

Philosophical reflections on Argument Strength and Gradual Acceptability

Henry Prakken^{1,2}[0000–0002–3431–7757]

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

² Faculty of Law, University of Groningen, Groningen, The Netherlands

Abstract. This paper proposes a classification of three aspects of argument strength based on philosophical insights, in particular Aristotle’s distinction between logic, dialectic and rhetoric. It is then argued that when developing or evaluating gradual accounts of argument strength it is essential to be explicit about which aspect of argument strength is modelled and about the adopted interpretation of the arguments and their relations in abstract or bipolar argumentation frameworks. The underlying aim is to encourage a principled development and evaluation of (principles for) gradual argumentation semantics.

Keywords: Computational argument · Argument strength · graduality

1 Introduction

A recent trend in the formal study of argumentation is the development of gradual notions of argument acceptability. These notions are proposed as alternatives to extension-based notions that are defined on top of the theory of abstract [16] or bipolar [13] argumentation frameworks. The gradual notions are often motivated by a discontent with the fact that extension-based notions of acceptability only allow for rather coarse distinctions between degrees of acceptability. The current developments arguably go back to [12] and really took off with publications like [29] and [1] (although largely ignored is that Pollock [35] already proposed a formalisation of gradual acceptability). In this body of work, the gradual nature of argumentation can have various sources: different base strengths of arguments, different sets or numbers of attackers and/or supporters, and attack or support relations that hold to varying degrees.

Although the new developments are very interesting and the formal achievements have been impressive, there are also reasons to take a step back. To start with, there is a need to reflect on which notions or aspects of argument acceptability, or argument strength, are modelled, and why proposed semantics or proposed sets of principles for those semantics are good. What is needed is a conceptual or philosophical underpinning of the formal ideas and constructs.

Furthermore, almost all work builds on abstract or bipolar argumentation frameworks and thus does not give explicit formal accounts of the nature of arguments and their relations, while yet this may be relevant when evaluating the formal proposals. Moreover, the arguments in abstract and bipolar frameworks are increasingly not seen as genuine arguments in the sense of inferential structures (as in Dung’s seminal paper and much initial follow-up work) but as statements that can be true or false. In this paper I will argue that some proposed (principles for) gradual semantics, while making sense when arguments are interpreted as statements, make less sense when arguments are regarded as inferential structures and hence cannot be regarded as general accounts of argumentation semantics.

More generally, in this paper I aim to make two contributions. First, I will propose a classification of three aspects of argument strength based on philosophical insights, in particular Aristotle’s distinction between logic, dialectic and rhetoric. I then argue that when developing or evaluating gradual accounts of argument strength it is essential to be explicit about which aspect of argument strength is modelled and about the adopted interpretation of the arguments and their relations. The underlying hope is that this paper will encourage a principled and focused development and evaluation of (principles for) gradual argumentation semantics. The discussion will largely proceed in terms of examples of principles and semantics, since given the fast growing literature a comprehensive discussion and analysis is outside the scope of this paper.

The paper is organised as follows. First in Section 2 I informally sketch the assumed formal background, and then in Section 3 I present the three aspects of logical, dialectical and rhetorical strength. In Section 4 I discuss why it is important to be explicit about which of these aspects is modelled and about the nature of arguments and their relations. Finally, in Section 5 I conclude.

2 Background

In this section I informally summarise the formal background assumed in this paper. The aim of my paper is not to carry out formal investigations but to offer a conceptual framework that can guide the development and evaluation of specific formal proposals. I will therefore use as little formal notation as possible and assume that the reader is familiar with the basics of formal argumentation, in particular of the theory of abstract [5] and bipolar [13] argumentation frameworks and of the main structured accounts of argumentation [27].

2.1 Arguments as statement or as inferential structures

In this subsection I explain in more detail the recent trend to view arguments as statements and how this is relevant for the evaluation of gradual argumentation semantics. Abstract argumentation frameworks as introduced by Dung [16] consist of a set of arguments with a binary relation of attack. Nothing is assumed about the structure of the arguments and the nature of the attack relations. Bipolar frameworks add to abstract argumentation frameworks a binary

support relation between arguments, sometimes but not always assumed to be disjoint from the attack relation. Various semantics have been proposed for evaluating the arguments in an abstract or bipolar argumentation framework. For present purposes their details do not matter.

