The Role of Knowledge in Intelligent Agents

Ingo J. Timm, Lukas Reuter

Trier Research Laboratory for Simulation
Center for Informatics Research and Technology (CIRT)
Trier University, Behringstraße 21, 54296 Trier, Germany
{itimm,reuter}@uni-trier.de

Abstract. Artificial intelligence systems embedded in a world of big data, enriched by learning can theoretically extend their knowledge infinitely. However, the computational efficiency of an artificial intelligence system highly depends on the size of its knowledge base as well as the complexity of its inferencing mechanisms. Bounded capacity of the mind is a well-known problem in human cognition and psychology, as well. In an interdisciplinary research program, intentional forgetting as mean for knowledge dynamics is in question to address these limitations.
The AdaptPRO project focuses on intentional forgetting in context of human and multiagent teams. In this paper we investigate the specific aspects of knowledge dynamics and their implications on intentional forgetting in multiagent systems, i.e., in intelligent agents as their constituting elements. Therefore, the knowledge structure within and knowledge distribution between agents are discussed. While the dynamics of increasing the amount of available knowledge or for updating information are widely researched, reducing the amount of available knowledge is rarely discussed. We analyze the requirements for integrating intentional forgetting within intelligent agents and propose a first specification with focus on action- and plan-related knowledge.

Keywords: Knowledge Dynamics · Intelligent Agents · Intentional Forgetting

1 Introduction

Humans, embedded in the *real-world* have access to almost unlimited amount of information through perception of their surrounding environment. While perceiving the environment, humans process huge amounts of audio-visual perceptions and derive an internal world-model as a basis for their cognitive decision-making. Efficient mechanisms for handling this information have been established in human evolution, e.g., integration, aggregation, abstraction, focusing, learning, and forgetting. Information overload or cognitive overload is in discussion since the early 70ies as an effect of *receiving too much information* in organizations. Eppler and Mengis identify the following causes for information overlaod: *"information itself (its quantity, frequency or itensity, and quality or general characteristics), the person receiving, processing or communicating information, the tasks or processes which need to be completed by a person, team* *or organization, the organizational design (i.e., the formal and informal work structures), and the information technology that is used (and how it is used) in a company."* [8]. Tushmann and Nadler (1978) introduce two main variables for influencing information overload: the information processing capacity (IPC) and the information processing requirements (IPR) [34]. While the IPC is mainly determined by personal characteristics of the information processing person, e.g., level of expertise and experience, the IPR is affected by tasks or processes, i.e., organizational design [11]. In order to cope with the limited capacity of the human brain as well as to improve performance issues in case of information overload, humans adapt their knowledge and delete, override, suppress, or sort out outdated information, i.e., they *forget* [3]. Information overload challenges organizational design as the configuration of teams and integration of roles on an individual level are key variables for team's efficiency and reliability. The significance of this problem highly depends on the level of dynamics and standardization. So called knowledge-intense processes, i.e., processes that require specific knowledge in execution, are characterized by frequent autonomous decisions and have a high degree of freedom [14]. Hence, psychologists analyze forgetting as an *intentional* process as well as team members' and team's capacity [6].

In distributed artificial intelligence (DAI), various research on multiagent systems (MAS) and intelligent agents has been conducted. Particularly for implementing such knowledge-intense processes MAS is a promising technology. Agents are associated with autonomous and reactive as well as proactive behavior [22,41,32]. Especially in dynamic environments, agents benefit from their autonomous behavior and can execute actions according to environment changes [17]. In contrast to purely reactive design, Russel and Norvig state that *"A system is autonomous to the extent that its behavior is determined by its own experience "* [22]. Consequently, they define that a *truly* autonomous agent is able to perform actions in a variety of environments with enough time to adapt [22]. When dealing with knowledge adaptation in dynamic environments, a dominant strategy lies in extending agents knowledge bases by learning new facts and rules. However, by extending the knowledge bases, the inference mechanism becomes more inefficient, too. This resembles a state of IO in agents. Therefore, the consideration of presuming or reducing a knowledge base to relevant aspects of an adapted environment becomes an important approach for intelligent, i.e., autonomous agents with respect to their computational limitations. In real-world applications, learning agents could also suffer from information overload. Thus the question arise if theoretical and methodological approaches from psychology can be transferred to agents.

