

Argumentation in artificial intelligence

T.J.M. Bench-Capon, Paul E. Dunne*

Department of Computer Science, University of Liverpool, Liverpool, United Kingdom

Received 27 April 2007; received in revised form 27 April 2007; accepted 1 May 2007

Available online 10 May 2007

Abstract

Over the last ten years, argumentation has come to be increasingly central as a core study within Artificial Intelligence (AI). The articles forming this volume reflect a variety of important trends, developments, and applications covering a range of current topics relating to the theory and applications of argumentation. Our aims in this introduction are, firstly, to place these contributions in the context of the historical foundations of argumentation in AI and, subsequently, to discuss a number of themes that have emerged in recent years resulting in a significant broadening of the areas in which argumentation based methods are used. We begin by presenting a brief overview of the issues of interest within the classical study of argumentation: in particular, its relationship—in terms of both similarities and important differences—to traditional concepts of logical reasoning and mathematical proof. We continue by outlining how a number of foundational contributions provided the basis for the formulation of argumentation models and their promotion in AI related settings and then consider a number of new themes that have emerged in recent years, many of which provide the principal topics of the research presented in this volume.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Argumentation models; Dialogue processes; Argument diagrams and schemes; Agent-based negotiation; Practical reasoning

1. Introduction

In its classical treatment within philosophy, the study of argumentation may, informally, be considered as concerned with how assertions are proposed, discussed, and resolved in the context of issues upon which several diverging opinions may be held. Thus philosophical investigations of argumentation, from Aristotle to the present day, have addressed such themes as: the mechanisms by which “legitimate” argumentation in support of a claim may be distinguished from “flawed” argumentation; analyses of the typical structures that constitute argument components and argumentation development; the processes by which participants engaging in debate may advance their respective positions and undermine contrary stances and arguments, etc; and the contexts in which these questions are decided. The importance of such philosophical theories to so-called everyday reasoning has a long and distinguished history in AI, and contributions from contemporary philosophical analyses continue to play a major role in the evolution of effective computational exploitation of argumentation technology.

Within the simplified overview of argumentation outlined in the preceding paragraph, one can, already, identify a number of themes whose elements embody issues of a computational nature in the following:

* Corresponding author.

E-mail address: ped@csc.liv.ac.uk (P.E. Dunne).

- Defining the component parts of an argument and their interaction.
- Identifying rules and protocols describing argumentation processes.
- Distinguishing legitimate from invalid arguments.
- Determining conditions under which further discussion is redundant.

It is, of course, the case that similar issues underpin one well-established and highly-developed theory: that of formal logic and mathematical proof. It is no coincidence that much of the formal computational treatment of argumentation has its roots in ideas developed from AI inspired contributions to logic and deductive reasoning. So one finds in mathematical proof theory core concepts such as: precisely defined means for expressing assertions (e.g. formulae in a given logical language); accepted bases on which to build theorems (e.g. collections of axioms); procedures prescribing the means by which further theorems may be derived from existing theorems and axioms (e.g. templates for inference rules); and precise concepts of termination (e.g. a sentential form is derivable as a theorem, “true”; or is logically invalid, “false”).

While the structural elements presented in this view of mathematical reasoning have proven to be a useful basis in the development of argumentation-based models in AI, the formal apparatus and methods of mathematical reasoning are, ultimately, radically different in nature to those of importance when considering the concept of argumentation as it is familiar from everyday contexts, e.g. as it might occur in political debate, the discussion of ethical principles, deliberation in judicial settings, etc. While there are, of course, parallels that can be made,—e.g. that those engaged in debate have some collection of accepted premises on which there is agreement, possibly, even, some recognition of when contributions to a discussion are “unreasonable” or flawed, etc.—there are, however, a number of fundamental distinctions between the concepts “ P is a formal proof that T holds” and “ P is a persuasive argument for accepting T ”. Thus, in mathematical reasoning,

- The premises can, ultimately, be explicitly defined in terms of closed concepts, e.g. the axioms of Euclidean geometry, the Zermelo–Frankel basis for set theory (ZF). Furthermore classical mathematical reasoning is based on an assumption that such premises are, collectively, *consistent*.¹
- Reasoning and analysis takes place within a closed, tightly defined context, i.e. there is no notion of “incomplete” or “uncertain” information.
- Conclusions are final and definite: if P is a correct proof that T , then T is, *ipso facto* valid and this status does not admit subsequent qualification or amendment, let alone retraction.
- Reasoning and conclusions are entirely *objective*, not susceptible to *rational* dispute on the basis of subjective views and prejudices.² Proof is demonstration whereas argument is persuasion.

In argument and discussion as encountered in everyday contexts, it is rare that any, let alone all, of these apply: the premises upon which debates may build are often presupposed as forming part of the background assumptions common to all parties involved; the information and knowledge brought to bear in the course of discussion will often be incomplete, vague, or uncertain. The remaining two aspects, in many ways, highlight the most significant differences between “logical proof” and “persuasive argument”. Arguments are *defeasible*: the reasoning that formed a persuasive case for T , in the light of changes in viewpoint or awareness of information not previously available, may subsequently fail to convince. This defeasibility is never removed: an argument may cease to be challenged and so accepted, but the *possibility* of challenge remains. Finally, the extent to which an argued case is accepted is *subjective*, dependent on the views, attitudes, and prejudices of the audiences to which it is directed. The same case may convince some people but, equally, fail to convince others.

¹ We note that in a number of systems, consistency *cannot* be formally proven, cf. [95] and so, in such cases, consistency is, indeed, an *assumption*.

² Some clarification of this claim may be in order. Suppose Δ is a derivation of φ within a theory $\langle A, R \rangle$ (with axioms A and inference rules R). Within the *same* theory, the proof Δ admits no rational, objective basis for dispute: criticisms that “ φ is ‘inconvenient’ or ‘counter-intuitive’” are subjective, and entirely irrelevant to the status of φ within the theory $\langle A, R \rangle$. In order to give rational grounds for not accepting φ it is necessary to endorse an alternative theory within which φ *cannot* be derived. As a concrete example, consider the axiomatic basis ZF extended by the so-called “Axiom of Choice” (ZF + AC): although widely adopted in modern theory this conflicts with Intuitionist principles which disqualify AC as an axiom so that theorems dependent on AC are (rationally) not accepted by Intuitionists.

One can summarise the distinction between argumentation and proof by the observation that the object of argumentation is to *persuade* (to acceptance of a given claim; to performance of a desired action, and so on). Unlike the concept of “proof”—at the level of deriving a sentential representation of an assertion—whether an argument is “correct” is not a factor, and, indeed, “correctness” may not even be sensibly defined. In contrast, mathematical reasoning, in order to have any value, must be correct where “correctness” has a strict, formal definition: beyond this requirement, however, notions of “persuasiveness” are unimportant.

In summary, the importation of elements from logic and formal deductive reasoning has provided a powerful basis for modelling and analysing argumentation in computational settings of AI. As we shall discuss later, these continue to form an important strand of contemporary work. It is also the case, however, that a number of significant directions pursued in recent years, have broadened the scope and concerns of argumentation in AI beyond this earlier logic driven motivation. As a consequence, one has a shift of emphasis within the developed treatment of argumentation in AI progressing from formalisms rooted in classical deductive reasoning through models handling concepts of incomplete information and uncertainty, to precise semantics for capturing defeasibility, and, within recent work, propounding computational bases to account for subjective aspects of argumentation, often using the notion of “audience” introduced by Perelman [145]. One consequence of such analyses has been the growth of work dealing with computational procedures and issues of resource-boundedness in implementing these, e.g. as discussed in Loui [118]. All of these features make argumentation particularly attractive to applications requiring distributed intelligence, autonomous components and synchronous interaction.

To conclude this overview, it is worth noting one further historical and continuing tradition. In common with many established areas of AI, the computational theory of argumentation has benefited from contributions and ideas originating in many diverse disciplines, so that a number of fundamental themes draw on the earlier work of, for example, philosophers, logicians, and legal theorists. In presenting concrete realisations of these theories, work on argumentation in AI has in its turn informed further research in these fields. Argumentation is thus an excellent example of interdisciplinary interchange and the mutual benefits that can stem from this.

2. Foundations of argumentation in AI

A discussion of early influences on the development of argumentation models in AI may be found in the comprehensive survey of Chesnevar, Maguitman, and Loui [57], so we will be content merely to outline a few significant aspects, referring the reader to [57] for a more detailed exposition. We concentrate on three important influential themes,

- 2.1 Origins in non-classical logic.
- 2.2 Models of argumentation as dialogue process.
- 2.3 Diagrammatic views of argument structure.

2.1. Influence of non-classical logics on argumentation in AI

Early studies using argumentation inspired methods in AI contexts can be found in the work of Birnbaum, Flowers, and McGuire, [42] in which a structural model of argument embracing notions of *support* and *attack* within a graph-theoretic base comprising propositional forms, is applied to textual reasoning; and Alvarado and Dyer’s approaches, [4,5], to the analysis of editorial presentation.

Undoubtedly, the important early motivations that brought argumentation theory into use in AI arose from the issues of reasoning and explanation in the presence of incomplete and uncertain information. The failings of classical propositional logic as a means to address these had been delineated in the influential work of Reiter [165], and, a pressing concern of work throughout most of the 1980s and early 1990s was to build on the proliferation of treatments of *non-monotonic* logics within AI. This state is succinctly summarised by [57, pp. 337–338].

Within AI, several non-monotonic reasoning formalisms emerged . . . In these formalisms, conclusions drawn may be later withdrawn when additional information is obtained. Formal logics of argument emerged as one style of formalising non-monotonic reasoning. The literature on non-monotonic reasoning dominated AI’s journals in the mid 1980s.

Thus argumentation was initially adopted as a possible supporting approach with which to effect a formal treatment of non-monotonic reasoning, rather than as a paradigm whose study might be of independent interest in itself. The engagement of philosophers and legal theorists with reasoning and argumentation in AI marked a key stage in the move towards computationally grounded models of argument. Particularly notable is the impact of Pollock’s work on defeasible reasoning and justification: originally promoted in specialist philosophical literature, e.g. [146–148] its relevance and significance to AI was recognised following Pollock’s dissemination of these ideas in [149–151].

In parallel with this development of the formal logical theory—in which context the significance of argumentation techniques with respect to non-classical logic was further emphasised in the contributions of Simari and Loui [173] and Brewka [44] (ideas in the latter being subsequently developed in [45])—the early 1990s saw important uses of argumentation techniques in the computational treatment of legal reasoning: notably in Rissland and Ashley’s treatment of legal argumentation from Case Law, [14,167,174] and its later extension by Aleven [3]; Prakken’s analyses in [152]; Sartor’s models of legal reasoning as described in [169,170]; the use of argumentation techniques in explaining complex legislation from Bench-Capon, Coenen, and Orton [29], etc.

The technical treatment evident in AI contributions to non-monotonic logics and the argumentation-based methodologies offered in the field of legal reasoning found some degree of common ground in the exploitation of logic programming paradigms and knowledge-based systems. It was in this context, building on argument-based treatments of “negation-as-failure” of Kakas, Kowalski, and Toni [108], together with Eshghi and Kowalski’s work on abductive interpretation [89], that the watershed contribution of Dung [72,73] appeared: the model of argumentation described in [73] is now recognised as providing an important bridge between argumentation theory as a supporting analytic tool for non-monotonic reasoning and the independent exploitation of argumentation models in wider AI contexts.

Two important ideas are put forward and expanded in [73]:

- (A) The reduction of argumentation about a given issue to a completely abstract setting consisting of a set of “atomic” arguments, \mathcal{X} , and a binary relation over these, $\mathcal{A} \subseteq \mathcal{X} \times \mathcal{X}$, with $\langle x, y \rangle \in \mathcal{A}$ interpreted as “the argument x attacks the argument y ”.
- (B) The proposal that intuitive notions of “collection of justified arguments” can be formally described by that of *extension-based semantics*: that is, through various properties of subsets, S of \mathcal{X} within an *argumentation framework* (AF), $(\mathcal{X}, \mathcal{A})$.³

The effect of (A) is that neither the structure of an argument nor the nature and semantics underpinning “ x attacks y ” need explicit consideration within the abstract framework. Thus an “argument”, x , may be a simple atomic proposition, p ; or a (defeasible) rule, e.g. $p \leftarrow q \wedge r$; or an instantiation of a richer, more particular, perhaps even domain specific, argument scheme. That x attacks y may be on account of reasons varying in form from “ x promotes a claim logically equivalent to the negation of that promoted by y ”, e.g. $x : p$ and $y : \neg p$; or “ x promotes a claim incompatible with the premises supporting the claim in y ”, e.g. $x : \neg p$ and $y : q \leftarrow p \wedge r$, and so on to the extent that “attacks” disputing the applicability of a given inference scheme and more complex structures are represented entirely abstractly in a single binary relation.

