

On Validating Theories of Abstract Argumentation Frameworks: the Case of Bipolar Argumentation Frameworks

Henry Prakken¹²

¹ Department of Information and Computing Sciences, Utrecht University, The Netherlands h.prakken@uu.nl

² Faculty of Law, University of Groningen, The Netherlands

Abstract. This paper discusses the validation of abstract formal or computational theories of argumentation as normative theories of argumentation. Three validation approaches are discussed: instantiation with a more concrete theory of argumentation (theory-based validation), validating with intuitions about concrete examples (intuition-based validation) and comparing the theory with how humans actually argue (empirical validation). It is argued that intuition-based validation can be useful for validating structured but not for validating abstract accounts of argumentation, that empirical validation can be used at all levels of abstraction but at best as a partial validation method, and that a full validation of abstract accounts of argumentation should include theory-based validation. A case study of the ‘standard’ theory of bipolar frameworks reveals that it is to a large extent still awaiting validation as a good theory of rational argumentation.

Keywords: Validating abstract argumentation theories · Bipolar argumentation frameworks

1 Introduction

When a formal or computational theory of argumentation is proposed, an important question is why it is a good theory of argumentation. In this paper I discuss several ways of answering this question for abstract theories of argumentation. By ‘abstract’ I mean theories that leave the nature of arguments and their relations unspecified. With ‘goodness’ I do not mean whether the theory is empirically adequate as a theory of how humans *actually* argue. Rather I mean whether it is a good theory of *rational* reasoning and argumentation, that is, of how people (or machines) *should* argue. I will in particular focus on bipolar argumentation frameworks (*BAFs*, frameworks with both attack and support relations between arguments), since these have recently been the subject of two detailed validation studies [4, 11]. I will critically analyse these studies with two aims in mind: a specific aim to see what these studies imply about whether *BAFs* are a good normative account of argumentation, and a general aim to develop guidelines for evaluating abstract theories of argumentation.

The first attempt to validate an abstract theory of argumentation was made by Dung in his seminal paper [5] on the theory of abstract argumentation frameworks (*AFs*, with just attack relations between arguments). It is worth quoting Dung in full [5, pp. 324-5].

Our main goal is to give an analysis of the nature of human argumentation in its full generality. This is done in two steps. In the first step, a formal, abstract but simple theory of argumentation is developed to capture the notion of acceptability of arguments. In the next step, we demonstrate the “correctness” (or “appropriateness”) of our theory. It is clear that the “correctness” of our theory cannot be “proved” formally. The only way to accomplish this task is to provide relevant and convincing examples. Two “examples” are provided. The first one shows how our theory can be used to investigate the logical structure of many human economic and social problems. The second one shows that many major approaches to nonmonotonic reasoning in AI and logic programming (...) are in fact different forms of our theory of argumentation.

Interestingly, as regards reasoning and argumentation, Dung did not instantiate his abstract theory with concrete natural-language examples but with logical systems, namely, default logic [19], Pollock’s [12] argumentation system and two logic-programming semantics. In this paper I argue that this is the best way of validating an abstract theory of argumentation, i.e., instantiating it with more concrete theories that have been independently validated as “correct” normative accounts of argumentation. I will call this *theory-based* validation. As a case study I will investigate whether there are such theory-based “proofs of correctness” for the ‘standard’ semantics of *BAFs*, and argue that these only partially exist. Then I will review two alternative ways in which *BAFs* have been validated, validating with intuitions in concrete examples (*intuition-based validation*) and Polberg & Hunter’s [11] recent experimental comparison with how humans argue (*empirical validation*). For *BAFs* I will argue that these validation attempts are inconclusive. More generally I will argue that intuition-based validation can at best be a partial validation method at more concrete levels of abstraction, that empirical validation can be used at all levels of abstraction but at best as a partial validation method, and that a full validation of abstract accounts of argumentation should include theory-based validation.

The style of my paper will be largely informal, paraphrasing rather than stating formal results and observations, and assuming that the reader is familiar with the discussed works. Nevertheless, for ease of reference I will start with a brief recapitulation of the main definitions of the theories of *AFs* and *BAFs*.

2 Formal preliminaries

In this section the basics of Dung’s theory of abstract argumentation frameworks and the main theories of bipolar frameworks are reviewed.

