Towards Measuring Sharedness of Team Mental Models by Compositional Means

Catholijn M. Jonker[1], Birna van Riemsdijk[1], Iris C. van de Kieft[1], Maria Gini[2]
[1]Delft University of Technology, Man-Machine Interaction Group, Faculty of EEEMCS, Mekelweg 4, 2628 CD, Delft, The Netherlands
[2]University of Minnesota, Department of CS&E, 4-192 Keller Hall, 200 Union St SE, Minneapolis, MN 55455, USA
{c m.jonker, m.b.vanriemsdijk, i.c.vandekieft}@tudelft.nl, gini@cs.umn.edu

Abstract

The better the team mental model, the better the teamwork. An important aspect of what determines a good team model is the extent to which the model is shared by the team members. This paper presents suggestions for measuring the extent to which teams have a shared mental model and describes how these measures are related to team performance. The most promising measures of sharedness proposed so far rely on using a compositional approach for team modeling and on a situation-sensitive relevance relation that indicates to what extent components contribute to team performance. A case study illustrates the approach and initial results on measuring performance when teams use different levels of sharedness.

Introduction

A Shared Team Mental Model (STMM) is considered beneficial for teamwork, see e.g., (Bolstad and Endsley, 1999; Cannon-Bowers et al., 1993; Lim and Klein, 2006; Mathieu et al., 2000). Several aspects are relevant for working in a team. In this paper the distinction made in (Cannon-Bowers et al., 1993; Mathieu et al., 2000) is taken as a point of departure: equipment, task, team interaction, and team members. The better the shared mental model the better a team is capable of performing its tasks, even in unforeseen circumstances. However, so far, it is unclear how and when to measure the extent to which teams have a shared mental model and even more unclear is how sharedness relates to team performance. Therefore, the notions of team model, measures of sharedness, and team performance need formalization.

In (Jonker et al., 2010), the authors introduce a formalization of shared mental models. Definitions are presented that help to determine the extent of similarity of models: subject overlap, compatibility, agreement, and the extent of similarity. These definitions are formulated in terms of the set of questions that are relevant for the goal for which the models are to be used.

The existing measure (Jonker et al., 2010) defines a measure in terms of the percentage of questions answered in the same way by the agents. This paper shows that this measure cannot be simply extended to cover all aspects of a team mental model. As a way to resolve this issue a compositional approach is presented to construct models that underlie teamwork. This compositional approach enables the formulation of a compositional measure of sharedness of team models. We identify the following components as relevant to shared mental models for teamwork: domain model, competence/capabilities model, and organizational model. Each of these can, in turn, be composed of smaller models as explained in the article.

A relevance relation is associated with the composition relation, to indicate which to what extent the components contribute to team performance. This relevance relation depends on the current circumstances, team task, the domain, and the performance criteria. The relevance relation is used to formulate a compositional measurement of sharedness of mental models. Experiments are done within the context of a case study in the Blocks World For Teams (BW4T) as introduced in (Johnson et al., 2009) to illustrate the formalizations and test the adequacy of the compositional approach and the proposed measures of sharedness.

This paper is organized as follows. After a brief introduction to shared mental models, we focus on the components of team mental models and their compositional structure. This is followed by an introduction to the testbed (BW4T) we use to measure the performance of different types of agents that share different types of team mental models and show how the sharing affects performance.

Shared Mental Model

The measure of sharedness of (Jonker et al., 2010) is defined in terms of a definition of model subject overlap and a definition of model agreement.

Subject overlap provides a measure for the extent to which models provide answers to the set of relevant questions Q. These answers may be different, but at least an answer should be given ("unknown" is not considered to be an answer). For example, posing a question about bicycles to a model of birds would typically not yield an answer. Also, answers are assumed to be individually consistent.

Definition 1 Model Subject Overlap (Jonker et al., 2010). Let the set of questions for which the models provide answers (not necessarily similar answers) be OverAns(M1,M2,Q) = {q ∈ Q | ∃a1, a2 : M1 |- answer(a1, q) and M2 |- answer(a2, q)}. The level of subject overlap between the model M1 and M2 with respect to set of questions Q is defined as

SO(M1,M2,Q) = |OverAns(M1,M2,Q) | / |Q|.