With respect to the close dependency of psychology and computer science on information overload, a new priority research program "Intentional Forgetting in Organizations" (DFG-SPP 1921) including tandem projects from computer scientists and psychologists has been initiated. Our project AdaptPRO [1] as part of this program, addresses organizational, i.e., team aspects of limited capacity, information overload, and intentional forgetting from an organizational

[1] `adaptpro.uni-trier.de`

psychology and DAI perspective. We aim at identifying necessary requirements for formal agent architectures to implement intentional forgetting. Therefore, we focus on the interdisciplinary model, analyzing knowledge structures, knowledge dynamics, and adaptive behavior in intelligent agents and MAS. Using a formal architecture for intelligent agents (Discourse Agents) based on the well-known BDI-paradigm, we illustrate how to extend the formal model to allow for forgetting of action- and task-related knowledge.

The remainder of this paper is structured as follows. We introduce our interdisciplinary theoretical approach to intentional forgetting as a basis for balancing efficiency and reliability in Section 2. In Section 3, knowledge structures and knowledge dynamics in agents are analyzed and shortcomings for adopting intentional forgetting are identified. Finally, Section 4 presents a first step approach for extending a *conventional* deliberative agent architecture with means of intentional forgetting with focus on action- and plan-related knowledge.

2 Adaption by Intentional Forgetting

Intelligent agents allow for modeling of autonomous and dynamic behavior. When dealing with autonomous behavior, knowledge and process capacity of agents are limited. For humans, working with limited capacities is a well-known problem. In the project *Adaptation of Roles and Processes in Organizations (AdaptPRO)*, we focus on these aspects by adopting intentional forgetting in a team from psychology and artificial intelligence research. The next paragraphs give an overview on how knowledge can be organized in individuals as well as teams from a psychological perspective and address the (dis-)advantages.

Team members memorize knowledge required for their tasks, they specialize on particular areas of expertise, or they share knowledge and information with each other [15]. These various approaches to the organization of team knowledge are known as *team cognitions* (see Figure 1) [23]. Team cognitions describe the structure in which knowledge important to team functioning is mentally organized, represented, and distributed within the team and allows team member to anticipate and execute actions [5,6]. Therefore, they are particularly suitable as a theoretical concept for describing, modeling, and analyzing knowledge configuration approaches in collaborative work processes. Team cognition, as an emergent state, are conceptualized as (1) shared or (2) distributed team knowledge [6].

When working together, it is important for team members to share their knowledge about task and team relevant information with each other in the form of *team mental models* to facilitate successful cooperation and coordination [5,6]. On the one hand, this generates trust and increases coordination and the robustness of the work process against disturbances by means of information exchange and the acquisition of group knowledge [39,18]. On the other hand, sharing of the entire knowledge among all team members results in an increased amount of information that needs to be processed by each individual which can lead to information overload [7,8]. Information overload endangers the effectiveness and efficiency of the team as its members struggle to focus on specific current

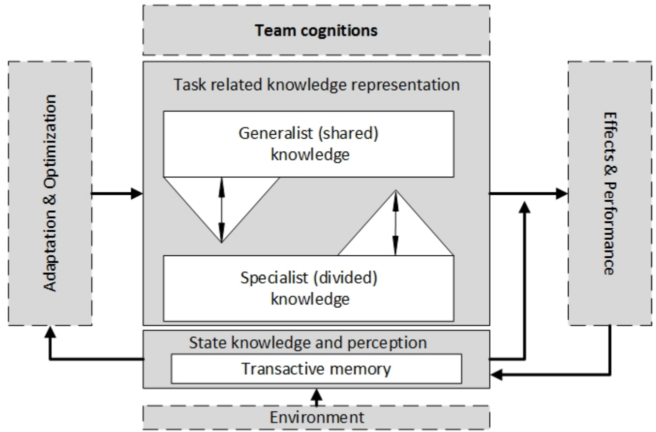

Fig. 1. Interdisciplinary Forgetting Model

tasks when constantly switching between different contexts [6]. Contrastingly, specializing on particular areas of competence reduces the cognitive load faced by members of a team [12]. That is, each team member can focus on its specific expertise which reduces the load of information being processed. This distribution of knowledge in specialized teams increases the overall knowledge capacity of the whole team since individual members only have to memorize and process information which is relevant to their areas of expertise [13]. However, this potentially makes the team as a system more fragile as it lacks the required redundancy of knowledge to avoid confusion, conflicts, and failures [36].

