Domain-specific Trust for Context-aware BDI Agents
Preliminary Work

Arthur Casals[1], Eduardo Fermé[2] and Anarosa A. F. Brandão[1]
[1]*Laboratório de Técnicas Inteligentes, EP/USP, Av. Prof. Luciano Gualberto, 158 - trav. 3, São Paulo, SP, Brazil*
[2]*Universidade da Madeira, Campus Universitário da Penteada, 9020-105, Funchal, Portugal*

Keywords: Context-aware Systems, Multiagent Systems, BDI, Trust, Contextual Planning, Experience Sharing, Learning.

Abstract: Context-aware systems are capable of perceiving the physical environment where they are deployed and adapt their behavior accordingly. Multiagent systems based on the BDI architecture can be used to process contextual information in the form of beliefs. Contextual information can be divided and structured in the form of information domains. Information and experience sharing enables a single agent to receive data on different information domains from another agent. In this scenario, establishing a trust model between agents can take into account the relative perceptions each agent has of the others, as well as different trust degrees for different information domains. The objective of this work is to adapt an epistemic model to be used by agents with their belief revision in order to establish a mechanism of domain-specific relative trust attribution. Such mechanism will allow for each agent to possess different trust degrees associated with other agents regarding different information domains.

1 INTRODUCTION

Context-aware systems are capable of capturing information from the environment and using it to adapt their functions accordingly (Hong et al., 2009). The information captured from the environment itself can be referred to as context (Abowd et al., 1999), and it can be represented and used in different ways – depending on which aspects are relevant to the context-aware system using it (Kim and Chung, 2014; Nalepa and Bobek, 2014). Intelligent agent architectures are among the ones that can be used by context-aware systems (Kwon and Sadeh, 2004). The belief-desire-intention (BDI) agent architecture (Rao and Georgeff, 1991) is of particular interest due to its inherent use of environmental information (context) in the form of beliefs. Beliefs are used to determine what the agent has chosen to do (its desires), and how committed it is to that choice (its intentions) (Cohen and Levesque, 1990). Systems composed of multiple intelligent agents are called Multiagent Systems (MAS) (Wooldridge, 2009). The intelligent agents that compose a MAS may interact among themselves, collaborating and exchanging information in order to achieve their objectives.

The information exchanged between agents can be originated from the environment in which each agent is situated, or it can be a representation of experiences derived from each agent's actions over time. Nevertheless, the information exchanged and used by intelligent agents is aggregated to any information that each agent already possesses. In particular, BDI agents use this information as part of their belief revision process - which, in a general sense, refers to the process of altering beliefs to take into account new acquired information (Gärdenfors, 2003). When this new information is conveyed by other agents, questions related to trust and expertise may arise: environmental observations usually can be trusted (acquired by direct observation), but other agents' experiences may be associated with individual trust or reputation metrics (Huynh et al., 2006). At the same time, reputation and trust are subject to pertinence: while a high degree of trust may be associated with a given agent, the experiences it possess can be more or less relevant depending on the subject they refer to.

From an agent's perspective, experience from other agents may be more or less relevant. Since agents in different MAS can use different subsets of all contextual information available, their experience relevance is also bound to the information they used in their reasoning process. Therefore, trust and reliability on experience exchange between agents may also be associated with which subsets - or domains -

Casals, A., Fermé, E. and Brandão, A.
Domain-specific Trust for Context-aware BDI Agents - Preliminary Work.
DOI: 10.5220/0006600602440249
In *Proceedings of the 10th International Conference on Agents and Artificial Intelligence (ICAART 2018) - Volume 1*, pages 244-249
ISBN: 978-989-758-275-2
Copyright © 2018 by SCITEPRESS – Science and Technology Publications, Lda. All rights reserved

of the existing information are actually used by each agent. This situation is similar to what we observe in the real world: information provided by doctors on medical matters and sports are subject to different credibility degrees. While medical matters fall within the domain of expertise of doctors, sports have no correlation to their expertise domain. Therefore, health diagnostics provided by a doctor are highly reliable, but opinions related to sports may not receive the same degree of reliability (since the subject is not pertinent to a doctor's experience). We will use this example along the text to illustrate different concepts and aspects related to the present work.