An **abstract argumentation framework** (AF) is a pair $\langle \mathcal{A}, \mathcal{C} \rangle$, where \mathcal{A} is a set of arguments and $\mathcal{C} \subseteq \mathcal{A} \times \mathcal{A}$ is a relation of attack. Dung’s theory of AF s identifies sets of arguments (called *extensions*) which are internally coherent and defend themselves against attackers. An argument $A \in \mathcal{A}$ is *defended* by a set by $S \subseteq \mathcal{A}$ if for all $B \in \mathcal{A}$: if B attacks A , then some $C \in S$ attacks B . Then relative to a given AF , $E \subseteq \mathcal{A}$ is *admissible* if E is conflict-free and defends all its members; E is a *complete extension* if E is admissible and $A \in E$ iff A is defended by E ; E is a *preferred extension* if E is a \subseteq -maximal admissible set; E is a *stable extension* if E is admissible and attacks all arguments outside it; and $E \subseteq \mathcal{A}$ is the *grounded extension* if E is the least fixpoint of operator F , where $F(S)$ returns all arguments defended by S . It holds that any preferred, stable or grounded extension is a complete extension. Finally, for $T \in \{\text{complete, preferred, grounded, stable}\}$, X is *sceptically* or *credulously* justified under the T semantics if X belongs to all, respectively at least one, T extension.

Bipolar frameworks add a binary support relation \mathcal{S} to AF s. Thus BAF s are a triple $(\mathcal{A}, \mathcal{C}, \mathcal{S})$. In [2] a *sequence of supports* for argument B by argument A is a sequence ASB_1, \dots, SB_nSB (it is said that A *supports* B). Often the semantics of BAF s is defined in terms of constraints on the attack relation given sets of attack and support relations between arguments, specifying which attack relations should also hold given other attack and support relations. Arguments are then evaluated with a given semantics for AF s applied to BAF s that are closed under these constraints. The following constraints are the most important ones that have been considered in the literature (the formulations below are adapted from [11]). Accordingly, I will call them the ‘standard semantics’ for BAF s. A semantics of BAF s can use any subset of these constraints.

- There is a *supported attack* from A to B iff there exists an argument C such that there is a sequence of supports from A to C and C attacks B .
- There is a *secondary attack* from A to B iff there exists an argument C such that there is a sequence of supports from C to B and A attacks C .
- There is an *extended attack* from A to B iff there exists an argument C such that there is a sequence of supports from C to A and C attacks B .
- There is a *mediated attack* from A to B iff there exists an argument C such that there is a sequence of supports from B to C and A attacks C .

3 Dung-style validation attempts of theories of bipolar argumentation frameworks

To the best of my knowledge, only two works following Dung’s validation strategy for bipolar argumentation frameworks exist: my own preliminary research in [15] and the recent excellent comprehensive study of Cohen et al. in [4]. Below I will mainly discuss the latter paper, since it includes the main positive result of [15] as a special case.

Cohen et al. take $ASPIC^+$ [9] as the basis for their investigations. $ASPIC^+$ instantiates AF s by giving definitions of the structure of arguments and the

nature of attack. Informally, in $ASPIC^+$ arguments are constructed from logical knowledge bases by chaining applications of two kinds of inference rules, and can be displayed as acyclic inference graphs (which are trees if no premise is used more than once). Arguments have *subarguments*. Informally, every argument B corresponding to subgraph of an argument A (viewed as a graph) that is also an argument (so takes all its premises from the premises of A) is a subargument of A . Moreover, if B does not equal A , then B is a *proper subargument* of A . Arguments can be attacked on their premises, except on those that are declared unattackable (*undermining attack*), on the conclusions of their defeasible inferences ([12]’s *rebutting attacks*) and on their defeasible inferences themselves by arguing that it has an exception ([12]’s *undercutting attacks*). These attack relations can be *indirect* in that if an argument is attacked, also all arguments that have this argument as a proper subargument are attacked. Thus $ASPIC^+$ ’s attack relation satisfies closure under secondary attacks. Attack relations can be resolved into so-called *defeat* relations by applying a binary preference relation on arguments to attacks. A set of $ASPIC^+$ arguments and their defeat relations then induces an AF by letting the AF attack relations correspond to the $ASPIC^+$ defeat relations. In this paper I will assume for convenience that all attacks succeed as defeats, which is equivalent to assuming that there are no strict preference relations. This allows me to speak of attack relations only. However, my observations will generalise to any defeat relation.