Regarding the knowledge dimensions of sharing and dividing knowledge, we define *intentional forgetting* as a reorganization of knowledge in teams to enhance a team's knowledge capacity. As the previous paragraphs have shown, knowledge distributions have different impacts on team performance and human cognition which resembles capacity issues of knowledge bases and inference mechanism in intelligent agents. The goal is to find a knowledge structure by intentional forgetting which allows for efficient and resilient teamwork.

To integrate this model into MAS it is essential to cover functions for adaption and optimization, task-related and state knowledge, as well as perception. Adaptation at runtime can be performed by (a) modifying an agent's knowledge and (b) situation dependent deliberation on goals, plans, and actions. In the next section, knowledge structure and knowledge dynamics as prerequisites for intentional forgetting are analyzed.

3 Agent's Knowledge

Agents are the key concept of DAI. Even if agents are knowledge processing entities with specific properties like autonomy, reactive, and proactive behavior as well as social deliberation, agents also provide a platform to implement almost any AI technique, e.g., machine learning, planing, informed search, and constraints satisfaction propagation. So they are similar to expert or knowledge-based systems. An agent implements an agent function mapping perceptions to actions. The decision-making within this mapping is determined by the agent architecture [25]. The architecture separates inference knowledge from the domain and task knowledge, i.e, a specific agent architecture could be embedded in different environments. Due to the theoretical considerations of this paper we focus on architectural aspects and knowledge relevant for deliberation rather than discussing specific knowledge elements from a domain.

In this section, we analyze requirements as well as typical elements of an agent's knowledge. As agents are expected to deal with dynamic elements of the environment and non-deterministic actions from other agents, knowledge revision is relevant with respect to intentional forgetting, too. Accompanying the discussion of knowledge structures and dynamics, we introduce a possible formalization within BDI-Logic [21]. The so-called *Discourse Agent* [30] are specified with a multi-modal logic which integrates the approaches VSK-Logic [42] for inter-agent behavior and LORA (Logic Of Rational Agents) [40] for deliberative agent behavior. Whereas, LORA allows to represent and reason about beliefs, desires, intentions, and actions within an agent and how these change over time [40], VSK-Logic takes into account that different agents create different pictures of the world, which may be only partly visible in general and especially for the single agent. Formally a (global) visibility function (visibility) and an agent depending perception function (see) are introduced in addition to local states (knowledge) in terms of (multi-)modal sorted first order logic including the possible worlds semantics.

In this approach, BDI is used as an example as it is the dominant architectural approach to intelligent agents and Discourse Agents are chosen because of their sophisticated specification of the concepts for intentions and plans. Even if the computational complexity prevents a straightforward implementation of such multi modal logics they have proven of use for steering implementation in conventional programming languages.

3.1 Knowledge Structure

Knowledge is widely used in agent design and implementation. From an organizational perspective knowledge can be found at the level of individual agents, agent interaction or in the MAS at a whole. These perspectives will be analyzed in the following paragraphs.

Knowledge in MAS The representation of teams is assumed to be on the level of MAS, i.e., one MAS is intended to represent one team. Of course there are various approaches on groups and teams in MAS, but the creation of the team is not in scope for the underlying psychological conceptualization [40]. MAS are supposed to support information hiding, i.e., knowledge is usually stored in agents and not as generally available information. However, it is important to share common concepts in communication to enable teamwork or information exchange. Furthermore, knowledge on interaction processes and norms for excepted behavior have to be shared, too. In many DAI approaches there is no explicit discussion on where or how to share common knowledge. Common knowledge can be distributed as local copies within each agent or can be part of the environment visible to any agent within the system. The decision on distributing or centralizing knowledge heavily depends on its dynamics: if common knowledge changes regularly, knowledge distribution could suffer by frequent knowledge propagations through the system. Nevertheless centralized knowledge could also become a bottleneck in a high-scale MAS. Next to agent-related knowledge, information on the environment could be part of common knowledge. In MAS, it is assumed that external and internal knowledge representation of agents can differ with respect to communication and reasoning efficiency [33]. Thus, even in the case of centralized knowledge it is plausible that agents also represent this knowledge internally. Knowledge in MAS can be summarized as follows: conceptual knowledge, e.g., shared ontologies [28], normative knowledge, e.g., norms and obligations [35], interaction knowledge, i.e., communication protocols [2], and environmental knowledge. In contrast to the theoretical consideration of the previous section, there is no explicit definition of the MAS capabilities.