The objective of this work is to extend an existing epistemic model to be used by context-aware BDI agents in conjunction with their belief revision process. This would allow for a single agent to possess different trust degrees associated with other agents regarding different information domains, in a manner that different credibility parameters can be attributed to received experiences. In that manner, a mechanism of trust attribution can be used in conjunction with experience processing and multiple information domains. Trust models can then be incorporated into the agent's planning process through the attribution of different trust indicators to experiences received from the same agent.

This paper is organized as follows: Section 2 details the general concepts used in this work. An existing epistemic model for multi-source belief revision is presented in Section 3, along with the modifications made to accommodate information domains. Discussions about the proposed model and related work are presented in Section 4. In Section 5 we present our considerations on the proposed model and future work.

2 GENERAL CONCEPTS

Since the trust model is intended to be used in conjunction with context-aware BDI agents, it is important to understand a few concepts involving context and information domains. The process of physically gathering contextual information is not part of the scope of this work.

2.1 Contextual Information

Contextual information can be collected and distributed differently. It can also be detailed and organized in different levels, depending on its intended use. Different constraints can also determine its distribution model, such as interoperability with a pre-existent communication model or bandwidth limitations. Contextual information can be organized in different ways, depending on its purpose. This organization can also differ across different information dimensions. Mobile devices, for example, use different sensors to gather data mostly related to physical environment aspects, such as localization and acceleration. Information on the social information dimension is limited or non-existent. On the other hand, identification cards can retain organizational data – such as role in the company, unique identification record, and clearance level. In this case, the social information dimension is more detailed than in the previous one, while the physical information dimension is almost non-existent.

As a term, ”information domain” is broadly used to refer to different aspects and purposes of information organization (Hjørland, 2002). Generally speaking, information domains can be used to represent deterministic sets of information that are different among themselves in both content and organization (Hennessy, 1991). In the example previously presented, ”soccer” and ”medicine” are examples of two different information domains. Depending on how the information is organized, the content - or what is being represented - can be determined by its own representation. An ontology, for example, can be defined as a set of terms of interest in an information domain, along with the relationships among these terms (Mena et al., 1998).

Therefore, we consider that contextual information can be composed of different information domains. Sensor data gathered and aggregated by an internal sensor network, for example, can be considered as an information domain within a given environment. Another information domain could be represented by user preferences stored and organized in a mobile device. When the user is in the environment, the contextual information is composed by both information domains - which can be used by an agent in its reasoning process (Figure 1).

Information domains can be also used to structure information exchanged between two or more agents - such as past experiences. Different experiences can contain information related to different information domains. When we consider the information domains ”soccer” and ”medicine”, for example, past experiences can describe how effective a player can be according to the weather (soccer), or perceived physical symptoms that may lead to a specific prognosis (medicine). Using information domains, however, does not change the fact that one single agent can receive conflicting information from different agents. In that case, the agent receiving the conflicting informa-

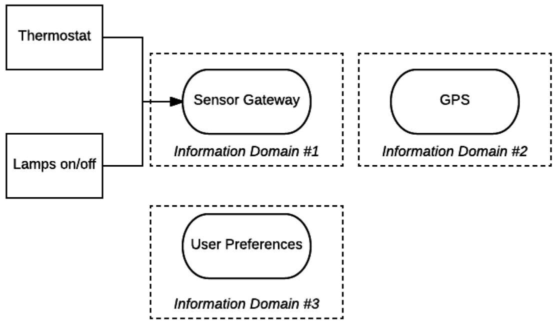

Figure 1: Example of different information domains within the contextual information.

tion must take into account factors such as how it can reach a consistent conclusion, or if all involved agents can be equally trusted, for example. This process is called belief revision, and it is explained in more detail in the next paragraphs.

2.2 Belief Revision in MAS

Belief revision is the process of changing beliefs to take into account a new piece of information. This is a non-trivial process, since there are several different ways to revise current knowledge - taking into account a different number of factors, and following different set of rules. The logical formalization of belief revision is discussed in different fields of research, including philosophy, databases, and artificial intelligence (mostly for the design of rational agents) (Gärdenfors, 2003).

According to Gärdenfors (Gärdenfors, 2003), there are three main methodological questions to be settled when trying to solve belief revisions in a logical manner: (i) the representation of beliefs; (ii) the relationship between the represented beliefs (explicit) and any other beliefs that can be derived from those (implicit); and (iii) the logic behind discarding and retaining existing beliefs in the revision process. In order to solve these questions, there are a number of integrity constraints in place that will not be detailed in the present work. It is important, however, to highlight how existing knowledge can be represented in order to establish a logical belief revision process.