Model agreement defines the extent to which models provide equivalent answers to questions. In the definition T refers to a background theory used for interpreting answers, and $\equiv_T$ stands for equivalence with respect to T. For example, the answers "1 meter" and "100 centimeter" are equivalent with respect to the usual definitions of units of length.

Definition 2 Model agreement (Jonker et al., 2010). Let the set of questions for which the models agree be AgrAns(M1,M2,Q) = {q ∈ Q | ∃a1, a2 : M1 |- answer(a1, q) and M2 |- answer(a2, q) and a1 $\equiv_T$ a2}. Then, the level of agreement between the model M1 and M2 with respect to set of questions Q is defined as:

A(M1,M2,Q) = |AgrAns(M1,M2,Q) | / |Q|.

Definition 3 Shared Mental Model (Jonker et al., 2010). A model M is a mental model that is shared to the extent θ by agents A1 and A2 with respect to a set of questions Q iff there is a mental model M1 of A1 and M2 of A2, both with respect to Q, such that

1. SO(M,M1,Q) = 1, and SO(M,M2,Q) = 1
2. A(M,M1,Q) ≥ θ, and A(M,M2,Q) ≥ θ

In the remainder of the paper Definition 3 is used to determine the sharedness of models. However, what should be part of the team mental model and should be shared is described in the next section. In order for the definitions to work, it is assumed that the model corresponds to a finite ground model. If not, then in the definitions dividing by the number of relevant questions causes problems.

Team Mental Models

Team models should at least distinguish equipment, task, team interaction, and team members (Cannon-Bowers et al., 1993; Mathieu et al., 2000). This section adds some components, and then discusses possible compositional structures of team mental models in relation to the influence of the situation on the relevance of the components for team performance.

Components of Team Mental Models

Equipment, task, team interaction, and team members should be part of the team mental model according to (Cannon-Bowers et al., 1993; Mathieu et al., 2000). This paper proposes a restructure of the team mental model and to add information and knowledge about the domain (world and agents), and to add the underlying ontologies for the whole team mental model as elements of the team mental model. The usual components that refer to task, team interaction, and team members will be part of the organizational specification.

Ontologies: referring to the domain ontology and ontologies for all parts of the domain model as described below.

World state model: the relevant aspects of the world state should be known to the team members, and thus part of the team mental model. This part of the model also contains knowledge that would improve performance if shared, e.g., strategies and/or procedures useful for this domain / task. For hunting teams, one could think of exploration patterns and hunting strategies.

Agent models: what are the relevant agents in the domain, what does the one that has the agent models know about the capabilities of agents. Furthermore, a related concept, but still different: the model should describe to what extent agents are competent with respect to their capabilities. For example, there is a difference between the capability of retrieving a block and the competence in retrieving, e.g., one agent can retrieve much faster than another agent. The next sections show that sharing information about the intentions of team members (and even non-team member agents) can contribute significantly to team performance. The same holds for information about their personalities, preferences, and habits.

Organizational specification: contains the models about task, team members, and team interaction. We refer to existing methods to specify organizations, such as MOISE (Hübner et al., 2002), OperA (Dignum, V., 2003), AGR (Ferber and Gutknecht, 1998). Regarding team members, the model should describe which agents are team members, and what role the agents play in the team. The role specifications will refer to the role behavior and the relevant capabilities and competences for that role. To describe capabilities and competences use is made of the ontologies for the domain related to tasks and world state elements. The team task is also part of the organizational specification: it specifies when a task is finished successfully. The team interaction refers to the expected and allowed communications and the timing thereof. Team interaction further refers to e.g., rules of engagement, doctrine, norms, and role interaction. The organizational specification might also indicate which submodels need to be shared, and which need not to be shared. An example in BW4T is two agents A and B having different roles, A retrieves the blue blocks, B the red ones. The agents need not have a fully shared mental model of the location of the blue and red blocks.

