of required Colour in Room.
(2) pickedUp(Sender, Colour, Block): Sender has picked up Block of Colour.
(3) visited(Sender, Room): Sender has visited Room. This message is sent irrespective of whether room contains required blocks.

3.4 Using Shared Mental Models

Agents employ the following policies to SMM to choose their activities such that it prevents potential conflicts with the activities of others. The following outlines how the agents use the components of the shared mental model. We chose a straightforward use of each intention and world knowledge, which was sufficient to test the effect of the component on the team's performance and avoids side-effects that would have been introduced because of using more complex mechanisms. The intentions are used as follows:

(1) An agent will not adopt a goal to go to a particular room if another agent has the goal of going to that room. For reciprocal task, this logic applies when both agents are delivering blocks from the same colour set, that is in a 4-agent team and not in a 2-agent team.
(2) An agent will not adopt a goal to hold a block that has been delivered.
(3) An agent will not adopt a goal to hold a block/colour that another agent has the goal of holding — *unless* the block is heavy (both agents need to lift it together). For reciprocal tasks, this logic is applicable in a 4-agent team.

World knowledge is used as follows:

(1) An agent will not search for a colour if this been found by another agent.
(2) An agent will search for the next colour if the currently required colour has been picked up.
(3) An agent will not search a room that another agent has already searched.

Agents employ the above policies to SMM to reduce interference and duplication of effort. However, the agents have their own decision processes and may make decisions simultaneously. This may result in instances where the agents may adopt similar goals, for example to look for the same colour. Like Wei et al. [17], we simply implement a "first-come first-served" policy instead of implementing detailed negotiation mechanisms to assist agents resolve these issues.

4 Experiment Design

We ran a series of simulation experiments, measuring the following:

(1) Completion Time: Time it takes the team to complete the task. We used this measure as a proxy for team performance.
(2) Number of messages: We measured the total number of messages exchanged by the agents. We also counted the number of messages per component. These measures are indicative of the communication cost.
(3) Sharedness: We measured the sharedness of the agents' mental models. This is a compositional measure (see Sect. 2.1) and was calculated at the time any block was delivered to the drop zone. When one agent drops off a correct colour in drop zone, all agents log their belief and goal bases. These logs are then analysed to find the overlapping content, which is used to compute the sharedness values. The two components had a weight of 0.5 and each of the three sub-components had a weight of 0.33. In experiments where only one component was measured, the weight of the component was set to 1, and only questions related to that component were asked.

In case of sub-teams, we also measured the number of messages and sharedness of the agents with each sub-team.

Independent Variable. The independent variable is the component of the SMM. This variable has three values (see Sect. 3.3): (*1*) World Knowledge (WK); (*2*) Intentions (INT); and (*3*) World Knowledge and Intentions (ALL).

Table 1. Experimental setups (S1 – S6) for each type of task interdependence.

	Set 1				Set 2	
Team size	2 agents		4 agents		4 agents	
Map	1	2	1	2	1	2
Setup	S1	S2	S3	S4	S5	S6

Setup. We used two different maps, one for each task outlined in Sect. 3.1. Variations of each task gave us three different task interdependence types. We refer to Task 1 (Fig. 3a) as Map 1 and Task 2 as Map 2. The setups are as shown in Table 1. We had two sets. In set 1, we had four setups (S1-S4) (both maps combined with two team compositions) for each of the three types of task interdependence giving us 12 combinations.

Set 2 has two setups, S5 and S6, representing reciprocal and sequential tasks with joint actions respectively. Here the sub-teams were reciprocally or sequentially interdependent and were required to lift heavy blocks. We tested the effect of the SMM components on completion times by employing three communication strategies: (*1*) ALL-ALL: where agents exchanged the two components with every other agent. (*2*) WK-Within: where agents shared world knowledge within each sub-team but shared intentions with all agents. (*3*) INT-Within: where agents shared intentions within each sub-team but the world knowledge with all agents. For these two setups, we only used a 4-agent team because a 2-agent team would not have enabled us to fully test the effects of the two components. For example, we needed to have at least 2 agents in each sub-team to be able to test the effect of sharing a component within the sub-team. Combining S5 and S6 with the two types of task interdependence (sequential and reciprocal) gave us further 4 combinations, and a total of 16 combinations.

Combining each of the 12 combinations from Set 1 with the three components of the SMM (ALL, INT, WK) and the 4 combinations from Set 2 with the three communication strategies (ALL-ALL, WK-Within, INT-Within) resulted in 48 combinations in total. Each combination was run 30 times resulting in 1440 runs. Each map had 25 blocks pre-allocated to rooms and further 10 blocks were randomly generated giving a total of 35 blocks for each run. Each map had 9 rooms, 1 drop zone and 6 blocks in the joint task. Statistical significance tests were conducted using Wilcoxon rank-sum (WRS) and Kruskal-Wallis (KW) tests.