When an epistemic model is used to represent an agent's beliefs in a given moment, the epistemic state is an idealization of the psychological concept and represents the cognitive state of the agent in a given moment (Gärdenfors, 1988). The belief revision process involves the study of how the agent's knowledge base is changed in the presence of new information. Since it must be performed in a computational environment, it is necessary to establish a belief representation which can be processed in conjunction with logical operators. For that purpose, the AGM model (Alchourrón et al., 1985) is used to represent epistemic states in the form of belief sets, while each belief is represented in the form of a single sentence. This model predicts three different changes that can be made to the knowledge base: (i) expansion (addition of new information), (ii) contraction (removal of existing information), and (iii) revision (addition of new information while preserving consistency of the knowledge base).

In order to formally represent the beliefs, we will use a language ζ that assumes a finite set of atomic propositions closed under truth-functional operations. Each element of ζ is a sentence denoted by lowercase Greek letters. Arbitrary tautologies and contradictions are represented by $\top$ an $\perp$, respectively. A consequence operator *Cn* is used, taking sets of sentences to sets of sentences. This operator satisfies the Tarskian properties of *inclusion*, *monotony*, and *iteration*. *Cn* is also compact and satisfies the deduction theorem and supraclassicality. These properties can be referred to as *AGM-assumptions* (Ribeiro and Wassermann, 2009).

Following this formalism, the change operations recognized by the AGM model can be described as follows:

- **Expansion:** a sentence α is added to the belief set ϕ. This operation is represented by $\phi + \alpha$;
- **Contraction:** a sentence α is removed from the belief set ϕ. This operation is represented by ϕ - α;
- **Revision:** a sentence α is added to the belief set ϕ, while other sentences are removed from the same belief set in order to preserve its consistency. This operation is represented by $\phi * \alpha$.

When a MAS is considered, the belief revision process can be used either to maintain the consistency of a single agent's epistemic state or to achieve collective goals (involving multiple agents) - which implies maintaining the consistency of a shared knowledge base. The first process is referred as *Multi-Source Belief Revision* (MSBR), and the second one is called *Multi-Agent Belief Revision* (MABR) (Tamargo, 2012). The MSBR process - which involves processing multiple sources of information from a single agent's perspective - will be used as a basis for the present work.

3 EXTENDED EPISTEMIC MODEL FOR MSBR

As previously mentioned, we will use an existing epistemic model for MSBR as a basis for the present work (Tamargo, 2012). This model was developed specifically to represent knowledge in a MAS, and its formalism was intended to be used in conjunction with a MSBR mechanism. Therefore, there are already formal support for belief-based operations (expansion, contraction, and revision). The revision process also uses a credibility order to deal with eventual inconsistencies. Multiple existing agents are considered as information sources (informants), and each different agent's knowledge base if represented through the use of belief bases with additional (meta) information. The formalism details present in the original epistemic model will be explained in the next paragraphs, along with the extensions made to the original model in order for it to accommodate the formal use of information domains.

When interacting among themselves, each agent incorporate the other agents' knowledge through the use of *information objects*. Each information object associates a sentence with an agent. All agents are identified in a finite set of agents, denoted as $\mathbb{A} = \{A_1, A_2,..., A_n\}$. Also, in order to consider different information domains in its belief revision process, it is necessary that this information is properly formalized and included in the original epistemic model. For that purpose, we define the following formalism:

- **Information Domain:** An information domain is a tuple $D = (G, M)$, where G is a grammar that describes the information structure contained in the domain and M represents all metadata associated with the domain. While neither the grammar nor the metadata need to be detailed for the purpose of this work, it is important to recognize them since they are associated with an unique identifier to each information domain.
- **Context:** A context $\mathbb{D} = \{D_1, D_2,...,D_n\}$ is defined as a finite set of different information domains that compose the contextual information, where each information domain D_i ($1 \leqq i \leqq n$) is unique.