Dung’s introduction of various extension-based semantics has, as we shall discuss in Section 3.1, had a profound influence on subsequent analyses of the concept of “collection of justified arguments”. In extremely informal terms, an extension semantics, \mathcal{E} , can be thought of as describing properties that a subset of arguments within a given framework must satisfy in order to be deemed collectively justified, i.e. $\mathcal{E} : \langle \mathcal{X}, \mathcal{A} \rangle \times 2^{\mathcal{X}} \rightarrow \{\top, \perp\}$. Dung demonstrates how different choices of \mathcal{E} may be used to colour varying degrees of an argument’s “acceptability” ranging from very liberal (so-called credulous) conditions through to extremely restrictive (so-called sceptical) requirements. The elements of Dung’s original set-theoretic semantics are reviewed in a number of articles in this issue and for further technical exposition we direct the reader to the article in this volume by Baroni and Giacomin [22].

The past 5–7 years have witnessed an intensive study of mechanisms with the common aim of developing Dung’s ideas in various directions. For a detailed comparative critique of abstract argumentation techniques we refer the reader to the valuable perspective provided by Vreeswijk [182].

³ The significance of this work is enhanced since approaches which include additional information, such as preferences, may do so in such a way that the evaluation of argument status remains in terms of an underlying abstract framework.

Subsequent work [43] of Dung in conjunction with Bondarenko, Kowalski and Toni, makes explicit the link between abstract argumentation and uniform treatment of non-classical logics. The *Assumption-based frameworks* (ABF) of [43] consider deductive theories— $(\mathcal{L}, \mathcal{R})$ —(with \mathcal{L} a formal language such as the language of well-formed propositional sentences, and \mathcal{R} a (countable) set of inference rules) augmented by a triple $\langle T, Ab, \bar{\cdot} \rangle$ in which $T \subseteq \mathcal{L}$ is a set of *beliefs*, $Ab \subseteq \mathcal{L}$ a (non-empty) set of *assumptions* and $\bar{\cdot}: Ab \rightarrow \mathcal{L}$ maps assumptions to their *contrary* in \mathcal{L} .⁴ Such frameworks are shown to be applicable as a generic approach to describing a wide range of non-classical logics⁵ including: Reiter’s *Default Logic* [165], Moore’s *Autoepistemic Logic* [131], logic programming, and divers other non-monotonic reasoning formalisms.

While ABF structures may on first inspection seem unrelated to the abstract argumentation frameworks of [73], these can be presented as AFS by building the attack relation from the “contrary” mapping. A fuller overview of this approach may be found in, e.g. the paper of Dung et al. in this volume [76, Section 2.2]. One feature of importance in this abstraction of ABFs is that the resulting structure will typically describe an infinite graph. Informally, extension-based semantics for ABFs are introduced as subsets of Ab whose union with the belief set T constitute a consistent theory. Fuller technical descriptions may be found in [43, pp. 70–71]. Such links between abstract argumentation frameworks and the deductive bases underpinning assumption-based schemes bring two powerful analytic approaches to bear in algorithmic studies of extension-based semantics for argumentation: combinatorial and algorithmic graph theory have been usefully applied in the former case; whereas technology developed for deductive reasoning and formal logic has provided insight into the latter. We expand on such computational and algorithmic issues in Section 3.1.

2.2. Argumentation and dialogue processes

The perspective of argumentation presented in Section 2.1 is strongly biased to a view wherein the overall aim of argumentation is in deciding the status of some claim and in presenting a justification for it: thus, an assertion, p , is established in the light of available information, but recognised as potentially defeated should new data emerge; the nature of “justification” often being through some logical reasoning process. In total such a view treats argumentation in support of a claim as a somewhat one-sided process in which a single party merely presents a reasoned justification. In many applications such an abstraction is, of course, a natural analogue to use, e.g. in explanation-driven systems such as [29], or in the context of decision-support processes.

An objection to such treatments, however, is that they fail to embrace the dialectical nature of argument, discourse, and debate as encountered in everyday contexts: here argumentation is rarely a matter of a single party presenting a case but is more commonly an informed *exchange* of ideas and positions involving several contributors: in other words, argumentation concerning an issue, typically, arises as a *dialogical* process. Given this it is, perhaps, surprising that significant computational exploitation of the established treatments of dialogue within philosophical, rhetorical, and linguistic analyses, has been a comparatively recent phenomenon. Although originally explored to a limited extent as a means of interacting with expert systems, the significant factor motivating contemporary computational use of dialogue methods can be found in supporting multi-agent system applications, a topic that we review in Section 3.2. In this section we review a number of foundational contributions and the preliminary AI motivated developments of these.

As with the developments discussed in Section 2.1 many of the ideas within computational treatments of dialogue build on contributions originating from philosophical analysis. One established concern of such study is the notion of “fallacy”: a key aspect of which is the view that so-called “fallacious argument” encompasses a much wider collection of issues than simply what may be (more accurately) termed “erroneous (mathematical or logical) reasoning”. Thus, argument employing fallacious reasoning (in this wider philosophical and rhetorical sense) is not *ipso facto* “wrong” nor easily dismissed merely by the action of highlighting occurrences of fallacy.⁶ It is the case, however—and here one finds a basis for the interaction between argumentation and dialogue processes—that particular fallacies occasion

⁴ Some treatments of ABFs omit explicit specification of a belief base T , e.g. [75,76].

⁵ Choosing $(\mathcal{L}, \mathcal{R})$ to be the language and standard inference structures of classical propositional logic, together with $\bar{\varphi} = \neg\varphi$ —that is the contrary is simply logical negation—the resulting ABF structures recover standard propositional reasoning.

⁶ Although it is, of course, true that in regarding errors of logic as a form of fallacy, indicating such incidences could suffice as an attack on the argument in whose support they are used.

potential attacks on a specific line of argument. For example, “*argumentum ad verecundiam*” whereby a claim is unreasonably justified through appeal to the opinion of an authority, is regarded in rhetorical theories of debate as fallacious. The use of such support cannot be properly deemed invalid merely by signalling its occurrence. Such fallacies are, however, open to attack by processes that invite further discussion. Thus the argument, “ φ is the case since X has stated as much” can be attacked in several ways, e.g. by disputing the authority of X in matters relating to the domain of φ , by challenging the assertion that X has made any statement regarding φ , etc. We, thus, find two important themes emerging: the classification of distinct types of fallacy; and the nature of possible attacking arguments. Categories of fallacious reasoning have long been a topic of interest in rhetorical analysis and the treatment of Hamblin [100] has had some influence on computational ideas. An important contribution to the second issue—the nature of attacks on fallacious argument—is found in Walton’s formulation of *Critical Question* [186] which has been adopted in several computational treatments of persuasive argument, e.g. [17,18,155,188]. An influential contribution in which a number of key ideas that have played a significant role in computational realisations of dialogue machinery, also arose in modelling and detecting one specific type of fallacious reasoning: MacKenzie’s *dialogue game*, DC, aimed at exposing uses of *petitio principii* in argumentation, [121]. A striking feature of MacKenzie’s analysis,⁷ evident from [121, Appendix, pp. 129–132], is the wealth of computational ideas that are introduced and their expression in an operational form. Thus, the concepts of commitment store, dialogue rules, locutions, etc. have all been adapted and extended in current dialogue-based applications. Among the earliest argumentation-based presentations of MacKenzie’s ideas beyond their use in recognising a specific class of fallacious reasoning, are those of Moore [130] and Bench-Capon, Dunne, and Leng [31].

Analyses of fallacy in debate contexts had, typically, not differentiated the range of styles and aims to which dialogue processes might be directed. This concept—that dialogues could be distinguished by their intentions—is expanded in the seminal work of Walton and Krabbe [189]. The taxonomy of dialogue types in [189] is neither intended to be nor presented as a definitive, complete catalogue of dialogue forms,⁸ but it has had considerable influence on the treatment of dialogue in multi-agent systems. The central significance of [189] to later work on argumentation in AI is in promoting an awareness that the *purposes* of dialogue encompass a number of different aims and, therefore, the appropriate procedural mechanisms (for example, as might be defined from MacKenzie’s model) employed in computational use have distinctive requirements, e.g. the operational specification of dialogue processes geared to negotiating agreements are unlikely to be best-suited to use in dialogues whose purpose is to elicit information. Dialogue game approaches figure in several influential contributions dating from the mid-late 1990s: Gordon’s *Pleadings Game*, [96]; Lodder’s study of legal justification (*Dialaw*) in [113]; and Loui’s use of dialectic approaches to non-monotonic reasoning in which one of the first considerations of computational limits is presented [118]. A currently active area in which these ideas have proven to be highly relevant is the exploitation of argumentation in multi-agent systems applications.

We conclude this discussion of dialogue processes by reviewing one further aspect that has been fruitfully adapted to argumentation in AI: analysing argument justifiability via dialogue games.

Interpretations of mathematical reasoning as a dialogue process have been advocated, from the early 1960s, in work of Lorenz and Lorenzen [115–117]. An approach that has provided a useful abstraction within the more general context of argumentation considers discussion over a disputed argument, p , as involving two participants conventionally denoted PRO (who argues in favour of p) and OPP (who objects to p). A generic dialogue game building on Dung’s argumentation model is presented in Jakobovits and Vermeir [105] (see also [104, Section IV]) and this approach has been used as the basis of a number of later studies, e.g. [30,51,52,70,71]. This view of argument justification as resolved by a dialogue implicitly underpins a number of formal ideas that have been adopted in algorithmic methods, e.g. the concepts of argumentation lines and the generalisation of such as argumentation proof-trees. Analysis of the properties of such structures has proven useful in examining a number of issues in argumentation and we discuss such approaches further in Section 3.1.

⁷ It should be noted that [121] appeared almost 30 years ago in 1978.

⁸ Indeed, independently Dunne, Doutre, and Bench-Capon [85] and Walton, himself, in [187] have analysed one dialogue form not presented in [189].

2.3. Diagrammatic treatments of argument structure

In the presentation above we have made a distinction between the concepts of *argument* and *argumentation*: the latter being understood as the processes by which given arguments are analysed and evaluated. An argument may, informally, be considered as the basic supporting case behind a given assertion. Thus an argument, “for p ”, in this sense, may itself give rise to a variety of distinct structures ranging in complexity from simple statements of fact (“ p is an accepted fact”), through to deductive templates (e.g. “ p follows from q and q is the case”) to even more intricate structures (or *argument schemes*) that themselves may rely on further (sub)-arguments. A variety of argument schemes have been proposed and studied, in work by legal theorists and philosophers, e.g. [145,186], however, the bulk of our preceding review has largely addressed the issue of argumentation and its evolving application within AI with the notion of “argument” itself only briefly discussed.

It has been seen that Dung’s fundamental model, as described in [73], abstracts away such internal structure from individual arguments in order to focus on the manner in which arguments interact via the defined attack relationship. In unfolding the exact nature of “the argument x attacks the argument y ”, however, the *reason* why such an attack is present needs to be considered in terms of those *structural schema* underlying the arguments x and y from which the attack arises. Such an interpretation, therefore, raises issues that concern the form an argument might take, i.e. issues regarding the components and representation of *arguments* rather than the process and outcome of the *argumentation* involved.

Questions regarding argument form and uniform treatments of these have assumed increasing importance in recent years, especially with respect to multi-agent exploitation of argumentation methods. A key aspect of this work has been the extent to which diagrammatic models of argument structure have been adopted.

Early diagram-based models of argument were intended to aid in illustrative hand construction and analysis of argument, with the resulting schemes being static depictions. An important example of such an approach is provided by *Wigmore diagrams* [190], which although used as a method of describing legal arguments have only recently been rediscovered and promoted in AI contexts, e.g. [41,164]. More widely known is the highly influential model of argument promoted by Toulmin in the 1950s [178]. Toulmin’s structural interpretation treats an argument as consisting of five sub-components: the Claim advanced (which could be qualified by a modal operator to describe concepts such as “normally”); the (factual, evidential etc.) Data supporting this Claim; a Warrant providing a licence to infer the Claim from the Data, together with Backing for this Warrant; and, to encapsulate exceptional cases, Rebuttal conditions. Although—in common with Wigmore’s scheme—Toulmin diagrams were originally presented as a static representation of the totality of an argument, they have proven a flexible approach in AI treatments of argumentation. Thus, Bench-Capon et al. [32] describe a dynamically evolving extension of Toulmin’s schema and its use in a dialogue game. Later work of Bench-Capon [26] develops the dialogue game of [32] providing a complete move repertoire and operational semantics for it.

The exploitation of such argument diagram techniques offers an important basis for a number of contemporary ideas among which are: argument visualisation methods, e.g. as might be used in decision support and explanation; argument construction from source material; the specification of methods for interchanging arguments between distinct parties; and in providing a unifying link between informal argument descriptions and formal abstract approaches such as [73]. Current work in these areas will be discussed in Section 3.4.