The formal specification of the agents' environment usually considers the set of environmental states E which is partially observable by Agents Ag. Limitations in the *visibility* of the environment are explicitly represented by $vis : E \times Ag \rightarrow E$ for each agent. Any of the above mentioned knowledge structures are formally part of the environmental state. However, norms and obligations could implicitly be implemented in means of the state transformer function τ, preventing agents to reason about them. The environment with $e_0 \in E$ as its initial state is specified as $Env = \langle E, Act_1, ..., Act_n, vis_1, .., vis_n, \tau, e_0 \rangle$. In terms of this specification there is no knowledge on the MAS level except for knowledge included in the environmental state: $MAS = \langle Env, Ag_1, ..., Ag_n, \phi \rangle$.

Knowledge in agent interaction A key feature of MAS is dynamic coordination. By means of communication, agents dynamically negotiate on roles and processes such that the organization emerges [10]. The persistence of such emerging organizations range from one touch interaction to establishment of a virtual organization. Usually, agents are coordinated by interaction protocols solving a specific task [16]. Alternatively, a team or a group of agent is formed dynamically if an agent is unable or *unwilling* to complete the task by itself [40]. Two different decision problems are addressed by this: coalition formation, and team formation. In coalition formation agents with homogeneous capabilities are put together to increase efficiency or overcome computational boundaries [24]. Such approaches are researched in context of distributed rational decision-making and mechanism design. In contrast coalition formation, in team formation agents with heterogeneous capabilities are combined to solve a particular task. various approaches are discussed within distributed problem solving [19]. These two decision problems are highly relevant to the psychological team cognition process. In context of analyzing the knowledge structure in MAS, specific elements for implementing cooperation or coordination are available by various authors, e. g., mutual beliefs and joint intentions [40]. However those elements would be part of the agent's knowledge base rather than a separate structure on group level. Distributed or shared memories for teams can easily be implemented by use of the visibility concept introduced on the multiagent level. Thus, no additional specification in context of the Discourse Agents are required, here.

Agent's knowledge The core agent function has to transform perceptions to actions on the environment. Obviously this is a strict simplification as properties like proactiveness and social deliberation are independent from external triggers or happening without – in case of direct communication– environmental manipulation. However both aspects can formally be modeled with the initial assumption. There is almost no limitation to types of knowledge representation, acquisition, and inference within an agent when the expertise model is in question, e.g., the domain dependent or problem-solving knowledge base. These parts of the knowledge base are out of scope, here. More specifically, within the agent, knowledge on the domain and tasks are encapsulated [22] which is deeply connected to action and behavior [19]. The agent's knowledge base is often specified as a local state. There are many variations of internal knowledge representation or local states in agents as specific agent functions require specific structures for knowledge or internal states. In context of BDI, the local state covers knowledge on the current situation, i.e, internal representation of the environment, experience with other agents, self-awareness (beliefs), persistent goals (desires), and committed goals (intentions). Following the Discourse Agent architecture, plans and a relevance function are also part of the local state. By integrating plans as a structure by its own, advance concepts of reflection on this knowledge become possible. The relevance function is an implicit representation of users knowledge on specific goals. From a formal perspective the local state is specified by a 5-tuple : $L = \langle B, D, I, Plan, \gamma \rangle$ where B is a set of beliefs, D is the set of desires, I is the set of intentions, and $Plan$ is the set of plans. γ evaluates the relevance of desires. Intentions are representing goals which an agent has committed to pursue ranging from a simple action sequence, the desire itself, to a tuple consisting of a plan as well as the goal [31,27]. In Discourse Agents the intention is specified as follows $I \subset D \times Plan$. Additional knowledge structures as the priorly mentioned shared ontologies for enabling communication are assumed to be part of B. With respect to the psychological approach of team cognition, task relevant knowledge, e.g., actions and plans, representing skills and capabilities as well as forming roles are in focus.

