



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Alshehri, A;Miller, T;Sonenberg, L

Title:

Modeling communication of collaborative multiagent system under epistemic planning

Date:

2021-10-01

Citation:

Alshehri, A., Miller, T. & Sonenberg, L. (2021). Modeling communication of collaborative multiagent system under epistemic planning. *International Journal of Intelligent Systems*, 36 (10), pp.5959-5980. <https://doi.org/10.1002/int.22536>.

Persistent Link:

<https://hdl.handle.net/11343/298649>

RESEARCH ARTICLE**International Journal of Intelligent Systems**

Modeling Communication of Collaborative Multi-Agent System under Epistemic Planning

Abeer Alshehri^{1,2*} | Tim Miller^{2†} | Liz Sonenberg^{2‡}

¹Department of Computer Science and Information Systems, King Khalid University, Asir, Abha, Saudi Arabia

²School of Computing and Information Systems, University of Melbourne, Melbourne, Victoria, Australia

Correspondence

Abeer Alshehri, School of Computing and Information Systems, University of Melbourne, Melbourne, Victoria, 3057, Australia
Email: aalshehri@student.unimelb.edu.au

Funding information

In most multi-agent applications, communication is essential among agents to coordinate their actions and achieve their goals. However, communication often has a related cost that affects overall system performance. In this paper, we draw inspiration from epistemic planning studies to develop a communication model for agents that allows them to cooperate and make communication decisions effectively within a planning task. The proposed model treats a communication process as an action that modifies the epistemic state of the team. We evaluate whether agents can cooperate effectively and achieve higher performance using communication protocol modeled in our epistemic planning framework in two simulated tasks. Based on an empirical study conducted using search and rescue tasks with different scenarios, our results show that the proposed model improved team performance across all scenarios than baseline models.

KEYWORDS

Epistemic Planning, Epistemic Logic, Communication Process, Multi-agent system, Team Performance, Collaboration, Shared Mental Model

1 | INTRODUCTION

In many applications, multiple autonomous agents working as a team have potential advantages relating to the effective execution of tasks beyond the capabilities of a single agent. With interdependent tasks, effective collaboration and coordination are the key foundation of a higher-performing system. An agent has to be transparent to other agents to be able to coordinate their actions¹.

1

This is the author manuscript accepted for publication and has undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the [Version of Record](#). Please cite this article as doi: [10.1002/int.22536](https://doi.org/10.1002/int.22536)

One of the main requirements of effective collaboration is the existence of a communication process. In critical domains, such as disaster response and healthcare, the communication process among team members is considered an essential requirement for better team cooperation. In a multi-agent system (MAS), agents communicate to develop and maintain a shared mental model (SMM) on a task. An SMM is an internal representation of an environment that enables agents to form accurate explanations and expectations regarding the task and coordinate their actions to meet the demands of that task². Although communication is a natural way to increase the SMM of a team, excessive communication, which may arise due to computational requirements, has been costly and hinders system performance³. Besides, the information obtained may not be valuable for the agent in making their decision. When faced with collaborative tasks, effective coordination through communicating at the right time with the right information becomes crucial to achieve a higher system performance.

Recently, studies have discussed different methods for improving multi-agent systems' performance, namely, the ability to reason regarding the exchanged information, and when and what the agent should communicate. By gaining insight into how humans cooperate effectively, research shows that an intelligent human interaction is a result of their ability to reasoning about others' mental states —beliefs and knowledge— and thus *selectively* communicating based on the needs of team members^{4;5;6}.

Most works in developing communication considered it to be in a separate model from the planning process; thus, it assumes teams work simultaneously to make planning and communication decisions at each step^{7;3}. Since planning does not typically consider communication as an involved action, we believe that using an *epistemic planning* framework leads to a planning process that can consider communication action.

There have been recent works on planning using *epistemic logic*^{8;9;10;11}. Since it is concerned with representing knowledge/beliefs rather than simple facts about an environment, it forms a strong basis from which to reason in the absence of complete knowledge. By enabling agents to reason in such a way, they have a compact class of planning that models communication as a natural action, allowing a planner to fold in communication planning during task planning¹. These actions are carried out if newly observed information is inconsistent with the agents' current beliefs and useful for other agents in achieving their tasks. This model could ultimately maintain better performance by enabling agents to communicate and keep track of SMMs in one compact model.

In this paper, we address the following research question: *Does reasoning about knowledge of others (epistemic reasoning) in the presence of communication as an action improve the cooperation of a multi-agent system performing in a dynamic environment?* In simulated tasks motivated by disaster response and search and rescue, we evaluate whether agents can cooperate effectively and achieve higher performance using communication protocol modeled in our epistemic planning framework.

The paper is structured as follows. Section 2 provides the relevant background and discusses related work. We set out the proposed model in detail in Section 3. In section 4, we present the experiments that tested our approach. Then, section 5 discusses the results of the conducted experiments. Finally, our paper concludes with our work's contribution, along with some future directions of research.

2 | BACKGROUND

In this section, we overview relevant prior work on multi-agent systems and epistemic planning framework related to effective team performance.

2.1 | Shared Mental Models (SMM)

In the context of teamwork, the psychology literature has introduced the concept of a shared mental model (SMM) to explain how a human team works. It represents the knowledge of team members for the world around them that helps to understand and

predict others' behavior and, in turn, adapt their actions to task demands². For instance, a blind pass between members of a basketball team shows that each member can predict others' position on the court. Extensive teamwork experiments have shown that higher similarity of team mental models results in higher team performance^{12;13;14}.

The ability to predict others' behavior allows anticipating actions to their expected behavior along with team tasks. Team implicit coordination relies on having common knowledge that assists in predicting team behavior and needs, including beliefs and goals along with a sharing understanding of the team task^{15;16}.

Maintaining SMMs developed in the psychological literature can be applied to human-agent or agent-agent systems to improve their work. Having a shared mental model about the team functioning will improve its performance^{4;5}. Although the word "shared" suggests identical mental models, this is not the case where autonomous agents are distributed in a dynamic environment. Instead, studies have proposed similarity measures to investigate the degree of similarity for good team performance or good enough performance to achieve specific tasks^{17;4;5;13;18}. In this work, maintaining an SMM enables agents to share information efficiently and achieve a better team performance.

2.2 | Communication Strategies

The essential process for maintaining and developing SMMs is the communication between team members. However, excessive exchanging of information results in a communication overhead, which will negatively affect team performance. Due to an agent's limited view of an environment, it might be unwanted or impractical to share all knowledge all the time. Communication has a related cost associated with the required resources or bandwidth. Thus, effective team performance results from evaluating the trade-off between communication cost and the value of information received. That leads to more efficient communication by which a common ground (i.e., SMM) is developed among team members to better coordinating implicitly¹⁹.