5 Results

This study was aimed at identifying the components that contributed most to team performance across different forms of task interdependence. Recall that going from sequential to team tasks represents increasing levels of dependence between agents as well as coordination requirements. For simplicity, we collapse the results of the two tasks (shown in Fig. 3) and report the averages.

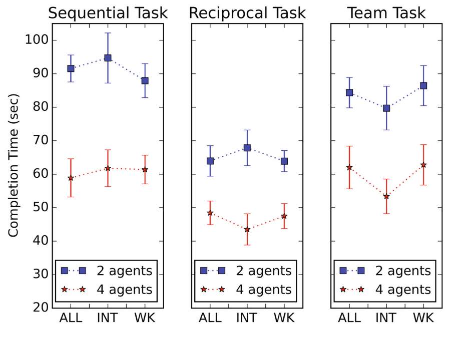

(a) Performance

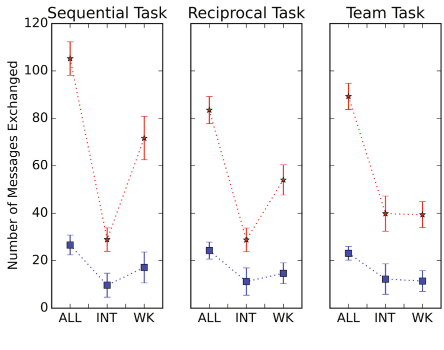

(b) Communication Cost

Fig. 5. Performance and communication cost for different forms of task interdependence. The communication cost is expressed as the average number of messages exchanged by all team members. Error bars represent on standard deviation.

5.1 SMM Components and Team Performance

Figure 5a shows the average task completion times for the 2-agent and 4-agent teams performing different tasks. These results are for experiments resulting from setups S1-S4. Recall that a 4-agent team comprises 2 sub-teams of 2 agents each. For team tasks, the intentions contributed more to team performance than world knowledge. This finding is significant at 5% for all except two combinations and consistent for both team compositions. In the team task, some blocks were heavy and the agents could pick any colour. In such scenarios, knowing the intentions of team members allows agents to avoid duplicating their activities, therefore reducing interference. These results are in line with Li et al. [10], who reported that with joint actions, exchanging goals results in improved completion times.

However, for sequential and reciprocal tasks, different trends have been observed between 2-agent and 4-agent teams. For sequential tasks and 2-agent team, the world knowledge contributed significantly more ($p < 0.05$) than intentions in terms of task completion times. In this task setting, the first agent delivered first three blocks while the remaining three by the other agent. Because agents had separate sub-tasks, exchanging world knowledge helped the other agent find it's required blocks faster. However, for reciprocal tasks, this difference was less pronounced. We discuss this more later.

In 4-agent teams performing sequential tasks, no significant difference in terms of completion times were noted between the two components. However, it is worth noting that moving from 2-agent to 4-agent team, the importance of intentions increases. A similar trend occurs for reciprocal and team tasks. In these team settings, the agents within each sub-team could choose conflicting goals, for example, choosing the same block to deliver. By exchanging intentions, agents within sub-teams avoided duplicating their activities, therefore improving the completion times.

To make these trends clearer, we computed *component influence* (CI) for each task. CI is computed based on the difference between the completion times achieved when communicating both components and any one of the two components. To normalise the difference between completion times across different experiments, we used the *tanh* function. The CI for component c is:

$$CI_c = tanh(CompletionTime_{all} - CompletionTime_c)$$

The resulting values were normalised to between 0 and 1 using $CI_{normalised} = (CI - min(CI))/(max(CI) - min(CI))$. Figure 6 for 2-agent teams show that with increasing dependence between agents, that is, going from sequential to team interdependence, the importance of intentions increases while the importance of world knowledge decreases. For 4-agent team, the intentions were almost always more important than world knowledge.

SMM and Joint Actions. Results of experiments relating to setups S5 and S6 indicated that the difference between completion times of WK-Within and INT-Within is significant (p-value = 0.009) in favour of INT-Within. This indicates that sharing intentions within sub-teams and world knowledge with everyone

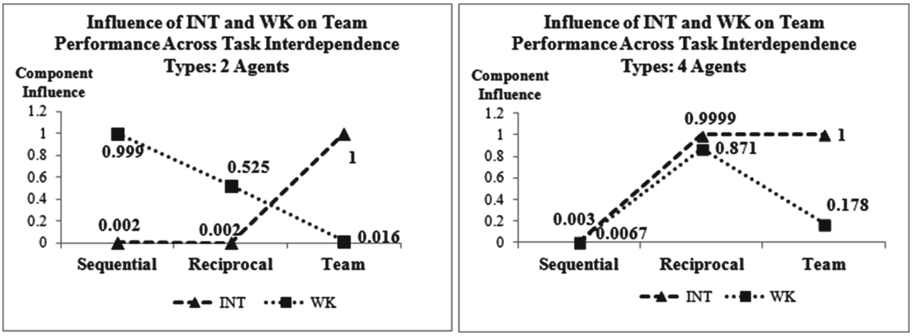

Fig. 6. The two graphs (2 agents and 4 agents) show that intentions become more important more as the level of interdependence increases and as the number of agents in each sub-team increases.

achieves the best team performance. This is consistent with our earlier findings that intentions and team tasks are positively correlated. Also, world knowledge and sequential and reciprocal tasks are positively correlated.