3.2 Knowledge Dynamics

Agents are inherently dealing with dynamics in knowledge. The environment is assumed to change during agent execution such that an agent continuously has to update its internal world model as well as to analyze if in the current situation another course of action or another goal is desirable. The first part of an agent function consists of perceiving the environment and updating the internal knowledge. As a dominant assumption in agent design it is assumed that an agent can be *wrong* about the environment on its perception or another agent could deliver irrelevant or false information. For a formal foundation of agents, mechanisms to handle inconsistent information have to be implemented. Therefore, many approaches represent the agent's knowledge in a set of beliefs rather than a knowledge base filled up with facts. In consequence logics are required which can differentiate states with respect to time. Multi modal logic with a possible world semantics, e.g., KD45 [37] is often used in DAI. This approach uses a consistent current date but allows for inconsistency over time. Processing the perceptions of the environment does also include the transformation of external information to internal one, widely researched as belief revision [9,1,29]. Even if belief revision is an explicit task within agents deliberation, the adaptation of knowledge by adding, updating, or even contracting (forgetting) information and knowledge is performed unintentionally and do not allow for modeling experience of task-related knowledge.

The update function is not related to the capacity or computational boundaries of the agent. Of course agent implementations can apply the same representation internally (local state) as well as externally (environmental state). After updating its beliefs, an agent has to check if the intentions or other action-related concepts are still valid. In BDI, this step is called intention reconsideration and it evaluates if an intention should be still pursued or dropped. Schut and Wooldridge (2001) analyzed intention reconsideration strategies and point out that in a MAS world intentions can represent commitments between agents such that local decisions have to be propagated to other agents [26]. Additional knowledge revision functions can be designed, e.g., plans, desires, strategy-revision in Discourse Agents. In the next step of the agent function, the decision on the course of action takes place. In BDI, this is associated with reasoning about desires, intentions, plans, and the agent's beliefs. Usually this does not involve the adaptation of knowledge structures except from instantiation of additional intentions or plans. After deciding on the next objectives, an agent is executing the planned actions. In learning agents, the result of action or the result of the deliberation or the execution process can be analyzed and knowledge used by the selection mechanism can be adapted for continuous improvement of the agents behavior. Therefore Russel and Norvig proposed to integrate a critique element into the agent which evaluate agents performance [22]. In context of the Discourse Agent a markov decision process is used as the basis for learning and optimizing interaction behavior.

In our formal framework plans are an important knowledge structure in the local state as they do not only representing the capabilities of agent but also allowing for adaptation: $Plan = \langle \varphi_{pre}, \varphi_{post}, Actions, status, select \rangle$. In analogy of the definition of AI planning, e.g., STRIPS [4] a plan itself is specified with precondition, defining its applicability, a post condition, a set of possible actions and two functions. The functions are hiding the internal behavior of a plan. The complexity of the selection function can vary from simple action selection, from a predefined sequence to continuous planning. In case of markov decision process, experience of interactions is preserved. To forget such experience is possible by resetting the transition probability to defaults.

4 Towards Intentional Forgetting in Multiagent Systems

In the previous section, we have analyzed key concepts for structuring and adapting knowledge in agents. It can be assumed that *intelligent* agents implement those or similar concepts. Thus the question arise if such concepts are sufficient to implement intentional forgetting in MAS. Transferring the approach of intentional forgetting in organization (cf. section 2) to MAS require mechanisms for reconfiguring knowledge distribution and role assignment. The focus of knowledge distribution in the interdisciplinary model lies on action- and task-related knowledge. The concepts of actions and plans correspond to the psychological concepts of skills and capabilities. The actions or skills determine IPC. However, skill improvement or intentional forgetting on skills is not in focus in DAI research. Nevertheless a naive approach to intentional forgetting could (temporarily) delete skills, i.e., restricting the set of available actions. As the amount of actions is determining the branching factor for planning, the computational complexity could drastically be reduced by such an approach. In contrast to the psychological phenomena this does not include improvement or deprecation of skills as a computer program does not forget efficient routines by not applying them on a regular base. If an agent is able to improve its efficiency with respect to a specific behavior this would not be represented as an action but as a plan.

