The Knowledge Engineering Review, Vol. 26:1, 35–43. © Cambridge University Press, 2011
doi:10.1017/S0269888910000391

On the evolving relation between Belief Revision and Argumentation

MARCELO A. FALAPPA[1], ALEJANDRO J. GARCÍA[1], GABRIELE KERN-ISBERNER[2] and GUILLERMO R. SIMARI[1]

[1]*Artificial Intelligence Research and Development Laboratory, Department of Computer Science and Engineering, Universidad Nacional del Sur, Consejo Nacional de Investigaciones Científicas y Técnicas, Av. L.N. Alem 1253, B8000CPB, Bahía Blanca, Argentina;*
e-mail: mfalappa@cs.uns.edu.ar; ajg@cs.uns.edu.ar; grs@cs.uns.edu.ar;
[2]*Department of Computer Science, University of Dortmund, Otto-Hahn-Str. 16, 44227 Dortmund, Germany;*
e-mail: gabriele.kern-isberner@cs.uni-dortmund.de

Abstract

Research on the relation between Belief Revision and Argumentation has always been fruitful in both directions: some argumentation formalisms can be used to define belief change operators, and belief change techniques have also been used for modeling the dynamics of beliefs in argumentation formalisms. In this paper, we give a historical perspective on how belief revision has evolved in the last three decades, and how it has been combined with argumentation. First, we will recall the foundational works concerning the links between both areas. On the basis of such insights, we will present a conceptual view on this topic and some further developments. We offer a glimpse into the future of research in this area based on the understanding of argumentation and belief revision as complementary, mutually useful disciplines.

1 Introduction

The past quarter century has seen the development of the interaction between two increasingly significant areas of research in Knowledge Representation and Reasoning: *Belief Revision* and *Argumentation*. Both areas are concerned with establishing the status of the beliefs contained in an agent's knowledge base.

In belief revision, investigation is centered on the process of changing beliefs in order to adapt the agent's epistemic state when considering the addition of a new piece of information to its knowledge base. Argumentation research is concerned with deciding the epistemic state of the beliefs in the agent's knowledge base by considering the support that each belief has in the form of arguments that can be constructed from the knowledge base. Each belief takes the role of the claim of an argument with a set of premises, and the decision of accepting the belief is made after considering all the arguments against and in favor of the argument supporting the claim.

Each area contributes tools that model forms of human commonsense reasoning as it is applied in everyday situations. As an example, consider an agent's environment. Belief revision techniques could provide ways in which the agent is supposed to change its beliefs when new information arrives, or when changes in the world are observed. Complementarily, argumentation could provide support for monological (i.e. internal) reasoning for deciding how the revision of the knowledge base should proceed, or could give support for dialogical (i.e. external) reasoning, backing the strategies used in the agent's dialogs with other agents.

This article is organized as follows. First, in Section 2, we will introduce a brief overview of the evolution of Belief Revision. Then, in Section 3, we will give a perspective on how the ideas of

Argumentation are contributing to new developments in Belief Revision. After that, in Section 4, we will advance some possible new connections between Belief Revision and Argumentation. We will finish this short article in Section 5, presenting our main conclusions.

2 A brief summary of the evolution of Belief Revision

The origins of belief change theory go back to the work of Isaac Levi (Levi, 1977), who discussed the problems concerning this field of research and William Harper's proposal of a rational way to interrelate some belief change operators (Harper, 1975). However, the main advances on belief change theory came during the 1980's when Carlos Alchourrón and David Makinson studied changes in legal codes (Alchourrón & Makinson, 1981), and Peter Gärdenfors's introduced rational postulates for change operators (Gärdenfors, 1982). After that, the three authors wrote a foundational paper containing what became known as the *AGM* model (Alchourrón *et al.*, 1985), named after the three originators: Alchourrón, Gärdenfors, and Makinson. The extraordinary contribution of the AGM model was the presentation of a new and more general formal framework for the studies of belief change. Today, this work is considered as the cornerstone from which belief change theory evolved.