There are some factors for effective communication—what information to communicate, when, and to whom.^{20 20} argued that there is a condition where the environment may affect the agents' mental states if they work using a static communication protocol. Handling the observability and environmental changes is required for cooperative agents. They also argued how an agent evaluates information before sharing it to determine its usefulness for both the task and other agents in order to avoid the cost of consuming communication resources. In addition,²¹ and¹² observed that communicating goals (intentions) contributes more to team efficiency in human-agent teams than sharing beliefs (world knowledge) to build a strong SMM. Regarding certain classes of knowledge that need to be communicated, it is better to consider what beliefs should be shared and whether they are sufficient for the team.

Other work has concentrated on the relationship between task structure, anticipatory information sharing, and team performance.²² proposed a set of hypotheses to identify these relationships. They evaluated them through empirical study, including multiple teams of four people in a collaborative search-and-deliver task within a simulation environment. They analyzed the strategies for effective team communication in tasks with different complexity levels. They observed higher rates of implicit coordination among the best performing team rather than explicit coordination. This was true even with a high-complexity task as a result of training teams before starting experiments. Furthermore, they identified that proactive communication between team members regarding the next goal reduces communication costs and improves team performance. However, these results may not cover the case where teams have not worked together before or lack understanding of the task.

³ developed a communication policy, ConTaCT, which enables agents to decide whether to share information with others or not during time-critical collaborative tasks within a simulated task motivated by rescue operations. They evaluated its performance within a deterministic domain (partially observable states) with a partially-known initial state. The algorithm enables an agent to maintain other agents' knowledge about the environment and then compare the benefits of this decision against communication cost. They found that ConTaCT reduces communication overhead within artificial agent teams while achieving a comparable task performance, reducing more than 60%. The reason is that some information is unnecessary for other members and benefits only

the local agent.

²³ worked on communication problems related to the question of why and when to communicate. They introduced communication to the interactive, partially observable Markov decision processes (I-POMDPs). They presented a communicative, interactive POMDPs (Com-I-POMDPs) by integrated costly communication actions into (I-POMDPs) framework. Their results show a graphical model of the framework that consumes less time and space to solve a problem in uncertain multi-agent settings. More recently,²⁴ built on the I-POMDP framework and included the communication process to formulate Communicative I-POMDPs (C-IPOMDPs). Similar to our work, communication is treated as an action. They use the Bellman optimality principle as decision-theoretic planning to deciding whether they are communicative acts, just as the other actions. However, their framework has not been tested to verify their findings.

2.3 | Evaluation of Team Performance and Communication

Several studies have shown a critical need for a solid understanding of what influences team interactions and performance^{25;26}. Here we describe standard methods that are used to evaluate team performance and communication. Analyzing these methods would help in identifying ways to improve team performance and communication.

Methods of evaluating team performance used in prior research have been tailored to specific scenarios and involved objective measures such as a total number of accomplished tasks and accuracy rates, or subjective measures such as task uncertainty and the workload assigned to each agent^{27;12;3;26;18}. In this paper, team performance is measured based on the total number of actions it takes to complete the task (completion cost) and task completion time.

Measuring communication behavior in prior studies was based on the total number of communications per minute or per task^{16;12;3;18}.²⁸ measured communication efficiency based on anticipation ratio, which is the ratio of total communications to communication requests. The higher the team shows the anticipation ratio, the more efficient communication, and implicit coordination. We measured communication behaviors based on the total number of messages exchanged by the agents over a task in our work.

SMMs have been used to predict and understand team performance. Agents communicate as soon as they obtain new observations in order to develop an SMM for a task. Measuring the degree (level) of sharedness of mental models that ensure good team performance is important, that is, which beliefs of mental models should be shared?^{13;29} proposed similarity measures that can be applied to intelligent agents.²⁹ defined what is called (subject overlap) of agents' mental models to determine the extent of their similarity. It is a measure of the percentage of answers agents were provided to a set of queries related to a task. Based on this, we measured the sharedness of mental model contents (i.e., agents' beliefs) over a task.

2.4 | Dynamic Epistemic Logic (DEL)

Epistemic logic (EL) is a means of handling knowledge, a concept that was first published by³⁰. Having knowledge something is true (false) means either being true (false) in every state that is considered to be possible or being uncertain about it (i.e., true in some states and false in others). A key merit of EL is that the state of knowledge for several agents can be represented using a Kripke structure³¹. Given a set of propositional variables, Kripke structure forms a set of possible states (epistemic alternatives), a set of accessibility relations (binary relations on the world states), and a set of propositional facts (possibilities for the truth of propositional variables in each world state)³².

Extending epistemic logic with dynamic aspects has been studied over recent decades under the term dynamic epistemic logic (DEL)^{31;33;34;35}. In DEL, changing knowledge is modeled by transforming Kripke structures. These transformations usually involve factual (or ontic) changes and changes in the agents' mental states. This means that agents' accessibility relations have to change, and accordingly, the set of possible states may change as well³².

2.5 | The DEL Planning Framework

To achieve a specific goal in a particular situation, we need to organize our actions in a plan which represents how we should act. In a planning task, otherwise known as a planning problem, three components are given: an initial state, a set of available actions, and a set of goal states. A planning solution is a sequence of actions that converts the initial state into one of the goal states³⁶.

A method for generating plans works by predicting the outcome of actions that lead to the goal. There are many approaches to planning systems, mostly connected to propositional logic³⁷. Epistemic planning is one such approach, which uses epistemic logic as the underlying formalism^{8;38;39} - we refer to⁴⁰ for an overview of DEL-based epistemic planning. Generally, epistemic planning considers how to get from a current state of knowledge to a goal state. The three components of the DEL-based planning approach are an epistemic model as the initial state, a set of event models (or action models) to describe the actions, and an epistemic formula to describe a goal. The solution for such a planning problem is a sequence of event models that uses product update operators to produce a model that meets the goal formula^{8;33}.

Planning in partially observable environments is computationally challenging. Partially observable Markov decision processes (POMDP) are a common approach to planning in such an environment^{41;42;43}. However, the POMDP approach has limited applicability to planning under incomplete knowledge compared to epistemic logical approaches⁴⁴. The latter approaches allow for more compact modeling of planning tasks about *nested knowledge* than POMDPs; that is, knowledge about what other agents know. In the DEL-based planning approach, actions (event models or action models) are considered a feature that encodes partial observability and non-determinism³⁹. Thus, observability will be based on executing actions, and there is no need to specify an observation function as an action description.

Furthermore, the DEL-based planning approach allows us to formalize reasoning about knowledge or belief changes (including higher-order knowledge or beliefs about other agents). Such reasoning is required for many AI applications that involve multiple autonomous agents interacting to accomplish collaborative or competitive tasks.