Compositional Structures and Relevance

Not in all situations all possible components of team mental models are relevant for team performance. For example, the agents described in the next section don't have the capability to reason about capabilities of themselves or other agents. As a result requiring the agents to have a shared mental model about these competences will not improve performance.

Different hierarchical composition relations can be constructed that respect to the previous section. The point that this paper would like to make is that a hierarchical composition relation is more transparent than having one big unstructured model.

Note that the compositional structure of the model can be reflected in a compositional structure of the set of questions Q, e.g., the subset of questions on the current situation (domain model), the subset of questions on the organizational specification. Note that if there is no organizational specification from the start, then the team members have to determine their organizational structure for themselves. For this the agents can use communication on the topics normally addressed in organizational specifications.

Team Performance

To measure team performance requires performance criteria. Criteria typically used are: how long did it take to complete the task, to what extent was the task successfully completed (quality), and the way the task was performed. This last category can encompass a variety of elements, such as, did the team work have unintended and / or unexpected side effects, under what circumstances was the task performed, how was the interaction amongst team members (atmosphere, team spirit, amount of communication needed). Furthermore, in case of repetitive tasks, one might consider changes in performance with respect to previous times the task was performed. How well is the team capable of dealing with unexpected events (robustness)?

As the aim of this paper is to show that the predictive power of sharedness measures for team performance is situation dependent, not all possible aspects of team performance need to be considered, and the same holds for team model components. Therefore, in the experiments discussed in the remainder of this paper, only time to finish the task is taken as the measure for team performance. Similarly, the team mental model is reduced to world state and intention state.

The Team Member Agents Tested in BW4T

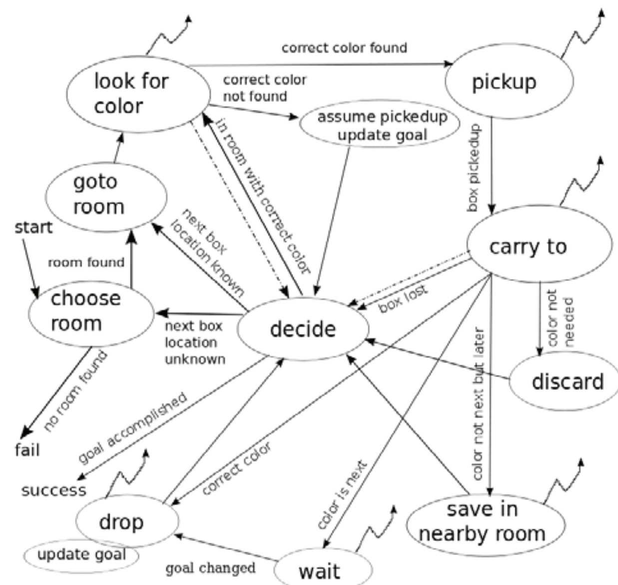

Figure 1: Abstract Decision Cycle of the Agents

To measure team performance under different levels of sharedness, agents were constructed in GOAL (Hindriks, 2009) for the relatively simple BW4T domain (Johnson et al., 2009). BW4T is an extension of the classic blocks world that is used to research joint activity of teams in a controlled manner. A team of agents has to deliver colored blocks from a number of rooms to the so-called drop zone in a certain color sequence. The agents can communicate with each other but their visual range is rather limited. They can see where they are in the world, but they can only see objects not other agents. Furthermore, objects can only be seen by agents that occupy the room that the objects are in. That means that if an agent drops a block in the hall ways, then this block is lost to all agents. Finally, the world restricts access to rooms to one agent at a time. In the version of the BW4T used in this paper, communication and sensing can be assumed to be reliable, and the world is only changed by the actions of the agents.