With these considerations in mind, the original epistemic model definitions were extended to consider information domains in its structure. Using the formalism described above, we could consider three agents - an engineer (A_E), a doctor (A_D) and a soccer player (A_P) - talking about different subjects. Each subject refers to a specific information domain, including "soccer" (D_S) and "medicine" (D_M). Whenever the conversation goes towards each of these subjects, it is expected that the engineer gives more or less credibility to either of the other people. If we also take into consideration the credibility labels defined by "not credible", "plausible", and "very credible" (following a strict total order C_o), opinions from the doctor on medicine topics tend to be perceived by "very credible" by the engineer, while opinions on the same subject from the soccer player tend to be "plausible" at most. This situation represents a credibility order that can be represented by $A_S <_{C_o}^{(A_E,D_M)} A_M$ and $A_M <_{C_o}^{(A_E,D_S)} A_S$ (considering that the doctor will never be as credible as the soccer player in soccer-related subjects, and the opposite in medicine-related subjects).

- **Information Object:** An information object is a tuple $I = (\alpha, A_i, D_j)$, where α is a sentence of a propositional language ζ, $A_i \in \mathbb{A}$, and $D_j \in \mathbb{D}$. Information objects are used to represent an agent's belief base, and can be used to associate a given sentence to a specific agent. This allows for the identification of the source of each information received by an agent. The extension proposed allows for information objects to associate a given sentence to a specific agent and a specific domain within a context $\mathbb{D}$.
- **Belief Base:** A belief base of a given agent A_i ($1 \leqq i \leqq n$) is a set $K_{A_i} = \{I_1, I_2,...,I_q\}$ that contains information objects (α, A_p, D_j) ($1 \leqq p \leqq q$) received from other agents ($p \neq i$) and proper beliefs ($p = i$) regarding different domains in the context $\mathbb{D}$. Thus, the set $\kappa = 2^{\,\zeta \times A \times D}$ represents all the belief bases for all information domains within the context. As an example, we can consider the finite set of agents given by $\mathbb{A} = A_1, A_2, A_3, A_4$ and the belief base for agent A_1: $K_{A1} = (\beta, A_1, D_2), (\alpha, A_2, D_3), (\alpha, A_3, D_1)$.
- **Sentence Function:** The sentence function *Sen* ($Sen : K \rightarrow (2^{\zeta}, \mathbb{D})$) over a belief base $K \in \kappa$ is defined as $Sen(K) = \{\alpha : (\alpha, A_i, D_j) \in K\}$. For a given agent A_i, its belief base is consistent for a given domain if $C_n(Sen(K_A))$ is consistent for the same domain. Considering the example in the belief base definition, $Sen(K_{A1}) = (\alpha, D_1), (\alpha, D_3), (\beta, D_2)$.
- **Agent Identifier Function:** An agent identifier function *Ag* ($Ag : \kappa \rightarrow (2^{\mathbb{A}}, \mathbb{D})$) establishes a relationship between a belief set and a finite set of agents for different information domains within a context $\mathbb{D}$, allowing for the identification of agents that are referenced within a given belief set $K \in \kappa$. This function is defined by: $Ag(K) = \{A_i : (\alpha, A_i, D_j) \in K\}$. Considering the example in the belief base definition, $Ag(K_{A1}) = (A_1, D_2), (A_2, D_3), (A_3, D_1)$.
- **Assessment:** In order to represent the credibility that one agent assigns to other agents for every domain in $\mathbb{D}$, an assessment function is used. Credibility, as a value assigned to an specific agent,

can be represented by a finite set of credibility values (or labels) $\mathbb{C} = c_1,...,c_k$ common to all agents. The credibility values follow a strict total order. Therefore, for a finite set of agents $\mathbb{A}$ and a credibility set $\mathbb{C}$, an assessment c_{A_i,D_j} for an agent A_i regarding a domain D_j is a function $c_{Ai} : (\mathbb{A},\mathbb{D}) \rightarrow \mathbb{C}$ that assigns a credibility value from $\mathbb{C}$ to each agent $A_j \in \mathbb{A}$, regarding each domain D_j in $\mathbb{D}$. Since the credibility set is common to all agents within $\mathbb{A}$, each agent possess comparable credibility values assigned by other agents. On the other hand, since the credibility values are conditioned to specific information domains, credibility values regarding different information domains for the same given agent cannot be compared. Similarly as in the original model, different credibility values can be assigned to the same agent by different other agents for the same information domain. For example: considering $\mathbb{A} = A_1, A_2, A_3$, $\mathbb{D} = D_1, D_2$, and $\mathbb{C} = c_1, c_2, c_3$, the credibility values assigned to A_3 regarding the domain D_1 can be different: $c_{A_1}(A_3, D_1) = c_2$ and $c_{A_2}(A_3, D_1) = c_3$. At the same time, the credibility values for the same agent regarding different information domains can also be different: $c_{A_1}(A_2, D_1) = c_2$ and $c_{A_1}(A_2, D_2) = c_3$.