Cohen et al. first define four ways in which $ASPIC^+$ arguments can support each other. The first is the $ASPIC^+$ **proper subargument relation** between arguments. The second notion of support is a notion of argument accrual: two different arguments A and B **conclusion-support** each other if they have the same final conclusion. A third notion is premise support. Argument A **premise-supports** another argument B if A ’s final conclusion is a premise of B . Fourth, a variant of conclusion support is intermediate support: if A conclusion-supports a proper subargument of another argument B that does not equal a premise of B , then A **intermediate-supports** B .

Cohen et al. then consider three semantics for $BAFs$ in terms of the AFs generated by sets of attack constraints. A BAF , is for **general support** iff its attack relation is closed under supported and secondary attack. As in [3] a BAF is for **deductive support** iff its attack relation is closed under supported and mediated attacks. And as in [10] a BAF is said to be for **necessary support** if its attack relation is closed under secondary and extended attacks and, moreover, its support relation is irreflexive and transitive.

Cohen et al. then investigate whether their four notions of support in $ASPIC^+$ can be related to these three BAF semantics. They do this for each of the three $ASPIC^+$ attack relations separately. Like in the present paper they assume that all $ASPIC^+$ attacks succeed as defeats. For each $AF = (\mathcal{A}, \mathcal{C})$ induced by an $ASPIC^+$ instantiation and a particular $ASPIC^+$ attack relation (undermining, rebutting or undercutting attack), they first consider the $BAF = (\mathcal{A}, \mathcal{C}, \mathcal{S}_s)$ where \mathcal{S}_s is the support relation on \mathcal{A} according to $ASPIC^+$ support type s (proper-subargument, conclusion, premise or intermediate sup-

port). Then for each of the three *BAF* semantics x (general, deductive or necessary) they consider $BAF_x = (\mathcal{A}, \mathcal{C}_x, \mathcal{S}_s)$, where \mathcal{C}_x is the closure of \mathcal{C} under the attack constraints of semantics x . They then compare for each *ASPIC*⁺-induced $BAF = (\mathcal{A}, \mathcal{C}, \mathcal{S}_s)$ and each corresponding $BAF_x = (\mathcal{A}, \mathcal{C}_x, \mathcal{S}_s)$ the sets \mathcal{C} and \mathcal{C}_x . The semantics x is an abstraction of support type s just in case $\mathcal{C} = \mathcal{C}_x$.

Cohen et al.’s findings on this question are largely negative. They identify only one full correspondence, between *ASPIC*⁺ proper-subargument support and *BAFs* for necessary support, regardless of the type of *ASPIC*⁺ attack. A similar result was earlier proven by me in [15], except that I only assumed closure of attack under secondary attack relations. Cohen et al.’s result implies that also assuming closure under extended attacks does not affect this result.

Cohen et al. remain neutral on the methodological implications of their findings but in my opinion they provide strong support for the claim that the ‘standard’ semantics for *BAFs* has not yet been validated as a good theory of rational argumentation. Of course, this conclusion presupposes that *ASPIC*⁺ is itself a good structured account of rational argumentation, although it does *not* presuppose that *ASPIC*⁺ is the *only* such good account; nor does it presuppose that the only theory-based way to validate a theory of *BAFs* is by instantiating it with *ASPIC*⁺. In Section 5 I will return to these issues.

4 Alternative validation attempts of theories of bipolar argumentation frameworks

In the present section I discuss two main alternative validation approaches that have been employed in the literature: appeals to intuition in concrete examples and comparisons with how humans actually argue.

4.1 Validation with intuitions in concrete examples

An often applied strategy in validating the normative adequacy of theories of reasoning is applying them to concrete examples and checking whether they validate one’s intuitions about these examples. In the literature on computational argumentation this is occasionally used but not in a very systematic way. For structured accounts of argumentation (which specify the structure of arguments and the nature of their relations) it can be a useful aspect of validation but it should be used with care. One problem is that it is not obvious whose intuitions should count [21, Ch. I.1]. Those of logicians or argumentation researchers are seriously biased by overexposure to formalism. So should we ask the ‘average language user’, hoping that they are not infected by theoretical bias? Then the problem often arises that their answers reveal a lacking understanding of the reasoning patterns; however, teaching them about these reasoning patterns infects them with the theoretical bias we were hoping to avoid. Also, whether we ask researchers or lay people, intuitions of different persons often conflict.

In validating theories of defeasible reasoning another problem is that it is often hard to distinguish between intuitions that a particular reasoning pattern is

invalid and implicitly assumed additional information that invalidates inferences that are defeasibly valid in the absence of that information [13].