3. Recent trends and concerns

Section 2 has offered a, necessarily condensed, summary of influences on argumentation in AI covering up to the start of the present century. As we turn now to more recent developments, that is subsequent to those contributions discussed in [57,182] a number of trends becomes apparent: the continuing enrichment of the formal theory of argumentation building on [43,73,105]; the growth of argumentation-based methodologies in multi-agent systems applications; new computational treatments of argument diagramming and visualisation; the exploitation of argumentation in novel specialist domains; and the development of theoretical bases embracing subjectivity in argumentation and concepts of practical reasoning.

Overall one finds in such themes a broadening of the scope of argumentation in AI beyond its earlier traditional uses in realisations of non-classical logic scenarios. In this section, we discuss some of these themes in greater depth with particular reference to the articles contributing to this volume.

3.1. Development of the Dung-style model of argumentation

The graph-theoretic model of argumentation framework in [73] and the deductive schema supporting the assumption-based frameworks of [43] have given rise to an extensive body of research with particular concentration on the following,

- (a) Extension based semantics of argumentation.
- (b) Algorithmic and complexity issues in argumentation.
- (c) Dialogue processes for deciding acceptability.

Extension-based semantics of argumentation

Each of the extension-based semantics presented in [73] builds on sets of arguments, S , that are *conflict-free*: that is, no argument within S attacks another in S . Conflict-freeness, as observed by Baroni and Giacomin [22] in their study of evaluative criteria for extension semantics, is viewed as a minimal requirement to be satisfied within any computationally sensible notion of “collection of justified arguments”. Conflict-freeness, however, is too weak a condition, in itself, to be applied as a reasonable guarantor that a set of arguments S is “collectively acceptable”: for example, such a set could be attacked by arguments not among its members.

It is in the approaches posited to form additional conditions on (conflict-free) subsets of arguments that complications become apparent, and out of the divers methods proposed to resolve such complications that the current, to use Guillermo Simari’s phrase, “plethora of argumentation semantics”⁹ has emerged.

The three principal extension-based semantics introduced in [73]—the so-called *Grounded*, *Preferred* and *Stable* semantics—can exhibit a variety of problematic aspects.¹⁰

- (P1) *Emptiness*: although an extension satisfying the prescribed conditions always exists, there are AFs for which the only such extension is the empty set. This can arise with both the grounded and preferred semantics of [73].
- (P2) *Non-existence*: an extension, when it exists is never empty, but there are frameworks for which no extension meeting the required criteria exists. This can occur, for example, with Dung’s stable semantics.
- (P3) *Multiplicity*: in an AF there may be several “incompatible” extensions, i.e. sets S_1 and S_2 which are well-defined extensions of $\langle \mathcal{X}, \mathcal{A} \rangle$ but with $S_1 \cup S_2$ failing to be so. While Dung’s grounded semantics does not suffer from this problem, frameworks are easily constructed in which both the preferred and stable semantics exhibit this phenomenon.

A number of approaches have been proposed in order to address these and other perceived drawbacks. Thus, Cayrol and Lagasque-Schiex [55] define concepts of “graduality” in order to evaluate classes of acceptable arguments, Caminada [46] introduces *semi-stable semantics*; Dung, Mancarella and Toni [75, p. 151] develop *ideal* semantics and their paper in this volume [76] presents further analyses concerning the computation of ideal extensions in ABFs. Baroni et al. [21,23] define various extension-based semantics for an AF building from the strongly-connected component (SCC) decomposition of its directed graph: of the resulting SCC-recursive semantics, CF2-semantics have been examined in depth in [23]. In [62], Coste-Marquis, Devred and Marquis consider a refinement of the concept of “conflict-free set” in order to exclude “*controversial arguments*” [73, p. 332], i.e. arguments $\{x, y\}$ such that, although $\langle x, y \rangle \notin \mathcal{A}$ there is an “indirect attack” by x on y : the resulting approach gives rise to the *prudent* semantics of [62].

A number of extension-based semantics have been proposed motivated by new interpretations of the interactions between arguments that should be considered: thus the basic binary attack relation of [73] is developed. Important contributions of this type include the articles by Cayrol et al.—[50,53,54]—wherein the relation “the argument x supports the argument y ” is introduced leading to the formulation of *bipolar* argumentation frameworks. In such frameworks each of the existing extension semantics can be qualified through bipolarity, e.g. [50] considers bipolar prudent semantics. Other developments of Dung’s attack structure are offered in work of Nielsen and Parsons [134]

⁹ During the presentation of [129] at COMMA 2006, 12th September, 2006.

¹⁰ It should be noted that although our description is given in terms of AFs exactly the same issues arise in the analogous semantics within ABFs.

using an approach predicated on the idea that a *binary* attack relation is not always appropriate and thus this should be defined in terms of $\mathcal{A} \subseteq 2^{\mathcal{X}} \times \mathcal{X}$, i.e. each $x \in \mathcal{X}$ has an associated set of *subsets* of \mathcal{X} that attack it.¹¹

A different treatment of \mathcal{A} forms the basis of Amgoud and Cayrol’s approach in [7] to the interpretative issues created by the presence of multiple preferred extensions. In this any AF is augmented by a *preference* relation over arguments defining an attack in \mathcal{A} . In the resulting *preference-based* AF (PAF)¹² an attack $\langle x, y \rangle$ is relevant only if the argument y is *not* preferred to the argument x . By imposing suitable restrictions the effect of specifying preferences within a given AF is to reduce it to an *acyclic* graph: for such frameworks, [73] has shown that the grounded, preferred, and stable semantics coincide in a unique extension.

The *value-based* argumentation frameworks (VAF) of Bench-Capon [27,28] also stem from attempts to provide a formal basis on which to rationalise choices between several preferred extensions. The basic elements of VAFs are described in the paper by Dunne [80, Section 8]. In common with the preference-based approach, resolving choices in VAFs can be interpreted in terms of consistently removing attacks (using “value” orderings) so that the resulting framework is, again, one in which all three basic extension semantics coincide. The philosophical rationale underpinning VAFs ultimately derives from Perelman [145] and is treated in more depth in Section 3.3. A detailed comparative discussion of preference and value-based methods may be found in [30, Section 7.1].

The article by Baroni and Giacomin [22], makes a powerful case for re-examining the proliferation of new semantics:

“In fact, various kinds of motivations have been used to support the introduction of new semantics with respect to “classical” proposals ... These motivations range from the desire to formalise some high-level intuition, not captured by other proposals, to the need to achieve the “correct” treatment of a particular example (or family of examples) regarded as particularly significant. ...

Clearly, these kinds of heterogeneous intuitions hardly lend themselves to systematic comparisons. Given this situation it is not surprising that comparisons are quite often carried out using specific problematic examples, often ingeniously devised so as to bring to light patently different behaviours exhibited by the semantics under discussion.” ([22, Intro.])

Extension-based semantics in AFS continues to be an extremely active topic for argumentation models in AI and a number of specialised technical questions remain unresolved.¹³ Important as such questions are, it may well be the case, however, that treatments of extension-based semantics will come to focus less on the construction of novel specialised forms and more on consolidation theories such as the evaluative principles of [22] or the complementary approach applied to arguments with a particular structure of [47,48]: just as the attempts to construct a notional “definitive” non-monotonic logic from the disparate alternatives proposed in the 1980s are now recognised as ill-fated, such is likely to be the outcome of efforts to build an “ultimate” extension-based semantics.

Algorithmic and complexity issues

While the preponderance of formal theoretical study into computational issues arising from [43,73] has addressed semantic concerns within these abstract frameworks, there is a significant core of results relating to algorithmics and computational complexity.

Early work of Dimopoulos and Torres [69] derived exact complexity classifications for a number of decision problems involving extension-based semantics in AFS¹⁴ and a summary of these results may be found in [80, Table 1(a–d)]. In [81], Dunne and Bench-Capon further develop complexity-theoretic analysis of Dung’s model in deriving exact bounds on the computational complexity of two questions (neither of which is considered in [69]): that of deciding if a given argument is justified under the most restrictive semantics defined in [73] (so-called *sceptical acceptance*);

¹¹ For details see the paper by Nielsen and Parsons [135] in this volume.

¹² These should not be confused with the *partial* argumentation frameworks (also denoted PAF) described in the article by Coste-Marquis et al. [64] in this volume.

¹³ Among which are issues such as conditions under which particular extension-based semantics coincide, existence properties etc.

¹⁴ The analyses of [69] are not presented in the context of Dung’s frameworks from [73] but may be readily translated into this. A discussion of the links between [69] and [73] may be found in Dunne and Bench-Capon [81, pp. 188–189].

and the problem of determining whether an AF satisfies Dung’s concept of *coherence* [73, Definition 31(1), p. 332] whereby every set of arguments defining a preferred extension also defines a stable extension.

Both [69] and [81] relate to decision problems in AFs: the basis for hardness results is via suitable directed graph mechanisms, cf. the two core constructions described in [80, Secn. 3]. In an important series of articles—[66–68]—Dimopoulos, Nebel and Toni consider analogous questions within various instantiations of assumption-based frameworks. A significant achievement of this work is in characterising the computational complexity of decision questions in ABFs with respect to that of testing “*derivability*” (i.e. of φ from a given base Δ) within the associated logic modelled by the ABF: derivability being central in determining the existence of attacks between arguments. In consequence a number of instantiations of ABFs describing particular non-classical logics exhibit a significant increase in complexity compared with the NP or Π_2^P -completeness of related problems in AFs.

The contributions of [66–69,81] primarily focus on purely complexity-theoretic analyses. Efficient algorithmic methods are introduced in [73] for special classes of AF (directed acyclic graphs—DAGs) and, more recently, in work of Coste-Marquis et al. [63] (for *symmetric* frameworks). The extent to which such graph-theoretic conditions can ameliorate complexity issues forms the central topic of Dunne’s paper [80] in this volume. The treatment of “graduality” in [55] includes a number of algorithmic elements and other useful work has emerged from modelling of argument justification via dialogue games, e.g. [51,52,184].

One collection of methods which have received increasing attention over the past five years concern enumerative techniques for constructing all extensions (of a particular form) within a given AF, key contributions being the work of Doutre and Mengin [70] and Verheij’s labelling approach to generating all *stable* extensions of an AF described in [181]. More recent work includes algorithms of Nielsen and Parsons [134] and Vreeswijk [183], the first of these relative to the authors’ set-theoretic notion of attack mentioned earlier. An alternative slant on the question of enumerating preferred sets is offered in [77] which considers the following issue: under the assumption that an enumeration, \mathcal{S} , has already been produced, to what extent can it be represented compactly with the representation allowing $S \in ?\mathcal{S}$ to be decided efficiently for any subset of arguments of \mathcal{S} ?¹⁵ While [77] presents indications that a number of computational questions remain difficult even with significant additional information provided, more positive results indicate that concise encodings of preferred sets which *can* be efficiently queried *may* be possible.¹⁶

The algorithmic analysis of ABFs has been rather less advanced than that within AFs: one major factor accounting for this is, of course, the formidable complexity-theoretic issues raised in [68]. Nevertheless, promising dialogue based techniques are presented in work of Dung, Kowalski and Toni [74]. The paper by Dung et al. in this volume [76] offers a further example of dialogue approaches by adapting these to the computation of ideal extensions in ABFs. In addition, recent work of Egly and Woltran [88] is of some interest: this proposes approaches building on translations to quantified Boolean formulae and the subsequent exploitation of highly-tuned QBF solvers to resolve decision questions.¹⁷

The article by Coste-Marquis et al. in this volume [64] introduces an important topic that appears to have been largely neglected in previous studies: given a number of distinct AFs, describing, say the views of a number of observers regarding a specific issue, how should these be merged into a “sensible” unified framework that “fairly” reflects individual viewpoints? The techniques in [64] contribute both to semantic and algorithmic aspects of this question.

For the developments of [73] represented by PAFs and VAFs, only the latter gives rise to non-trivial algorithmic and complexity issues. Treatments of these, including exact complexity classifications of the principal decision questions together with algorithmic approaches may be found in the series of papers by Dunne et al. [30,71,83,84]; Dunne’s paper in this volume indicates that a number of non-trivial issues remain to be resolved in the algorithmic treatment of VAFs.

Dialogue-based approaches to deciding argument acceptability

The view of reasoning as a dialogue mechanism has been widely adopted in formal algorithmic approaches to determining the acceptance status of arguments within both AF and ABF models. Such a view has also featured

¹⁵ While the AF itself provides a compact encoding of its preferred sets, in view of [69] it is unlikely that this would satisfy the “efficient querying” criterion.

¹⁶ Of course, *finding* such encodings given an AF is another matter.

¹⁷ Treatments of non-classical logic via propositional encoding had also been proposed in earlier work of Ben-Eliyahu and Dechter [25]; similar techniques, for AFs, are discussed in [81, p. 202].

significantly in models of argument methods building on deductive reasoning templates. Key ideas underpinning these techniques include *argument line* (a chain $\langle x_0, x_1, \dots, x_k \rangle$ in which the argument x_i attacks the argument x_{i-1} for $i > 0$) and the concept of (partial) *proof tree*, which, in informal terms, can be interpreted as combining a number of distinct argument lines concerning a common initial argument. In pursuing such approaches a number of basic questions arise: rules and strategies affecting the selection of arguments with which to continue a dialogue; termination properties; demonstrating soundness and completeness of procedures intended to establish acceptability of arguments in particular semantics; approaches to assessing the “efficiency” of dialogue methods, and so on.