2.6 | Epistemic Planning for a Cooperative Multi-agent System

In multi-agent situations, epistemic planning can be used for decision making in the case of distributed knowledge. It is essential for successful collaboration and coordination that agents can reason about other agents' knowledge, capabilities, and uncertainty. Such a planning framework integrates non-determinism and partial observability in multi-agent tasks⁴⁵.

Multi-agent planning tasks can be modeled using Kripke structures as world states and event models as actions. We have to consider some additional design choices for such tasks, such as whether agents are collaborating or competing; or communicating arbitrarily and committing to a joint plan. In a decentralized setting,^{46;10} proposed implicitly coordinated plans that help to solve tasks with joint goals where agents at the time of planning are not allowed to commit to a joint policy.^{47;48} modeled epistemic planning problems with bounded nesting of knowledge (or belief) and compiled this into a form that classical planning can handle. Such a model can plan from a single agent's viewpoint with the possibility for goals and actions having nested beliefs, co-present observations, and taking other agents' beliefs into account while reasoning. In our work, we evaluate the efficiency of using an epistemic planning framework to model the communication process using an epistemic planner proposed by⁴⁷.

3 | METHOD

In this section, we describe the objective of this work. We describe an epistemic planning task of multi-agent settings. Then we lay out the proposed model (Selective communication Model) in detail.

3.1 | Objective

Our model's objective is to incorporate a cooperative team with varying tasks, in which the planning task is categorized under the epistemic planning framework. This will enable an agent to model their beliefs and other beliefs; thus, aligning SMM to keep track of consistent beliefs. Based on this, agents will communicate *selectively* only in case of inconsistent beliefs that are newly observed from a current state and valuable for other members to achieve their tasks. We aim to investigate the impact of applying such a framework on team performance and compare it with two generated baseline strategies.

Hypothesis: We hypothesize that a team can communicate more *selectively* with minimal loss of performance if they represent its SMM compared with baseline communication models. Having a model of others' beliefs will help them to collaborate effectively.

3.2 | Multi-agent Planning Task

A model consists of \mathcal{P} a finite set of epistemic literals, $\mathcal{A}g$ a finite set of cooperative agents, and \mathcal{A} a finite set of actions. We can define a multi-agent epistemic planning (MEP) task \mathcal{D} as follows:

\mathcal{D} is a tuple of the form $(\mathcal{P}, \mathcal{A}, Ag, \mathcal{G})$, where \mathcal{P} , Ag and \mathcal{A} are as above, and \mathcal{G} is the goal condition. Each action $a \in \mathcal{A}$ is assumed to be of the form $precondition(a), effect(a)$, where both are finite sets over \mathcal{P} . An action a is said to be applicable in a state if $precondition^+(a) \subseteq \mathcal{P}$ and $effect^-(a) \cap \theta = \emptyset$. The set of formulae, \mathcal{L} , for multi-agent epistemic logic is given by the following grammar⁴⁷:

$$\phi ::= p \mid \phi \wedge \phi \mid \neg\phi \mid B_i\phi \mid [\alpha]\phi$$

where $p \in \mathcal{P}$, $i \in Ag$, $B_i\phi$ is that agent i believes ϕ , and $[\alpha]\phi$ is believed ϕ after applying action α . The epistemic planning process has been approached by compilation to classical planning that deals with bounded nesting of knowledge/belief. For more details, we refer to⁴⁷.

We designed two tasks for Gridworld and Blocks World for Teams (BW4T) maps to test the effect of modeling communication as a natural action under an epistemic planning framework on team performance. A variation of task scenarios has been used to realize the validity of the proposed model in some different scenarios motivated by search and rescue operations. We assume that actions' effects are deterministic and commonly known for these tasks, but the environment is often initially unknown. Common knowledge of the planning behavior, initial state, and goal state is assumed to be known by the collaborator agents.

3.2.1 | Selective Communication Model

Communication process is modelled as an action $a = \{precondition(a), effect(a)\} \in \mathcal{A}$, in which the $precondition(a)$ and $effect(a)$ are with respect to belief. The Selective communication model for a cooperative system allows each agent to communicate during a task as a part of the planning process. The following example will describe the model in details:

In a mission of searching for survivors on a map in the Gridworld domain (refer to section 4.1.1 for more details), agents are moving around to find the survivor locations and update the team's mental model about it. The decision of updating others' mental state involves the following three parts: (1) *what to communicate*; (2) *to whom should be communicated*; and (3) *when to communicate*. Communication action $a \in \mathcal{A}$ is modeled to answer them as the following:

- $precondition(a) = B_i s_1$, an agent i believes that there is a survivor s at location 1.
- $effect(a) = \{\forall n \in Ag, B_n s_1\}$, for all other agents they believes that there is a survivor s at location 1.

In the action's precondition, we model (*what to communicate*) part of the communication decision by determining a proposition that needs to be communicated - one communication action for each proposition; that is here, the survivor location on the map. In the action's effect, we model (*to whom should be communicated*) part by selecting the receiver, either all agents or specific agents -based on the task scenario, as we will see later (refer to section 4). As the planner models the effect of actions in \mathcal{A}_g , that allows performing epistemic reasoning. Besides, it is planning from the perspective of a single agent (root agent) to form a coordinated plan accomplished by multiple agents⁴⁷. Therefore, deciding (*when to communicate*) is the planner ability in which the information is communicated only just as needed to maintain consistent beliefs (SMM) essential for them to achieve their task.

| Example

Here is the communication action *commsurvivor*, to communicate survivor location, coded in PDDL for Gridworld domain where the agents are required to search positions of the three survivors on the map. We use the types $?p$ for positions, $?a$ for agents and $?s$ for survivors. We introduce fluents (`at ?a ?p`) to represent that agent $?a$ is at position $?p$, and `[?a] (survivorat ?s ?p)` to represent that the agent $?a$ believes that survivor $?s$ is at position $?p$. The effect of that action is to change the belief state of other agents $?g$ to all believe that survivor $?s$ is at position $?p$:

```
(:action commsurvivor
  :derive-condition always
  :parameters      (?p - pos ?a - agent ?s
                   - survivor)
  :precondition    (and (at ?a ?p) [?a]
                       (survivorat ?s ?p))
  :effect          (and (forall ?g - agent
                       [?g] (survivorat ?s ?p
                               )))
)
; ...
```

3.2.2 | Constructing a Shared Mental Model (SMM)

A mental model of a particular agent, perhaps a root agent who perceives the environment, represents a state of the world; The planner's reasoning is from the perspective of that agent, what its true and false beliefs⁴⁷. Action execution, then, is grounded on these hold beliefs (action's precondition). A Shared Mental Model (SMM) construction represents the overlapping content of agents' mental models; that is, how many propositions are always known among agents. It is updated based on the actions' effects. Thus, communication action's effect will inform others about what has been observed and the communicator agent's current location on the map. Additionally, the SMM can be updated if more than one agent observes the same state; that is, they all have the same beliefs (a common knowledge) about that state.