Next to forgetting of skills, reconfiguration of knowledge distribution can be performed by (re-) assigning roles in the system to improve IPR with respect to team capacity. From a psychological perspective a team where any roles assigned to any team has a maximal fault tolerance as any team member can substitute any other team member. In the opposite scenario, each team member is strictly specialized as it has only one assigned role which should lead to maximum efficiency. Roles can be specified by a collection of capabilities which can be represented in MAS as plans. Capabilities is an important topic in multiagent research even if there are only few results available which specifically deal with a sophisticated concept of capability. Many approaches focus on the combination or match making of capabilities to fulfill services. As an early work on this topic Padgham and Lambrix are formalizing capabilities in context of BDI to enable reflection and reasoning about capabilities [20]. White, Tate and Rovatsos propose explicit capability representation for improving plan reliability [38]. Both approaches have relevant elements to be considered in our approach but aim at different research objectives. Thus implementation of dynamic capability sets within an agent to establish an organizational frame is not pursued there. The plan concept of the Discourse Agent can serve as a starting point for representing or transferring capabilities to MAS. However in analogy to White, Tate and Rovatsos, experience about capability usage have to be explicitly included. The proposed confidence on capability execution is related to the adaptive plan with markov decision processes of our Discourse Agent approach [31,38]. Both approaches, suffer from a statistical representation of experience: if in a reorganization process an agent *looses* a role, it should forget about experience related to this role. In a business application this could be a specific customer. However, if the experience with customers is stored in confidence intervals or transition probabilities it cannot be selectively removed.

Thus we propose a different approach for the representation of capabilities: $Plan = \langle \varphi_{pre}, \varphi_{post}, Actions, status, select, history \rangle$, where $history = \langle \varphi_{pre}, \varphi_{post}, Act_1, ..., Act_n, evaluation \rangle$. The main difference lies in an explicit set of experience, where role dependency of experience is preserved. Doing so, it can be modified when the agent looses a role such that the learning algorithm rebuilds its knowledge automatically. The integration of adaptive skill sets, i.e., limitation of available actions, does also require slight modifications in formalization: $Ag = \langle L, Act, Act_{ltd}, see, reflect, decide, execute, l_0 \rangle$. The secondary set of actions Act_{ltd} represents the role dependent limitation of available actions.

5 Conclusion

In this paper we have introduced the information overload problem in knowledge intensive processes and presented an interdisciplinary approach of intentional forgetting as a counter strategy. Analyzing knowledge structures in MAS, groups of agents, and within agents we have discussed typical elements, like actions, plans, and goals, which are significant in this context. Following BDI as a representative architecture for intelligent agents, knowledge dynamics within the agent's deliberation cycle have been derived. The contribution of this paper is the identification of two missing concepts: (a) fixed action set, (b) no explicit connection of roles to experience in adaptive plans. In the last section we specified extensions to the formal structure to overcome this limitation in context of the Discourse Agent architecture.

However, next to the formal structures, behaviors are required for implementing intentional forgetting in MAS. While the effects of role dependent experience has to be handled in specific types of adaptive plans, restricting or extending available actions set with respect to role assignment could be developed on a more general level. Both approaches still have to proof their benefits even if they are psychologically plausible. As a next step, we plan to implement the proposed concepts in a multiagent-based simulation system and evaluate dynamic role assignment in contrast to static role assignment.

Acknowledgments.

The project *AdaptPRO: Adaptive Process and Role design in Organisations* (TI 548/-1) is funded by the German Research Foundation (DFG) within the Priority Program "Intentional Forgetting in Organisations" (SPP 1921). We would also like to thank Jan Ole Berndt, Anna-Sophie Ulfert, Thomas Ellwart, and Conny Antoni for their fruitful discussion