- **Credibility Order among Agents:** Since different credibility values (following a strict total order) can be assigned by a single agent A_i for the same domain D_j, a credibility order for the other agents can be established. In the previous example, considering A_1 and the credibility values over A_2 and A_3 regarding the information domain D_1 respectively given by $c_{A_1}(A_2, D_1)$ and $c_{A_1}(A_3, D_1)$. If $c_{A_1}(A_2, D_1) \prec c_{A_1}(A_3, D_1)$ or $c_{A_1}(A_2, D_1) = c_{A_1}(A_3, D_1)$, it means that - according to A_1 - A_3 is at least as credible as A_2 regarding the information domain D_1. This relationship is represented by $A_j \leq_{C_o}^{(A_i, D_p)} A_k$, meaning that - according to A_i - A_k is at least as credible as A_j regarding an information domain D_p. Similarly, the strict relationship $A_j <_{C_o}^{(A_i, D_p)} A_k$ can be defined, meaning that A_k is strictly more credible than A_j regarding the information domain D_p. Likewise, $A_j =_{C_o}^{(A_i, D_p)} A_k$ means that A_k is as credible as A_j for the information domain D_p. Therefore, the relationship $\leq_{C_o}^{(A_i, D_p)}$ is a total order over $\mathbb{A}$. It is important to notice, however, that credibility order among agents is confined to specific information domains, and that the credibility order between two agents can change for different information domains.

4 DISCUSSION AND RELATED WORK

The present work aims at serving as an initial step towards solving the problem of the use of different trust degrees associated with different information domains in belief revision processes. While the original epistemic model proposes a formalism that can be used in multi-source belief revision processes, it does not take into account the existence of different information domains. Therefore, an agent can only receive a global credibility (or trustworthiness) degree from any other agent.

Also, from a learning perspective, a global credibility degree can have a negative impact in the overall belief revision process. If a given agent possess extremely accurate experiences regarding a specific activity but no consistent or useful information about anything else, it can be perceived as ”unreliable” by other agents in the system. While this perception would help in the case of the unreliable information transmitted to the other agents, it would diminish or even void the benefits that processing the accurate information could bring to the revised beliefs.

There are different publications on formalisms related to both multi-agent and multi-source belief revisions ((Kfir-Dahav and Tennenholtz, 1996; Liu and Williams, 2001; Cantwell, 1998; Dragoni and Puliti, 1994)). Our decision to extend an existing belief revision formalism was based on the fact that the epistemic model used as a reference also took into consideration aspects of trust and reputation of agents in a distributed environment. Aspects such as plausibility, reputation maintenance and information retransmission in a multiagent system were addressed. Since our long-term goal is to study these same aspects under a multi-expertise perspective, it was appropriate to adapt an existing formalism and revisit the already studied aspects in conjunction with the new research.

As previously mentioned, our studies are also related to existing work on agent argumentation schemes associated with expertise ((Melo et al., 2016)), as well as trust and reliability aspects on MAS ((Wang and Singh, 2007; Tamargo, 2012)). In this work, we focus on the formalism of relative association of trust (credibility) regarding different information domains. Our intent is to explore the concept of both relative and localized trust, and how it can impact the experience and information sharing process in a multiagent system.

It is also important to notice that as a preliminary work, there are more complex questions and challenges that are not addressed yet. Different intelligent agents in different MAS can possess different primitives, for example, which would require mechanisms in place for experience or meaning interpretation - which is not part of the scope of the present study. Alternative methods for belief revision, argumentation mechanisms, or credibility evaluation processes are also not part of the scope of the present work.

5 CONCLUSIONS AND FUTURE WORK

The contribution of this paper resides on the extension of an existing epistemic model in order to allow its use by context-aware BDI agents in their belief revision process. Using an epistemic model in conjunction with the concept of information domains provides the formalization necessary for a multi-source belief revision process based on contextual information. The use of such model allows for a single agent to possess different trust degrees associated with other agents regarding different information domains.

While more complex problems are not addressed in this work, we intend to use the extended epistemic model presented here as a basis for future research. This will include further development of the extended epistemic model and the implementation of a MSBR mechanism to be used by context-aware agents, along with trust calculation and conflict solving mechanisms that can benefit from this model.