In the 25 years since its publication, different frameworks for belief dynamics and their respective epistemic models have been proposed. The epistemic model corresponds to the formalism in which beliefs are represented, providing the framework in which different kinds of operators can be defined. The AGM model is conceived as an idealistic theory of rational change in which epistemic states are represented by belief sets, which are sets of sentences closed under logical consequence, and the epistemic input is represented by a sentence. The AGM approach studies the changes at the *knowledge level*, whereas some other approaches are concerned with the changes at the *symbolic level*. The distinction about knowledge and symbolic level was proposed by Newell (1982). According to Newell, the knowledge level lies above the symbolic level where all of the knowledge in the system is represented. An implicit assumption in Artificial Intelligence is that a knowledge-level system must contain a symbol system. Following the Newell's unified theory of cognition, the symbol level corresponds to the cognitive part and the knowledge level to the rational part. Moreover, some belief bases with different symbolic representations may represent the same knowledge representations.

EXAMPLE 1. Suppose that we are using a propositional language $\mathcal{L}$ with a classical consequence operator *Cn*. *Cn* takes sets of sentences in $\mathcal{L}$ and produces new sets of sentences. The operator *Cn* satisfies *inclusion* ($A \subseteq Cn(A)$), *idempotence* ($Cn(A) = Cn(Cn(A))$), and *monotony* (if $A \subseteq B$ then $Cn(A) \subseteq Cn(B)$). Then, suppose that p and q are logically independent propositions, and $K_1 = \{p, q\}$ and $K_2 = \{p, p \rightarrow q\}$ are two (different) belief bases. Since their closure is the same, $Cn(K_1) = Cn(K_2) = Cn(\{p, q\})$, they represent the same beliefs at the knowledge level. Although they are statically equivalent (they represent the same beliefs), they could be dynamically different: changes in K_1 are probably different to changes in K_2 because p and q are totally independent in K_1 and they are interrelated in K_2 by the sentence $p \rightarrow q$.

In the AGM model, three basic change operators are defined: *expansion*, *contraction*, and *revision*. Expansion is the simplest and best known form of theory change, in which a new belief, hopefully consistent with the given epistemic state, is added. Two other operations are considered: contraction removes a particular belief from the state, and revision adds a new belief in such a way that a new consistent epistemic state that contains the belief is obtained. More precisely, suppose that **K** is a belief set (a set of sentences closed under logical consequence, usually noted in boldface). Then, three basic changes can be defined for **K**:

- **Expansions**: the result of expanding **K** by a sentence α is a larger set which infers α;
- **Contractions**: the result of contracting **K** by α is a smaller set which usually does not infer α;

- **Revisions**: the result of revising $\mathbf{K}$ by α is a set that neither extends nor is part of the set K. In general, if $\mathbf{K}$ infers $\neg\alpha$ then α is consistently inferred from the revision of $\mathbf{K}$ by α.

The great importance of AGM comes from providing axiomatic characterizations of contraction and revision in terms of rationality postulates. Such rationality postulates regard the operators as black boxes, characterizing *what* they do, but not explaining *how* they do it. In other words, their behavior is constrained with regard to inputs in basic cases, without describing the internal mechanisms used for achieving that behavior, so it is crucial to say that both types of operators can also be obtained via more constructive approaches. Suppose that $\mathbf{K}$ represents the belief set of some agent and α represents a new piece of information, and that a contraction of $\mathbf{K}$ by α, noted as $\mathbf{K} \div \alpha$, is to be performed. According to Alchourrón *et al.* (1985) and Gärdenfors (1988), the following are the basic postulates for contraction:

- **Closure**: $\mathbf{K} \div \alpha = Cn(\mathbf{K} \div \alpha)$.
- **Inclusion**: $\mathbf{K} \div \alpha \subseteq \mathbf{K}$.
- **Success**: If $\nvdash \alpha$ then $\alpha \notin \mathbf{K} \div \alpha$.
- **Vacuity**: If $\alpha \notin \mathbf{K}$ then $\mathbf{K} \div \alpha = \mathbf{K}$.
- **Recovery**: $\mathbf{K} \subseteq (\mathbf{K} \div \alpha) + \alpha$.
- **Extensionality**: If $\vdash \alpha \leftrightarrow \beta$ then $\mathbf{K} \div \alpha = \mathbf{K} \div \beta$.