As noted in the introduction, the arguments in abstract or bipolar frameworks are in applications increasingly not seen as genuine arguments in the sense of inferential structures but as statements that can be true or false. For evaluating (principles for) gradual argumentation semantics it is crucial how the nodes in an argument graph are interpreted, so the adopted interpretation should be made explicit. Sometimes this is indeed done, e.g. by Baroni et al. [7] and Rago et al. [39], who present their model as a formalisation of the IBIS model [28]. Then the formalism can be evaluated on its adequacy for what it was explicitly meant for. However, often the ‘statement’ interpretation of the nodes in argumentation frameworks is not made explicit but has to be inferred from the informal text and from the examples that are given. For instance, in [42] an example is discussed with the following (and some other) arguments:

S_1 : *We should buy an SUV; it's the right choice for us*
 C : *SUVs are very safe, safety is very important to us*

where C supports S_1 . This support relation in fact expresses an argument with premises *SUVs are safe* and *Safety is important to us* and the conclusion *We should buy an SUV*, which readers familiar with the theory of argument schemes will recognise as an instance of the argument scheme from good consequences. In a formalisation where arguments are inferential structures this argument would appear as a single node in the argument graph, while here it is spread out over a subgraph of a bipolar argumentation framework with nodes C and S_1 .

The difference between the inferential and statement interpretations of the nodes in an abstract or bipolar argumentation framework is important for the design and choice of abstract formalisms. While bipolar argumentation frameworks model support as a relation *between* arguments, when support expresses some kind of inference, it can also be modelled *inside* arguments, at least if arguments are interpreted as inferential structures. Then this can be done by using abstract argumentation frameworks 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 [34]. 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 [44], Defeasible Logic Programming [19] and *ASPIC*⁺ [37]. While this approach is possible when arguments are interpreted as inferential structures, this is different when they are interpreted as statements; then support relations *between* arguments (viewed as statements) are needed, as in bipolar argumentation frameworks. Relations between arguments then become what philosophers call *reasons* [34, 24]. In this approach, arguments-as-inferential-structures are not nodes in but subgraphs of the argument graph, as illustrated by the above example. Moreover, since in general

statements are supported by sets of statements, we need support relations from *sets* of arguments to arguments, which in ‘standard’ bipolar frameworks cannot be expressed. So how the arguments in abstract and bipolar frameworks are interpreted greatly matters for what abstract formalism is needed.

2.2 Basic concepts of argument structure and relations

I next informally sketch what I mean in this paper by arguments and their relations, trying to remain as close as possible to the formal (e.g. [27]) and informal (e.g. [18]) literature on argument structure. Basic arguments have a set of *premises* and a *conclusion* (statements that can be true or false) and an *inference* from the premises to the conclusion licensed by an *inference rule*. Basic arguments can be combined into complex arguments by letting their conclusion be among the premises of another argument. Arguments can be informally represented as directed acyclic hypergraphs, in a way similar to the usual visualisation methods in argumentation theory [40], with the nodes corresponding to premises or intermediate or final conclusions and the links from sets of nodes to nodes corresponding to inferences (see e.g. Figures 1 and 2 below).

Both premises and inference rules can be *attackable* or *non-attackable*, so arguments can also be *attackable* or *non-attackable*. I will call attackable and non-attackable inference rules *defeasible*, respectively, *deductive*. An attackable argument can be *attacked* in three ways: it can be *undermined* by an argument of which the conclusion is incompatible with one of its premises, it can be *rebutted* by an argument of which the conclusion is incompatible with an intermediate or its conclusion, and it can be *undercut* by an argument of which the conclusion says that some inference rule applied in the attacked argument does not apply. Attacks can be *allowed* or *not allowed* depending on further constraints on these informal definitions. Finally, allowed attacks can succeed or not succeed as *defeats*. Henceforth, when I say that argument *A attacks* argument *B* I assume that the attack from *A* on *B* is allowed, while when I say that argument *A defeats* argument *B* I assume that the attack from *A* on *B* succeeds as defeat.

This informal sketch could be regarded as an abstraction of the *ASPIC*⁺ framework. Depending on the precise formal definitions of arguments, of incompatibility of statements and of the constraints on allowed attacks and defeats, one variant or another of this framework can be obtained, or a variant of some related approach like assumption-based argumentation or defeasible logic programming. For present purposes the precise design choices and the differences between these formal frameworks do not matter.

3 Logical, dialectical and rhetorical argument strength

In classifying aspects of argument strength it is natural to take Aristotle’s famous distinction between logic, dialectic and rhetoric as starting point. Very briefly, *logic* concerns the validity of arguments given their form, *dialectic* is the art of testing ideas through critical discussion and *rhetoric* deals with the principles of

effective persuasion [17, Section 1.4]. Accordingly, I distinguish between logical, dialectical and rhetorical argument strength.

Logical argument strength divides into two aspects: inferential and contextual argument strength.

