Applications of Argumentation Schemes

Chris Reed
Department of Applied Computing
University of Dundee,
Dundee DD1 4HN Scotland
Tel. (+44)(0)1382 348083
Fax (+44)(0)1382 345509
chris@computing.dundee.ac.uk
www.computing.dundee.ac.uk/staff/creed/

Doug Walton
Department of Philosophy
University of Winnipeg,
Winnipeg, Manitoba R3B 2E9 Canada
Tel. (204) 786 9878

d.walton@uwinnipeg.ca
www.uwinnipeg.ca/academic/as/philosophy/walton.htm

Abstract

Argumentation schemes capture common, stereotypical patterns of reasoning which are nondeductive and nonmonotonic. As interest in understanding these schemes from a theoretical point of view grows, so too does an awareness within computational work that these schemes might yield powerful techniques in a range of domains. This paper aims to perform two functions. First, to briefly review the literature on argumentation schemes, including the key works by Hastings, Walton and Kienpointner, and to set it in a broader context, adducing concerns from deductivism and presumptive, nonmonotonic reasoning. The second is to consider the various roles argumentation schemes might play in Artificial Intelligence, and in particular to consider (i) how schemes might be characterized as planning operators in domains such as natural language generation, with operationalized schemes representing means of achieving specific (e.g. persuasive) goals; (ii) how schemes might be exploited in AI domains typically characterized by a deductive (if argumentation theoretic) basis for communication, such as communication between intelligent agents; (iii) how schemes might be used in teaching critical thinking, and how they might be characterized in pedagogical software; (iv) more broadly, the representation tasks that can be facilitated using argumentation schemes (drawing particularly upon the wide–ranging discussions held at the *Symposium on Argument and Computation* held in Scotland in July 2000); and (v) the role that critical questions associated with schemes have in generative, representational and pedagogical aspects of argumentation in AI. The paper concludes with desiderata for a theoretical understanding of argumentation schemes, motivated by the demands of the AI applications discussed, with the aim of stimulating and outlining further foundational research in the area.

Introduction

Argumentation schemes offer a means of characterising stereotypical patterns of reasoning. Walton (1996:81) employs the following example:

A PhD student, Susan, has spent more than five years trying to finish her thesis, but there are problems. Her advisers keep leaving town, and delays are continued. She contemplates going to law school, where you can get a degree in a definite period. But then she thinks: "Well, I have put so much work into this thing. It would be a pity to give up now."

This is not just an instance of a presumptive, defeasible argument, but an instantiation of a particular pattern of reasoning, a particular scheme – in this case, the *argument from waste*. By taking this classificatory step, it becomes possible to probe in a straightforward, deterministic way, for missing premises, and to consider a set of predetermined 'critical questions' in carrying out an analysis and evaluation of the argument.

Argumentation schemes have recently been attracting increasing interest for several reasons. The first is their contribution to fallacy theory. As Walton has pointed out in a series of monographs, arguments which fit into traditional categories of fallacies seem, under the right circumstances to be appropriate, acceptable and persuasive. (Walton, 1996) posits that schemes offer one way of tackling this apparent contradiction. Argumentation schemes are also very attractive on pedagogical grounds. Schemes can be used as a way of providing students with additional structure and analytic tools with which to both analyse natural arguments, and to evaluate them critically.

Finally, schemes also hold great potential for tackling a variety of problems in artificial intelligence (AI). The real world represents an immense challenge to artificial agents. Even if we focus only upon reasoning capabilities, and leave to one side the physical aspects of interacting with the world, an agent must deal with two fundamental problems: uncertainty and incompleteness. Not only will an agent not know everything, it cannot even be sure of the things that it does know – and this demands a complete shift from more traditional approaches to software design. Recent work (see (Carbogim et al., 2000; Reed, 1997; Walton, 2000) for reviews) has shown that argumentation offers a powerful means of tackling these problems by moving away from purely deductive, monotonic approaches to reasoning, towards presumptive, defeasible techniques. Presumptive reasoning supports inference under conditions of incompleteness by allowing unknown data to be presumed, and defeasible conclusions can be withdrawn or modified if known (but uncertain) data turns out to be flawed (Fox & Das, 2000).

Typically, however, such reasoning systems will have to interact not only with the world, but also with humans. To do this places further demands on these systems: not only must this reasoning be carried out, but it must also be presented, perhaps dialogically, in a form that is appropriate for human consumption. Once again, it has been demonstrated that theories of argumentation offer flexible, realistic, and, crucially, implementable techniques, (Grasso *et al.*, 2000; Reed, 1998).

Thus the two requirements in AI – of reasoning and the subsequent expression of that reasoning in language – mirror the twofold view of argumentation within philosophy. In the same way that argumentation schemes have a role to play in both views of argument, so it has been recognised that they may also have a role to play in both facets of AI. At a recent symposium focusing upon the overlap between argumentation and AI, collaborative efforts in diverse areas such as computational linguistics, legal reasoning and rhetoric all independently identified argumentation schemes as having the potential for playing a central role in future computational work (Reed, Norman & Gabbay, *to appear*). In this paper, several computational applications of argumentation schemes are explored in more detail, in an attempt both to understand the role that schemes play in AI, and to develop desiderata for foundational theoretical work upon which researchers in AI can build.