The generic formalism for describing dialogue games within Dung’s model of argument introduced in [105] was discussed, briefly, earlier. A significant subsequent development is found in the methods presented by Vreeswijk and Prakken in [184]. This describes the structure of *Two-part Immediate Response* (TPI) disputes. Adopting the two player PRO and OPP convention for debate over an argument x , among the features of TPI-disputes is the requirement for each player to attack the most recently played argument of their opponent whenever it is *possible* (within the game’s rules) to do so. Several examples in [184] establish that both players require moves allowing back-tracking to a defined earlier point in a dialogue. The resulting game is shown to be sound and complete for so-called *credulous reasoning*, i.e. where the aim is to decide if x is a member of *at least one* preferred extension. Thus for any AF and argument x within it, TPI-disputes are guaranteed to terminate and correctly to determine whether x is justified under Dung’s credulous preferred semantics (PRO wins) or x cannot be so justified (OPP wins). A variant of this game provides sound and complete methods for *sceptical reasoning* in *coherent* AFS i.e. where PRO wins if and only if x belongs to *every* preferred extension.

Vreeswijk and Prakken’s results in [184] were instrumental in motivating one of the first systematic studies concerning formal concepts of “efficiency” of dialogue games in argumentation: the definition and analysis of *dispute complexity* presented by Dunne and Bench-Capon in [82]. Informally, the dispute complexity of a dialogue game is measured in terms of the (worst-case) number of moves that might be required in order to resolve the status of a given argument in an AF. One significant contribution of [82] is its positioning of such dialogue games within an established body of work regarding the relative efficiency of propositional proof methods via the concept of dispute complexity, i.e. the basis provided in Cook and Reckhow [61]. Thus, [82] not only demonstrates that TPI-disputes occasion a propositional proof method but also, adopting the comparative criteria for such systems presented in [61], further show that the resulting system is equivalent to the CUT-free sequent calculus of Gentzen [92]. In consequence, via results of Urquhart [180], one may construct (a family of) AFS and arguments within these— $\langle \text{AF}_n, \varphi \rangle$ such that resolving the status of φ requires exponentially long TPI-disputes.

Important treatments combining elements of MacKenzie’s dialogue model [121] with the formal approach of [105]—for example, locutions, utterances, rules for dialogue continuation, termination—are found in McBurney, Parsons, and others, e.g. [106,122–124,141,142]. Much of the emphasis of this work is directed towards providing a basis for dialogue exploitation in multi-agent system contexts—for example Torroni’s analysis of termination properties in negotiation dialogues [177]—as discussed in Section 3.2.

Treatments of proof-theoretic techniques via dialogue methods using the concepts of argumentation line and partial proof-tree have been considered in a number of recent papers. An important issue in this context concerns the design of heuristics that reduce the search space thereby obviating the requirement to consider all expansions of each argument line. Dung, Kowalski and Toni [74] propose a novel “backward reasoning” approach to the construction of proof-trees in ABFS. Recent work of Chesnevar and Simari [58] deals with sceptical argumentation via a lattice-theoretic encoding of the relevant search space. A related question in implementing dialogue mechanisms is that of deciding which (from a range of available options) is the “best” continuation for a participant to contribute. There are, of course, many interpretations of “best” that may be applicable from loosely defined intuitive qualitative notions (e.g. most “persuasive” or “convincing”) to quantitative ideas, e.g. guaranteed to terminate debate in the fewest possible moves. In [86,87] Dunne and McBurney consider one formalisation of this problem that allows it to be related to the “literal selection” problem examined by Liberatore [112].

A final collection of issues, concerning which computational treatment has only recently been initiated, addresses questions arising from wider considerations of the motives of participants. Thus, recognising that contributors to a discussion may have rational bases to obstruct its development or be anxious to avoid revealing information regarding their pursuit of an issue. In [91], Gabbay and Woods examine the use of so-called “stone-walling” tactics as one means of impeding the progress of information-seeking dialogues, while Dunne [78] considers settings in which one participant seeks to prolong discussion and reviews such approaches against a variety of legal applications. Informa-

tion hiding strategies are examined in a number of recent papers relating to multi-agent systems, e.g. Otterloo [139], Paruchiri et. al [144]. A very preliminary study of computational elements relating to the concept of “hidden agenda”¹⁸ is initiated in [79].

3.2. Exploiting argumentation techniques in multi-agent systems

If the significant interests moving argumentation into AI during the 1980s and early 1990s were driven by its application to non-classical logics, within the last 10 years argumentation technology has increasingly been widely adopted in driving the development of another computational field, also of importance in AI: the paradigm of autonomous agent computing, e.g. as described in Sycara [176], Wooldridge [192].

It is not difficult to account for this interest, the basis for which lies in the key paper of Sycara [175] in which a system for multi-agent negotiation—PERSUADER—is described: negotiation in [175] being presented and treated as a persuasive argumentation process. Automated negotiation mechanisms have long formed a central concern of work in multi-agent systems and the article by Ramchurn et al. [159] in this volume marks a further development of this field. We refer the reader to [159, Section 7] for a more detailed discussion of subsequent work building on [175] in particular the important contributions of Kraus, Sycara and Evenchik [110] and Parsons, Sierra and Jennings [140].

A number of important themes have emerged from such treatments of inter-agent negotiation as an argumentation driven persuasive dialogue: the rationalisation of individual agent contributions as stages in a goal-directed plan; the study of logic-based language formalisms in terms of both syntactic (e.g. the manner in which agents represent contributions to debate, proposals, goals they seek to bring about, etc.) and semantic (e.g. how an agent’s perspectives are affected by particular contributions as negotiation progresses) aspects; the development and analysis of formal agent oriented dialogue games; the consideration of comparative criteria for differentiating and classifying dialogue mechanisms, etc. As should be evident from the brief survey presented in Section 2.2 such themes interact heavily with the core body of philosophical and rhetorical theories of dialogue, most notably in the contributions of MacKenzie [121] and Walton and Krabbe [189]. Treatments of such topics forms the object of study in work of Parsons et al. [141–143], Amgoud et al. [8], McBurney et al. [125,126]. A collection of recent articles on the theme of inter-agent communication languages may be found in Dignum [65].

One further feature of argumentation has been influential in encouraging its adoption as an enabling technology for multi-agent systems developments. A central idea advanced in agent-based approaches is that of *autonomy*: agents act as individual entities, often, but not always, attempting to cooperate and coordinate with others. In such environments, however, the actions that an individual agent wishes to perform may conflict with the actions attempted by other agents, e.g. in seeking access to particular limited resources. An agent’s understanding and knowledge of its environment (and of its awareness of the perspective of other agents) is likely to be incomplete and uncertain, thereby subject to continual revision. These elements of incomplete knowledge, uncertain information, and the potential for conclusions initially formed subsequently to be rejected are, as we have seen, fundamental to the nature of argumentation. In sum multi-agent applications provide a natural arena for formalisms building on the study of argumentation directed towards justifying *actions* (as opposed to beliefs) as the mechanism by which an agent seeks to bring about particular desired *goals*—that is, the study of so-called *practical* reasoning which we review in the next section. Argumentation as an approach for autonomous agents to make decisions is discussed in work of Kakas and Moraitis [109] and the paper by Oren et al. [137] in this volume deals with important aspects regarding reaching agreement in the presence of uncertain (and potentially unreliable) evidence. Observing that individual agents seeking to come to a shared understanding about aspects of their environment require methods by which conflicting evidence concerning this may be assessed and resolved, [137] discuss those factors which must be considered and present a solution approach whose logical framework builds on Josang’s Subjective Logic [107].

3.3. Practical reasoning

Arguments are often thought of a set of reasons for a claim, and the claim is typically thought of as being propositional, that such and such is the case. Many arguments are, however, not about whether some belief is true, but

¹⁸ The term “hidden agenda” is attributed to Barsky [24] by Silverman [172].

about whether some action should or should not be performed. Many dialogue types such as persuasion, deliberation and negotiation can concern what should be *done* in a given situation, rather than what is *true*. Reasoning about what should be done has been seen as a specific topic in Philosophy since the time of Aristotle, and is there termed *practical reasoning*. For a collection of philosophical essays on the topic, see [160].

Practical reasoning has a number of important differences from theoretical reasoning. There may be many ways to achieve a given goal, and so sufficiency of an action to achieve a goal is not enough—it needs to be the best way. Side effects—both beneficial and harmful—need to be considered. Also—and this is an important distinction between practical reasoning and traditional planning systems—it cannot simply be taken as given that the goal itself is worth achieving, or that it should be pursued at the expense of other goals which might be adopted instead. Practical reasoning involves the *selection* of goals as well as their realisation.

Important too is the notion of *direction of fit* [171]. In theoretical reasoning agents are trying to make their beliefs fit the world and, since reality is thought to be the same for all, perfectly rational agents with complete information should be able to come to agreement. When there is disagreement, one agent will be right and the other will be wrong. In contrast, practical reasoning attempts to make the world fit what the agent wants it to be, the point of actions being to change the state of the world so that it is some respects more acceptable to the agent. Different agents may quite properly have different interests and aspirations, and so, even if perfectly rational and in possession of complete information, they may disagree as to what it is best to do. In practical reasoning there is no implication that one is right and one is wrong: each may be right, according to their own perspective.

Representation of such individual perspectives has been built on the notion of “audience” introduced by Perelman [145]. Perelman stressed that the purpose of an argument is to persuade, and whether an argument succeeds in persuading is a function not only of the argument itself but also of the audience to which it is being addressed. It is those hearing the argument that are important in this respect, not the speaker. For example, an argument to cut top rates of tax because it will enable high earners to keep more of their income might persuade captains of industry while leaving their work force entirely unmoved. To convince the latter an argument in terms of increased general prosperity in which all will share must be made. Perelman’s notion of audience has been used in AI in work such as [98] and [27]. Although it is particularly important in the context of practical reasoning, Perelman would argue that the notion of audience is important for all arguments, theoretical and practical. In the context of theoretical reasoning, different assumptions and reasoning capabilities may affect the acceptance of an argument. For an exploration of the notion of audiences with respect to theoretical reasoning, see [102] and [103].

Issues relating to practical reasoning have received far less attention than theoretical reasoning, but have become increasingly important as autonomous agents grow in popularity. The autonomy of agents requires them to be able to select, at least to some extent, the goals they will pursue: it is not enough that they find ways to satisfy goals given from outside, or built in at design time. One approach, advocated by Amgoud and Rahwan [157], uses a set of beliefs together with a base of desire rules to produce arguments for which desires to adopt, and then a set of planning rules to determine how to fulfil these desires given the beliefs. Huljstein and van der Torre [101] have also provided a framework for combining goal generation and planning. An alternative approach, represented by the paper by Atkinson and Bench-Capon in this volume [19], uses an ordered set of values to represent the interests and the aspirations of an agent, differences in agents being represented by differences in the value ordering. This value ordering is found to produce arguments by instantiating an argument scheme for practical reasoning [20], goals being selected by reason of their enhancement of some favoured value.

Autonomous agents, by their very nature need to select actions, and perhaps to justify them to others and to persuade and negotiate joint actions with other agents. Of course, theoretical reasoning is an important part of this, since an accurate understanding of the state of the world and the effects of their actions is essential. Practical reasoning is, however, also vitally important if this knowledge of the world is to be put to good use.

3.4. Informal logic: Argument diagrams and schemes

The arguments we typically encounter, both in the everyday contexts of newspaper editorials, and in scholarly contexts such as philosophical writing, are presented in a rather discursive manner. Examples are accumulated, principles stated, objections raised and countered, and claims made. Understanding such arguments is often helped by an analysis which attempts to establish the components of the argument and the relations between them. Once the claim has been precisely formulated, and the various supporting material identified it becomes much easier to see exactly

what the argument is intended to establish, and whether it does so successfully. Typically implicit assumptions will be identified, and often subtle shifts of meaning will be uncovered. The analysis of arguments in this way has long been the stock in trade of philosophers and critics, but over the last twenty-five to thirty years the study of methods to reformulate arguments has developed into the discipline of informal logic, given considerable impetus by the increasing importance given to the teaching of critical thinking, especially in the United States. There are now many textbooks devoted to this topic, of which [191] and [168] may serve as examples.

Although many of the analyses of arguments are expressed in natural language, there is also a tradition of using diagrams to explicate the relations between the components of the arguments. Early examples are Wigmore [190], who used diagrams to represent the elements of legal cases, and Toulmin [178] who made use of a simple diagrammatic structure to promote thinking about arguments critically.