5.2 Communication Performance

Figure 5b shows the communication cost (average number of messages exchanged). The number of intentions exchanged was significantly lower ($p < 0.05$) than world knowledge for about two-thirds of the combinations. This indicates that agents generally have more information to communicate about the world than their intentions. There was no correlation between the number of messages and team performance. More communication resulted in worst performance in some cases, particularly for larger teams. This is due to the two-block look-ahead policy. When agents exchange information about possible blocks, in larger teams this often results is agents trying to collect the same block/colour and this increases the completion time. When agents only exchange intentions, all agents are required to find the blocks themselves, and so search randomly, thus reducing the number of unnecessary runs for the same block/colour. While it is clear that a mechanism could be designed to improve this by using a different look-ahead policy, we believe our policy is reasonable. Importantly though, this result shows that simply throwing more information towards agents can result in worse performance if significant thought is not given to how that information is used.

5.3 Analysis of Sharedness

We computed the sharedness in relation to each component at the time a block was delivered to the dropzone. Generally, higher sharedness correlates with improved completion times. For simplicity we show the data for team task and

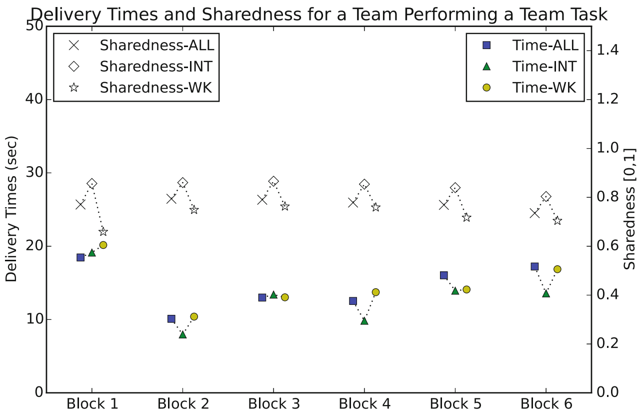

Fig. 7. Sharedness and delivery times for a 2-agent team engaged in a team task.

note that the results for sequential and reciprocal tasks are similar. For example, Fig. 7 shows the sharedness at the time each correct block is dropped off for team tasks. The plotted delivery times are the time differences between block deliveries. For team tasks, exchanging intentions achieved the best completion times and the sharedness was highest for this component. Notice that in Fig. 7, sharedness of intentions is highest across all six blocks and the delivery times when teams exchange intentions are fastest across most of the six blocks.

Sharedness and Sub Teams. We measured the sharedness of members within each sub-team for tasks solved by 4-agent teams. Sharing intentions resulted in the best completion times and the sharedness of intentions was highest for reciprocal and team tasks. For sequential tasks, we noted a significant increase in the importance of intentions compared to 2-agent team. This supports the finding that higher sharedness results in better completion times. The other consistent finding is that in situations where we may have members of sub-teams potentially duplicating their efforts, sharing intentions with each other helps avoid such conflicting actions and therefore, improves the completion times.

6 Discussion

We intended to identify the components contributing most to team performance across the different forms of task interdependence. Our results show that as the interdependence increases, the importance of intentions to team performance also increases. These results are in line with [5,10] who found that when team members exchanged intentions, the team performance improves. In [5], teams were engaged in sequential tasks and their team composition was similar to our 4-agent team while in [10], the authors introduced joint actions in sequentially interdependent tasks.

While our results are in line with the above works, we have observed that when team members can perform their sub-tasks independently, e.g. in 2-agent teams, exchanging world knowledge contributes more to team performance for sequential and reciprocal tasks. This makes sense intuitively: if other members provide potentially useful information, such as location of blocks that one is required to deliver, the team performance improves. This is a form of *soft interdependence* [7] where one team/member 'helps' another voluntarily. In case of 4-agent teams, we found that intentions contributed more to team performance across all forms of task interdependence. This indicates that team composition plays a role in which component is important to team performance.

Our findings that are partially consistent with [17] who found that for sequential tasks, beliefs contributed more to team performance. While this is consistent with the results of our 2-agent team, we noted a marked increase in the importance of intentions when 4-agent team was concerned. These differences may hinge on other factors, such as how effectively the agents use the information that it receives. This is an area of future work.