What are argumentation schemes?

Argumentation schemes are the forms of argument (structures of inference) that enable one to identify and evaluate common types of argumentation in everyday discourse. In (Walton, 1996), twenty–five argumentation schemes for presumptive reasoning are identified. Matching each argumentation scheme, a set of critical questions is given. The two things together, the argumentation scheme and the matching critical questions, are used to evaluate a given argument in a particular case, in relation to a context of dialogue in which the argument occurred. An argument used in a given case is evaluated by judging the weight of evidence on both sides at the given point in the case where the argument was used. If all the premises are supported by some weight of evidence, then that weight of acceptability is shifted towards the conclusion, subject to rebuttal by the asking of appropriate critical questions.

In deductive logic, we are used to working with forms of argument. Deductively valid forms of argument like *modus ponens* and disjunctive syllogism are used as formal structures to analyze and evaluate arguments. In a comparable way, inductive forms of argumentation of various sorts can be used to model probabilistic argumentation. The new tool we need is a set of argumentation schemes (form of argument) to model many kinds of plausibilistic (abductive, presumptive, defeasible) forms of argument. But the modeling with respect to this third type of argument needs to be pragmatic in nature. The argument is used to fulfill a probative function whereby probative weight is transferred from the premises to the conclusion. Probative weight is defeasible. Its function is to tilt a balance of considerations on an ultimate issue in a dialogue to one side or the other. An argumentation scheme is evaluated in a given case in light of appropriate critical questions in a dialogue. This approach sounds new and unusual to those traditionalists trained in deductive logic. But its utility is immediately apparent from a computational viewpoint, especially in realtion to recent work on defeasible reasoning in AI.

The list of presumptive argumentation schemes given in (Walton, 1996) is not complete, but identifies many of the most common forms of defeasible argumentation that should be the focus of research. Perelman and Olbrechts–Tyteca (1958) identified many distinctive kinds of arguments used to convince a respondent on a provisional basis. Arthur Hastings' Ph.D. thesis (1963) made an even more systematic taxonomy by listing some of these schemes, along with useful examples of them. Recently Kienpointner (1992) has produced an even more comprehensive outline of many argumentation schemes, stressing deductive and inductive forms. Among the presumptive argumentation schemes presented and analyzed in (Walton, 1996) are such familiar types of argumentation as argument from sign, argument from example, argument from commitment, argument from position to know, argument from expert opinion, argument from analogy, argument from precedent, argument from gradualism, and the slippery slope argument. Helpful examples of each type of argumentation are given and discussed. In other recent writings on argumentation, like van Eemeren and Grootendorst (1992), there is a good deal of stress laid on how important argumentation schemes are in any attempt to evaluate common arguments in everyday reasoning as correct or fallacious, acceptable or questionable. The existing formulations of the argumentation schemes are not very precise or systematic, perhaps because they have arisen out of practical concerns in dealing with real cases. New work is needed to refine, classify and formalize these schemes. To provide a more detailed introduction to what argumentation schemes are and how they work, two examples are presented below. One is called argument from position to know and the other is called argument from expert opinion.

The argument from position to know is a type of argument based on the presumption by a proponent that a respondent is a source that is privy to some information or knowledge that can be extracted from him (she, it) by questioning. The classic example is the dialogue in which someone lost in a foreign city asks a stranger where the Central Station is (or some other location or building). The questioner presumes, perhaps wrongly, that the person queried is familiar with the town. The following form of argument from position to know is that given in (Walton, 1996).

Argument from Position to Know

Major Premise: Source *a* is in a position to know about things in a certain subject domain *S* containing proposition *A*.

Minor Premise: *a* asserts that *A* (in Domain *S*) is true (false).

Conclusion: *A* is true (false).

Matching the argument from position to know, as indicated in (Walton, 1996), are the following three critical questions.

CQ1: Is *a* in a position to know whether *A* is true (false)?

CQ2: Is *a* an honest (trustworthy, reliable) source?

CQ3: Did *a* assert that *A* is true (false)?

As indicated above, argument from position to know is taken shifts a probative weight from the premises to the conclusion, thus tilting the balance of considerations in a dialogue more towards one side. But this outcome is only tentative, depending on what happens next in the dialogue. If an appropriate critical question is posed by the respondent, the probative weight shifts the balance of considerations to the other side. Only if the question is answered satisfactorily is the probative weight shifted back again.

Argument from expert opinion is a subtype of the more general argumentation scheme for argument from position to know. Argument from expert opinion is represented by the following argumentation scheme) in the analysis given in (Walton, 1997, p. 210).

Argument from Expert Opinion

Major Premise: Source *E* is an expert in subject domain *S* containing proposition *A*.

Minor Premise: *E* asserts that proposition *A* (in domain *S*) is true (false).

Conclusion: *A* may plausibly be taken to be true (false).