Inferential argument strength is about how well the premises support the conclusion if we only look at the arguments premises, inferences and conclusion(s). Example criteria for argument strength are that arguments with only deductive inferences are stronger than arguments with defeasible inferences, or that arguments with only non-attackable premises are stronger than arguments with attackable premises. Such criteria can be refined by combining them (for example, first looking at the type of inference and then for arguments with equally strong inferences looking at the types of premises), by defining preference relations on inference rules and/or premises, or by defining notions of strength (for example, probabilities) on inference rules and/or premises.

Contextual argument strength is about how well the conclusion of an argument is supported if we look at the context of all relevant arguments. Formal frameworks like Dung’s theory of abstract argumentation frameworks, assumption-based argumentation, *ASPIC*⁺ and defeasible logic programming formalise this kind of argument strength. The reader might wonder why this is not called dialectical strength, since after all, determining an argument’s contextual strength as defined here involves the comparison of argument and counterargument. Yet this is not truly dialectical, since the just-mentioned formalisms do not model principles of critical discussion but just define mathematical consequence notions on the basis of a given body of information; likewise [20, 30]. This even holds for the argument games proposed as proof theories for extension-based semantics [32]: these apply to a given framework of arguments and their relations, while principles of discussion allow for introduction of new arguments during a dispute.

Dialectical argument strength looks at how well defended an argument is in the context of an ongoing or terminated critical discussion. This context can be regulated by formal or informal principles of fair and effective disputes. Informal examples are legal procedures or rules of order for meetings. Formal examples are the many dialogue systems for argumentation proposed in philosophy and AI, e.g. [31, 20, 46, 36, 4]. Dialectical strength has both static and dynamic aspects. A static aspect is given by the outcome of a critical discussion: has the argument been successfully defended in the discussion? Dynamic aspects concern how well-defended or challengeable an argument is in a given state of the discussion. In [49] the latter is formulated as “a function of the (un)availability of permissible move sequences originating at the present dialogue stage, and ending in a discussant’s role-specified goal being achieved”.

To illustrate the idea of dialectical strength, I next suggest some possible criteria for determining dialectical argument strength, without claiming to be exhaustive. First, one might regard an argument as dialectically weaker the more attacks on it are allowed in the current state. Among other things, this may

imply that arguments are dialectically weaker the more attackable premises or inference rules they have. This idea is motivated by an underlying principle that many decision makers are aware of, namely, to justify one's decisions as sparsely as possible, in order to minimize the chance of successful appeal.

One might also look at how many attacks an argument has survived in a given state. For instance, if arguments are generated by argument schemes [47], one might regard an argument as dialectically stronger the more critical questions have been asked and successfully answered.

Yet another aspect of dialectical strength is to what extent it is possible to change the current contextual strength of the argument by moving to a new state. This aspect is arguably formalised by formal work on the dynamics of argumentation, in particular on so-called preservation, realisability and enforcement properties [8, 15]. Preservation is about the extent to which the current contextual status of arguments is preserved under change, while realisability and enforcement concern the extent to which particular outcomes can or will be obtained by changing the current state.

Rhetorical argument strength looks at how capable an argument is to persuade other participants in a discussion or an audience. Persuasiveness essentially is a psychological notion; although principles of persuasion may be formalised, their validation as principles of successful persuasion is ultimately psychological (as acknowledged in [26] and done in e.g. [23]). One way to formulate criteria for persuasiveness is in terms of agreement with shared background information or with information in a model of the other discussants or the audience (cf. [25]), possibly refined with probability distributions on what can be in these models. Such opponent models could also be used in game-theoretic investigations of optimal debate strategies [41]. Another way is to formulate heuristics about what is generally known or expected to be persuasive, such as the argumentation techniques of Perelman [33] or the argument schemes of Walton [45], or a technique such as *procatalepsis*, modelled by Bonzon et al. [11], which is the attempt of a speaker to strengthen their argument by dealing with possible counter-arguments before their audience can raise them.

Argument strength is multi-faceted. The classification proposed in this section shows that argument strength is a multi-faceted notion. Not only can we distinguish between logical, dialectical and rhetorical argument strength but each of these aspects of strength involves multiple criteria, which sometimes reinforce but sometimes oppose each other, and which have to be combined to provide an overall assessment of an argument's logical, dialectical or rhetorical strength. In fact, defining such overall notions seems a daunting task, and it may be better to focus on just one criterion or a small set of criteria for aspects of argument strength. This also holds for combining logical, dialectical and rhetorical strength into one overall notion for argument strength. Since the three aspects of argument strength serve different purposes, it may not be good to combine them into an overall notion. Another reason for this is that dialectical strength may *presuppose* contextual strength, since one aspect of dialectical strength is the extent to

which an argument's logical strength may be changed in the course of a dispute. In any case, even if logical, dialectical and rhetorical strength are combined into an overall notion of strength, they should first be separately defined, in order to make their combination a principled one.