AGM contractions can be realized by *partial meet contractions*, which are based on a selection among (maximal) subsets of $\mathbf{K}$ that do not imply α. Via the Levi identity (Gärdenfors, 1988), associated revision operations called *partial meet revisions* are obtained.

Another possible approach for contraction is based on a selection among the (minimal) subsets of $\mathbf{K}$ that contribute to make $\mathbf{K}$ imply α, as in *safe contraction* (Alchourrón & Makinson, 1985). A more general variant of the same approach, known as *kernel contraction*, was introduced later (Hansson, 1994). It has been shown that both safe contractions and kernel contractions are equivalent to partial meet contractions, and hence to the AGM approach to contraction (Hansson, 1994, 1999).

Gärdenfors made an important remark that started a new line of research: 'Certain pieces of our knowledge and beliefs about the world are more important than others when planning future actions, conducting scientific investigations, or reasoning in general. We will say that some sentences in a belief system have a higher degree of epistemic entrenchment than others…' (Gärdenfors, 1992). *Epistemic entrenchment* allows a contraction (or revision) to be based on a relation among sentences such that, when forced to choose between two beliefs, an agent will give up the less entrenched one. Moreover, it has been shown that a contraction based on an epistemic entrenchment relation is also equivalent to a partial meet contraction (Gärdenfors & Makinson, 1988), providing a more semantical view of AGM revision; (Rott, 1992) proposed a more intuitive entrenchment relation that does not satisfy *recovery* and, therefore, is not a partial meet contraction.

Likewise, Grove's Sphere Systems (Grove, 1988), which are very similar to the sphere semantics for counterfactuals proposed by Lewis (1973), may characterize AGM revisions in a semantic way. Concentric spheres allow the ordering of possible worlds according to some notion of similarity. Then, the result of revising by α can be semantically described by the set of α-models that are closest to the $\mathbf{K}$-models.

Other interesting types of change were proposed, such as *updating* (Katsuno & Mendelzon, 1992). Although in revision both old beliefs and the new information refer to the same situation, in updating the new information is about a possible change in the situation. The main goal of belief revision is to change beliefs preserving as much as possible the old information; the target of updating is just to change the worlds in which the change is possible.

All of the approaches summarized above are prioritized in the sense that when a change is performed, the new information has primacy over the beliefs contained in the original set when restoring consistency. A complete presentation of prioritized classical models of belief revision can be found in Hansson (1999). A general setting for non-prioritized revision can be found in Hansson (1997) and Hansson *et al.* (2001).

3 The contribution of Argumentation to Belief Revision

Although neither field was fully demarcated at the time, retrospectively Jon Doyle's *Truth Maintenance System*, or TMS, can be regarded as the first work relating Argumentation and Belief Revision (Doyle, 1979). A TMS offers a method for representing both beliefs and their dependencies, having the ability to restore consistency. It associates a special data structure, called a node, with each problem-solver datum that includes database entries, inference rules, and procedures. The TMS records arguments for potential inferred beliefs, so as to characterize the current set of beliefs. It manipulates *nodes* representing beliefs and *justifications* that represent reasons that support the beliefs. Later, the capabilities of a TMS are extended to deal with assumption sets (de Kleer, 1986). In an *Assumption-based TMS*, or ATMS, each datum is labeled with the sets of assumptions (representing the contexts) under which that datum holds. Doyle's seminal work on reason maintenance was expanded by explaining some similarities and differences between belief revision and argumentation (Doyle, 1992).