References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: Partial meet contraction and revision functions. The journal of symbolic logic **50**(2), 510–530 (1985)
2. Bellifemine, F., Poggi, A., Rimassa, G.: Developing multi-agent systems with a fipa-compliant agent framework. Software: Practice and Experience **31**(2), 103–128 (2001)
3. Bjork, E.L., Anderson, M.C.: Varieties of goal-directed forgetting. In: Golding, J.M., MacLeod, C. (eds.) Intentional Forgetting: Interdisciplinary Approaches, pp. 103–137. Lawrence Erlbaum: Mahwah, NJ (1998)
4. Bylander, T.: The computational complexity of propositional strips planning. Artificial Intelligence **69**(1-2), 165–204 (1994)
5. DeChurch, L.A., Mesmer-Magnus, J.R.: The cognitive underpinnings of effective teamwork: A meta-analysis. Journal of Applied Psychology **95**(1), 32–53 (2010)
6. Ellwart, T., Antoni, C.H.: Shared and distributed team cognition and information overload. evidence and approaches for team adaptation. In: Information and Communication Overload in the Digital Age. IGI Global (2017)
7. Ellwart, T., Happ, C., Gurtner, A., Rack, O.: Managing information overload in virtual teams: Effects of a structured online team adaptation on cognition and performance. European Journal of Work and Organizational Psychology **24**(5), 812–826 (2015)
8. Eppler, M.J., Mengis, J.: The concept of information overload: A review of literature from organization science, accounting, marketing, mis, and related disciplines. The Information Society **20**, 325–344 (2004)
9. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning about Knowledge. The MIT Press, Cambridge, Massachusetts (1995)
10. Ferber, J., Gutknecht, O., Michel, F.: From agents to organizations: An organizational view of multi-agent systems. In: Giorgini, P., Müller, J.P., Odell, J. (eds.) Agent-Oriented Software Engineering IV. pp. 214–230. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
11. Galbraith, J.R.: Organization design: An information processing view. Interfaces **4**(3), 28–36 (1974)
12. Hollingshead, A.B.: Communication, learning, and retrieval in transactive memory systems. Journal of Experimental Social Psychology **34**, 423–442 (1998)
13. Hollingshead, A.B., Gupta, N., Yoon, K., Brandon, D.P.: Transactive memory theory and teams: Past, present and future. In: Salas, E., Fiore, S.M., Letsky, M.P. (eds.) Theories of team cognition: Cross-disciplinary perspectives, pp. 421–455. Routledge, New York (2012)
14. Isik, Ö., Van den Bergh, J., Mertens, W.: Knowledge intensive business processes: An exploratory study. In: System Science (HICSS), 2012 45th Hawaii International Conference on. pp. 3817–3826. IEEE (2012)
15. Kozlowski, S.W.J., Ilgen, D.R.: Enhancing the effectiveness of work groups and teams. Psychological Science in the Public Interest **7**(3), 77–124 (2006)
16. Krempels, K.H., Spaniol, O., Scholz, T., Timm, I.J., Herzog, O.: Interaction design. In: Kirn, S., Herzog, O., Lockemann, P., Spaniol, O. (eds.) Multiagent Engineering: Theory and Applications in Enterprises, pp. 383–403. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
17. Müller, H.J.: Towards agent systems engineering. Data & Knowledge Engineering **23**(3), 217–245 (1997)