4 | EVALUATION

In this section, we first explain an experimental design for the two simulated tasks used to empirically evaluate our proposed model. Then we present a set of measurements proposed to measure team performance.

TABLE 1 Experimental setups for Gridworld domain

Communication model	Selective		No Communication		Communication All	
	3×3	4×3	3×3	4×3	3×3	4×3
Map size	3×3	4×3	3×3	4×3	3×3	4×3
Team size	3 agents	4 agents	3 agents	4 agents	3 agents	4 agents
Task scenarios	S1 .. S5	S1 .. S5	S1 .. S5	S1 .. S5	S1 .. S5	S1 .. S5

TABLE 2 Experimental setups for Blocks World for Teams BW4T domain

Communication model	Selective		No Communication		Communication All	
	3	6	3	6	3	6
Map size	3	6	3	6	3	6
Team size	3 agents	4 agents	3 agents	4 agents	3 agents	4 agents
Task scenarios	S1 .. S5	S1 .. S5	S1 .. S5	S1 .. S5	S1 .. S5	S1 .. S5


4.1 | Experimental Setup

We conducted a study involving two team compositions: a 3-agent team and a 4-agent team. Each team performed two tasks with five different scenarios (S1, S2, S3, S4, S5): epistemic goal, non-epistemic goal, commander (broadcast communication), commander (non-broadcast communication), and blocked cells. These tasks run within two simulated domains: Gridworld and Blocks World for Teams (BW4T). We used two map sizes for Gridworld domain: 3×3 and 4×3; two map sizes for Blocks World for Teams (BW4T) domain: three rooms and six rooms. The setups are as shown in Tables 1 and 2.

We employed three communication models: (1) Selective communication model where agents can communicate based on SMM reasoning; (2) No Communication baseline model where agents cannot communicate as we dropped that action from the domain file; and (3) Communication All baseline model where agents must communicate every new observation—that is, they do not track an SMM.

We ran five task scenarios for each communication model for two maps combined with two team sizes, giving us 60 combinations. All the experiments were performed in the epistemic planner proposed by⁴⁷. These instances were run on Intel Core i7-4770 CPU @ 3.40GHz.

To complete a task successfully, agents need to plan both their actions and communications. A sequential plan for such a task is a sequence of actions where executing them will lead to a goal. The planner is responsible for allocating tasks for agents. For simplicity, we use a sequential planner and assume that every agent takes one turn after another.

**FIGURE 1** The Gridworld (3×3) map.

4.1.1 | Gridworld Domain - Search Task

Gridworld is a simulated world, with a predetermined size and obstacles, consisting of a rectangular subset of the integer plane⁴⁹ (Figure 1). Each point is a position that is occupied by no more than one agent at a time. The agent can move onto any of the free neighbors of the currently occupied position. The distance between adjacent positions is equal to one unit. Each agent has a map of the world and can find their way to the goal with a planned path. Planning an agent's path is beyond the scope of this paper. As our tasks are motivated by search and rescue operations, we modified the Gridworld domain by adding a specified number of survivors (three survivors) in different positions. Thus agents are required to search positions of them on the grid. Agents do not have prior knowledge of the distribution of survivors, and thus, they must observe the grid and communicate effectively. At the outset, agents are distributed at known starting positions.

Each agent must perform a set of actions that we encoded within a domain file to accomplish the goal. To move in any direction, agents used the *move* action. Additionally, they used the *observe* action to observe a survivor's position and the *communicate* action to broadcast it to others and update their beliefs.

Epistemic Goal Scenario

In this scenario, we specified the goal in an epistemic form related to agents' beliefs. Therefore, the goal was to search the whole grid and make all agents know all survivors' positions.

Non-epistemic Goal Scenario

In this scenario, we specified the goal in a non-epistemic form related to simple facts about an environment. Therefore, the goal was only to search for the whole grid.

Commander (broadcast communication) Scenario

In this scenario, we specified one agent as a commander who always stays in one position and does not actively search. Therefore, the goal was to have agents check all positions and have only the commander be aware of all survivors' positions. This means that individual search agents will observe the survivors' positions, even if that is not a part of the goal, and then communicate them to the commander (broadcast communication).

Commander (non-broadcast communication) Scenario

Like the previous scenario, the only difference was that only one individual search agent could communicate all survivors' positions to the commander (non-broadcast communication).

Blocked Cells Scenario

In this scenario, we added another dimension of uncertainty by specifying one survivor position and two cells to be blocked. Agents do not have prior knowledge about the distribution of both the survivor and blocked cells. Therefore, the goal was for agents to reach a specific position while searching for the survivor's position and avoiding blocked cells.

4.1.2 | Blocks World For Teams (BW4T) Domain - Search and Retrieval Task

Blocks World For Teams (BW4T) is a simulated environment that consists of a number of rooms in which colored blocks are distributed randomly⁵⁰ (Figure 2). Agents must search, locate, and retrieve blocks from rooms and put them in a drop-zone in a specific order. An agent can only carry one block at one time. At the outset, agents are distributed at known starting locations. They operate within a potentially unknown environment. They know the locations of rooms (i.e., the map), but they do not know

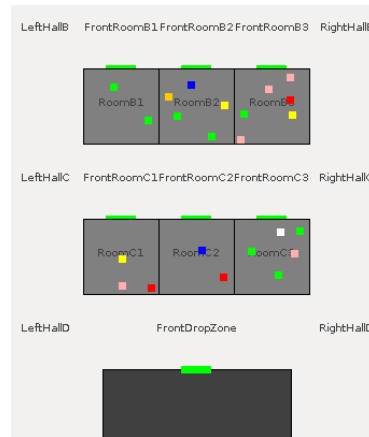


FIGURE 2 Blocks World for Teams BW4T (six rooms) map.

the blocks' location. This knowledge is obtained when the agent visits rooms that have these blocks and observe. Thus, agents need to communicate with each other to achieve this task effectively.

As our tasks are motivated by search and rescue operations, agents must search in rooms and retrieve two specific color blocks (higher priority elements to be rescued) and bring them to the drop-zone. There is a total of four blocks of various colors on the map. Each agent was required to perform a set of actions that we encoded within a domain file to accomplish the task, corresponding to the actions in the BW4T simulator. The agent uses the *goTo* action and the *goToDrop* action when they deliver a block to the drop-zone to move between rooms. To pick up the block, they use the *pickUp* action and *putDown* action to drop it at the drop-zone. They use the *observe* action to observe and believe survivor's positions. To broadcast it for other agents and update their beliefs, they use the *communicate* action.

Epistemic Goal Scenario

In this scenario, we specified the goal in an epistemic form related to agents' beliefs. Therefore, the goal was to search rooms and make all believe that required color blocks are brought to the drop-zone.

Non-epistemic Goal Scenario

In this scenario, we specified the goal in a non-epistemic form that relates to an environment. Therefore, the goal was to search in rooms and bring the required color blocks to the drop-zone.