In an approach oriented to the handling of inconsistency, (Benferhat *et al.*, 1995) presented research that dealt with the contradictions that the use of multiple sources of information would introduce. Each formula in a knowledge base and is associated with a certainty level corresponding to the layer to which it belongs; in this manner knowledge bases become stratified. The authors suggested that by using an argumentation-based inference mechanism conclusions can be derived and supported by the reasons that allow for their belief, independently of the consistency status of the knowledge base. Therefore, it is not necessary to restore consistency in order to make sensible inferences from an inconsistent knowledge base.

Following a 'derivational approach', (Pollock & Gillies, 2000) studied the dynamic of a belief revision system considering relations among beliefs and trying to obtain a theory of belief revision from a more concrete epistemological theory. According to them, one of the goals of belief revision is to generate a knowledge base in which each piece of information is justified (by perception) or warranted by arguments containing previously held beliefs. A difficulty with this approach is that the set of justified beliefs can exhibit all kinds of logical incoherences because it represents an intermediate stage in reasoning. Therefore, they propose a theory of belief revision concerned with warrant rather than justification.

In a framework in which epistemic states are represented as a set of undefeasible beliefs (called strict knowledge) and defeasible beliefs that support tentative (and potentially inconsistent) conclusions, (Falappa *et al.*, 2002) proposed a non-prioritized revision operator based on the use of explanations (arguments). The idea is that before accepting information inconsistent with its knowledge, an agent requests an explanation supporting it. They presented a framework oriented to defeasible reasoning that considers the explanation given to decide how to revise the agent's epistemic state. An interesting result of this work is the generation of defeasible conditionals from the revision process. This approach preserves consistency in the strict knowledge and it provides a mechanism to dynamically qualify the beliefs as strict or defeasible.

Putting argumentation and belief revision together in the same conceptual framework, (Paglieri & Castelfranchi, 2006), stressed the important role played by Toulmin's layout of argument in fostering such integration. They consider argumentation as 'persuasion to believe' and this perspective makes the connection with belief revision more explicit. They propose a model of belief dynamics alternative to the AGM approach: *data-oriented belief revision* (DBR). Two basic informational categories (data and beliefs) are used in their model to account for the distinction between pieces of information that are simply gathered and stored by the agent (*data*) and pieces of information that the agent considers (possibly up to a certain degree) truthful representations of states of the world (*beliefs*). Whenever a new piece of evidence is acquired through perception or communication, it directly affects the agent's data structure and only indirectly his beliefs. Belief revision is often triggered by information update either on a fact or on a source: the agent receives a new piece of information, rearranges his data structure accordingly, and possibly changes his beliefs.

In a position paper, Boella *et al.* (2008a) offered the view that a direct relation between argumentation and belief revision exists. They also consider argumentation as persuasion to believe, so that persuasion should be related to belief revision. In Boella *et al.* (2008b) the same authors elaborated on the previous presentation deepening the study of the interrelation between argumentation and belief revision in multi-agent systems, emphasizing that when an agent uses an argument to persuade another, she must consider not only the proposition supported by the argument, but also the overall impact of the argument on the beliefs of the addressee.

Another line of research where the dynamic nature of reasoning is captured through the application of belief revision concepts was introduced in (Rotstein *et al.*, 2008). The argumentation framework handles abstract arguments, but their structure is considered to a certain point allowing to apply the change operators to that structure. They define a *Dynamic Abstract Argumentation* theory including dialectical constraints, and then present argument revision techniques to describe the fluctuation of the set of active arguments (the ones that are considered by the inference process of the theory). In a separate work that shows the capabilities of the theoretical framework, Moguillansky *et al.* (2008) proposed an instantiation of their operators to *Defeasible Logic Programming*, *DeLP* (García and Simari, 2004).

The problem of the change in an abstract argumentation system by the addition of an argument was presented in a recent work (Cayrol *et al.*, 2010). The authors defined several change operations considering four possibilities: modifying the attack relation by adding or removing an attack, adding an argument together with a set of interactions, and removing an argument together with all interactions related to it, where an interaction is an attack received or effected by the argument involved. The work describes the impact on the set of extensions that such changes effect on the outcome of the argumentation, focusing on two particular semantics: the grounded and preferred semantics (Dung, 1995).

