International Journal of Approximate Reasoning 93 (2018) 395–423

Contents lists available at ScienceDirect

International Journal of Approximate Reasoning

www.elsevier.com/locate/ijar

An extension-based approach to belief revision in abstract argumentation

Martin Diller [a], Adrian Haret [a,*], Thomas Linsbichler [a,*], Stefan Rümmele [b], Stefan Woltran [a]

[a] *Institute of Information Systems, TU Wien, Vienna, Austria*
[b] *UNSW Sydney and University of Sydney, Australia*

ARTICLE INFO

Article history:
Received 21 March 2017
Received in revised form 24 November 2017
Accepted 29 November 2017
Available online xxxx

Keywords:
Abstract argumentation
Belief revision

ABSTRACT

Argumentation is an inherently dynamic process, and recent years have witnessed tremendous research efforts towards an understanding of how the seminal AGM theory of belief change can be applied to argumentation, in particular to Dung's abstract argumentation frameworks (AFs). However, none of the attempts have yet succeeded in solving the natural situation where the revision of an AF is guaranteed to be representable by a single AF. Here we present a solution to this problem, which applies to many prominent argumentation semantics. To prove a full representation theorem, we make use of recent advances in both areas of argumentation and belief change. In particular, we use the concept of realizability in argumentation and the concept of compliance as introduced in Horn revision. We also present a family of concrete belief change operators tailored specifically for AFs and analyze their computational complexity.

1. Introduction

Argumentation has emerged, over the last two decades, as a major research area in Artificial Intelligence (AI) [10,58]. This is due not just to the intrinsic interest of the topic and to its recent applications (see [50] and [2] for surveys emphasizing applications of argumentation in areas such as legal reasoning, medicine, and e-governance) but also because of fundamental connections between argumentation and other areas of AI, mainly non-monotonic reasoning.

The significant landmark in the consolidation of argumentation as a distinct field of AI has been the introduction of abstract argumentation frameworks (AFs) [33], which are directed graphs whose nodes represent arguments and where links correspond to attacks between arguments. To this day AFs remain the most widely used and investigated among the several argumentation formalisms. The study of AFs is mainly concerned with evaluating the acceptance of arguments when taking into consideration the structure encoded in the graph. A common approach to this is finding subsets of arguments (called *extensions*) that can all be accepted together. As a result, the argumentation literature offers a wide range of criteria (called *semantics of AFs*) for establishing which arguments are jointly acceptable [5].

Our work fits into the growing number of studies on the *dynamics* of argumentation frameworks [7,11–13,15,19,32,44,45,57,60]. This line of research is motivated by the realization that, as part of interactive reasoning processes, argumentation frameworks have to undergo change when new information becomes available. Particularly important in this respect is

* Corresponding authors.
E-mail addresses: diller@dbai.tuwien.ac.at (M. Diller), haret@dbai.tuwien.ac.at (A. Haret), linsbich@dbai.tuwien.ac.at (T. Linsbichler), s.rummele@unsw.edu.au (S. Rümmele), woltran@dbai.tuwien.ac.at (S. Woltran).

https://doi.org/10.1016/j.ijar.2017.11.013

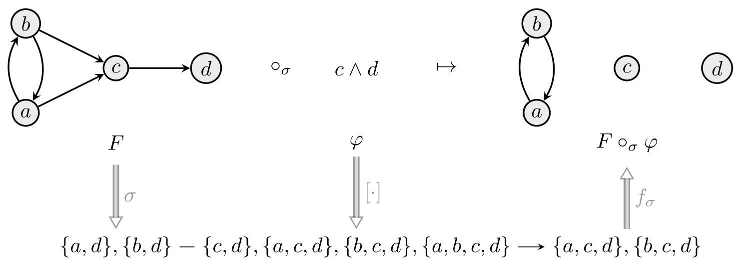

Fig. 1. F undergoes revision by φ.

change with respect to the acceptability of certain arguments: it is to be expected that increased knowledge of facts settles certain issues, with the effect that arguments pertaining to them have to either become part of, or be excluded from any extension of our AF. Thus, such increased knowledge must be reflected in a new AF which manages to preserve as much semantic information from the original one, while making sure that its extensions satisfy the added constraints. The main issue, in this setting, is to find appropriate ways of formalizing the notion of minimal change at the semantic level, with the understanding that the graph structure of the revised AF is then constructed around the semantic information. Settling on a specific graph structure for the revised AF is an interesting problem in its own right, though it is a separate issue from the one concerning us here, and left for future work.