Commander (broadcast communication) Scenario

In this scenario, we specified one agent as a commander who always stays in one location and does not actively search. Therefore, the goal was for agents to search all rooms, and the commander believes that all required color blocks are brought to the drop-zone. This means that individual search agents observe the required color blocks' locations and bring them to the drop-zone, even if that is not a part of the goal. Then they must communicate it to the commander (broadcast communication).

Commander (non-broadcast communication) Scenario

As in the previous scenario, the only difference was communicating with the commander as it has to be restricted to one known agent who allowed to share the commander what is being delivered (non-broadcast communication).

Blocked Cells Scenario

In this scenario, we added another dimension of uncertainty by specifying three-color locations and one blocked room where agents do not have prior knowledge about the distribution of both blocks and blocked room. Therefore, the goal was for agents to search in rooms and bring two required color blocks to the drop-zone while observing a blocked room.

4.1.3 | Baseline1: No Communication Model

No Communication baseline model is implemented by eliminating explicit communication action from a domain file in some scenarios, including epistemic and non-epistemic goal scenarios, or by restricting communication action to a place where a commander stays, that is, for commander scenarios. This baseline model is motivated by some prior work^{51,3}.

In this baseline, despite eliminating explicit communication process, agents can still obtain some information about their team members due to interference among agents in their shared environment. Namely, where agents have common knowledge at the beginning of a task and a shared plan, they will not go through the same state twice.

4.1.4 | Baseline2: Communication All Model

Communication All baseline model is implemented by forcing the planner to communicate each time an agent obtains a new observation. This would give us an upper bound for applying an epistemic framework to building an effective communication model. That is what would happen without having an SMM.

4.2 | Data Collection

We collected data on various items of simulation experiments, measuring the following:

- Completion Time: Time it takes a team to solve a planning task.
- Total number of actions: We measured the total number of actions it takes a team to complete a task.
- Total number of communications: We measured the total number of messages exchanged by agents on a task.
- Level of sharedness: We measured the consistency between agents' mental models over a task.

To illustrate how we measured SMM on a task using the concept of subject overlap (refer to Section 2.5), we used an example from the Gridworld domain (epistemic goal scenario). The map consists of positions and survivors; thus, we constructed some queries Q regarding, e.g., observed positions and observed survivors. For instance, does position 1 have a survivor? The answer to such queries will be compared with answers given by other models. Given that there are nine positions, for each position, whether it is observed or not, and whether it has a survivor or not, we formulated 36 queries related to the goal. Taking this set of queries Q , we calculated the number of queries answered by all agents' mental models, which was three out of 36. The subject overlap was $3/36 = 8.33\%$ (as shown in Table 3).

In some cases, SMM is articulated into sub-models; thus, we also measured SMM between pairs, as in commander task scenarios.

We used a completion time and a total number of actions collectively as a proxy for team performance. The total number of communications indicates the coordination level among agents. We measured the level of sharedness for all runs to see its effect on team performance, particularly in the case of the Selective communication model. Also, we measure an average ratio as a relative effectiveness measure between approaches.

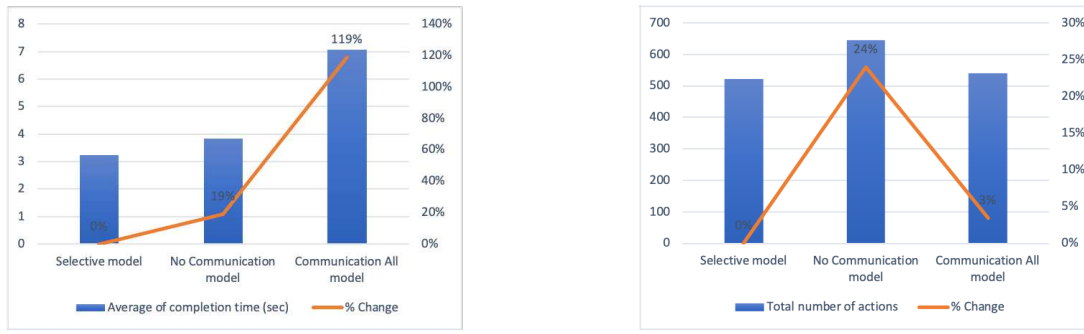


FIGURE 3 The team performance of Gridworld domain, higher is worse

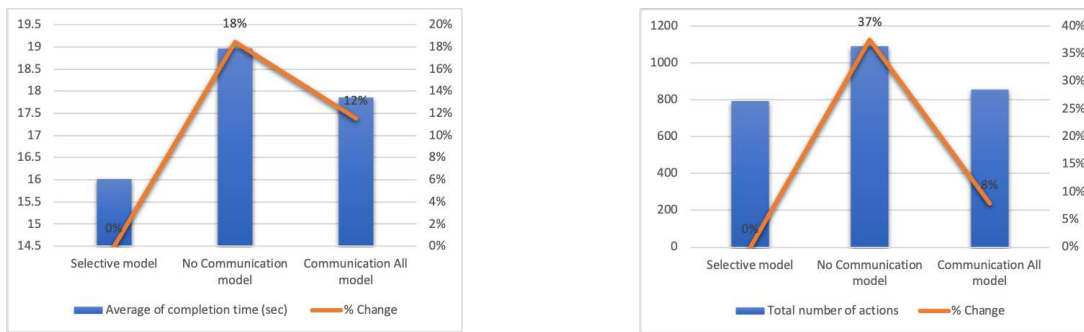


FIGURE 4 The team performance of Blocks World for Teams (BW4T) domain, higher is worse

5 | RESULTS

In this section, we first report our findings from our simulation experiments that evaluate the proposed model's efficacy against two generated baselines. Then we discuss the results of the case studies in detail. We end with a discussion of our findings.

The team's performance was assessed based on a number of measurements (refer to Section 4.2). The overall performance was measured in terms of completion time and the total number of actions collectively. We collapsed the results of the five task scenarios for each domain for simplicity and gave the average completion time and the total number of actions. The low completion time and the total number of actions indicate high team performance on the task. Overall, teams using the Selective communication model exhibited, on average, lower completion times across all task scenarios with shorter plans (total number of actions) compared with the two other baselines (shown in Figures 3 and 4). We show the relative effectiveness between models using a percentage change of the Selective-to-No Communication model and the Selective-to-Communication All model. A higher percentage indicates a lower team performance. In the Gridworld domain, as shown in Figure 3, the trend percentages reflect worse performance for baseline models compared to the Selective model by 19% and 119% respectively in terms of completion time, and by 24% and 3% respectively in terms of the total number of actions. For the BW4T domain, as shown in Figure 4, the trend percentages reflect worse performance for baseline models compared to the Selective model by 18% and 12% respectively in terms of completion time, and by 37% and 8% respectively in terms of the total number of actions.