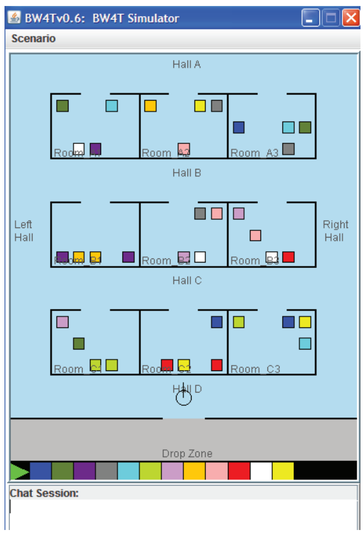

Figure 2: Snapshot of the BW4T environment

We consider two scenarios: one with two color repetitions (called *medium*), and another one with a target sequence of six unique colors (called *high*). Both scenarios are executed in the same environment consisting of 3 rows of 3 rooms each (A1, A2, A3, B1, B2, B3, C1, C2, C3) connected by corridors. In each scenario, 22 blocks (box_1, … box_22) of various colors were placed in the different rooms.

In the first scenario, the task was to bring blocks to the drop zone according to the following sequence: [red, white, white, yellow, darkBlue, darkBlue]. The relevant available blocks and their colors and locations are specified by the following:

DarkBlue: box_8, box_14.

Red: box_4.

White: box_11, box_19.

Yellow: box_9.

Room_A1: box_2, box_3.

Room_A2: box_1, box_4, Box_8, box_9.

Room_A3: box_5, box_6, box_7.

Room_B1: box_10.

Room_B2: box_13, box_12, box_11.

Room_B3: box_16, box_14, box_15.

Room_C1: box_18, box_17.

Room_C2: box_20, box_19.

Room_C3: box_21, box_22.

In the second scenario, the task sequence was [red, white, yellow, darkBlue, orange, lightGreen].

The agent programming to solve the BW4T tasks was based on the decision cycle presented in Figure 1. Note that the aim was not to produce the most efficient agent for this problem, but agents useful for measuring how different levels of sharedness affect team performance and the effect of using a compositional team model versus a flat team model. Futhermore, agents and example domain should be both simple and rich enough to demonstrate the impact of the set of questions Q underlying the team model. Therefore, the following choices were made.

One template agent was modeled, that, in principle, can solve the task alone. The agents spawned from this template are telling the truth and believe everything they are told. Agent capabilities were left out of consideration in the modeling of the agent, and out of the team models. Four homogeneous teams were formed using copies of the template agent.

Team A consisted of agents that *do not communicate* at all. The agents in Team B communicate only *world information* to each other. Team C agents only communicate their *intentions* to each other. Finally, the agents in team D communicate *world information and intentions* to each other. Each team consisted of three agents, called Bot0, Bot1, and Bot2. This fact was known to all agents, in all teams, in all scenarios. The decision cycle for dealing with the BW4T environment is depicted in Figure 1. To give an idea of how the GOAL code of the agents looks, we include here some examples of action rules for processing world information. The first rule processes the percept of a block of a certain color, by inserting this information into the agent's belief base and communicating it to the other agents. The second rule processes the corresponding message by inserting the content into the belief base. The third rule adopts a goal of holding a block and sends the information about this intention to the other agents. The last rule processes messages concerning the intention of another agent to deliver a block.

```
% If Block of Color is perceived, insert in belief base and send to others
forall bel(percept(color(Block,Color)))  do
  insert(color(Block,Color)) + send(allother, :color(Block,Color)) .

% received messages are inserted in belief base and message deleted
forall bel(received(Sender, color(Block,Color))) do
  insert(color(Block,Color)) + delete(received(Sender, color(Block,Color))).

% macro used as abbreviation of condition expressing which block has the
% color to be delivered next in sequence
#define nextColorLocKnown(Block)
  bel(nextColorToDeliver(Color), color(Block, Color), at(Block,Location)).
```

```
% If I know where a block of the next color to be delivered is, and I don't
%already have the goal of holding a block and I'm not already holding a
%block, then adopt the goal of holding the block and send this intention to
%others
if nextColorLocKnown(Block), not(a-goal(holding(Block2))),
   bel(color(Block, Color), not(in('Drop_Zone')), not(holding(_)))
   then adopt(holding(Block)) + send(allother, !deliver(Block, Color)).

% If a message is received that another agent wants to deliver a Block of
%Color, insert this in belief base and delete the message
forall bel(received(Sender, imp(deliver(Block, Color))))
   do insert(imp(Sender, deliver(Sender, Block, Color)))
   + delete(received(Sender, imp(deliver(Block, Color)))).
```

Pilot Experiment: First Results

In this section we discuss first results on the experiment that we are conducting in order to determine the relationship between different measures of sharedness and team performance.