We look at the problem through the lens of the semantic approach to propositional belief revision [42], where a knowledge base has a finite representation in a formal language, and this representation is used to encode a finite set of models. In our setup, AFs play the role of knowledge bases and their extensions under a certain semantics are the models. Thus, given a semantics σ, an AF F and a revision formula φ encoding desired changes in the status of some arguments, the task of a revision operator $\circ_\sigma$ is defined as follows: find an AF $F \circ_\sigma \varphi$ which manages to both satisfy φ and preserve as much useful information from F as possible. Example 1 illustrates the main steps in this process.

Example 1. Consider a propositional knowledge base $K = \{a \leftrightarrow \neg b, \neg c, d\}$, to undergo revision by $\varphi = c \wedge d$. A propositional revision operator $\circ$ would be expected to return a knowledge base $K \circ \varphi$ which implies φ: one can envision many ways to do this, but an approach based on minimizing information loss such as Dalal's operator (see [25] or Section 2) would pick the models of φ considered most plausible from the point of view of the models of K, and return the knowledge base $\{a \leftrightarrow \neg b, c, d\}$.

Consider, now, the AF F depicted in Fig. 1, where some semantics σ has singled out the extensions $\{a, d\}$ and $\{b, d\}$ as jointly acceptable: we think of these sets as the models of F. Suppose, next, that in light of new information (in the form of a propositional formula or another AF), we learn that c and d must be accepted. If F is to undergo revision by a formula $\varphi = c \wedge d$, this is interpreted as asking for an AF $F \circ_\sigma \varphi$ whose extensions satisfy certain constraints, e.g., they are models of φ. A strategy of minimizing information loss such as the one mentioned above would return $\{\{a, c, d\}, \{b, c, d\}\}$ as a suitable set of candidates. In the final step, a function f_σ constructs an AF $F \circ_\sigma \varphi$ with precisely this set of extensions.

In the paper we fill out this picture by formulating rationality constraints to guide the revision process, distinguishing between different forms which the new information can take (a formula or another AF), and making sure the resulting set of extensions can be represented by a single AF. The latter step turns out to be sensitive to the semantics used and poses non-trivial challenges. Remarkably, a *representation theorem* illuminates the problem, by showing that performing AF revision in accordance with some rationality postulates is equivalent to choosing among possible extensions of AFs, according to a particular type of rankings on extensions.

For the rationality constraints, we adapt a well-known core set of postulates from the literature on propositional revision [42]. In keeping with the different ways in which new information can be expressed, we study two types of revision operators. The first considers the new information represented as a propositional formula. This formula encodes, by its models, a set of extensions representing the change (in terms of extensions) to be induced in the original AF. The second type is revision by an AF, where new information is restricted in the sense that it can only stem from another AF's outcome. While the first type is similar to existing work [21], the latter assumes that new information stems from another agent's beliefs, and is in the form of an AF. This is more in line with work on Horn revision [28], where all involved formulas belong to some fragment of propositional logic. The two types of revision present interesting differences, particularly when considering the realizability of the result as an AF. Revision by a propositional formula is characterizable using standard revision postulates, as long as rankings on extensions satisfy a *compliance* restriction. Revision by an AF, on the other hand, turns out not to require compliance, but is only characterizable using an extra postulate called Acyc and what we call *proper I-maximal semantics*. Finally, we analyze the computational complexity of the main revision tasks.