We hypothesized that a team could communicate more *selectively* if the members represent an SMM with minimal loss of performance compared with baseline communication models. Generally, we found a significant effect of presenting the SMM on the Selective communication model's team performance. In the two domains, as shown in Tables 3 and 4, presenting SMMs helps agents to communicate *selectively* as the model shows better performance in terms of the total number of actions and communications compared with the two baseline models. That is also with a marginally better time completion of the

tasks. However, the fewer messages presented by the No Communication model comes at the cost of the total number of actions. Agents using the No Communication model need to observe each state to get information, resulting in redundant actions. The Communication All model shows a comparable performance in terms of the total number of actions compared with the Selective model. However, that comes at the cost of a higher number of communications and completion times. Agents using the Communication All model communicate each new piece of information, resulting in unnecessary communications, exchanging information not required for other agents to complete their tasks.

The Selective model's performance comes from agents' ability to communicate when necessary to keep track of SMMs, thus, improving the team's overall performance. These findings are in line with³, who found that communicating only valuable information in cooperative tasks helps to maintain comparable performance with fewer communications. However,³ implemented the communication process separately from the planning model. Our proposed model utilized the epistemic planner⁴⁷ to model communication as an action; thus, the planning and communication occur in one compact model.

5.1 | Scenario Results

Tables 3 and 4 show detailed results for each domain case study—Gridworld and BW4T—over five task scenarios (refer to sections 4.1.1 and 4.1.2 for task details).

5.1.1 | Epistemic Goal Scenario

In this scenario, it can be seen from Tables 3 and 4 that the performance of the Selective communication model is marginally better in terms of completion time compared with the Communication All model, which has roughly the same number of actions and communications. In the No Communication model, it is clear that despite lower completion times compared to other models, it is more costly in terms of the total number of actions.

5.1.2 | Non-epistemic Goal Scenario

In this scenario, Selective communication and No Communication models show a roughly similar performance. However, the latter is slightly better in completion time since there is no need to communicate as the goal is only to search. Agents do not need to tell each other the value of positions on the map to achieve the task, and if they do so, the team performance will worsen, as we see in the Communication All model.

5.1.3 | Commander (broadcast communication) Scenario

This scenario shows that the Selective model's team performance is better than the baseline models. Agents using the No Communication model show worse performance as they must return to the commander agent's position each time to update their beliefs.

5.1.4 | Commander (non-broadcast communication) Scenario

This scenario shows that the Selective model's team performance is better than the baseline models. With comparable total number of actions and communications, it is shown a better performance in terms of completion time.

TABLE 3 Experimental results of the Gridworld domain

	Map Size	Selective Model				No Communication Model				Communication All Model			
		3*3		4*3		3*3		4*3		3*3		4*3	
	Agents	3	4	3	4	3	4	3	4	3	4	3	4
Epistemic Goal Scenario	Completion time(sec)	1.19	1.91	2.03	3.02	1.13	1.47	1.93	2.69	1.42	2.09	2.27	3.51
	Total # actions	18	17	24	23	28	35	36	44	18	17	24	23
	Total # Communication	3	3	3	3	0	0	0	0	3	3	3	3
	Level of sharedness (%)	8.33%	8.33%	6.25%	6.25%	8.33%	8.33%	6.25%	6.25%	8.33%	8.33%	6.25%	6.25%
Non-epistemic Goal Scenario	Completion time(sec)	1.23	1.78	2.04	3.1	1.03	1.47	1.78	2.61	1.51	2.07	2.3	3.3
	Total # actions	15	14	21	20	15	14	21	20	18	17	24	23
	Total # Communication	0	0	0	0	0	0	0	0	3	3	3	3
	Level of sharedness (%)	0%	0%	0%	0%	0.00%	0.00%	0.00%	0.00%	8.33%	8.33%	6.25%	6.25%
Commander (broadcast communication) Scenario	Completion time(sec)	1.34	2.27	2.15	3.09	2.96	5.18	5.44	9.26	1.75	3.01	2.9	4.63
	Total # actions	19	19	25	25	21	24	28	31	19	19	25	25
	Total # Communication	3	3	3	3	3	3	3	3	3	3	3	3
	Level of sharedness (%)	7.69%	7.69%	5.88%	5.88%	7.69%	7.69%	5.88%	5.88%	7.69%	7.69%	5.88%	5.88%
Commander (non-broadcast communication) Scenario	Completion time(sec)	1.34	2.02	2.23	3.4	3.03	5.18	5.49	9.38	7.74	15.45	13.76	30.69
	Total # actions	23	23	30	30	33	34	42	42	23	23	30	30
	Total # Communication	3	3	3	3	3	3	3	3	3	3	3	3
	Level of sharedness (%)	7.69%	7.69%	5.88%	5.88%	7.69%	7.69%	5.88%	5.88%	7.69%	7.69%	5.88%	5.88%
Blocked Cells Scenario	Completion time(sec)	4.11	6.7	8.02	11.68	2.42	3.4	4.58	6.41	5.61	9.33	10.32	17.76
	Total # actions	33	41	47	54	37	40	49	52	36	37	54	54
	Total # Communication	3	4	3	5	0	0	0	0	8	8	11	11
	Level of sharedness (%)	5.56%	5.56%	5.21%	5.21%	5.56%	4.17%	4.17%	3.13%	11.11%	11.11%	11.46%	11.46%

5.1.5 | Blocked Cells Scenario

This scenario shows that the No Communication model achieves a marginally better performance than other models. Agents tend not to communicate due to having common knowledge and a shared plan at the beginning of the task, in which an agent will not follow another agent to the same state.

5.2 | Discussion

This study identified the effectiveness of the Selective communication model that contributed to team performance across different task scenarios compared to the No Communication and Communication All baseline models. In some task scenarios, the Selective model shows more sensible results for the BW4T domain than the Gridworld in terms of the total number of communications and level of sharedness. For instance, in the epistemic goal scenario, communication's selectivity is more apparent with different values across different team compositions and map sizes, unlike its results for the Gridworld domain. This may be because the task goal of the BW4T is to search and retrieve two specific color blocks from four blocks on the map. Thus, we determined that agents should communicate just what is relevant to the goal.