Whether analysed using text or diagrams, it is a striking feature that the supporting statements produced from the analysis often do not seem to entail the claim. On one view, *deductivism*, such arguments are claimed to be elliptical, with some premises that are hidden. This view leads its exponents to reconstruct such arguments, by treating the assumptions as implicit premises which when added yield a deductive argument. This view underlies a good deal of thinking about argument in AI: it is perhaps the standard approach to define an argument as a sequence of inference steps, either in a classical or a non-monotonic logic (see e.g. [47] for a recent definition of argument in these terms). In informal logic, however, deductivism is not universally accepted, on the grounds that such a restrictive model fails to do justice to the richness and variety of arguments found in natural language. As will be discussed below, the view that not all arguments can be reconstructed as deductive arguments has also received attention in AI.

An important tool in early informal logic was the study of *fallacies* as discussed in Section 2.2—patterns of poor reasoning which somehow mimic good reasoning—and a number of patterns of fallacious reasoning were identified. There is, however, another side to this: some of these patterns do indeed seem to present instances of good reasoning. The stereotypical patterns of good reasoning are often termed *argument schemes*. Consider “if p then q , and q , so p ”. At one level this is simply fallacious, the fallacy of affirming the antecedent. Often, however, this pattern is used to express a perfectly good argument based on “inference to the best explanation”. Of course, to be a good argument, not only does the antecedent have to be an explanation of the consequent, but all other potential explanations of the consequent have to be considered, and this particular antecedent chosen as the best explanation.

The deductivist approach would be to enumerate all explanations, and demonstrate the falsity of all explanations other than the favoured one.¹⁹ It is, however, often impractical to enumerate all the explanations, and even more so to eliminate all the alternative explanations. Instead it is possible to give a more procedural account in terms of *argument schemes*, e.g. [186]. Here the instantiation of an argument scheme such as “inference to the best explanation” gives a *presumptive* justification for the conclusion. The argument scheme is, however, subject to a number of critical questions characteristic of the scheme, such as, for inference to the best explanation, *is there an alternative explanation?*, which challenge this presumption. On this view, the presumptive conclusion will stand unless some alternative explanation is produced. Should such an alternative explanation be produced, the presumption can be defended by showing it to be false, or at least inferior to the original explanation. This has great appeal for modelling reasoning against a background of incomplete information and limited resources, both of which are common in AI. The conclusion is justified not by demonstration, but by withstanding the appropriate critical procedure in the particular circumstances. This then leads to the consideration of what argument schemes should be accepted, and what critical questions are associated with them. Ref. [186] lists more than twenty schemes, but the list is extensible. Another set of argument schemes can be found in [145].

In recent years, there has been important cross fertilisation between informal logic and computer science. The practise of argument diagramming is an area which can be effectively supported with computer tools, providing computer support, and argument diagrams offer an attractive way of presenting the fruits of automatic reasoning. During the 1990s, one main focus was on Toulmin’s diagrammatic argument scheme: [128] and [120] provide two early examples. Another very influential tool, not based on Toulmin, was gIBIS [60]. A more general tool in current use is Araucaria [162]. Araucaria gives a general way of building up argument diagrams, includes facilities for representing argument schemes, and allows for the translation between the original diagrams of Araucaria and diagrams in the

¹⁹ It is often said that the great fictional detective Sherlock Homes is wrongly called the “Master of Deduction” since he explains his method as “when you have eliminated the impossible, whatever remains, however improbable, must be the truth”, which is abduction. It is, however, a deductivist interpretation of abduction.

forms introduced by Toulmin and Wigmore. While use of Araucaria has tended to focus on specific arguments, other tools, such as Claimaker and Compendium developed at the Knowledge Media Institute of the Open University were intended to give an overview of an entire body of thought across a number of authors, such as the Turing machine intelligence debate (e.g. [111,179]).

The advance of the World Wide Web also made the vision of an on-line resource of modelled arguments being available to all, and of this resource being made extensible in the manner of Wikipedia. For an example of the traditional textual modelling of the arguments of Aristotle and Plato see the Archelagos Project of Scaltsas [11]; with respect to diagrams, many of the other tools mentioned above are also available to be used over the Web. As this vision spreads there has been an attempt to bring together various strands of this work and to produce a standard for representing and exchanging arguments. This has resulted in the Argument Interchange Format (AIF) [56]. A development of AIF designed to enable the provision of a large corpus of represented arguments on-line is the topic of the paper by Rahwan et al. [158] in this volume.

There has been an increasing interchange between informal logicians and argumentation people in computer science, and a number of interdisciplinary events have been organised, fostering collaboration, e.g. [161]. From these collaborations interest in argumentation schemes has flourished. The notion of presumptive arguments as the instantiation of argument schemes subject to critical questioning has particularly influenced two of the papers in this volume: Gordon et al. [97], emphasise the procedural nature of justification using a generalised argument scheme while Atkinson and Bench-Capon [19] offer a detailed exploration of a particular argument scheme and its critical questions.

3.5. Specialist domains and applications

There are a number of particular disciplines in which argumentation is central, and which have particular styles of argument associated with them. One such area which has been the subject of a good deal of argumentation related AI research is Law (see [166] for a survey of AI and Law research, and for overviews of argumentation in AI and Law see [33,34]). Argumentation is particularly important in law: a legal case typically centres on a conflict between two parties which is resolved by each side producing arguments in an effort to persuade the judge that their side is right. The judge then decides which party to favour, and publishes a decision in which he argues why his decision is justified. Modelling legal reasoning can then be seen, to a large extent, in terms of modelling argument, and so it is unsurprising that attempts to understand legal argumentation have been a key strand in AI and Law.

One of the earliest AI and Law projects was the Taxman project of Thorne McCarty ([127] provides a good summary). The idea there was to model the argument of the majority and minority opinions of the Supreme Court decision in a famous tax case, *Eisner vs Macomber*, which had centred on whether a particular share issue was income or not. That case is quite representative of legal argument: there are things to be said on both sides, and the verdict was in doubt and—as the existence of a minority decision testifies—it is even possible to justify finding for either party.

In the 1980s argumentation in law was pursued by the influential HYPO project of Rissland and Ashley [14], which modelled arguing with cases in the field of US Trade Secrets Law. Important features of this system were the modelling of adversarial argument as a three ply structure, where a case is first cited, then contested by the other party, before finally the original party is allowed a chance to rebut, and the modelling of facts as *dimensions* so as to allow for hypothetical cases to be created by strengthening or weakening cases in particular respects. This work broke into two strands, CABERET [174], which identified a set of argument moves and strategies, and CATO [3], which was designed to teach the HYPO style of reasoning with cases to law students. Another important paper in this tradition was [119], which extended the notion of argument moves to include the restatement of positions to uncover the rationale underlying the position, the better to identify the point at which to attack it.

This work, predominately from the US, concentrated on cases and arguments based on precedent cases. In Europe the focus had rather been on the representation of law, particularly statute law, as rules. A key problem with such representations was that the rules were often conflicting, and of uncertain application because they contained open textured terms. Although argumentation had long been suggested as an approach to resolving open texture [35], it was not until [73] offered a way of approaching conflict, defeasibility and reinstatement in terms of argumentation frameworks that argument became accepted as central to rule based approaches. Prakken, e.g. [153], was particularly influential in promoting the idea that the conflicting arguments generated from the rule based representation should be evaluated by being organised into a Dung-style argumentation framework in order to resolve conflict. Once argumentation had become appreciated by the rule based exponents of AI and Law, the divide between rule based and case

based argumentation narrowed, and the integration of the approaches was further facilitated by work which modelled HYPO style reasoning in rule based terms, such as [156].

Another important theme relevant to argumentation in AI and Law relates to modelling legal reasoning as a *dialogue game*. This approach was introduced into AI and Law by Gordon's *Pleadings Game* [96], which modelled the pre-trial process of pleading, designed to identify which aspects of a case were agreed and which were disputed. The key element here was that dialogue was used to model the *process* of a legal dispute, appealing to the notion of procedural justice, whereby a decision derives its validity from being the output of a properly conducted procedure. There followed a number of legally directed dialogue games, such as [26,113,114,181]. The emphasis on argumentation as a dialectical process is reflected in the papers by Gordon et al. [97] and Artikis et al. [13] in this volume. Another important domain in which argumentation has been applied is that of medicine and one early demonstration of its effectiveness in this context is provided in work of Fox et al. [90]. The paper by Mozina et al. [132] in this volume describes the use of argumentation to enhance a Machine Learning technique also relates to medicine, but an earlier application applied the same technique in the legal domain [133].

This brief discussion of argumentation in AI and Law shows it to be good example of how developments in argumentation techniques, such as case based techniques, dialogue and the argumentation frameworks of Dung, have been absorbed into a particular area of application, and developed within it, driven by particular problems and needs arising from that domain. These developments have then fed back into argumentation generally. Although none of the papers in this volume are specifically directed to law, several of the authors are stalwarts of that field (Sergot, Bench-Capon, Gordon and Prakken have all been Programme Chairs of the biennial International Conference on AI and Law, and many of the other authors have also published in AI and Law).

3.6. Other important trends

The themes discussed over the preceding sections reflect the principal trends of interest to the articles contributing to this volume. It would be inappropriate, however, to regard these as defining the entire scope of argumentation in AI. For this reason we now very briefly outline a select number of areas which, although not explicitly considered in subsequent articles, merit some discussion. Considerations of space prevent a more detailed analysis of these fields and interested readers are directed to the references indicated and their associated bibliographies for more extensive treatments.

A formal model of argumentation with some features in common with both AF and ABF methods is found in *Deductive Argument Frameworks*. These model an argument for a claim p as a pair $\langle S, p \rangle$ wherein S is a collection of (propositional) formulae (called the *support*) drawn from a knowledge-base Δ and whose collective acceptance (logically) entails the *conclusion* p . The usual concept of “attack” that is adopted is that of the conclusion of one argument being inconsistent with the support for another. Treatments of argumentation using this model are provided in a series of papers by Besnard and Hunter [39,40,102,103]; the relationship between deductive frameworks and AFS is examined in a recent paper of Wooldridge et al. [193]. A number of current issues within this model are the subject of Hunter's *Argumentation Factory* project [15].

The fact that information used in argumentation is often uncertain suggests probability theory as a natural analytic approach to adopt. In consequence a number of models of *probabilistic argumentation* have been put forward. The series of papers by Benferhat, Dubois, and Prade [36–38] and the, more recent, analyses of Amgoud and Prade [6,9, 10] provide excellent exemplars of such methods. Other related work builds on concepts of *possibilistic logic* and has been developed in recent papers of Alsinet et al. [1,2].

Treatments of *natural language* and argumentation date back to the foundational work of Birnbaum et al. [42] and Alvarado and Dyer [4,5]. An important recent trend in natural language studies has been in the direction of *generating* natural language arguments rather than the interpretative analysis that underpins earlier work. The articles by Green and Carberry [99], Walker et al. [185] and Carenini and Moore [49] provide a good overview of these recent trends in natural language approaches to argumentation.

We conclude this short summary by noting one area of activity which is only just beginning to be developed in the study of argumentation in AI. In our analysis of Dung's work [73], one aspect of this which was not considered is—as emphasised in the title of [73]—its relationship with classical game-theory as pioneered by von Neumann and Morgenstern [136]. The question of how extension-based semantics in AFS relate to classical concepts of solutions within n -player games, e.g. as described in Osborne and Rubinstein [138]—has been reviewed in work of Rombouts [163].

Of particular interest is the extent to which AI study of argumentation can benefit from (and, indeed, contribute to) the well-established theory of debate in game-theoretic economic models. A typical and accessible treatment of argumentation in such models is provided in Glazer and Rubinstein [94].

4. Summary

As we have seen, argumentation covers a wide range of approaches and concerns, and has drawn on influences from a number of sources. In this volume we have selected from the substantial number of good quality submissions a set of papers intended to reflect this diversity both of approach and concern.

Several of the papers build on abstract argumentation frameworks as introduced by Dung. Central to these argumentation based frameworks, are extension based semantics, but a number of competing semantics have been proposed. The paper by Dung et al. considers one such alternative—the ideal semantics—an underlying motivation of which is to provide an extension-based form that relaxes the extreme sceptical requirements of Dung’s grounded semantics while being less credulous than admissibility semantics. The paper by Baroni and Giacomin attempts to bring some principle and order to the treatment of extension-based semantics by proposing a number of criteria by which such semantics can be evaluated, and applying these criteria to a representative range of existing proposals. This systematic exploration of the strengths and weaknesses of the various different semantics provides an excellent overview of what has been proposed and the issues at stake. The need to identify extensions of various types in argumentation frameworks in order to evaluate the status of the arguments within them gives rise to a number of decision problems, several of which have been shown to be computationally intractable, although efficient decision procedures may exist if certain restrictions are imposed. Argumentation frameworks are naturally represented as graphs, and the paper by Dunne examines the complexity properties of argument frameworks which are subject to certain graph theoretic restrictions. These restrictions are shown to have positive results for some decision problems, although others remain intractable. Most work on argumentation frameworks has investigated a single structure. In many applications, however, argumentation frameworks may be developed independently. If several such frameworks are developed, there may be arguments that are included in some but not all, and there may be disagreement with respect to the attack relation. In order for the different perspectives to be pooled, the various frameworks need to be merged, and the differences between them reconciled to achieve some kind of consensus. The paper by Coste-Marquis et al. investigates issues relating the merging of Dung-style frameworks. They show that simple voting is not adequate, and offer a general framework in which the frameworks can be merged in a principled manner. The problem of extracting a consensus from a group of agents with different perspectives is also the topic of the paper by Nielsen and Parsons. In their treatment, however, each agent is equipped with a Bayesian network, and their approach is to provide an open framework in which the agents can use argumentation to arrive at an agreed network. In their framework the agents can, in a distributed fashion, explore the consequences of various compromises, and so judge which are acceptable.