4 Nascent connections between Belief Revision and Argumentation

Investigating the multifaceted relationship between Belief Revision and Argumentation requires considering cross-links between different aspects on either side while also considering their place in the higher context of reasoning. As the field advances, new elements and interconnections become necessary in order to extend this already rich and complex formalism.

Here, we present connections between Belief Revision and Argumentation based mainly on ideas contained in Falappa *et al.* (2009). To obtain a clear picture of the possible connections, we begin (as sketched in the introduction) by making more precise the basic reasoning steps, tracing the process from the reception of (new) information to the problem of coming up with adequate plausible beliefs on which a decision can be based. We will assume that the current epistemic state is given and represented within some appropriate framework that takes into account the following aspects of this setting:

- *Reception of new information*: The new information I may come in very different shapes and forms. In the simplest scenario, I is a propositional fact. Assuming the epistemic state of the agent is given as a belief set, this lies inside the scope of the basic AGM theory. However, I might be much more complex; for instance, it can be equipped with a degree of plausibility, come from a particular informant agent, have the form of a rule or even a complete argument, or consist of a set of such entities.
- *Evaluation of new information:* For the further processing of I, it is crucial for the agent to know its origin, as this knowledge will decisively influence her willingness to adopt I. For instance, if I is based on an observation made by the agent herself, usually she will be convinced of its truth. However, if I is conveyed to her by another agent, be it given as part of official news, in personal communication, or found as written material, the agent will require some justification for I. In any case, as a mandatory step for rational, critical thinking, she will evaluate in an argumentative manner both I and the justification possibly given on the basis of her own beliefs and decide whether or not I is to be incorporated into her stock of beliefs.

- *Changes of beliefs:* If the agent has decided to adopt I in the previous step, she must employ some strategy to incorporate I consistently into her beliefs. For this, she must use belief revision techniques that will allow her to change her epistemic state accordingly.
- *Inferences:* From her new epistemic state, the agent derives the (most) plausible beliefs that guide her behavior.

This scenario can also be partially applied, if I is not new information, but a query to which the agent is expected to reply. In most such cases, the change step would be unnecessary, and the evaluation and inference steps would collapse. However, if I is a conditional query such as 'Suppose α holds, would you believe β?', whatever α and β are, reasoning would include a hypothetical change process that the current epistemic state of the agent must undergo.

EXAMPLE 2. Suppose that the epistemic states and the epistemic input belong to a language that includes a subset of a first-order logic. Further assume that the *reception of new information* allows I to be not only a single sentence as in AGM model, but also one that is composed by an explanation (a set of beliefs with some restrictions) that supports the new information. Suppose that Michael believes that *all birds fly* ($\alpha \equiv (\forall x)(b(x) \rightarrow f(x))$) and that *Tweety is a bird* ($\beta \equiv b(t)$). Thus, he will believe that *Tweety flies* ($\delta \equiv f(t)$). Then, Johana tells him that *Tweety does not fly* ($\neg f(t)$). As a consequence, Michael will have to drop the belief in α or the belief in β forced by having to drop δ. However, it does not seem like a sensible attitude to incorporate any external belief without pondering it. Usually, an intelligent agent demands an explanation supporting the information provided, even more so if that information contradicts its own set of beliefs. In that case, Michael will demand an explanation for the belief $\neg\delta$. For instance, Johana accompanies her contention of *Tweety does not fly* ($\neg f(t)$) with the sentences *Tweety is an ostrich* ($o(t)$) and *ostrichs are birds but they do not fly* ($(\forall x)(o(x) \wedge b(x) \rightarrow \neg f(x))$). Then, after the *evaluation of the new information*, Michael would have to check his beliefs in order to determine whether he believes that *Tweety flies* or *Tweety does not fly*. Therefore, the set of *beliefs can be changed* and some *inferences* can be dropped from the new epistemic state.