4 Evaluating semantics and principles

In this section I discuss how the above classification into three aspects of argument strength, together with the interpretation of the nodes and links in abstract or bipolar argumentation frameworks, is relevant for developing and evaluating gradual accounts of argument strength. A complication here is that several recent accounts rely on the distinction between the base and overall score of an argument, but it is not obvious whether notions of argument strength can always be suitably defined in this format. For example, such accounts presuppose that the base and overall scores are of the same sort and can therefore be compared but this does not have to be the case. Consider, for example, gradual definitions of contextual strength in terms of extension-based [10] or labelling-based [48] semantics. A very simple definition of overall strength would be that (given a grounded labelling) being *in* is better than being *undecided*, which is better than being *out*. It is not obvious what the base score of arguments would be in such an approach. For these reasons I will below mainly discuss accounts with no distinction between base and overall scores.

4.1 Be explicit about which aspects of argument strength are modelled

It is important to be explicit about which aspects of argument strength are modelled (as in [11], who explicitly model two aspects of persuasiveness). The aspects serve different purposes, so principles or definitions that are good for one aspect may not be good for another aspect. Consider, for example, two arguments A and B where A defeasibly infers q from p while B first defeasibly infers r from p and then defeasibly infers q from r . Consider a definition of dialectical strength capturing that having fewer attackable elements is dialectically better and a definition of rhetorical strength that captures that a larger overlap of an argument's elements with the audience's beliefs is rhetorically better. Even without formalising these notions it is obvious that argument A is dialectically stronger than argument B , since A has one attackable element less than B . However, if the audience accepts that p defeasibly implies r and that r defeasibly implies q but not that p defeasibly implies q , then B is rhetorically stronger than A since it shares some elements with the background theory while A does not. This illustrates that while sparsely justifying one's claims or decisions may be dialectically good, it may at the same time make an argument less persuasive.

Another example is the phenomenon of procatalepsis, modelled in [11], which is the attempt to strengthen an argument by dealing with possible counter-arguments before the audience can raise them. As modelled by Bonzon et al.,

procatalepsis makes an argument rhetorically stronger when combined with an attacker and an attacker of that attacker than when presented alone. Bonzon et al. prove that procatalepsis is inconsistent with the principle of ‘void precedence’ [1], according to which an argument that has no attackers is more acceptable than an argument that has attackers, even if these attackers are counterattacked, and which is a key element of many current gradual argumentation semantics [9]. It is implied by Basic Idea 7 of [6] that a strictly larger set of attackers determine a lower strength. One interpretation of the void precedence principle is as an aspect of dialectical strength, in particular, as capturing that having fewer ways to attack an argument makes it dialectically less challengeable. On this interpretation, examples of procatalepsis are other cases where an argument can be rhetorically stronger but dialectically weaker than another argument.

Being explicit about which aspects of argument strength are modelled is not only important when formulating theories of argument strength but also when evaluating applications of computational argumentation. For instance, according to [43] the Debater system was evaluated by twenty human annotators who had to indicate to what extent they agreed with the statement ‘The first speaker is exemplifying a decent performance in this debate’. It is unclear which aspects the annotators had in mind when answering this question or even whether all annotators looked at the same aspects and applied the same criteria.

4.2 Be explicit about the interpretation of arguments and their relations

It is important to be explicit about whether the arguments in an abstract or bipolar framework are regarded as statements or as inferential structures. One reason is that trade-offs between attacks and supports, or even to regard supporters as strengthening the supported argument, may make sense in the statement interpretation but may not make sense in the inferential-structure interpretation. When the arguments are regarded as statements, then attack and support relations between arguments can hardly be interpreted as anything else but expressing reasons for or against the statement. For example, in a decision-making application [7, 39] attack and support relations are reasons to adopt or not adopt a given decision option. Then it makes sense to consider trade-offs between supporting and attacking arguments, as, for example, captured in the (Strict) Franklin, Weakening and Strengthening principles proposed in [3]. The same holds for Basic Idea 8 of [6], according to which (everything else being equal) a strictly larger set of supporters determines a higher strength. In the statement interpretation of arguments it makes sense to say that (in the absence of attackers) a statement for which there are reasons to believe or accept it is more acceptable than a statement for which there are no such reasons.