As discussed earlier, to be able to determine the relation of sharedness to performance the sharedness of the various components of the team model has to be considered at various time points while the team is at work. As observed in (Jehn and Mannix, 2001) changes that happen over time while the team is at work affect team performance.

The following timing options were considered: after each observation by some agent, after each communication by some agent, and after each block delivery. Measuring after each communication leads to a weak increase in sharedness (at least of some component). As agents of different teams have various communication strategies, measuring sharedness after each communication would make the measures hard to compare; in some teams no measurements would be taken, in some teams many measurements would be taken.

Measuring after each observation leads to a change in sharedness of the world component of the team model: a decrease if an agent observes something that no other agent has observed before without telling the others, and an increase if this is the last agent to learn of the facts it observed. As this depends on the communication strategy of the agents in the different teams, measuring sharedness after each observation would also make the measures hard to compare.

Measuring after each block delivery would provide the same number of measurements in all the teams (if the teams are successful). Furthermore, the number of measurements is limited to the length of the sequence of colors to be delivered. Therefore, we decided to measure sharedness of the mental models of the agents after every block delivery.

In the following, we present initial results of the experiment with respect to team performance and amount of communication between team members. In future work, we will relate these results to other measures of sharedness.

The performance of the different teams in the two scenarios is depicted in Figure 3. The figure shows for each team and scenario the time at which each block in the

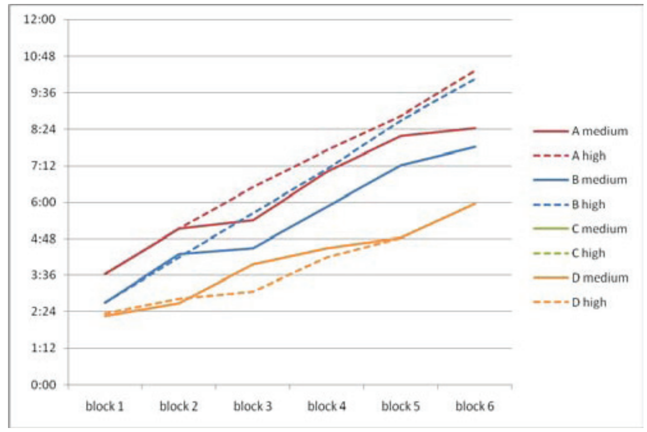

Figure 3: Performance: time to complete the team task

goal sequence was delivered. We can thus observe that the performance of team D (full communication) was the best, and there is not much difference between the two scenarios.

The results for team C are not depicted, but these are comparable to team B. This means that communicating world information and communicating intentions seem to have a similar effect on team performance. Also the effect of not communicating certain information seems to have a greater effect in the high scenario, i.e., the one with different colored blocks. This may be due to the fact that if two agents go for the same color due to lack of communication, this is more of a problem if there are no double colors in the goal sequence. If there are, it may still be useful to collect blocks of the same color. Also one can see that the difference in performance between team A (no communication) and teams B and C (either world information or intentions are communicated) is not very large. Only in case of full communication we can see a significant gain in performance.

Figure 4 depicts the amount of communication in the different teams and scenarios. Team A had no communications at all. We can see that the amount of communication is highest in team D, as is to be expected, but the difference with team B (world information) is small. This means that the performance gain from communication is relatively large for communication of only intentions, in comparison with communication of only world information. The communication of team C (intentions only) has the smallest communication requirements yet it produces a performance, shown in Figure 3, which is similar to that obtained by team B (world information) at a much smaller cost.