Dialogues, especially dialogues between agents, are an important way of exploiting argumentation in systems. In order for agents to engage in dialogue there must be some protocol which both parties will follow in order to make sense of the exchanges. The paper by Artikis et al. gives an example of this work by providing a specification in the action language $\mathcal{C}+$ [93] of a protocol based on a formal procedure for dispute resolution. This takes into account the physical capabilities of the agents concerned, the rules of the protocol itself, and importantly the normative consequences of these rules for the agents, the sanctions they incur for non-compliance, and the enforcement of these sanctions. In the context of agents, negotiation has been an important area in which argumentation techniques have been applied. In a multi-agent system, negotiation may have social aspects, in that a negotiation may often be part of a series of encounters. In addition to presenting a framework for argumentation based negotiation, the paper by Ramchurn et al. investigates the effect on negotiation of the opportunity to make promises offering rewards in future negotiations in return for concessions in the current negotiation. Their theoretical approach is reinforced by an empirical study which shows that this can improve both the efficiency of the negotiation and the utility of the deals. Another important feature of agents is that they take in information from their environment and need to form and adjust their beliefs on the basis of this evidence. Reasoning with evidence is the subject of the paper by Oren et al. who provide a framework for arguing about evidence based on Subjective Logic which allows for important factors such as accrual of evidence and burden of proof to be taken into account. The work is set in a dialogical context in which different agents may have different utilities associated with particular facts being accepted.

An important import into AI argumentation from informal logic is the notion of argument schemes. Oren et al. make use of several schemes appropriate to evidential reasoning. The paper by Atkinson and Bench-Capon provides an in-depth exploration of a single argument scheme, designed to justify choices about what should be done in particular situations. Of prime importance here is the notion of critical questions, the means by which the presumption given by the instantiation of the scheme can be challenged, and the mechanisms by which the subjective aspects of choice can be captured by taking into account the individual interests and aspirations of the agent making the choice. The notion of arguments schemes is also central to the paper of Gordon et al. which provides a general framework in which argument schemes and their associated critical questions can be represented. The framework emphasises the procedural nature of argumentation, the burdens of presentation and persuasion placed on those engaged in argumentation and the standards of proof required to discharge them, all of which are required to ascertain the dialectical status of the claims being advanced.

Also drawing inspiration from work on informal logic, the paper by Rahwan et al. describes a vision of a world wide web of argumentation, which can act as an extensible repository of represented arguments. Based on a proposed standard for argument interchange, they describe an open platform for representing arguments, and for building interlinked and dynamic argument networks to form a publicly available resource.

The final paper, by Mozina et al., presents an interesting example of how argumentation can be used to give a novel take on a traditional problem. They show how argumentation can be applied in Machine Learning by using arguments from an expert to guide the learning of concepts using an adapted rule induction algorithm. Argumentation is able both to improve efficiency by focusing search, and to improve the quality of the rules induced by making them closer to the terms in which they would be expressed by the expert.

Computational treatments of argumentation with reference to AI are now the subject of a number of well-established meetings, e.g. ArgMAS within the workshop programme of AAMAS, Computational Models of Natural Argument (CMNA) alternating between ECAI and IJCAI; a biennial conference (COMMA) [59] was inaugurated in 2006 and is intended to complement these workshops. Also of interest are the number of large-scale argumentation related research projects which are currently in progress. Among such are ASPIC [16] addressing issues arising in the provision of argumentation services and ARGUGRID [12] aiming to exploit argumentation technology as a foundation for Semantic Grid applications.

The contributions to this volume show the range of work currently being produced in argumentation in AI and we hope that these will encourage specialists from AI fields where such ideas have yet fully to be exploited to consider what argumentation techniques may have to offer them. From the diversity of contributions and their background as outlined within this introduction, a number of points are evident: that argumentation in AI now informs the development of many, historically core, AI topics; and that the computational treatment of argumentation has evolved from the abstract models pioneered in [43,73], through methodologies offering effective realisations of argumentation techniques, to practical implementations directed at concrete applications. Most significantly, however, that the body of theory, techniques, and applications we have discussed is very far from encompassing a final, definitive description of the scope and limits of what argumentation-based approaches can offer to the furtherance of AI as a scientific discipline: many questions remain unresolved, many avenues unexplored, and many applications offer a wealth of possibilities for future work. When people participate in reasoned debate they are engaging in argumentation not demonstration. Thus argumentation, rather than logical demonstration, should be seen as the core technique for justifying claims.

Acknowledgements

The 12 articles appearing in this volume were selected from a total of 30 papers submitted for publication in this special issue. The task of finalising this selection would have been impossible without the detailed analyses in the reports of more than 50 reviewers. We are happy to take this opportunity to thank the reviewers for their diligent and thorough efforts in contributing to this volume. In addition we thank Mike Wooldridge for providing many useful observations on an earlier version of the introduction. We are also extremely grateful to Anki Rune for her work on behalf of AIJ in tracking the progress of submitted articles and minimising the administrative complexities that would otherwise have arisen. Finally we thank Ray Perrault for his encouragement and support of this special issue.

References

- [1] T. Alsinet, C. Chesnevar, L. Godo, S. Sandri, Modeling defeasible argumentation within a possibilistic logic framework with fuzzy unification, in: Proc. 11th IPMU, 2006, pp. 1228–1235.
- [2] T. Alsinet, C. Chesnevar, L. Godo, S. Sandri, G. Simari, On the computation of warranted arguments within a possibilistic logic framework with fuzzy unification, in: Proc. 11th NMR, 2006, pp. 227–236.
- [3] V. Aleven, Teaching case based argumentation through an example and models, Ph.D. thesis, The University of Pittsburgh, 1997.
- [4] S. Alvarado, M. Dyer, Analogy recognition and comprehension in editorials, in: Proc. 7th Annual Conf. of the Cogn. Sci. Soc., 1985, pp. 228–235.
- [5] S. Alvarado, M. Dyer, Understanding analogies in editorials, in: Proc. 9th IJCAI, 1985, pp. 845–847.
- [6] L. Amgoud, J.-F. Bonnefon, H. Prade, An argumentation-based approach to multiple criteria decision, in: Proc. 8th ECSQARU, 2005, pp. 269–280.
- [7] L. Amgoud, C. Cayrol, A reasoning model based on the production of acceptable arguments, *Ann. Math. AI* 34 (2002) 197–215.
- [8] L. Amgoud, S. Parsons, N. Maudet, Argument, dialogue and negotiation, in: Proc. 14th ECAI, 2000, pp. 338–342.
- [9] L. Amgoud, H. Prade, Using arguments for making decisions: a possibilistic logic approach, in: Proc. 20th UAI, 2004, pp. 10–17.
- [10] L. Amgoud, H. Prade, Explaining qualitative decision under uncertainty by argumentation, in: Proc. 21st AAAI, 2006, pp. 219–224.
- [11] <http://www.archelogos.com/>, April, 2007.
- [12] <http://www.argugrid.org/>, September 2006.
- [13] A. Artikis, M. Sergot, J. Pitt, An executable specification of a formal argumentation protocol, *Artificial Intelligence*, this volume, 2007.
- [14] K. Ashley, *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*, MIT Press, Cambridge, MA, 1990.
- [15] <http://www.cs.ucl.ac.uk/staff/a.hunter/projects/af/>, Autumn 2006.
- [16] <http://www.argumentation.org/>.
- [17] K. Atkinson, What should we do?: Computational representation of persuasive argument in practical reasoning, Ph.D. thesis, Dept. of Comp. Sci., Univ. of Liverpool, 2005.
- [18] K. Atkinson, Value-based argumentation for democratic decision support, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument* (Proc. COMMA 2006), in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 47–58.
- [19] K. Atkinson, T.J.M. Bench-Capon, Practical reasoning as presumptive argumentation using action based alternating transition systems, *Artificial Intelligence*, this volume, 2007.
- [20] K. Atkinson, T. Bench-Capon, P. McBurney, Computational representation of practical argument, *Synthese* 152 (2) (2006) 157–206.
- [21] P. Baroni, M. Giacomin, Solving semantic problems with odd-length cycles in argumentation, in: Proc. 7th European Conf. on Symbolic and Quantitative Approaches to Reasoning With Uncertainty (ECSQARU), in: *Lecture Notes in Artificial Intelligence*, vol. 2711, Springer-Verlag, 2003, pp. 440–451.
- [22] P. Baroni, M. Giacomin, On principle-based evaluation of extension-based argumentation semantics, *Artificial Intelligence*, this volume, 2007.
- [23] P. Baroni, M. Giacomin, G. Guida, SCC-recursiveness: a general schema for argumentation semantics, *Artificial Intelligence* 168 (1–2) (2005) 162–210.
- [24] A.J. Barsky IIIrd., Hidden reasons some patients visit doctors, *Ann. Intern. Med.* 94 (1981) 492–498.
- [25] R. Ben-Eliyahu, R. Dechter, Default reasoning using classical logic, *Artificial Intelligence* 84 (1–2) (1996) 113–150.
- [26] T.J.M. Bench-Capon, Specification and implementation of Toulmin dialogue game, in: J.C. Hage, et al. (Eds.), *Legal Knowledge Based Systems*, 1998, pp. 5–20.
- [27] T.J.M. Bench-Capon, Persuasion in practical argument using value-based argumentation frameworks, *J. Logic Comput.* 13 (3) (2003) 429–448.
- [28] T.J.M. Bench-Capon, Agreeing to differ: modelling persuasive dialogue between parties with different values, *Informal Logic* 22 (3) (2003) 231–245.
- [29] T.J.M. Bench-Capon, F. Coenen, P. Orton, Argument based explanation of the British Nationality Act as a logic program, *Computers, Law and AI* 2 (1) (1993) 53–66.
- [30] T.J.M. Bench-Capon, S. Doutre, P.E. Dunne, Audiences in argumentation frameworks, *Artificial Intelligence* 171 (2007) 42–71.
- [31] T.J.M. Bench-Capon, P.E. Dunne, P.H. Leng, Interacting with knowledge-based systems through dialogue games, in: Proc. 11th Annual Conf. Expert Systems and their Applications, 1991, pp. 123–130.
- [32] T.J.M. Bench-Capon, P.E. Dunne, P.H. Leng, A dialogue game for dialectical interaction with expert systems, in: Proc. 12th Annual Conf. Expert Systems and their Applications, 1992, pp. 105–113.
- [33] T.J.M. Bench-Capon, J.B. Freeman, H. Hohmann, H. Prakken, Computational models, argumentation theories and legal practice, in: C. Reed, T.J. Norman (Eds.), *Argumentation Machines. New Frontiers in Argument and Computation*, Kluwer, 2004, pp. 85–120.
- [34] T.J.M. Bench-Capon, H. Prakken, Argumentation, in: A.R. Lodder, A. Oskamp (Eds.), *Information Technology & Lawyers: Advanced Technology in the Legal Domain, from challenges to daily routine*, Springer-Verlag, 2006, pp. 61–80.
- [35] T. Bench-Capon, M. Sergot, Towards a rule based representation of open texture in law, in: C. Walter (Ed.), *Computing Power and Legal Reasoning*, Greenwood Press, 1989, pp. 39–60.
- [36] S. Benferhat, D. Dubois, H. Prade, Argumentative inference in uncertain and inconsistent knowledge bases, in: Proc. 9th UAI, 1993, pp. 411–419.
- [37] S. Benferhat, D. Dubois, H. Prade, Non-monotonic reasoning, conditional objects and possibility theory, *Artificial Intelligence* 92 (1997) 259–276.

