

Compositionality of Team Mental Models in Relation to Sharedness and Team Performance

Catholijn M. Jonker¹, M. Birna van Riemsdijk¹, Iris C. van de Kieft¹, Maria Gini²

¹Delft University of Technology, Man-Machine Interaction Group, Faculty of EEEMCS,
Mekelweg 4, 2628 CD, Delft, The Netherlands

²University 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”, or so is said, in which better model refers to the extent to which the model is shared by the team members. This paper argues that according to circumstances, some components of that model are more relevant with respect to team performance than others. Circumstances change with the dynamics of the environment, the team composition and organization, its members, and the team task. Consequently, a compositional approach to measuring sharedness of team mental models is proposed. A case study illustrates the argument and the approach.

Keywords: Shared mental model, team model, measuring

1 Introduction

A Shared Team Mental Model (STMM) is considered beneficial for teamwork, see e.g., (Cannon-Bowers et al., 1993; Lim and Klein, 2006; Mathieu et al., 2000). Aspects to be taken into consideration are equipment, task, team interaction, and team members, see e.g., (Cannon-Bowers et al., 1993; Mathieu et al., 2000). Sharedness refers to the team members having equivalent mental models. For example, in American Football if the team does not share the offensive plan of the quarterback, chances are high that the other team will intercept the ball. The better the shared mental model the better a team is capable of performing its tasks, even in unforeseen circumstances, such as a ball that is not perfectly caught.

Simplistic reasoning would suggest that to improve team performance, the team members should share full knowledge of the state of the world and of what each team member is doing. In reality the communication overhead alone that this would require makes this totally impractical. Furthermore, for teams in which team members have highly individual expertise having full information on what another team member is

doing, would not help team performance; e.g., in incident management a fire fighter does not need to know the specifics of the injuries of a person being attended to by the ambulance personnel. This reasoning induces the *hypothesis* that the extent of sharedness of some aspects or components of team mental models is more important for team performance than that of other components, and in particular, that this extent changes with circumstances.

Testing such a hypothesis requires ways of measuring the extent to which teams have a shared mental model. The formalization of the notions of mental model, sharedness, and team mental model as presented in (Jonker et al., 2010), form a starting point. However, the formalizations do not take compositional models into account, nor do they account for the relative importance of each component.

This paper proposes a compositional approach to construct models that underlie teamwork, in which the following components are identified as potentially 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 this paper.

A relevance relation is associated with the composition relation, to indicate to what extent the components contribute to team performance. This relevance relation depends on the current circumstances such as the particular instance of the team task, the performance criteria, the dynamics of the environment (world, equipment, and non-team member actors), the composition and organization of the team, and the team members themselves. The relevance relation is used to formulate a compositional way of measuring sharedness of mental models.

Experiments are done within the context of a case study in the Blocks World For Teams (BW4T) domain, as introduced in (Johnson et al., 2009), to test our hypothesis, and to test the adequacy of our compositional approach. In particular we construct different teams of agents that have to operate in two different scenarios, and measure team performance by the time taken by each team for task completion and communication overhead.

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) which we use to measure the performance of different types of agents that share different types of team mental models. We show how sharing affects performance.

2 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, discusses compositional structures of team mental models, and defines sharedness of compositional team mental models in such a way that the relevance of the components for team performance is taken into account.

2.1 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 restructuring of the team mental model and the addition of information and knowledge about the domain (world and agents), and of 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 be part of the team mental model. This part of the model may also contain knowledge about, e.g., the status of equipment.

Agent models: what are the relevant agents in the domain, and what are their capabilities. 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 clearly 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) and OperA (Dignum, V., 2003). 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.