Finally, our findings are consistent with others (e.g. [5]) in terms the role SMM plays in improving team performance. Across all tasks and both team compositions, higher sharedness of SMM resulted in improved team performance.

7 Conclusions and Future Work

The four types of task interdependence form a hierarchy, from pooled to team, representing increasing levels of dependence between team members as well as increasing needs for coordination. We found that with increasing levels of interdependence, the importance of intentions increases as well. Team composition also plays a role in which component contributes more to team performance. In team compositions, where agents can perform their tasks independently, e.g. in sequential and reciprocal tasks, world knowledge contributed more to team performance. When multiple team members may be engaged in a single sub-task, the potential of interference increases and so does the importance of knowing the intentions of others.

A factor to investigate further is the reasoning capability of the agents; that is, how the agents reason with information that they receive from others. We also have not explicitly analysed the behavioural changes in the agents when agents switch from one task interdependence type to another, making this another opportunity for future investigation.

References

1. Cannon-Bowers, J.A., Salas, E., Converse, S.: Shared mental models in expert team decision making. In: Castellan, J., John, N. (eds.) Individual and Group Decision Making: Current Issues, pp. 221–246. Lawrence Erlbaum., Hillsdale, NJ (1993)
2. DeChurch, L.A., Mesmer-Magnus, J.R.: Measuring shared team mental models: A meta-analysis. Group Dyn. Theor. Res. Pract. **14**(1), 1 (2010)
3. Fan, X., Yen, J.: Modeling cognitive loads for evolving shared mental models in human-agent collaboration. IEEE Trans. Syst. Man Cybern. **41**(2), 354–367 (2011)

4. Harbers, M., Jonker, C.M., Van Riemsdijk, M.B.: Context-sensitive sharedness criteria for teamwork. In: Proceedings of AAMAS 2014, pp. 1507–1508. ACM (2014)
5. Harbers, M., Riemsdijk, M.V., Jonker, C.: Measuring sharedness of mental models and its relation to team performance. In: Proceedings of 14th Workshop on COIN, pp. 106–120. Verlag (2012)
6. Hindriks, K.: Programming rational agents in GOAL. In: Seghrouchni, A.E.F., Dix, J., Dastani, M., Bordini, R.H. (eds.) Multi-Agent Programming, pp. 119–157. Springer, US (2009)
7. Johnson, M., Bradshaw, J., Feltovich, P., Jonker, C., Riemsdijk, M.B.V.: Coactive design : Designing support for interdependence in joint activity. J. Hum. Robot Interact. **3**(1), 43–69 (2014)
8. Johnson, M., Jonker, C., 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). doi:10.1007/978-3-642-10203-5_26
9. Jonker, C.M., Riemsdijk, M.B., Vermeulen, B.: Shared mental models. In: Vos, M., Fornara, N., Pitt, J.V., Vouros, G. (eds.) COIN -2010. LNCS (LNAI), vol. 6541, pp. 132–151. Springer, Heidelberg (2011). doi:10.1007/978-3-642-21268-0_8
10. Li, S., Sun, W., Miller, T.: Communication in human-agent teams for tasks with joint action. In: Dignum, V., Noriega, P., Sensoy, M., Sichman, J.S.S. (eds.) COIN 2015. LNCS (LNAI), vol. 9628, pp. 224–241. Springer, Cham (2016). doi:10.1007/978-3-319-42691-4_13
11. Manner, M.D., Gini, M.: Improving agent team performance through helper agents. In: Dignum, F., Brom, C., Hindriks, K., Beer, M., Richards, D. (eds.) CAVE 2012. LNCS (LNAI), vol. 7764, pp. 89–105. Springer, Heidelberg (2013). doi:10.1007/978-3-642-36444-0_6
12. Mohammed, S., Ferzandi, L., Hamilton, K.: Metaphor no more: A 15-year review of the team mental model construct. J. Manag. **36**(4), 876–910 (2010)
13. Rouse, W.B., Morris, N.M.: On looking into the black box: Prospects and limits in the search for mental models. Psychol. Bullet. **100**(3), 349 (1986)
14. Saavedra, R., Earley, P.C., Van Dyne, L.: Complex interdependence in task-performing groups. J. Appl. Psychol. **78**(1), 61–72 (1993)
15. 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 (LNAI), vol. 8861, pp. 381–389. Springer, Cham (2014). doi:10.1007/978-3-319-13191-7_31
16. Smith-Jentsch, K.A.: On shifting from autonomous to interdependent work. Technical report, NASA, Houston, Texas (2015)
17. Wei, C., Hindriks, K., Jonker, C.M.: The role of communication in coordination protocols for cooperative robot teams. In: Proceedings of International Conference on Agents and Artificial Intelligence, pp. 28–39 (2014)