The above example uses the concept of explanation. The role of explanations in knowledge representation has been widely studied, for instance, in Falappa *et al.* (2002), Boutilier and Becher (1993), Gärdenfors (1988), and Williams *et al.* (1995).

EXAMPLE 3. Consider a scenario where an agent wants to travel to a ski resort and she knows that usually roads are closed when it is snowing. In order to decide whether to go or not, the agent can get more information from different sources: Internet, the Tourist Office, and some other agents. It is clear that one source can be more credible than other, and that the reliability of the incoming information can also be evaluated in terms of the explanation that is given. Assume that the agent obtains the following information: (a) the Internet weather forecast shows that the chances of snow are high due to a storm, (b) the Tourist Office advises that the route is always open because the ski resort has snow clearing machines, and (c) a ski resort employee said on the news that there is much snow on the road, and that the number of machines is insufficient to clear the way today. Then, the agent should be able to *evaluate the new information* from (a), (b), and (c) in order to properly revise its belief.

From this embedding into a complex reasoning process, the complementary characters of argumentation and belief revision become clear: while argumentation can make substantial contributions to the evaluation step, belief revision theory should be employed in the change of beliefs. But this is not the end of the story; evaluation might include hypothetical change processes, considering what would happen if the new information were to be believed, and belief change implicitly relies on logical links between pieces of information which can be represented by

arguments. In the end, plausible beliefs can be obtained both from argumentation processes and from belief revision processes, but both areas focus only on parts of the dynamic reasoning process while at the same time providing general and versatile frameworks. Revision operators can not only be applied to beliefs, but also to intentions, preferences, theories, ontologies, law codes, etc.

Argumentation can be used for *negotiation* (when the agents have conflicting interests and they try to make the best out of a deal), *inquiry* (when the agents have general ignorance in some subject and they try to find a proof or destroy one), *deliberation* (when the agents collaborate to decide what course of action to take), or *information seeking* (when the agents have personal ignorance on some subject and every agent seeks the answer to some questions from other agents). Although Paglieri and Castelfranchi (2006) and Boella *et al*. (2008b) proposed related ideas considering argumentation as persuasion, the view depicted above is more general.

EXAMPLE 4. Considering the Example 2, suppose that we are defining a change operator where: (a) the epistemic state is represented by a belief base, (b) the epistemic input is a statement together with an explanation, (c) the underlying language contains a subset of a first-order logic with a complete set of connectives, and (d) there exists a classical consequence operator Cn. Following the proposal of Falappa *et al*. (2002), we say that a set A is an *explanation* for a sentence α if the following conditions are satisfied: (i) deduction: $\alpha \in Cn(A)$, (ii) consistency: $\perp \notin Cn(A)$, (iii) minimality: there is no $A' \subseteq A$ such that $\alpha \in Cn(A')$, and (iv) informational content: $Cn(A) \not\subseteq Cn(\alpha)$. Then, argumentation can be used in the revision process in order to decide if the external explanation prevails or if some of the current beliefs are better and, therefore, the new information with its respective explanation is fully accepted, fully rejected, or partially accepted.

These considerations result in a more complex, highly interrelated view of argumentation and belief revision. Consequently, discussing links between both fields must reflect this complexity, taking into account that the proposed frameworks on either side might only implement some but not all aspects of the corresponding field. In particular, the seminal AGM theory (Alchourrón *et al*., 1985) in belief revision is more concerned with judging the results of change in a very abstract way than with the change process itself; moreover, argumentative evaluation of justifications to believe the given information in the sense described above is not at all a topic of AGM. That is why *success* is one of the basic postulates of the AGM theory, according to which the agent is forced to believe the new information. Here, it is implicitly assumed that evaluation has been done beforehand. Critics may focus on the point that it is not rational to accept an inconsistent belief, something inferred from success and consistency (in AGM theory, *consistency* establishes that the revision of $\mathbf{K}$ by α is consistent only when α is consistent) when the epistemic input is inconsistent. Hence, a comparison between AGM and some approach that makes use of a sophisticated argumentative evaluation of information to select beliefs, as has been developed, for example, in Paglieri and Castelfranchi (2006), is very likely to cast negative light on AGM. It is clear that AGM is no more about argumentative reasons for belief than Dung's framework (Dung, 1995) is about change processes. Both are highly abstract, declarative approaches in the respective field, and hence both are (necessarily) over-simplifying in some respects but nevertheless extremely valuable as reference points.