In addition, we found that epistemic planning helps more to solve problems that have goals in an epistemic form as it is the only way to solve such problems. However, it is also useful to solve problems whose goal is non-epistemic, such as blocked cell scenarios. Agents using the No Communication model achieved comparable performance in tasks having a non-epistemic goal. This could result from having common knowledge at the beginning of the task and a shared plan preventing an agent from following another agent to the same state. It is an important finding because it indicates that the Selective model may be particularly more beneficial for effective teamwork when presenting epistemic goals. One possible explanation for these results is that where there is no need to communicate (i.e., update other beliefs as a part of the goal), agents using the No Communication model will achieve comparable performance. This is because a team coordinates its actions implicitly based on the assumption

TABLE 4 Experimental results of the BW4T domain

	Map Size	Selective Model				No Communication Model				Communication All Model			
		3 rooms		6 rooms		3 rooms		6 rooms		3 rooms		6 rooms	
	Agents	3	4	3	4	3	4	3	4	3	4	3	4
Epistemic Goal Scenario	Completion time(sec)	4.88	7.55	16.69	24.88	4.84	6.81	14.98	23.98	5.18	7.83	16.84	26.4
	Total # actions	37	41	43	47	94	117	98	121	40	41	46	49
	Total # Communication	2	4	2	3	0	0	0	0	4	4	4	4
	Level of sharedness (%)	8%	15%	4%	6%	8%	8%	4%	4%	15%	15%	8%	8%
Non-epistemic Goal Scenario	Completion time(sec)	6.07	9.84	21.67	32.01	5.71	8.87	19.93	33.37	6.52	9.79	20.88	35.38
	Total # actions	31	29	37	35	31	29	37	35	38	36	44	42
	Total # Communication	0	0	0	0	0	0	0	0	2	2	2	2
	Level of sharedness (%)	0%	0%	0%	0%	0%	0%	0%	0%	8%	8%	4%	4%
Commander (broadcast communication) Scenario	Completion time(sec)	5.29	8.37	17.21	26.41	8.66	14.63	28.64	44.78	6.13	10.52	20.5	32.74
	Total # actions	37	37	42	42	48	54	45	45	40	38	46	44
	Total # Communication	2	2	2	2	2	2	2	2	4	4	4	4
	Level of sharedness (%)	8%	8%	4%	4%	8%	8%	4%	4%	15%	15%	8%	8%
Commander (non-broadcast communication) Scenario	Completion time(sec)	5.22	7.83	17.36	27.9	8.23	13.38	26.1	45.08	6.07	10.45	19.31	29.46
	Total # actions	43	57	42	51	48	50	45	55	45	42	46	51
	Total # Communication	2	2	2	2	2	2	2	2	4	4	4	4
	Level of sharedness (%)	8%	8%	4%	4%	8%	8%	4%	4%	15%	15%	8%	8%
Blocked Cells Scenario	Completion time(sec)	7.11	11.16	23.81	38.97	6.51	10.31	20.81	33.79	8.28	13.84	27.65	43.47
	Total # actions	30	26	45	41	28	26	43	41	36	30	54	48
	Total # Communication	2	0	2	0	0	0	0	0	4	4	7	7
	Level of sharedness (%)	5%	0%	3%	0%	0%	0%	0%	0%	11%	11%	10%	10%

that action effects are deterministic.¹⁰ stated that relying on an epistemic planner's ability to plan for each agent as if it were another can elegantly solve problems without the necessity of explicit coordination. We can clearly see the importance of such a planning model in domains where there is a risk of information being explicitly communicated, for instance, in military operations.

6 | CONCLUSION AND FUTURE WORK

This paper aims to evaluate the efficiency of using a dynamic epistemic logic (DEL)-based planning approach to model communication processes through simulated experiments of multi-agent cooperative situations. The focus is mainly on search problems motivated by applications, such as search and rescue teams. Simulation domains were given to test the effects of the proposed model (Selective communication model) on team performance in incomplete information about the domain. In order for team members to work effectively towards achieving the tasks, they aligned the SMM that keeps track of consistent beliefs. They thus cooperated and made communication decisions effectively.

This paper modeled a communication process to be a part of planning as a natural action, in which its effect is to update agents' beliefs about the world. Thus, the planner communicates and keeps track of the SMM in one compact model. Our first step was to empirically evaluate how the communication model impacts team performance by characterizing the relationship between the degree of shared knowledge, communication actions, and team performance. We then compared our proposed model, the Selective model, with two baseline models generated from the proposed model. We found that the Selective model contributed more to team performance across all task scenarios than baseline models. This study is a first step towards the empirical evaluation of modeling communication explicitly as a part of planning in cooperative tasks and assessing its impact on team performance.

However, there are some limitations to our study that could be a further future area of work. First, we considered only two relatively-simple environments; looking at domains in which there are many competing goals and agents who must choose these goals would be a valuable extension. A further extensive evaluation is required for our proposed model using some simulation environments such as the Blocks World for Team (BW4T) simulator⁵⁰. We could also benchmark it against prior work of modeling communication proposed by³.

As we only ran one planner, we could get different results with different planners. Although the planner that we have used tends to minimize the cost, it was not optimal. Furthermore, as the epistemic planner used can reason nested beliefs (i.e., levels of knowledge) where it is bounded to three⁴⁷, we only evaluated level one of knowledge, that is what agents believe about the world. Future work is needed to evaluate level two, and three of knowledge⁵², and assess its impact on modeling communication and performing tasks.

We mainly focused on the communication content in terms of beliefs. A mental model consists of different components, such as intentions, goals, and beliefs, each with its weight and power to team performance⁵³. Results from prior work of measuring team performance with sharedness levels of various mental model components proved their contribution to team performance¹³. We could further investigate the impact of communicating different mental model components, such as intentions and goals, on team performance.

Finally, we evaluated the Selective model on team performance based on the assumptions that agents are cooperative and commonly know the effects of actions. One avenue for further study would be to evaluate that model where agents are in a non-cooperative mode in which their goals are not shared.

CONFLICT OF INTEREST

We have no conflicts of interest to disclose.

REFERENCES

1. Tim Miller, Adrian R Pearce, Liz Sonenberg in *Foundations of trusted autonomy*, Springer, Cham, **2018**, pp. 67–86.
2. William B Rouse, Nancy M Morris, *Psychological bulletin* **1986**, *100* (3), 349.
3. Vaibhav V Unhelkar, Julie A Shah, in *AAAI*, **2016**, pp. 2544–2550.
4. Beng-Chong Lim, Katherine J Klein, *Journal of organizational behavior* **2006**, *27* (4), 403–418.
5. John E Mathieu, Tonia S Heffner, Gerald F Goodwin, Eduardo Salas, Janis A Cannon-Bowers, *Journal of applied psychology* **2000**, *85* (2), 273.
6. Rineke Verbrugge, *Journal of Philosophical Logic* **2009**, *38* (6), 649–680.
7. Graeme Best, Michael Forrai, Ramgopal R Mettu, Robert Fitch, in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, **2018**, pp. 1050–1057.
8. Thomas Bolander, Mikkel Birkegaard Andersen, *Journal of Applied Non-Classical Logics* **2011**, *21* (1), 9.
9. Liz Sonenberg, Tim Miller, Adrian Pearce, Paolo Felli, Christian Muise, Frank Dignum, *arXiv preprint arXiv:1602.06483* **2016**.
10. Thorsten Engesser, Thomas Bolander, Robert Mattmüller, Bernhard Nebel, *arXiv preprint arXiv:1703.02196* **2017**.
11. François Schwarzentruher, Ph.D. thesis, Habilitation Thesis, University of Rennes, <http://people.irisa.fr/Francois...>, **2019**.