SCIENCE AND TECHNOLOGY PUBLICATIONS

ACKNOWLEDGEMENTS

Eduardo Fermé is partially supported by FCT MCTES and NOVA LINCS UID/CEC/04516/2013, FCT SFRH/BSAB/127790/2016 and FAPESP 2016/13354-3. Arthur Casals is supported by CNPq, grant no. 142126/2017-9.

REFERENCES

Abowd, G. D., Dey, A. K., Brown, P. J., Davies, N., Smith, M., and Steggles, P. (1999). Towards a better understanding of context and context-awareness. In *International Symposium on Handheld and Ubiquitous Computing*, pages 304–307. Springer Berlin Heidelberg.

Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *The journal of symbolic logic*, 50(2):510–530.

Cantwell, J. (1998). Resolving conflicting information. *Journal of Logic, Language and Information*, 7(2):191–220.

Cohen, P. R. and Levesque, H. J. (1990). Intention is Choice with Commitment. *Artificial Intelligence*, 42(2–3):213–261.

Dragoni, A. F. and Puliti, P. (1994). Distributed belief revision versus distributed truth maintenance. In *Tools with Artificial Intelligence, 1994. Proceedings., Sixth International Conference on*, pages 499–505. IEEE.

Gärdenfors, P. (1988). *Knowledge in flux: Modeling the dynamics of epistemic states*. The MIT press.

Gärdenfors, P. (2003). *Belief revision*, volume 29. Cambridge University Press.

Hennessy, P. (1991). Information domains in CSCW. *Studies in Computer Supported Cooperative Work: Theory, Practice and Design, Eds. JM Bowers and SD Benford, Elsevier*.

Hjørland, B. (2002). Domain analysis in information science: eleven approaches–traditional as well as innovative. *Journal of documentation*, 58(4):422–462.

Hong, J.-y., Suh, E.-h., and Kim, S.-J. (2009). Context-aware systems: A literature review and classification. *Expert Systems with Applications*, 36(4):8509–8522.

Huynh, T. D., Jennings, N. R., and Shadbolt, N. R. (2006). An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 13(2):119–154.

Kfir-Dahav, N. E. and Tennenholtz, M. (1996). Multi-agent belief revision. In *Proceedings of the 6th conference on Theoretical aspects of rationality and knowledge*, pages 175–194. Morgan Kaufmann Publishers Inc.

Kim, J. and Chung, K.-Y. (2014). Ontology-based healthcare context information model to implement ubiquitous environment. *Multimedia Tools and Applications*, 71(2):873–888.

Kwon, O. B. and Sadeh, N. (2004). Applying case-based reasoning and multi-agent intelligent system to context-aware comparative shopping. *Decision Support Systems*, 37(2):199–213.

Liu, W. and Williams, M.-A. (2001). A framework for multi-agent belief revision. *Studia Logica*, 67(2):291–312.

Melo, V. S., Panisson, A. R., and Bordini, R. H. (2016). Trust on beliefs: Source, time and expertise. In *TRUST@ AAMAS*, pages 31–42.

Mena, E., Kashyap, V., Illarramendi, A., and Sheth, A. (1998). Domain specific ontologies for semantic information brokering on the global information infrastructure. In *Formal Ontology in Information Systems*, volume 46, pages 269–283. Amsterdam: IOS Press, MCB UP Ltd.

Nalepa, G. J. and Bobek, S. (2014). Rule-based solution for context-aware reasoning on mobile devices. *Computer Science and Information Systems*, 11(1):171–193.

Rao, A. S. and Georgeff, M. P. (1991). Modeling rational agents within a BDI-architecture. In Allen, J., Fikes, R., and Sandewall, E., editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 473–484. Morgan Kaufmann publishers Inc.: San Mateo, CA, USA.

Ribeiro, M. M. and Wassermann, R. (2009). Agm revision in description logics. *Proceedings of ARCOE*.

Tamargo, L. H. (2012). Knowledge dynamics in multi-agent systems: Plausibility, belief revision and forwarding information. *AI Communications*, 25(4):391–393.

Wang, Y. and Singh, M. P. (2007). Formal trust model for multiagent systems. In *IJCAI*, volume 7, pages 1551–1556.

Wooldridge, M. (2009). *An introduction to multiagent systems*. John Wiley & Sons.