- [38] S. Benferhat, D. Dubois, H. Prade, Possibilistic and standard probabilistic semantics of conditional knowledge bases, *J. Logic Comput.* 9 (6) (1999) 873–895.
- [39] P. Besnard, A. Hunter, A logic-based theory of deductive arguments, *Artificial Intelligence* 128 (1–2) (2001) 203–235.
- [40] P. Besnard, A. Hunter, Practical first-order argumentation, in: *Proc. 20th AAAI*, 2005, pp. 590–595.
- [41] F. Bexhe, H. Prakken, C. Reed, D. Walton, Towards a formal account of reasoning about evidence: argumentation schemes and generalisations, *Artificial Intelligence and Law* 11 (2–3) (2003) 125–165.
- [42] L. Birnbaum, M. Flowers, R. McGuire, Towards an artificial intelligence model of argumentation, in: *Proc. AAAI 1980*, 1980, pp. 313–315.
- [43] A. Bondarenko, P.M. Dung, R.A. Kowalski, F. Toni, An abstract, argumentation-theoretic approach to default reasoning, *Artificial Intelligence* 93 (1–2) (1997) 63–101.
- [44] G. Brewka, A reconstruction of Rescher’s theory of formal disputation based on default logic, in: *Proc. 11th ECAI*, 1994, pp. 366–370.
- [45] G. Brewka, Dynamic argument systems: a formal model of argumentation processes based on situation calculus, *J. Logic Comput.* 11 (2) (2001) 257–282.
- [46] M. Caminada, Semi-stable semantics, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument (Proc. COMMA 2006)*, in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 121–130.
- [47] M. Caminada, L. Amgoud, An axiomatic account of formal argumentation, in: *Proc. AAAI05*, 2005, pp. 608–613.
- [48] M. Caminada, L. Amgoud, On the evaluation of argumentation formalisms, *Artificial Intelligence* 171 (2007) 286–310.
- [49] G. Carenini, J.D. Moore, Generating and evaluating evaluative arguments, *Artificial Intelligence* 170 (11) (2006) 925–952.
- [50] C. Cayrol, C. Devred, M.C. Lagasque-Schiex, Handling controversial arguments in bipolar argumentation systems, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument (Proc. COMMA 2006)*, in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 261–272.
- [51] C. Cayrol, S. Doutre, J. Mengin, Dialectical proof theories for the credulous preferred semantics of argumentation frameworks, in: *Sixth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-2001)*, in: *Lecture Notes in Artificial Intelligence*, vol. 2143, Springer-Verlag, 2001, pp. 668–679.
- [52] C. Cayrol, S. Doutre, J. Mengin, On decision problems related to the preferred semantics for argumentation frameworks, *J. Logic Comput.* 13 (3) (2003) 377–403.
- [53] C. Cayrol, M.C. Lagasque-Schiex, Gradual valuation for bipolar argumentation frameworks, in: *8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-2005)*, in: *Lecture Notes in Artificial Intelligence*, vol. 3571, Springer-Verlag, 2005, pp. 366–377.
- [54] C. Cayrol, M.C. Lagasque-Schiex, On the acceptability of arguments in bipolar argumentation frameworks, in: *8th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-2005)*, in: *Lecture Notes in Artificial Intelligence*, vol. 3571, Springer-Verlag, 2005, pp. 378–389.
- [55] C. Cayrol, M.C. Lagasque-Schiex, Graduality in argumentation, *J. AI Res.* 23 (2005) 245–297.
- [56] C. Chesnevar, J. McGinnis, S. Modgil, I. Rahwan, C. Reed, G. Simari, M. South, G. Vreeswijk, S. Willmott, Towards an argument interchange format, *Knowledge Eng. Rev.* 21 (4) (2006) 293–316.
- [57] C.I. Chesnevar, A.G. Maguitman, R.P. Loui, Logical models of argument, *ACM Comput. Surv.* 22 (2000) 337–383.
- [58] C. Chesnevar, G. Simari, A lattice-based approach to computing warranted beliefs in skeptical argumentation frameworks, in: *Proc. IJCAI07*, 2007, pp. 280–285.
- [59] <http://www.comma-conf.org/>.
- [60] J. Conklin, M.L. Begeman, gIBIS: A hypertext tool for team design deliberation, in: *Hypertext*, 1987, pp. 247–251.
- [61] S.A. Cook, R.A. Reckhow, The relative complexity of propositional proof systems, *J. Symbolic Logic* 44 (1) (1979) 36–50.
- [62] S. Coste-Marquis, C. Devred, P. Marquis, Prudent semantics for argumentation frameworks, in: *Proc. 17th ICTAI*, 2005, pp. 568–572.
- [63] S. Coste-Marquis, C. Devred, P. Marquis, Symmetric argumentation frameworks, in: L. Godo (Ed.), *Proc. 8th European Conf. on Symbolic and Quantitative Approaches to Reasoning With Uncertainty (ECSQARU)*, in: *Lecture Notes in Artificial Intelligence*, vol. 3571, Springer-Verlag, 2005, pp. 317–328.
- [64] S. Coste-Marquis, C. Devred, S. Konieczny, M.-C. Lagasque-Schiex, P. Marquis, On the merging of Dung’s argumentation systems, *Artificial Intelligence*, this volume, 2007.
- [65] F. Dignum (Ed.), *Advances in Agent Communication*, *Lecture Notes in Artificial Intelligence*, vol. 2922, Springer-Verlag, 2004.
- [66] Y. Dimopoulos, B. Nebel, F. Toni, Preferred arguments are harder to compute than stable extensions, in: *Proc. 16th IJCAI*, 1999, pp. 36–43.
- [67] Y. Dimopoulos, B. Nebel, F. Toni, Finding admissible and preferred arguments can be very hard, in: A.G. Cohn, F. Giunchiglia, B. Selman (Eds.), *KR2000: Principles of Knowledge Representation and Reasoning*, 2000, pp. 53–61.
- [68] Y. Dimopoulos, B. Nebel, F. Toni, On the computational complexity of assumption-based argumentation by default reasoning, *Artificial Intelligence* 141 (2002) 57–78.
- [69] Y. Dimopoulos, A. Torres, Graph theoretical structures in logic programs and default theories, *Theoret. Comput. Sci.* 170 (1996) 209–244.
- [70] S. Doutre, J. Mengin, Preferred extensions of argumentation frameworks: query answering and computation, in: *First Intern. Joint Conf. Automated Reasoning (IJCAR 2001)*, in: *Lecture Notes in Artificial Intelligence*, vol. 2083, Springer-Verlag, 2001, pp. 272–288.
- [71] S. Doutre, T.J.M. Bench-Capon, P.E. Dunne, Explaining preferences with argument positions, in: *Proc. IJCAI’05*, 2005, pp. 1680–1681.
- [72] P.M. Dung, On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming, and *N*-person games, in: *Proc. IJCAI’93*, Chambéry, France, 1993, pp. 852–857.
- [73] P.M. Dung, On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming, and *N*-person games, *Artificial Intelligence* 77 (1995) 321–357.
- [74] P.M. Dung, R. Kowalski, F. Toni, Dialectic proof procedures for assumption-based, admissible argumentation, *Artificial Intelligence* 170 (2) (2006) 114–159.

- [75] P.M. Dung, P. Mancarella, F. Toni, A dialectic procedure for sceptical assumption-based argumentation, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument (Proc. COMMA 2006)*, in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 145–156.
- [76] P.M. Dung, P. Mancarella, F. Toni, Computing ideal sceptical argumentation, *Artificial Intelligence*, this volume, 2007.
- [77] P.E. Dunne, On concise encodings of preferred extensions, in: *Proc. 9th Non-monotonic Reasoning Workshop (NMR'2002)*, Toulouse, 2002, pp. 393–398.
- [78] P.E. Dunne, Prevarication in dispute protocols, in: *Proc. 9th ICAIL*, ACM Press, 2003, pp. 12–21.
- [79] P.E. Dunne, Suspicion of hidden agenda in persuasive argument, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument (Proc. COMMA 2006)*, in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 329–340.
- [80] P.E. Dunne, Computational properties of argument systems satisfying graph-theoretic constraints, *Artificial Intelligence*, this volume, 2007.
- [81] P.E. Dunne, T.J.M. Bench-Capon, Coherence in finite argument systems, *Artificial Intelligence* 141 (2002) 187–203.
- [82] P.E. Dunne, T.J.M. Bench-Capon, Two party immediate response disputes: properties and efficiency, *Artificial Intelligence* 149 (2003) 221–250.
- [83] P.E. Dunne, T.J.M. Bench-Capon, Complexity in value-based argument systems, in: *Proc. 9th European Conf. on Logics in Artificial Intelligence (JELIA'04)*, in: *Lecture Notes in Artificial Intelligence*, vol. 3229, Springer-Verlag, 2004, pp. 360–371.
- [84] P.E. Dunne, T.J.M. Bench-Capon, Identifying audience preferences in legal and social domains, in: *Proc. DEXA'04*, in: *Lecture Notes in Computer Science*, vol. 3180, Springer-Verlag, 2004, pp. 518–527.
- [85] P.E. Dunne, S. Doutre, T.J.M. Bench-Capon, Discovering inconsistency through examination dialogues, in: *Proc. IJCAI05*, 2005, pp. 1560–1561.
- [86] P.E. Dunne, P.J. McBurney, Optimal utterances in dialogue protocols, in: *Proc. AAMAS'03*, ACM Press, 2003, pp. 608–615.
- [87] P.E. Dunne, P.J. McBurney, Concepts of optimal utterance in dialogue: selection and complexity, in: F. Dignum (Ed.), *Advances in Agent Communication*, in: *Lecture Notes in Artificial Intelligence*, vol. 2922, Springer-Verlag, 2004, pp. 310–328.
- [88] U. Egly, S. Woltran, Reasoning in argumentation frameworks using quantified Boolean formulas, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument (Proc. COMMA 2006)*, in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 133–144.
- [89] K. Eshghi, R.A. Kowalski, Abduction compared with negation as failure, in: *Proc. ICLP-89*, MIT Press, 1989, pp. 234–254.
- [90] J. Fox, P. Krause, M. Elvang-Goransson, Argumentation as a general framework for uncertain reasoning, in: *Proc. 9th Conf. on Uncertainty in AI*, 1993, pp. 114–121.
- [91] D. Gabbay, J. Woods, More on non-cooperation in dialogue logic, *Logic J. IGPL* 9 (2001) 305–323.
- [92] G. Gentzen, Investigations into logical deductions, 1935, in: M.E. Szabo (Ed.), *The Collected Papers of Gerhard Gentzen*, North-Holland, 1969, pp. 68–131.
- [93] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, H. Turner, Non-monotonic causal theories, *Artificial Intelligence* 153 (1–2) (2004) 49–104.
- [94] J. Glazer, A. Rubinstein, Debates and decisions: On a rationale of argumentation rules, *Games and Economic Behavior* 36 (2) (2001) 158–173.
- [95] K. Gödel, Über formal unentscheidbare Sätze der Principia Mathematica und verwandter System I, *Monatshefte Math. Phys.* 38 (1931) 173–198.
- [96] T.F. Gordon, *The Pleadings Game: An Artificial Intelligence Model of Procedural Justice*, Kluwer Academic Publishers, Dordrecht, 1995.
- [97] T.F. Gordon, H. Prakken, D. Walton, The Carneades model of argument and burden of proof, *Artificial Intelligence*, this volume, 2007.
- [98] F. Grasso, A. Cawsey, R. Jones, Dialectical argumentation to solve conflicts in advice giving: a case study in the promotion of healthy nutrition, *Int. J. Hum.-Comput. Stud.* 53 (6) (2000) 1077–1115.
- [99] N. Green, S. Carberry, Interpreting and generating indirect answers, *Computational Linguistics* 25 (3) (1999) 389–435.
- [100] C.L. Hamblin, *Fallacies*, Methuen, London, 1970.
- [101] J. Hulstijn, L.W.N. van der Torre, Combining goal generation and planning in an argumentation framework, in: *Proc. NMR 2004*, 2004, pp. 212–218.
- [102] A. Hunter, Making argumentation more believable, in: *Proc. 19th AAAI*, MIT Press, 2004, pp. 269–274.
- [103] A. Hunter, Towards higher impact argumentation, in: *Proc. 19th AAAI*, MIT Press, 2004, pp. 275–280.
- [104] H. Jakobovits, On the theory of argumentation frameworks, Ph.D. thesis, Dept. of Comp. Sci., Vrije Univ. Brussels, 2000.
- [105] H. Jakobovits, D. Vermeir, Dialectic semantics for argumentation frameworks, in: *Proc. 7th ICAIL*, 1999, pp. 53–62.
- [106] M.W. Johnson, P. McBurney, S. Parsons, A mathematical model of dialog, *Elec. Notes in Theor. Comp. Sci.* 141 (5) (2005) 33–48.
- [107] A. Josang, A logic for uncertain probabilities, *J. Approx. Reason.* 9 (2001) 279–311.
- [108] A.C. Kakas, R.A. Kowalski, F. Toni, Abductive logic programming, *J. Logic Comput.* 2 (6) (1992) 719–770.
- [109] A.C. Kakas, P. Moraitis, Argumentation based decision making for autonomous agents, in: *Proc. AAMAS'03*, 2003, pp. 883–890.
- [110] S. Kraus, K. Sycara, A. Evenchik, Reaching agreements through argumentation: a logical model and implementation, *Artificial Intelligence* 104 (1998) 1–69.
- [111] G. Li, V.S. Uren, E. Motta, S. Buckingham Shum, J. Domingue, Claimaker: weaving a semantic web of research papers, in: *International Semantic Web Conference*, 2002, pp. 436–441.
- [112] P. Liberatore, On the complexity of choosing the branching literal in DPLL, *Artificial Intelligence* 116 (2000) 315–326.
- [113] A.R. Lodder, *Dialaw: On legal justification and dialogue games*, Ph.D. thesis, Univ. of Maastricht, 1998.
- [114] A.R. Lodder, A. Herczog, DialLaw: A dialogical framework for modeling legal reasoning, in: *Proc. 5th ICAIL*, 1995, pp. 146–155.
- [115] K. Lorenz, *Arithmetik und Logik als Spiele*, Ph.D. thesis, Kiel, 1961.
- [116] P. Lorenzen, *Dialectical foundations of logical calculi*, in: *Constructive Philosophy*, Univ. of Massachusetts Press, 1987, K.R. Pavlovic (transl.).