When validating *abstract* theories argumentation there is the problem that if the theory is validated with natural-language examples of argumentation that are directly translated into *AFs* instead of through a theory of the nature of arguments and attacks, then the resulting modellings may be ad-hoc, so that the observations made about these modellings have no general validity. Prakken & Winter [18] illustrate this problem with several examples from the literature on probabilistic abstract argumentation frameworks and bipolar frameworks.

In conclusion, intuition-based validation can, if used with care, be useful for validating structured accounts of argumentation but they should not be used for validating abstract accounts of argumentation, since directly encoding natural-language examples in abstract frameworks, without guidance by theories of the nature of arguments and their relations, risks in yielding ad-hoc modellings.

4.2 Empirical validation

Polberg & Hunter [11] empirically evaluate both *AFs* and *BAFs* in experiments with human test subjects, who were instructed to identify support and attack relations between the arguments in two natural-language persuasion dialogues. They call their experiments “explorative” and warn that their results are merely “indicative” and should serve as “a basis for further studies, rather than as an “indisputable proof for or against a given argumentation approach”. I will not discuss the ‘strength’ of their findings in this sense but rather whether their experiments can in principle support conclusions of the kind they draw.

Their overall conclusion is that the data from their experiments support the need for bipolar approaches. The question arises: the need for bipolar approaches for what purpose? Are they needed to specify how people *should* argue (the topic of the present paper) or are they needed to model how people *actually* argue? These things are quite different. It is not entirely clear from their paper which of these questions Polberg & Hunter are interested in. Yet it is important to disambiguate, since if their claims are interpreted as being about normative adequacy, then the question arises how empirical findings on how people actually argue can be relevant at all for a normative theory on how they should argue.

Polberg & Hunter write “Without empirical evidence, we can accidentally increase the gap between applying argumentation and *successfully* applying argumentation in real life situations” (their emphasis, HP). In my opinion, what this hints at is an empirical constraint on normative theories of reasoning that they should refer to the concepts that people use when reasoning or arguing, on the penalty of not being applicable by people. The principles formulated by the theory may still deviate from how people actually reason, as long as they are stated in terms that are natural to people. Otherwise it may happen that a theory, even if it is normatively adequate, is so foreign to people that they are not able to apply the theory, so that they are left without any useful guidance for their reasoning. It may be that Polberg & Hunter have this kind of empirical naturalness constraint on the usefulness of normative theories of argumentation

in mind. When interpreted thus, their results do not inform us about whether a semantics of *BAFs* is normatively correct but about whether it is stated in terms that are natural to people. Henceforth I will call this kind of validation *quasi-empirical validation* of (the human naturalness of) normative theories.

Let us see whether Polberg & Hunter’s results can be seen as quasi-empirical validation of *BAFs* in this sense. In my opinion, this is not the case, since Polberg & Hunter state their claims in absolute terms while their data at best support a relative conclusion, namely, that using *BAFs* is more quasi-empirically adequate than using *AFs*, since their experiments are designed to compare *AFs* and *BAFs*. However, there are other alternatives, namely, using *AFs* in combination with a theory of the structure of arguments and the nature of attack. There are quite a few such theories, dating back to the seminal work of Pollock in [12]. For a historical overview see my [16]. All these theories allow for support relations that are not *between* but *inside* arguments, namely as inferential relations between (sets of) *statements* in some logical language. Examples are assumption-based argumentation, Defeasible Logic Programming [7] and *ASPIC*⁺. The arguments of [11] for their claims are thus instances of the fallacy of the excluded middle.

What is quasi-empirically more adequate: modelling support relations primarily as relations between arguments or primarily as relations between statements? Polberg & Hunter’s experiments are inherently unsuitable for answering this question, since they did not instruct their test subjects to identify support relations between statements in the sense of a structured account of argumentation, so between premises and conclusions *inside* arguments. In fact, Polberg & Hunter’s use of the terms ‘statement’ and ‘argument’ is sometimes ambiguous. They refer to the elements of their persuasion dialogues as “statements”. However, most of these “statements” are what argumentation theorists would call arguments, with statements as premises, another statement as a conclusion and a claimed relation of support between premises and conclusion. For example:

Hospital staff members are exposed to the flu virus a lot. *Therefore*, it would be good for them to receive flu shots in order to stay healthy (my emphasis, HP).

On the other hand, some of their ‘statement’ are just claims. For example:

Hospital members do not need to receive flu shots.