However, when arguments are regarded as inferential structures, then multiple interpretations of the support relation are possible and their differences matter. In the context of the *ASPIC*⁺ framework, Cohen et al. [14] define four kinds of support (visualised in Figure 1). Subargument support corresponds to

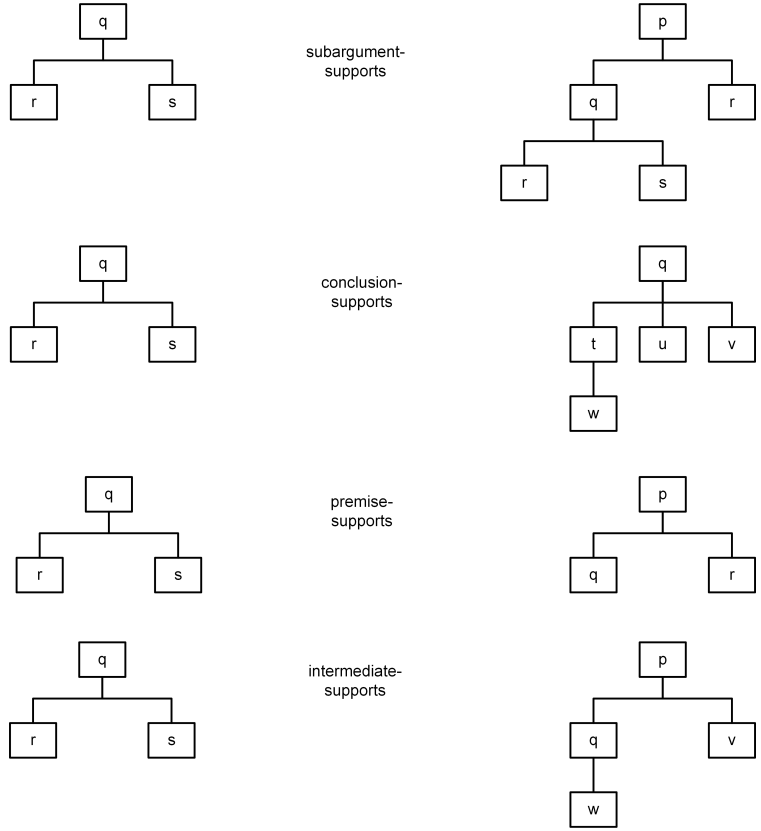


Fig. 1. Cohen et al.’s (2018) four kinds of support in *ASPIC*⁺.

the *ASPIC*⁺ proper-subargument relation. Informally, every argument *B* corresponding to subgraph of an argument *A* (viewed as a hypergraph) that is also an argument (so takes all its premises from the premises of *A*) is a subargument of *A*. A subargument of an argument *A* is not separate from *A* but is contained in *A* as part of it, as visualised in the top row of Figure 1. Clearly, when support corresponds to (proper) subargument support, it makes no sense to consider trade-offs between attacks and support or to regard supporters as strengthening the supported argument. Logically, the number of supports of an argument is then just a measure of its inferential complexity while dialectically, having more supporters may make an argument more vulnerable to attack and thus weaker.

For conclusion support (argument *A* supports *B* whenever they have the same conclusion) these ideas may make more sense but for premise support they are again questionable. An argument *A* premise-supports an argument *B* iff the final conclusion of *A* is equal to a premise of *B* (intermediate support is a variant of premise support in which not a premise but an intermediate conclusion

is supported; I will therefore not discuss it separately). Premise support may be useful in debate contexts, where there usually is no global knowledge base from which the debaters construct their arguments. Then if one debate participant presents an argument for a premise of another participant’s argument, it may be less natural to say that the supporting argument is part of the supported argument as captured in the subargument relation. Instead, the arguments may be said to remain separate, as depicted in the third row of Figure 1.

Consider an example from [38], in which John argues “Nixon was a pacifist since he was a Quaker and Quakers are usually pacifists” (argument *A*). Now suppose Mary supports John’s argument by saying “Nixon regularly attended service in a Quaker church, people who are regularly seen in a Quaker church usually are a Quaker, so Nixon was a Quaker” (argument *B*). Consider contextual strength. Can we say that Mary’s supporting argument makes John’s argument contextually stronger? If so, then a successful attack on Mary’s argument should intuitively also weaken John’s argument. Suppose that Bob attacks Mary’s argument by arguing that Nixon only attended service in a Quaker church to please his wife, who was a Quaker (argument *C*). It can be argued that this does not knock down John’s argument, since why should John be blamed for Mary’s flawed attempt to support his argument? However, it is still unsatisfactory that there is no logical relation at all between attacking a supporter and the status of a supported argument. If support means anything at all, then surely attacking a supporter should have some effect on the status of an argument supported by it (note that for subargument support this is automatic). The solution adopted in [38] is that whether *C*’s attack on *B* also weakens *A* is conditional on what the audience accepts as given. If the audience accepts that Nixon was a quaker without further support, then argument *C* has no effect on the acceptability of *A*, while if the audience wants further support for this premise, then argument *C* reduces *A*’s acceptability. This approach implies that whether trade-offs between premise-supporters and attackers should be considered, or whether a premise-supporter can strengthen the supported argument at all, cannot be determined in general but depends on the context, in particular on the audience’s beliefs.