- [117] P. Lorenzen, K. Lorenz, *Dialogische Logik*, Wissenschaftliche Buchgesellschaft, Darmstadt, 1978.
- [118] R.P. Loui, Process and policy: resource-bounded non-demonstrative reasoning, *COMPINT: Computational Intelligence: An International Journal* 14 (1) (1998) 1–38.
- [119] R.P. Loui, J. Norman, Rationales and argument moves, *Artificial Intelligence and Law* 3 (1995) 159–189.
- [120] L.S. Lutomski, The design of an attorney's statistical consultant, in: *Proc. 2nd ICAIL*, 1989, pp. 224–233.
- [121] J.D. MacKenzie, Question-begging in non-cumulative systems, *J. Philos. Logic* 8 (1) (1978) 117–133.
- [122] P. McBurney, S. Parsons, Representing epistemic uncertainty by means of dialectical argumentation, *Ann. Math. AI* 32 (1–4) (2001) 125–169.
- [123] P. McBurney, S. Parsons, Games that agents play: A formal framework for dialogues between autonomous agents, *J. Logic Language Inform.* 11 (2002) 315–334.
- [124] P. McBurney, D. Hitchcock, S. Parsons, The eightfold way of deliberation dialogue, *Int. J. Intell. Sys.* 22 (1) (2007) 95–132.
- [125] P. McBurney, S. Parsons, M. Wooldridge, Desiderata for agent argumentation protocols, in: *Proc. First Intern. Joint Conf. on Autonomous Agents and Multi-agent Systems (AAMAS'02)*, ACM Press, 2002, pp. 402–409.
- [126] P. McBurney, R.M. van Eijk, S. Parsons, L. Amgoud, A dialogue-game protocol for agent purchase negotiations, *J. Autonomous Agents Multi-Agent Syst.* 7 (3) (2003) 235–273.
- [127] L.T. McCarty, An implementation of Eisner v. Macomber, in: *Proc. 5th ICAIL*, 1995, pp. 276–286.
- [128] C.C. Marshall, Representing the structure of a legal argument, in: *Proc. 2nd ICAIL*, 1989, pp. 121–127.
- [129] D.C. Martinez, A.J. Garcia, G.R. Simari, On acceptability in abstract argumentation frameworks with an extended defeat relation, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument (Proc. COMMA 2006)*, in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 273–278.
- [130] D. Moore, Dialogue game theory for intelligent tutoring systems, Ph.D. thesis, Leeds Metropolitan Univ., 1993.
- [131] R.C. Moore, Semantical considerations on non-monotonic logic, *Artificial Intelligence* 25 (1985) 75–94.
- [132] M. Mozina, J. Zabkar, I. Bratko, Argument based machine learning, *Artificial Intelligence*, this volume, 2007.
- [133] M. Mozina, J. Zabkar, T. Bench-Capon, I. Bratko, Argument based machine learning applied to law, *Artificial Intelligence* 13 (1) (2005) 53–73.
- [134] S.H. Nielsen, S. Parsons, Computing preferred extensions for argumentation systems with sets of attacking arguments, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument (Proc. COMMA 2006)*, in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 97–108.
- [135] S.H. Nielsen, S. Parsons, An application of formal argumentation: fusing Bayes nets in multi-agent systems, *Artificial Intelligence*, this volume, 2007.
- [136] J. von Neumann, O. Morgenstern, *Theory of Games and Economic Behaviour*, Princeton University Press, Princeton, NJ, 1944.
- [137] N. Oren, T. Norman, A. Preece, Subjective logic and arguing with evidence, *Artificial Intelligence*, this volume, 2007.
- [138] M.J. Osborne, A. Rubinstein, *A Course in Game Theory*, The MIT Press, Cambridge, MA, 1994.
- [139] S. van Otterloo, The value of privacy, in: *Proc. AAMAS'05*, 2005, pp. 1015–1022.
- [140] S. Parsons, C.A. Sierra, N.R. Jennings, Agents that reason and negotiate by arguing, *J. Logic Comput.* 8 (3) (1998) 261–292.
- [141] S. Parsons, P. McBurney, E. Sklar, M. Wooldridge, On the relevance of utterances in formal inter-agent dialogues, in: *Proc. 6th AAMAS*, 2007.
- [142] S. Parsons, M. Wooldridge, L. Amgoud, An analysis of formal inter-agent dialogues, in: *Proc. AAMAS'02*, ACM Press, 2002, pp. 394–401.
- [143] S. Parsons, M. Wooldridge, L. Amgoud, Properties and complexity of formal inter-agent dialogues, *J. Logic Comput.* 13 (3) (2003) 347–376.
- [144] P. Paruchiri, M. Tambe, F. Ordóñez, S. Kraus, Security in multi-agent systems by policy randomization, in: *Proc. AAMAS'06*, 2006, pp. 273–280.
- [145] C. Perelman, L. Olbrechts-Tyteca, *The New Rhetoric: A Treatise on Argumentation*, University of Notre Dame Press, Notre Dame, 1969.
- [146] J.L. Pollock, Criteria and our knowledge of the material world, *Philos. Rev.* 76 (1967) 28–60.
- [147] J.L. Pollock, The structure of epistemic justification, in: *Amer. Philos. Quarterly Monograph Series*, vol. 4, Blackwell, 1970, pp. 62–78.
- [148] J.L. Pollock, *Knowledge and Justification*, Princeton Univ. Press, Princeton, NJ, 1974.
- [149] J.L. Pollock, Defeasible reasoning, *Cognitive Sci.* 11 (1987) 481–581.
- [150] J.L. Pollock, A theory of defeasible reasoning, *Int. J. Intell. Sys.* 6 (1991) 33–54.
- [151] J.L. Pollock, How to reason defeasibly, *Artificial Intelligence* 57 (1) (1992) 1–42.
- [152] H. Prakken, Logical tools for modelling legal argument, Ph.D. thesis, Vrije Univ., Amsterdam, 1993 (published as [154]).
- [153] H. Prakken, From logic to dialectics in legal argument, in: *Proc. 5th ICAIL*, 1995, pp. 165–174.
- [154] H. Prakken, *Logical Tools for Modelling Legal Argument*, Kluwer Academic Publishers, 1997.
- [155] H. Prakken, Analysing and reasoning about evidence with formal models of argumentation, *Law, Probability and Risk* 3 (2004) 33–50.
- [156] H. Prakken, G. Sartor, Modelling reasoning with precedents in a formal dialogue game, *Artificial Intelligence* 6 (1998) 231–287.
- [157] I. Rahwan, L. Amgoud, An argumentation based approach for practical reasoning, in: *Proc. AAMAS06*, 2006, pp. 347–354.
- [158] I. Rahwan, F. Zablith, C. Reed, Laying the foundations for a world wide argument web, *Artificial Intelligence*, this volume, 2007.
- [159] S.D. Ramchurn, C. Sierra, L. Godo, N.R. Jennings, Negotiating using rewards, *Artificial Intelligence*, this volume, 2007.
- [160] J. Raz (Ed.), *Practical Reasoning*, Oxford University Press, Oxford, UK, 1978.
- [161] C. Reed, T. Norman, *Argumentation Machines*, Kluwer Academic Publishers, Dordrecht, 2003.
- [162] C. Reed, G. Rowe, Araucaria: Software for argument analysis, diagramming and representation, *Int. J. AI Tools* 13 (4) (2004) 961–980.
- [163] T. Rombouts, Solutions of argumentation in cooperative game theory, Master's thesis, Cognitive Artificial Intelligence, Univ. of Utrecht, August 2004.
- [164] G. Rowe, C. Reed, Translating Wigmore diagrams, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument (Proc. COMMA 2006)*, in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 171–182.

- [165] R. Reiter, A logic for default reasoning, *Artificial Intelligence* 13 (1980) 81–132.
- [166] E.L. Rissland, K.D. Ashley, R.P. Loui, AI and Law: a fruitful synergy, *Artificial Intelligence* 150 (1–2) (2003) 1–15.
- [167] E.L. Rissland, D.B. Skalak, M. Friedman, Bankxx: a program to generate argument through case-based search, in: *Proc. ICAIL'93*, 1993, pp. 117–124.
- [168] V.R. Ruggiero, *Thinking Critically about Ethical Issues*, sixth ed., McGraw-Hill, 2003.
- [169] G. Sartor, A simple computation model for non-monotonic and adversarial legal reasoning, in: *Proc. ICAIL'93*, 1993, pp. 192–201.
- [170] G. Sartor, A formal model of legal argumentation, *Ratio Juris* 7 (1994) 212–226.
- [171] J.R. Searle, *Rationality in Action*, MIT Press, Cambridge, MA, 2001.
- [172] J. Silverman, Hidden agendas and how to uncover them, *Medicine* 33 (2) (2005) 27–29.
- [173] G.R. Simari, R.P. Loui, A mathematical treatment of defeasible reasoning and its implementation, *Artificial Intelligence* 53 (2–3) (1992) 125–157.
- [174] D.B. Skalak, E.L. Rissland, Arguments and cases: an inevitable intertwining, *Artificial Intelligence and Law* 1 (1992) 3–44.
- [175] K. Sycara, Persuasive argumentation in negotiation, *Theory and Decision* 28 (1990) 203–242.
- [176] K. Sycara, Multi-agent systems, *AI Magazine* (Summer 1998) 79–92.
- [177] P. Torroni, A study on the termination of negotiation dialogues, in: *Proc. AAMAS'02*, 2002, pp. 1223–1230.
- [178] S. Toulmin, *The Uses of Argument*, Cambridge University Press, 1959.
- [179] V.S. Uren, S. Buckingham Shum, M. Bachler, G. Li, Sensemaking tools for understanding research literatures: design, implementation, and user evaluation, *Int. J. Man-Machine Stud.* 64 (5) (2006) 420–445.
- [180] A. Urquhart, The complexity of Gentzen systems for propositional logic, *Theoret. Comput. Sci.* 66 (1) (1989) 87–97.
- [181] B. Verheij, Artificial argument assistants for defeasible argumentation, *Artificial Intelligence* 150 (1–2) (2003) 291–324.
- [182] G. Vreeswijk, *Abstract Argumentation Systems*, *Artificial Intelligence* 90 (1997) 225–279.
- [183] G. Vreeswijk, An algorithm to compute minimally grounded and admissible defence sets in argument systems, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Computational Models of Argument* (*Proc. COMMA 2006*), in: *Frontiers in Artificial Intelligence and Applications*, vol. 144, IOS Press, 2006, pp. 109–120.
- [184] G. Vreeswijk, H. Prakken, Credulous and sceptical argument games for preferred semantics, in: *Proceedings of JELIA'2000*, The 7th European Workshop on Logic for Artificial Intelligence, Berlin, 2000, pp. 224–238.
- [185] M.A. Walker, S.J. Whittaker, A. Stent, P. Maloor, J.D. Moore, M. Johnston, G. Vasireddy, Generation and evaluation of user-tailored responses in multi-modal dialogue, *Cognitive Sci.* 28 (2004) 811–840.
- [186] D.N. Walton, *Argumentation Schemes for Presumptive Reasoning*, Erlbaum, Mahwah, NJ, 1996.
- [187] D.N. Walton, Examination dialogue: an argumentation framework for critically questioning an expert opinion, *J. Pragmatics* 38 (2006) 745–777.
- [188] D.N. Walton, T. Gordon, Critical questions in computational models of legal argument, in: P.E. Dunne, T.J.M. Bench-Capon (Eds.), *Argumentation in Artificial Intelligence and Law*, in: *IAAIL Workshop Series*, vol. 2, Wolf Legal Publishers, Tilburg, 2005, pp. 103–111.
- [189] D.N. Walton, E.C.W. Krabbe, *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, Univ. of New York Press, 1995.
- [190] J.H. Wigmore, *The Principles of Judicial Proof*, second ed., Little, Brown and Co., 1931.
- [191] J. Woods, A. Irvine, D. Walton, *Argument: Critical Thinking, Logic and the Fallacies*, Prentice-Hall, 2004.
- [192] M. Wooldridge, *An Introduction to Multi-agent Systems*, John Wiley & Sons, Chichester, UK, 2002.
- [193] M. Wooldridge, P.E. Dunne, S. Parsons, On the complexity of linking deductive and abstract argument systems, in: *Proc. 21st AAAI*, 2006.