Therefore, an investigation of the connections between belief revision and argumentation theory that is to do justice to both areas must go beneath the surface of abstract frameworks. Such an investigation must study methods, and also the rationale underlying these methods, as well as purposes and intentions guiding the application of related techniques.

5 Conclusions

In the past 20 years, research in the area of Knowledge Representation and Reasoning (KR&R) has blossomed. Within KR&R, Belief Revision and Argumentation have inspired two fertile areas of research with two most separate communities. Starting from the more relevant works done in

Belief Revision, we have explored the research that has begun to build bridges between Belief Revision and Argumentation. We developed a conceptual view of Belief Revision and Argumentation based on the understanding that both areas offer complementary contributions to the process of reasoning. Although Argumentation seems to be more appropriate for the evaluation of a belief's support, Belief Revision has proven useful for handling the dynamics of beliefs. Hence, they should not be regarded as competing alternatives that may replace each other. Combining Belief Revision and Argumentation allows for a richer modeling of the reasoning process, achieving greater variety, and complexity of applications than each area could alone. Looking ahead, we have pointed out various future lines of research that look promising.

Acknowledgements

This work is partially supported by Universidad Nacional del Sur (UNS-Argentina), Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET-Argentina), Ministerio de Ciencia, Tecnología e Innovación Productiva (MINCyT-Argentina), and Deutscher Akademischer Austausch Dienst (DAAD-Germany).

References

Alchourrón, C. & Makinson, D. 1981. Hierarchies of regulation and their logic. In Hilpinen, R. (ed.), *New Studies in Deontic Logic*, 125–148.

Alchourrón, C. & Makinson, D. 1985. On the logic of theory change: safe contraction. *Studia Logica*, D. Reidel Publishing Company **44**, 405–422.

Alchourrón, C., Gärdenfors, P. & Makinson, D. 1985. On the logic of theory change: partial meet contraction and revision functions. *The Journal of Symbolic Logic* **50**, 510–530.

Benferhat, S., Dubois, D. & Prade, H. 1995. How to infer from inconsistent beliefs without revising. In *Proceedings of IJCAI'95*. Montreal, Quebec, Canada, 1449–1455.

Boella, G., da Costa Perera, C., Tettamanzi, A. & van der Torre, L. 2008a. Dung argumentation and AGM belief revision (*position paper*). In *Fifth International Workshop on Argumentation in Multi-Agent Systems (ArgMAS)*, Estoril, Portugal.

Boella, G., da Costa Perera, C., Tettamanzi, A. & van der Torre, L. 2008b. Making others believe what they want. In *IFIP 20th World Computer Congress, TC 12 IFIP AI 2008*, Milan, Italy, 215–224.

Boutilier, C. & Becher, V. 1993. Abduction as belief revision: a model of preferred explanations. In *AAAI*, Washington, DC, USA, 642–648.

Cayrol, C., de Saint-Cyr, F. D. & Lagasquie-Schiex, M.-C. 2010. Change in abstract argumentation frameworks: adding an argument. *Journal of Artificial Intelligence Research* **84**, 49–84.

de Kleer, J. 1986. An assumption-based TMS. *Artificial Intelligence* **28**(2), 127–162.

Doyle, J. 1979. A truth maintenance system. *Artificial Intelligence* **12**, 231–272.