2.2 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 do not have the capability to reason about their own capabilities or those of 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 of the components of the team mental model. An example of a basic hierarchical structure is depicted in Figure 1. To have a fully shared team mental model, each team member has to have this model within his own mind. As part of this team mental model, each team mem-

ber maintains an agent model for each of his team mates. For example, the team mental model that agent Arnie maintains, contains an agent model of his team mate Bernie. That agent model that Arnie has of Bernie describes not only the beliefs that Arnie has about the availability, capability, competences, and so on of Bernie, it also describes what Arnie thinks that Bernie believes about the team, in other words, it contains Arnie’s idea of Bernie’s mental model of the team. In Figure 1, for reasons of simplicity the organizational specification component is not decomposed further, and neither is the recursive “Other agent’s team mental model”.

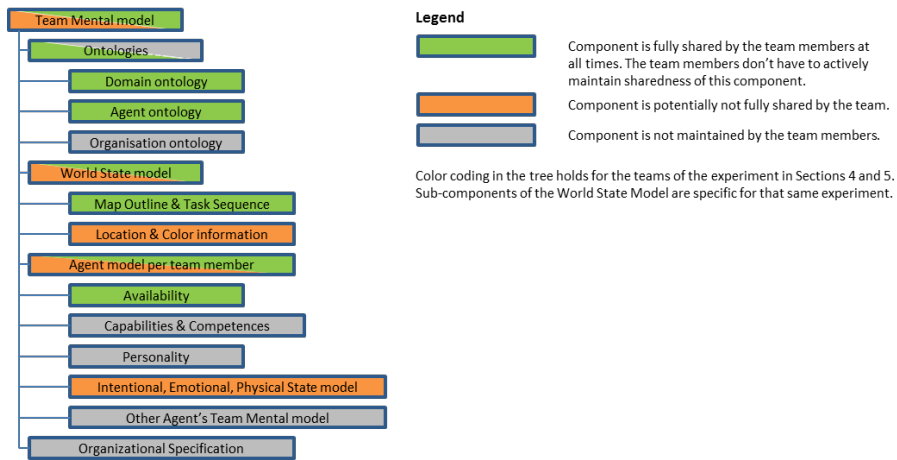


Fig. 1. Hierarchical composition of a team mental model

As mentioned, not all possible components of a team mental model are important in every situation. For example, a team consisting of agents incapable of reasoning about the competencies of other team members does not need to maintain that information. Such components are colored grey in Figure 1. Having sharedness of any of these components would not make the team perform better. Also if all team members are guaranteed to have the same mental model for a component (the green components in Figure 1), then it does not have to be considered as a factor affecting the performance of the teams. Therefore, when measuring the sharedness of a team mental model, only those components should be taken into account that are relevant (colored orange in Figure 1). The color coding is an indication of the relevance relation on components for measuring sharedness of team mental models in such a way that the measure would be predictive of team performance.

3 The Experimental Set Up

An experiment was set up in order to test the hypothesis derived in the introduction:

Hypothesis The extent to which sharedness of some components of team mental models is more important for team performance than that of other components is situation dependent.

In the following we describe the experimental set up we used to test our hypothesis. In these experiments, aspects of team performance considered are time to finish the task, and communication overhead. The team mental model is kept small and consists only of the intentions of team members and of the world state model. Changes that happen over time while the team is at work affect team performance, see e.g. (Jehn and Mannix, 2001). Thus team performance is monitored over time. Furthermore, we built teams of agents for which the structure of the team mental models is the same, but for which the sharedness differs. We constructed two scenarios, and show that the extent to which sharedness of components contributes to team performance differs per scenario.

The experimental environment is the BW4T domain (Johnson et al., 2009). BW4T is an extension of the classic blocks world that is used to research joint activity of teams. 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, see Figure 2 (right). The agents can communicate with each other but cannot see other agents. Furthermore, blocks can only be seen by agents that occupy the room that the blocks are in. If an agent drops a block in a hall way, 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.

The two scenarios differ in the color sequence that has to be delivered: 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 second task sequence was [red, white, yellow, darkBlue, orange, lightGreen].

The agent teams that have comparable structures in their team mental models but different sharedness of these models were programmed using the GOAL agent programming language (Hindriks, 2009). The agent programming to solve the BW4T tasks was based on the decision cycle presented in Figure 2 (left). Note that we need not produce the most efficient agent for this problem! 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 on the basis of the template agent.