While this is one approach, there may also be reasons to always consider trade-offs between premise-supporters and attackers and to regard a premise-supporter as, everything else being equal, strengthening the supported argument. However, even then the nature of the arguments and their relations matters. Consider Figure 2, with two bipolar frameworks in the top row and two instantiations of these frameworks in the bottom row (in BAF1 and BAF2 the dashed arguments depict support relations between arguments). According to the principle that, everything else being equal, having more supporters is better (e.g. the Cardinality Preference axiom of [2]), argument *C* on the top right is better supported than argument *A* on the top left since *C* has two premise-supporters while *A* has just one. However, as shown in the bottom row, all of *A*’s premises (namely, *q*) are supported while only one of *C*’s two premises is supported, so dialectically and perhaps also rhetorically *A* might just as well be regarded as better supported than *C*. Or imagine that *D* does not premise-support *C* on *u*

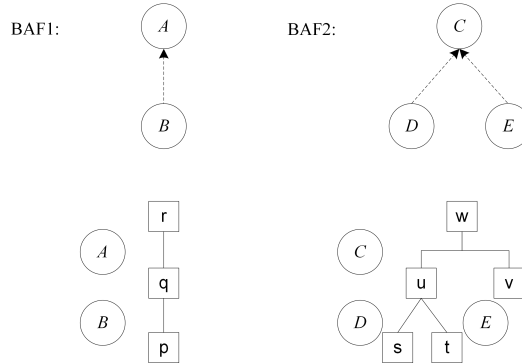


Fig. 2. Is having more supporters better?

but on v : then both A and C have all their premises supported, so there seems no reason to prefer C over A . Consider next Basic Idea 8 of [6], which says that a strictly larger set of supporters (w.r.t. to set inclusion) determines a higher strength, and suppose that argument A has no premise supporters while argument C only has supporters for premise u . Since any set includes the empty set, Basic Idea 8 implies that C is stronger than A . It is not obvious why this should be, given that they both have the same number of unsupported premises. This becomes even less obvious if C is changed to have more than two premises, of which only one is supported. Concluding, even in applications in which it makes sense to regard premise-supporters as, everything else being equal, strengthening the supported argument, it is important to take the structure of arguments and the nature of their relations into account.

So far I have illustrated the importance of being explicit about whether arguments are statements or inferential structures, and in the latter case of being explicit about the structure of arguments and the nature of the support relation. I next illustrate the importance of being explicit about whether an argument is attackable or not (a distinction made in some of the main structured approaches to argumentation, such as assumption-based argumentation, defeasible logic programming and *ASPIC⁺*). Consider the Cardinality Precedence principle that having fewer attackers makes an argument stronger [1, 9] and consider AF1 with A being attacked by B and AF2 with C being attacked by D and E (Figure 3; in the AFs in Figures 3 and 4 the solid arrows depict attack relations). According to Cardinality Precedence argument A is stronger than argument C . However, if B is not attackable while D and E are attackable, then it is not obvious why this should be the case. For example, from the point of view of dialectical strength C is arguably dialectically stronger than A since C can still be made in by adding new arguments and attacks while for A this cannot happen.

Another example of why the distinction between attackable and non-attackable arguments matters concerns the principle that having more defenders makes an




Fig. 3. Is having fewer attackers better?

argument stronger. Consider the *AF*s displayed in Figure 4. According to the




Fig. 4. Is having more defenders better?

gradual semantics of [21, 22], A_2 is justified to a higher degree than A_1 , since A_2 has two defenders (C_2 and D_2) while A_1 has only one defender (C_1). However, if C_1 is unattackable while C_2 and D_2 are attackable then it is not obvious why this has to be so, whatever aspect of argument strength is modelled.

5 Conclusion

In this paper I proposed a classification of three aspects of argument strength, namely, logical, dialectical and rhetorical strength. I then showed with several examples that when developing or evaluating gradual accounts of argument strength it is essential to be explicit about which aspect of argument strength is modelled, since some principles or semantics may be suitable for one aspect but not for another. Likewise, I showed that it is important to be explicit about the adopted interpretation of the arguments and their relations in abstract or bipolar argumentation frameworks. For example, it matters whether the arguments are interpreted as statements or as inferential structures, how support is defined, and how the structure of arguments is defined.