Doyle, J. 1992. Reason maintenance and belief revision: foundations versus coherence theories. In *Belief Revision*, Gärdenfors, P. (ed.). Cambridge University Press, 29–51.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* **77**, 321–357.

Falappa, M. A., Kern-Isberner, G. & Simari, G. R. 2002. Belief revision, explanations and defeasible reasoning. *Artificial Intelligence Journal* **141**, 1–28.

Falappa, M. A., Kern-Isberner, G. & Simari, G. R. 2009. Belief revision and argumentation theory. In *Argumentation in Artificial Intelligence*, Rahwan, I. & Simari, G. R. (eds). chapter 17, Springer, 341–360.

García, A. J. & Simari, G. R. 2004. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming* **4**(1), 95–138.

Gärdenfors, P. 1982. Rule for rational changes of belief. In *Philosophical Essay Dediccated To Lennart Åqvist on his Fiftieth Birthday*, Pauli, T. (ed.), University of Uppsala, Sweden, 88–101.

Gärdenfors, P. 1988. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. The MIT Press, Bradford Books.

Gärdenfors, P. 1992. *Belief Revision*. Cambridge University Press, 1–28.

Gärdenfors, P. & Makinson, D. 1988. Revisions of knowledge systems using epistemic entrenchment. In *Second Conference on Theoretical Aspects of Reasoning About Knowledge*. Asilomar, California, USA, 83–95.

Grove, A. 1988. Two modellings for theory change. *The Journal of Philosophical Logic* **17**, 157–170.

Hansson, S. O. 1994. Kernel contraction. *The Journal of Symbolic Logic* **59**, 845–859.

Hansson, S. O. 1997. *Theoria: Special Issue on Non-Prioritized Belief Revision*. Department of Philosophy, Uppsala University.

Hansson, S. O. 1999. *A Textbook of Belief Dymanics: Theory Change and Database Updating*. Kluwer Academic Publishers.

Hansson, S. O., Fermé, E., Cantwell, J. & Falappa, M. 2001. Credibility limited revision. *The Journal of Symbolic Logic* **66**(4), 1581–1596.

Harper, W. 1975. Rational belief change, popper functions and counterfactuals. *Synthese* **30**, 221–262.

Katsuno, H. & Mendelzon, A. 1992. On the difference between updating a knowledge database and revising it. In *Belief Revision*, Gärdenfors, P. (ed.), Cambridge University Press, 183–203.

Levi, I. 1977. Subjunctives, dispositions and chances. *Synthese* **34**(4), 423–455.

Lewis, D. 1973. *Counterfactuals*. Harvard University Press.

Moguillansky, M. O., Rotstein, N. D., Falappa, M. A., García, A. J. & Simari, G. R. 2008. Argument theory change applied to defeasible logic programming. In *Proceedings of The Twenty-Third Conference on Artificial Intelligence, AAAI'2008*. Chicago, Illinois, USA, 132–137.

Newell, A. 1982. The knowledge level. *Artificial Intelligence* **18**, 87–127.

Paglieri, F. & Castelfranchi, C. 2006. The Toulmin test: Framing argumentation within belief revision theories. In *Arguing on the Toulmin model*, Hitchcock, D. & Verheij, B. (eds), Springer, 359–377.

Pollock, J. L. & Gillies, A. S. 2000. Belief revision and epistemology. *Synthese* **122**(1–2), 69–92.

Rotstein, N. D., Moguillansky, M. O., Falappa, M. A., García, A. J. & Simari, G. R. 2008. Argument theory change: revision upon warrant. In *Proceedings of the Computational Models of Argument*, *COMMA'2008*. Toulouse, France, 336–347.

Rott, H. 1992. Preferential belief change using generalized epistemic entrenchment. *The Journal of Logic, Language and Information* **1**, 45–78.

Williams, M.-A., Pagnucco, M., Foo, N. & Sims, B. 1995. Determining explanations using transmutations. In *Proceedings of IJCAI-95*. Montreal, Quebec, Canada, 822–830.