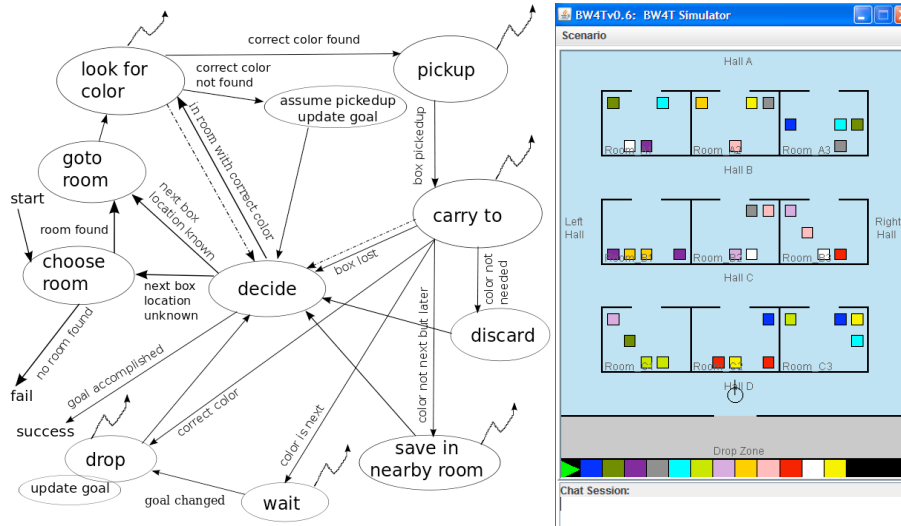


Fig. 2. Agent Decision Cycle (left). Snapshot of the BW4T environment (right).