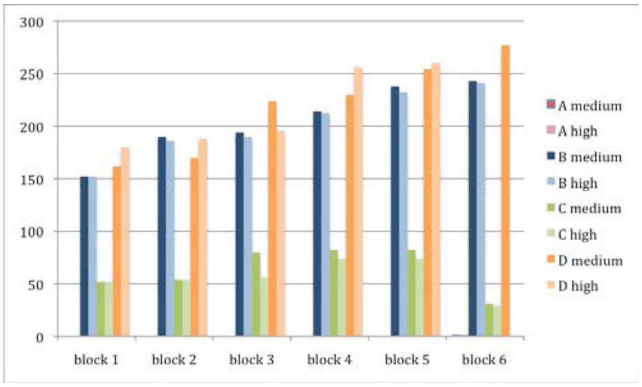

Figure 4: Communication overhead

Conclusions

In this paper we have presented some criteria for measuring sharedness of Team Mental Models. Properties of the criteria are that they are compositional, situation dependent, constrained, and avoid an infinity of terms in the questions / answers used to measure sharedness.

We have also presented results of initial experiments where we compared different types of simulated agents in the domain of the BW4T. The main assumption made is that agents can solve the task individually. Since the objective of the experiment was to measure if and how sharing mental models improves performance, agents do not need to be sophisticated or to use an optimal strategy. The domain, task, and team members are given up front and are relatively simple to enable assessing improvements without excessive complexity. For the same reasons, sensing and actions are assumed to be correct. The results show that the amount of sharedness affects team performance.

Future work will focus on additional performance measures, using more sophisticated agents capable of reasoning about their team mates and including errors in sensing and actions. More general questions will also be addressed, such as how the compositional structure of the model can be reflected in a compositional structure of the set of questions used to measure sharedness.

References

Bolstad, C., and Endsley. M., 1999. Shared mental models and shared displays: An empirical evaluation of team performance. *Human Factors and Ergonomics Society Annual Meeting Proceedings*, 43(3):213–217.

Cannon-Bowers, J. A., Salas, E., and, Converse, S., 1993. Shared mental models in expert team decision making. In: N. J. Castellan, editor, *Individual and group decision making*, Lawrence Erlbaum Associates, pp. 221–245.

Dignum, V., 2003. A model for organizational interaction: based on agents, founded in logic. Dissertation, Universiteit Utrecht.

Ferber, J., and, Gutknecht, O., 1998. A meta-model for the analysis and design of organisations in multi-agent systems. In: *Proceedings of the Third International Conference on Multi-Agent Systems* (ICMAS'98), IEEE Computer Society Press, pp. 128-135.

Hindriks, K. V., 2009. Programming rational agents in GOAL. In R. H. Bordini,M. Dastani, J. Dix, and A. El Fallah Seghrouchni, (eds.), *Multi-Agent Programming: Languages, Tools and Applications*. Springer, Berlin, pp. 119-157.

Hübner, J.F., Sichman, J.S., and, Boissier, O., 2002. A model for the structural, functional, and deontic specification of organizations in multiagent systems. In: *Advances in Artificial Intelligence*, 16th Brazilian Symposium on Artificial Intelligence, (SBIA'02), volume 2507 of LNCS, Springer, pp. 118–128.

Jehn, K.A., and Mannix, E. A., 2001. The dynamic nature of conflict: a longitudinal study of intragroup conflict and group performance. *The Academy of Management Journal,* 44(2):238–251.

Johnson, M., Jonker, C.M., Riemsdijk, M. B. van, Feltovich, P.J., and, Bradshaw, J.M., 2009. Joint activity testbed: Blocks world for teams (BW4T). In: *Proceedings of the Tenth International Workshop on Engineering Societies in the Agents' World* (ESAW'09), volume 5881 of LNAI, Springer. pp. 254–256.

Jonker, C.M., Riemsdijk, M.B. van, and, Vermeulen, B., 2010. Shared Mental Models: A Conceptual Analysis. In: *Proceedings of the 9th International Workshop on Coordination, Organization, Institutions and Norms in Multi-Agent Systems*, COIN@AAMAS2010, Springer LNCS series, to appear.

Lim, B., and, Klein, K., 2006. Team mental models and team performance: A field study of the effects of team mental model similarity and accuracy. *Journal of Organizational Behavior*, 27(4):403.

Mathieu, E., Heffner, T. S., Goodwin, G., Salas, E., and, Cannon-Bowers, J., 2000. The influence of shared mental models on team process and performance. *Journal of Applied Psychology*,85(2): 273–283.