12. Maaike Harbers, Catholijn Jonker, Birna Van Riemsdijk, in *Proceedings of the 4th Annual Human-Agent-Robot Teamwork Workshop*, **2012**, pp. 1–2.
13. Maaike Harbers, MB van Riemsdijk, CM Jonker, in *Proceedings 14th International Workshop on Coordination, Organisations, Institutions and Norms*, **2012**, pp. 106–120.
14. Susan Mohammed, Lori Ferzandi, Katherine Hamilton, *Journal of Management* **2010**, 36 (4), 876–910.
15. Sharolyn Converse, JA Cannon-Bowers, E Salas, *Individual and group decision making: Current issues* **1993**, 221.
16. Elliot E Entin, Daniel Serfaty, *Human factors* **1999**, 41 (2), 312–325.
17. Cheryl A Bolstad, Mica R Endsley, in *proceedings of the human factors and ergonomics society annual meeting*, SAGE Publications Sage CA: Los Angeles, CA, **1999**, pp. 213–217.
18. Ronal Singh, Liz Sonenberg, Tim Miller, in *International Conference on Autonomous Agents and Multiagent Systems*, Springer, **2016**, pp. 163–179.
19. Tim Miller, Jens Pfau, Liz Sonenberg, Yoshi Kashima, *Journal of Artificial Intelligence Research* **2017**, 58, 859–904.
20. Michael Kinney, Costas Tsatsoulis, *Applied intelligence* **1998**, 9 (1), 71–91.
21. Sirui Li, Weixing Sun, Tim Miller, in *International Workshop on Coordination, Organizations, Institutions, and Norms in Agent Systems*, Springer, **2015**, pp. 224–241.
22. Abhizna Butchibabu, Christopher Sparano-Huiban, Liz Sonenberg, Julie Shah, *Human factors* **2016**, 58 (4), 595–610.
23. Ling Zhou, Jian Luo, in *2012 7th International Conference on Computer Science & Education (ICCSE)*, IEEE, **2012**, pp. 169–174.
24. Piotr Gmytrasiewicz, Sarit Adhikari, in *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, International Foundation for Autonomous Agents and Multiagent Systems, **2019**, pp. 1985–1987.
25. E Hall, W Regian, *AFOSR Laboratory Research Initiation Request* **1996**.
26. Eileen B Entin, Elliot E Entin, in *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, SAGE Publications Sage CA: Los Angeles, CA, **2000**, pp. 73–76.
27. Claudia V Goldman, Shlomo Zilberstein, in *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, ACM, **2003**, pp. 137–144.
28. Jean MacMillan, Elliot E Entin, Daniel Serfaty, *Team cognition: Process and performance at the interand intra-individual level. American Psychological Association, Washington, DC. Available at http://www.aptima.com/publications/2004_MacMillan_EntinEE_Serfaty.pdf* **2004**.
29. Catholijn Jonker, M van Riemsdijk, Bas Vermeulen, *Coordination, organizations, institutions, and norms in agent systems vi* **2011**, 132–151.
30. Jaakko Hintikka, *Knowledge and belief : an introduction to the logic of the two notions.*, of *Contemporary philosophy*, Cornell University Press, **1962**.
31. R Fagin, JY Halpern, Y Moses, *MY Vardi Reasoning About Knowledge*, **1995**.
32. Hans van Ditmarsch, Wiebe van der Hoek, Barteld Kooi, *Dynamic Epistemic Logic* **2007**, 11–42.
33. Alexandru Baltag, Lawrence S Moss, Sławomir Solecki in *Readings in Formal Epistemology*, Springer, **2016**, pp. 773–812.
34. Horacio Arló Costa, *Studia Logica* **2002**, 71 (1), 87–118.

35. Jing Li, *Journal of Economic Theory* **2009**, 144 (3), 977–993.
36. Stephen M Majercik, Michael L Littman, *Artificial Intelligence* **2003**, 147 (1-2), 119–162.
37. Hector Geffner, Blai Bonet, *Synthesis Lectures on Artificial Intelligence and Machine Learning* **2013**, 8 (1), 1–141.
38. Benedikt Löwe, Eric Pacuit, Andreas Witzel, et al., *LORI* **2011**, 6953, 179–192.
39. Mikkel Andersen, Thomas Bolander, Martin Jensen, *Logics in Artificial Intelligence* **2012**, 94–106.
40. Thomas Bolander, *arXiv preprint arXiv:1703.02192* **2017**.
41. Leslie Pack Kaelbling, Michael L Littman, Anthony R Cassandra, *Artificial intelligence* **1998**, 101 (1), 99–134.
42. Francisco S Melo, MT Spaan, in *Proc. 1st Eur. Workshop on Multiagent Systems*, **2011**.
43. Christopher Amato, George Konidaris, Shayegan Omidshafiei, Ali-akbar Agha-mohammadi, Jonathan P How, Leslie P Kaelbling, in *2015 AAAI Fall Symposium Series*, **2015**.
44. Andreas Herzig, Jerome Lang, Pierre Marquis, in *IJCAI*, **2003**, pp. 1067–1072.
45. Guillaume Aucher, Thomas Bolander, in *IJCAI-International Joint Conference in Artificial Intelligence-2013*, **2013**.
46. Thomas Bolander, Thorsten Engesser, Bernhard Nebel, *Distributed and Multi-Agent Planning (DMAP-16)* **2016**, 42–49.
47. Christian J Muise, Vaishak Belle, Paolo Felli, Sheila A McIlraith, Tim Miller, Adrian R Pearce, Liz Sonenberg, in *AAAI*, **2015**, pp. 3327–3334.
48. Filippos Komonis, Hector Geffner, in *ICAPS*, **2015**, pp. 147–155.
49. Mark Boddy, Thomas L Dean, *Solving time-dependent planning problems*, Brown University, Department of Computer Science, **1989**.
50. Matthew Johnson, Catholijn Jonker, Birna Van Riemsdijk, Paul J Feltovich, Jeffrey M Bradshaw, in *International Workshop on Engineering Societies in the Agents World*, Springer, **2009**, pp. 254–256.
51. Changyun Wei, Koen V Hindriks, Catholijn M Jonker, in *ICAART (2)*, **2014**, pp. 28–39.
52. Harmen De Weerd, Rineke Verbrugge, Bart Verheij, in *ECSI*, **2014**, pp. 128–136.
53. Catholijn M Jonker, Birna van Riemsdijk, Iris van de Kieft, Maria L Gini, in *AAAI Fall Symposium: Robot-Human Teamwork in Dynamic Adverse Environment*, **2011**.