The underlying aim of this paper was to encourage a principled development and evaluation of (principles for) gradual argumentation semantics. As such, this was just an initial attempt. More comprehensive formal investigations should yield more systematic insights into the purposes for which (principles for) gradual semantics are suitable and into the assumptions on which they depend. This paper has aimed to lay the conceptual foundations for such investigations.

References

1. Amgoud, L., Ben-Naim, J.: Ranking-based semantics for argumentation frameworks. In: Liu, W., Subrahmanian, V., Wijsen, J. (eds.) *Scalable Uncertainty Management. SUM 2013*. pp. 134–147. No. 8078 in Springer Lecture Notes in Computer Science, Springer Verlag, Berlin (2013)
2. Amgoud, L., Ben-Naim, J.: Evaluation of arguments from support relations: axioms and semantics. In: *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16)*. pp. 900–906 (2016)
3. Amgoud, L., Ben-Naim, J.: Weighted bipolar argument graphs: axioms and semantics. In: *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-18)*. pp. 5194–5198 (2018)
4. Atkinson, K., Bench-Capon, T., McBurney, P.: A dialogue game protocol for multi-agent argument over proposals for action. *Journal of Autonomous Agents and Multi-Agent Systems* **11**, 153–171 (2005)
5. Baroni, P., Caminada, M., Giacomin, M.: An introduction to argumentation semantics. *The Knowledge Engineering Review* **26**, 365–410 (2011)
6. Baroni, P., Rago, A., Toni, F.: How many properties do we need for gradual argumentation? In: *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI 2018)*. pp. 1736–1743 (2018)
7. Baroni, P., Romano, M., Toni, F., Aurisicchio, M., Bertanza, G.: Automatic evaluation of design alternatives with quantitative argumentation. *Argument and Computation* **6**, 24–49 (2015)
8. Baumann, R.: What does it take to enforce an argument? minimal change in abstract argumentation. In: *Proceedings of the 20th European Conference on Artificial Intelligence*. pp. 127–132 (2012)
9. Bonzon, E., Delobelle, J., Konieczny, S., Maudet, N.: A comparative study of ranking-based semantics for abstract argumentation. In: *Proceedings of the 30st AAAI Conference on Artificial Intelligence (AAAI 2016)*. pp. 914–920 (2016)
10. Bonzon, E., Delobelle, J., Konieczny, S., Maudet, N.: Combining extension-based semantics and ranking-based semantics for abstract argumentation. In: *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference*. pp. 118–127. AAAI Press (2018)
11. Bonzon, E., Delobelle, J., Konieczny, S., Maudet, N.: A parametrized ranking-based semantics compatible with persuasion principles. *Argument and Computation* **12**, 49–85 (2021)
12. Cayrol, C., Lagasque-Schiex, M.C.: Graduality in argumentation. *Journal of Artificial Intelligence Research* **23**, 245–297 (2005)
13. Cayrol, C., Lagasque-Schiex, M.C.: Bipolar abstract argumentation systems. In: Rahwan, I., Simari, G. (eds.) *Argumentation in Artificial Intelligence*, pp. 65–84. Springer, Berlin (2009)
14. 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)
15. Doutre, S., Mailly, J.G.: Constraints and changes: A survey of abstract argumentation dynamics. *Argument and Computation* **9**, 223–248 (2018)
16. 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)