18. Oertel, R., Antoni, C.H.: Phase-specific relationships between team learning processes and transactive memory development. European Journal of Work and Organizational Psychology **24**(5), 726–741 (2015)
19. O'Hare, G.M., Jennings, N.: Foundations of distributed artificial intelligence, vol. 9. John Wiley & Sons (1996)
20. Padgham, L., Lambrix, P.: Formalisations of capabilities for bdi-agents. Autonomous Agents and Multi-Agent Systems **10**(3), 249–271 (May 2005). https://doi.org/10.1007/s10458-004-4345-2, `https://doi.org/10.1007/s10458-004-4345-2`
21. Rao, S., Georgeff, M.P.: Bdi agents: From theory to practice. In: Lesser, V.R., Gasser, L. (eds.) Proceedings of the First International Conference on MultiAgent Systems (ICMAS 1995). pp. 312–319. The MIT Press, Boston (1995)
22. Russell, S.J., Norvig, P., Davis, E.: Artificial intelligence: A modern approach. Prentice Hall series in artificial intelligence, Prentice Hall, Upper Saddle River and NJ, 3rd ed edn. (2010)
23. Salas, E., Fiore, S.M., Letsky, M.P. (eds.): Theories of team cognition – Cross-disciplinary perspectives. Taylor & Francis Group, New York (2012)
24. Sandholm, T., Lesser, V.R.: Coalitions among computationally bounded agents. Artificial Intelligence **94**(1), 99–137 (1997), special Issue on Principles of Multiagent Systems
25. Scholz, T., Timm, I.J., Herzog, O., Görz, G., Schiemann, B.: Semantics for agents. In: Kirn, S., Herzog, O., Lockemann, P., Spaniol, O. (eds.) Multiagent Engineering: Theory and Applications in Enterprises, pp. 431–464. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
26. Schut, M., Wooldridge, M.: The control of reasoning in resource-bounded agents. The Knowledge Engineering Review **16**(3), 215–240 (2001)
27. Simari, G.I., Parsons, S.: On the relationship between mdps and the bdi architecture. In: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems. pp. 1041–1048. ACM (2006)
28. Stuckenschmidt, H., Timm, I.J.: Adapting communication vocabularies using shared ontologies. In: Crane-field, S. (ed.) Proc. of the 2nd Intl. Workshop on Ontologies in Agent Systems, Workshop at 1st International Conference on Autonomous Agents and Multi-Agent Systems. pp. 6–12. Bologna (2002)
29. Tamargo, L.H., García, A.J., Falappa, M.A., Simari, G.R.: Modeling knowledge dynamics in multi-agent systems based on informants. The Knowledge Engineering Review **27**(1), 87–114 (2012). https://doi.org/10.1017/S0269888912000021
30. Timm, I.J., Scholz, T., Herzog, O.: Emerging capabilities in intelligent agents for flexible production control. Advanced Engineering Informatics Journal. Special Issue on Emergent Synthesis **20**(3), 247–259 (2006)
31. Timm, I.J.: Dynamisches Konfliktmanagement als Verhaltenssteuerung Intelligenter Agenten. Dissertation in der Künstlichen Intelligenz (DISKI) Aka Verlagsgruppe (2004)
32. Timm, I.J., Knirsch, P., Kreowski, H.J., Timm-Giel, A.: Autonomy in software systems. In: Understanding Autonomous Cooperation and Control in Logistics, pp. 255–273. Springer (2007)
33. Timm, I.J., Scholz, T., Herzog, O., Krempels, K.H., Spaniol, O.: From Agents to Multiagent Systems. In: Kirn, S., Herzog, O., Lockemann, P., Spaniol, O. (eds.) Multiagent Engineering: Theory and Applications in Enterprises, pp. 35–51. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
34. Tushman, M.L., Nadler, D.A.: Information processing as an integrating concept in organizational design. Academy of management review **3**(3), 613–624 (1978)

35. Vázquez-Salceda, J., Aldewereld, H., Dignum, F.: Norms in multiagent systems: from theory to practice. International Journal of Computer Systems Science & Engineering **20**(4), 225–236 (2005)
36. Wegner, D.M.: Transactive memory: a contemporary analysis of the group mind. In: Mullen, B., Goethals, G.R. (eds.) Theories of group behavior, pp. 185–208. Springer, New York (1987)
37. Werner, E.: Logical foundations of distributed artificial intelligence. pp. 57–117. John Wiley & Sons, Inc., New York (1996)
38. White, A., Tate, A., Rovatsos, M.: Improving plan execution robustness through capability aware maintenance of plans by bdi agents. International Journal of Agent-Oriented Software Engineering **5**(4), 306–335 (2017)
39. Wiedow, A., Konradt, U., Ellwart, T., Steenfatt, C.: Direct and indirect effects of team learning on team outcomes: A multiple mediator analysis. Group Dynamics: Theory, Research, and Practice **17**(4), 232–251 (2013)
40. Wooldridge, M.: Reasoning about Rational Agents. The MIT Press, Cambridge (2000)
41. Wooldridge, M.: An Introduction to MultiAgent Systems. John Wiley & Sons, 2nd edn. (2009)
42. Wooldridge, M., Lomuscio, A.: Multi-agent VSK logic. In: European Workshop on Logics in Artificial Intelligence. pp. 300–312. Springer (2000)