Since Polberg & Hunter do not systematically distinguish between statements as full arguments or as elements of arguments, they also do not systematically distinguish between support relations *between* and *inside* arguments. In any case, no conclusion can be drawn from their experiments on the relative quasi-empirical adequacy of *BAF*-style approaches and structured approaches to the modelling of support relations in argumentation.

5 Validation of the instantiations

As noted above at the end of Section 3, and as also acknowledged by [11, p. 488], the success of theory-based validation depends on whether the more con-

crete theories with which abstract theories are instantiated are themselves valid as normative theories of reasoning or argumentation (or at least accepted as valid by a substantial part of the research community). A full treatment of this issue is beyond the scope of this paper but some observations can be made. Dung instantiated his theory with one of the best known nonmonotonic logics, default logic, with two semantics for logic programming, which is a well-established computational theory, and with a system for defeasible argumentation (Pollock’s system) with foundations in epistemology which has been very influential in the study of formal argumentation. Others have instantiated Dung’s theory with, for instance, assumption-based argumentation [6] and the *ASPIC*⁺ framework [14]. Both of these frameworks have in turn been validated to some extent. Assumption-based argumentation has been used for reconstructing several well-known nonmonotonic logics and it has been successfully used in many applications [20], which contributes to its quasi-empirical validation. *ASPIC*⁺ has been used for reconstructing assumption-based argumentation [14] and several variants of classical-logic argumentation [8], which is a Dung-style validation of *ASPIC*⁺. Moreover, it has received quasi-empirical validation in several applications or case studies, as summarised in Section 6.3 of [9]. Of course, opinions may differ on the extent to which validation studies have been successful but in my opinion most of this other work has been sufficiently validated so that Dung’s theory of *AFs* has received theory-based validation to a very high degree.

6 Conclusion

In this paper I have argued that the best way of validating the ‘correctness’ of an abstract normative theory of argumentation is theory-based validation, that is, instantiation with more concrete theories of reasoning and argumentation that have themselves been sufficiently validated in other ways. As a case study I have investigated how current bipolar approaches, which extend Dung’s abstract argumentation approaches with abstract support relations, have been validated in this way. It turned out that the only two papers that, to my knowledge, apply theory-based validation to bipolar approaches, [15] and [4], yield negative results for *BAF* semantics for general and deductive support but positive results for necessary support, which has been shown to be an abstraction of *ASPIC*⁺’s subargument relation.

The largely negative results of [15] and [4] are telling. They should warn us that developing a theory of argumentation without taking the structure of arguments and the nature of attack into account is dangerous: it increases the risk of developing a piece of mathematics that is elegant and seductive but has no clear relations to real argumentation phenomena. Moreover, a ‘social’ risk is that it sends the wrong message to young researchers that this is the way research in argumentation should be done, namely, by ignoring the structure of arguments. It may make that not enough effort is devoted to the study of genuine argumentation phenomena. I therefore repeat my proposal from [16] that any newly proposed abstract formalism for argumentation should be accompanied

with at least one meaningful instantiation. In other words, we should follow the example that Dung set in his classic 1995 paper.

After applying Dung’s validation strategy to bipolar approaches, I investigated two common alternative validation strategies of abstract approaches. As regards validation of *BAF* semantics by intuitions in concrete examples, I repeated and summarised observations in the literature that this validation strategy is ad-hoc and unreliable and that modellings of concrete examples should be guided by accounts of the structure of arguments and the nature of their relations. As regards validation of *BAFs* by empirical comparison with how people actually argue I have made the following contributions. I first distinguished several senses of validation and argued that Polberg & Hunter’s experiments in [11] can best be interpreted not as validation of these theories as normative theories of rational argumentation but as validation of their naturalness to human arguers. The benefit of this kind of quasi-empirical validation is that it guides researchers in developing normative theories of argumentation that are not only rationally well-founded but also applicable by humans. Such theories can still deviate from the way humans actually argue but since they are stated in terms that are natural to humans, they have a higher chance of being applied by humans than possibly better but less natural normative theories.

Next, I argued that Polberg & Hunter’s results do not contribute to this kind of quasi-empirical validation, since their paper is silent about the vast body of work in AI on the structure of arguments and the nature of attack and support. I agree with them that support relations are important in argumentation but this has been known for a long time. All informal definitions of argumentation in the argumentation-theoretic literature and all formal definitions of the structure of arguments in the AI literature are variations on the theme that argumentation involves supporting claims with grounds. The question is how such support relations can best be modelled. One main issue here is whether support can best be modelled primarily as a relation *between* arguments or primarily as a relation between grounds and claims *inside* arguments. Polberg & Hunter do not answer to this question, since they do not consider the latter alternative.