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. To give an idea of the GOAL code of the agents, we include 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 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,
% then 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)))).

```

4 Results of the Pilot Experiment

In the following, we present the results of a pilot experiment with the set up described in the previous section. The results focus on team performance in terms of time to complete the task and communication overhead.




Fig. 3. Team Performance: Time (left) and Communication overhead (right)

The time to complete the task by the different teams in the two scenarios is depicted in Figure 3 (left). The figure shows for each team and scenario the time at which each block in the goal sequence was delivered. The right side of Figure 3 shows the amount of communication required from one block delivery to the next. As more of the sequence is delivered, clearly more information on locations and colors is communicated. Timewise, 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 comparable to those of team B. This means that communicating world information and communicating intentions seem to have a similar effect on the required time. Also the effect of not communicating certain information seems to have a greater effect in the high scenario, i.e., the scenario 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 clear improvement in the amount of time required to finish the task.

Our hypothesis is supported both by the findings of our pilot experiment and by practical experience. When considering team performance, team members do not need to have a fully shared team mental model consisting of all possible elements that literature discusses. In fact, practical evidence shows that overhead required to maintain full sharedness of all aspects that might belong to team mental models is counter-productive due to the communication overhead and information overload (Eppler and Mengis, 2004). This suggests that the selection of the components for which the team members should aim for sharedness depends on the circumstances, i.e., task, environment, team composition and organization, and team members.

It is important to realize that knowing that some component is, given the circumstances, probably relevant for team performance, is not the same as knowing to which extent such a component is relevant. Figure 3 shows that teams B and C outperform team A in the medium scenario, but that B and C perform equally well. The difference between these three teams is in the communication about the orange model components in Figure 1, in particular in team B about the locations of blocks and agents, and the color of blocks, and in team C about the intentions of agents. This suggests that the component that team B communicates about is approximately equally important as the component that team C communicates about. However, in the high scenario, neither team B, nor team C performs better than team A. If the results on team D are not taken into account, two possible combinations of these components can be made: AND, but also OR. If team D would have performed equally well as teams B and C, then, for the circumstances in our experimental set up, teams should communicate on intentions only (see Figure 3, right). Now that team D outperforms teams B and C, for high performance on time, teams should share information on BOTH locations, colors, and intentions. If the communication on locations and colors could not be reduced further, then the weight of the component on intentions should be increased, until a balance is found with respect to the weight of communication overhead in the overall team performance.

5 Measuring Sharedness of Compositional Models

This section presents a formal definition of measuring sharedness of compositional models with respect to a relevance relation in terms of weights. It builds on the measure of sharedness, subject overlap, and model agreement as defined in (Jonker et al., 2010). These notions are all defined relative to a set of questions Q , which can be seen as those questions that determine the purpose of the model. As team members do not have to be clones of each other with respect to all the information and knowledge in their minds, the set of questions Q is to be used to focus the team mental model on those subjects that matter. Therefore, the subject overlap of the team mental models in

the minds of the team members should be complete (denoted by 1 in the definition below). Furthermore, agreement means that these models provide the same answers (at least to extent θ), and thus provide sharedness to at least that extent. The definition of sharedness provided by (Jonker et al., 2010) is then used to determine the sharedness of models that are not composed of submodels.

Definition 1 Shared Mental Model (Jonker et al., 2010). A model M is a mental model that is shared to the extent θ by a set of agents A with respect to a set of questions Q , denoted by $\text{Sh}(M, A, \theta)$ iff there is a mental model M_A for each agent $A \in A$ with respect to Q , such that

1. $\text{SO}(M, M_A, Q) = 1$,
2. $\text{Agr}(M, M_A, Q) \geq \theta$,

We can define the sharedness of the compositional model in terms of a function that takes the sharedness per component to the extent that it is relevant, and follows that same composition relation as the team mental model itself. A compositional linear additive function suffices for the case study, using weights to model the relevance of each component. Following our hypothesis, we expect that then also the answer to the question which compositional function for measuring sharedness is indicative of team performance, should be situation dependent.

Definition 2 Given a set of agents A , a model M , and a set of questions Q , then

$$U_Q(A, M) = \max \{ \theta \mid \text{Sh}(M, A, \theta) \}, \quad \text{if } M \text{ is not composed.}$$

$$U_Q(A, M) = \sum_{m \in M} w_m U_Q(A, m), \quad \text{if } M \text{ is composed.}$$

Where w_m is the weight of component m , and $\sum_{m \in M} w_m = 1$.

6 Conclusions and Future Work

This paper shows that according to circumstances some of the components of team mental models are more relevant for team performance than others. Circumstances change with the dynamics of the environment (i.e., world, equipment, non-member actors), the team composition and organization, its members, and the team task.

Results of a pilot experiment in which different teams operate in the domain of the BW4T support our compositional approach of team mental models and enable measuring their sharedness. By comparing the team performance of the teams over two different scenarios it becomes clear that of the two components under consideration should be combined differently in the measure of sharedness over the two scenarios.

Finally, the paper proposes a compositional approach to measuring sharedness of team mental models, and suggests a linear additive function for the case study. A formal definition is provided that shows how such functions can be used to compute sharedness of composition models. In combination with the more qualitative analysis approach that can be associated with the color coding presented in Figure 1, such functions make this work generally applicable.

Future work will focus on running full experiments, determining precise sets of questions Q to underpin the formalizations presented in Section 5, and computing sharedness of the team mental models in various ways to test the practical value of such computations. Furthermore, additional performance measures will be taken into account and more sophisticated agents will be used, e.g., capable of reasoning about their team mates and including errors in sensing and actions.

Acknowledgement. The authors thank the anonymous reviewers for their comments. This work is subsidized by the Dutch Technology Foundation STW, applied science division of NWO and the Technology Program of the Ministry of Economic Affairs; the Pocket Negotiator project with grant number VICI-project 08075.

7 References

1. 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.
2. Dignum, V., 2003. A model for organizational interaction: based on agents, founded in logic. Dissertation, Universiteit Utrecht.
3. Eppler, M.J. and Mengis, J. 2004. The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines. *The Information Society*, 20, 5, 325-344.
4. 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.
5. 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.
6. 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.
7. 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.
8. 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.
9. 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.
10. 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.