17. van Eemeren, F., Garssen, B., Krabbe, E., Henkemans, A.S., Verheij, B., Wage-
mans, J.: *Handbook of Argumentation Theory*. Springer, Dordrecht (2014)
18. Freeman, J.: *Dialectics and the Macrostructure of Arguments. A Theory of Argu-
ment Structure*. Fors/de Gruyter, Berlin-New York (1991)
19. Garcia, A., Simari, G.: Defeasible logic programming: An argumentative approach.
Theory and Practice of Logic Programming **4**, 95–138 (2004)
20. Gordon, T.: The Pleadings Game: an exercise in computational dialectics. *Artificial
Intelligence and Law* **2**, 239–292 (1994)
21. Grossi, D., Modgil, S.: On the graded acceptability of arguments. In: *Proceedings
of the 24th International Joint Conference on Artificial Intelligence*. pp. 868–874
(2015)
22. Grossi, D., Modgil, S.: On the graded acceptability of arguments in abstract and
instantiated argumentation. *Artificial Intelligence* **275**, 138–173 (2019)
23. Hadoux, E., Hunter, A.: Comfort or safety? gathering and using the concerns of a
participant for better persuasion. *Argument and Computation* **10**, 113–147 (2019)
24. Horty, J.: *Reasons as Defaults*. Oxford University Press, Oxford (2012)
25. Hunter, A.: Making arguments more believable. In: *Proceedings of the 19th National
Conference on Artificial Intelligence*. pp. 6269–274 (2004)
26. Hunter, A.: Towards a framework for computational persuasion with applications
in behaviour change. *Argument and Computation* **9**, 15–40 (2018)
27. Hunter, A. (ed.): *Argument and Computation*, vol. 5 (2014), special issue with
Tutorials on Structured Argumentation
28. Kunz, W., Rittel, H.: *Issues as elements of information systems*. Working Paper
No. 131, Institute of Urban and Regional Development, University of California,
Berkeley, Calif. (1970)
29. Leite, J., Martins: Social abstract argumentation. In: *Proceedings of the 22nd
International Joint Conference on Artificial Intelligence (IJCAI-11)*. pp. 2287–2292
(2011)
30. Loui, R.: Process and policy: resource-bounded non-demonstrative reasoning. *Com-
putational Intelligence* **14**, 1–38 (1998)
31. Mackenzie, J.: Question-begging in non-cumulative systems. *Journal of Philosoph-
ical Logic* **8**, 117–133 (1979)
32. Modgil, S., Caminada, M.: Proof theories and algorithms for abstract argumen-
tation frameworks. In: Rahwan, I., Simari, G. (eds.) *Argumentation in Artificial
Intelligence*, pp. 105–129. Springer, Berlin (2009)
33. Perelman, C., Olbrechts-Tyteca, L.: *The New Rhetoric. A Treatise on Argumen-
tation*. University of Notre Dame Press, Notre Dame, Indiana (1969)
34. Pollock, J.: Defeasible reasoning. *Cognitive Science* **11**, 481–518 (1987)
35. Pollock, J.: Defeasible reasoning with variable degrees of justification. *Artificial
Intelligence* **133**, 233–282 (2002)
36. Prakken, H.: Coherence and flexibility in dialogue games for argumentation. *Jour-
nal of Logic and Computation* **15**, 1009–1040 (2005)
37. Prakken, H.: An abstract framework for argumentation with structured arguments.
Argument and Computation **1**, 93–124 (2010)
38. Prakken, H.: Modelling support relations between arguments in debates. In:
Chesñevar, C., Falappa, M.A., Ferm, E., Garca, A.J., Maguitman, A.G., Martínez,
D.C., Martínez, M.V., Rodríguez, R.O., Simari, G.I. (eds.) *Argumentation-based
Proofs of Endearment. Essays in Honor of Guillermo R. Simari on the Occasion of
his 70th Birthday*, pp. 349–365. College Publications, London (2018)

39. Rago, A., Toni, F., Aurisicchio, M., Baroni, P.: Discontinuity-free decision support with quantitative argumentation debates. In: *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifteenth International Conference*. pp. 63–72. AAAI Press (2016)
40. Reed, C., Walton, D., Macagno, F.: Argument diagramming in logic, law and artificial intelligence. *The Knowledge Engineering Review* **22**, 87–109 (2007)
41. Riveret, R., Prakken, H., Rotolo, A., Sartor, G.: Heuristics in argumentation: a game-theoretical investigation. In: Besnard, P., Doutre, S., Hunter, A. (eds.) *Computational Models of Argument*. Proceedings of COMMA 2008, pp. 324–335. IOS Press, Amsterdam etc (2008)
42. Rosenfeld, A., Kraus, S.: Providing arguments in discussions based on the prediction of human argumentative behavior. In: *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 2015)*. pp. 1320–1327 (2015)
43. Slonim, N., Bilu, Y., Alzate, C.: An autonomous debating system. *Nature* **591**, 397–384 (2021)
44. Toni, F.: A tutorial on assumption-based argumentation. *Argument and Computation* **5**, 89–117 (2014)
45. Walton, D.: *Argumentation Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, NJ (1996)
46. Walton, D., Krabbe, E.: *Commitment in Dialogue. Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany, NY (1995)
47. Walton, D., Reed, C., Macagno, F.: *Argumentation Schemes*. Cambridge University Press, Cambridge (2008)
48. Wu, Y., Caminada, M.: A labelling-based justification status of arguments. *Studies in Logic* **3**, 12–29 (2010)
49. Zenker, F., Debowska-Kozłowska, K., Godden, D., Selinger, M., Wells, S.: Five approaches to argument strength: probabilistic, dialectical, structural, empirical, and computational. In: *Proceedings of the 3rd European Conference on Argumentation*. pp. 653–674. College Publications, London (2020)