Finally, an interesting question is what is the role of postulate-based approaches (e.g. [1]) in validating abstract theories of argumentation. In my opinion, this depends on knowing what abstract theories abstract from, since otherwise it is hard to say whether a postulate makes sense. So in this sense postulates for abstract accounts also require theory-based validation. In [17] I made this point in more detail for abstract accounts of probabilistic argumentation.

To conclude, although the the present paper by no means offers a comprehensive account of validating abstract accounts of argumentation, it provides some arguably useful distinctions between various kinds of validation and some guidelines for how to conduct validation attempts.

References

1. Baroni, P., Rago, A., Toni, F.: How many properties do we need for gradual argumentation? In: Proceedings of AAAI 2018. pp. 1736–1743 (2018)

2. Cayrol, C., Lagasquie-Schiex, M.C.: Bipolar abstract argumentation systems. In: Rahwan, I., Simari, G. (eds.) *Argumentation in Artificial Intelligence*, pp. 65–84. Springer, Berlin (2009)
3. Cayrol, C., Lagasquie-Schiex, M.C.: Bipolarity in argumentation graphs: Towards a better understanding. *International Journal of Approximate Reasoning* **54**, 876–899 (2013)
4. Cohen, A., Parsons, S., Sklar, E., McBurney, P.: A characterization of types of support between structured arguments and their relationship with support in abstract argumentation. *International Journal of Approximate Reasoning* **94**, 76–104 (2018)
5. Dung, P.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n -person games. *Artificial Intelligence* **77**, 321–357 (1995)
6. Dung, P., Mancarella, P., Toni, F.: Computing ideal sceptical argumentation. *Artificial Intelligence* **171**, 642–674 (2007)
7. Garcia, A., Simari, G.: Defeasible logic programming: An argumentative approach. *Theory and Practice of Logic Programming* **4**, 95–138 (2004)
8. Modgil, S., Prakken, H.: A general account of argumentation with preferences. *Artificial Intelligence* **195**, 361–397 (2013)
9. Modgil, S., Prakken, H.: Abstract rule-based argumentation. In: Baroni, P., Gabbay, D., Giacomin, M., van der Torre, L. (eds.) *Handbook of Formal Argumentation*, vol. 1, pp. 286–361. College Publications, London (2018)
10. Nouioua, F., Risch, V.: Argumentation frameworks with necessities. In: *Proceedings of the 4th International Conference on Scalable Uncertainty Management (SUM’11)*. pp. 163–176. No. 6929 in *Springer Lecture Notes in AI*, Springer Verlag, Berlin (2011)
11. Polberg, S., Hunter, A.: Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *International Journal of Approximate Reasoning* **93**, 487–543 (2018)
12. Pollock, J.: Defeasible reasoning. *Cognitive Science* **11**, 481–518 (1987)
13. Prakken, H.: Intuitions and the modelling of defeasible reasoning: some case studies. In: *Proceedings of the Ninth International Workshop on Nonmonotonic Reasoning*. pp. 91–99. Toulouse, France (2002)
14. Prakken, H.: An abstract framework for argumentation with structured arguments. *Argument and Computation* **1**, 93–124 (2010)
15. Prakken, H.: On support relations in abstract argumentation as abstractions of inferential relations. In: *Proceedings of ECAI 2014*. pp. 735–740 (2014)
16. Prakken, H.: Historical overview of formal argumentation. In: Baroni, P., Gabbay, D., Giacomin, M., van der Torre, L. (eds.) *Handbook of Formal Argumentation*, vol. 1, pp. 73–141. College Publications, London (2018)
17. Prakken, H.: Probabilistic strength of arguments with structure. In: *Proceedings of KR 2018*. pp. 158–167 (2018)
18. Prakken, H., de Winter, M.: Abstraction in argumentation: necessary but dangerous. In: Modgil, S., Budzynska, K., Lawrence, J. (eds.) *Computational Models of Argument*. *Proceedings of COMMA 2018*, pp. 85–96. IOS Press, Amsterdam etc (2018)
19. Reiter, R.: A logic for default reasoning. *Artificial Intelligence* **13**, 81–132 (1980)
20. Toni, F.: A tutorial on assumption-based argumentation. *Argument and Computation* **5**, 89–117 (2014)
21. Veltman, F.: *Logics for Conditionals*. Doctoral dissertation, Department of Philosophy, University of Amsterdam (1985)