Appeal to expert opinion is a fallible form of argument that often carries probative weight. But, as the logic textbooks have rightly emphasized in the past so often, there is a tendency to defer to experts, sometimes too easily, and this tendency gives rise to fallacious appeals to expert opinion. Pressing ahead to aggressively, or "browbeating" the respondent, is associated with many cases of fallacious appeal to expert opinion. The sophistical tactic often used is for the proponent to try to trade on the respondent's respect for expert opinion by suppressing the respondent's legitimate critical questions in the dialogue. The following six basic critical questions for the appeal to expert opinion are presented in (Walton, 1997, p. 223) are indicated as appropriate.

1. *Expertise Question*: How credible is *E* as an expert source?
2. *Field Question*: Is *E* an expert in the field that *A* is in?
3. *Opinion Question*: What did *E* assert that implies *A*?
4. *Trustworthiness Question*: Is *E* personally reliable as a source?
5. *Consistency Question*: Is *A* consistent with what other experts assert?
6. *Backup Evidence Question*: Is *A*'s assertion based on evidence?

Someone attempting to analyze or critically evaluate a given argument can use the two devices of the argumentation scheme and the matching critical questions. The scheme identified the form of the argument and its premises. Once the premises have been identified they can then be questioned to see if there is support for them. The critical questions indicate other ways in which the argument can be questioned or criticized by indicating key assumptions that the worth of the argument depends on. The asking of a critical question throws doubt on the structural link between the premises and the conclusion.

Critical questions

One of the key features of argumentation schemes is their list of associated critical questions – questions that can be asked (or assumptions that are held) by which a non–deductive argument based on a scheme might be judged to be (or presented as being) good or fallacious. The critical questions form a vital part of the definition of a scheme, and are one of the benefits of adopting a scheme–based approach. One crucial aspect, then, of developing applications of argumentation schemes – computational and otherwise – is to capture these critical questions in an appropriate way.

The pattern of most argumentation schemes is similar to Modus Ponens, typically with something defeasible acting as the major premise (Walton, *in press*). In a standard diagramming approach, a Modus Ponens argument would be analysed as having a single conclusion supported by two linked premises. That the premises are linked rather than convergent might be demonstrated using Freeman's (1991) approach of considering a reconstructed dialogue in which an imaginary interlocutor asks a specific question after the presentation of the minor premise. So, in the example (Ex1),

(Ex1) *This computer does have an accumulator. It is built on the von Neumann architecture.*

the interlocutor might ask, "Why is that [the second sentence] relevant?", eliciting, the major premise that *having a von Neumann architecture implies having an accumulator*. (Of course, in a real dialogue a question such as 'Why is that relevant?' is more likely to elicit, through Gricean maxims, not just the linked major premise, but also further arguments in support of that premise). It is Freeman's question of relevance, rather than of ground adequacy ("Can you give me another reason?"), which moves the argument from one premise to the other in this case, suggesting that they are linked.

Viewing Modus Ponens as an example of a linked argument structure, it is possible to see argumentation schemes in the same way: a conclusion supported by two, or sometimes more, linked premises. Crucially, many of these premises are often left implicit. In a Modus Ponens it is usual to leave implicit the major premise – so usual, in fact that the enthymematic form has been analysed as a separate argument form entirely, the Modus Brevis (Sadock, 1977). Including the major premise of a Modus Ponens usually leads to hopelessly cumbersome text – though in certain extreme situations (high levels of audience scepticism or cognitive load), it may be appropriate to make all three components of the syllogism explicit (Reed, 1999). In argumentation schemes, however, it is not just the (defeasible) major premises that are left implicit. In many cases there are a range of assumptions each of which can be seen as acting as implicit linked premises. So for example, recall the scheme capturing *Argument from Position to Know* introduced above:

(P1) *a* is in a position to know whether *A* is true (false).

(P2) *a* asserts that *A* is true (false).

(C) Therefore, *A* is true (false).

(CQ1) Is *a* in a position to know whether *A* is true (false)?

(CQ2) Is *a* an honest (trustworthy, reliable) source?

(CQ3) Did *a* assert that *A* is true (false)?

In a canonical use of this scheme, the second premise, P2, is asserted explicitly, as is the conclusion. Premise P1 is left implicit (and as Walton points out, is probably assumed by the hearer by Grice's Principle of Charity, by which an assumption of honesty and relevance is made). The argument thus has its conclusion C, supported by the two linked premises P1 and P2 (if either premise fails then the argument falls down, just as with the minor and major premises of a Modus Ponens).

In addition, however, the propositional content of the critical questions also form necessary assumptions if the argument scheme is to successfully carry the burden of proof. Thus for an argument from position to know to be successful, it is necessary for an audience to accept that *a* is honest – in addition to the premises P1 and P2. Furthermore, this additional premise is also linked: if *a* were believed to be dishonest, then the entire argument would fall down.

The complete set of linked premises employed in a scheme is thus the union of those given as premises and (the propositional content of) those listed as critical questions. In the current fluidity of active work on argumentation schemes, the distinction and overlap between premises and critical questions may be unclear. Thus, for example, in argument from position to know, CQ1 and P1 are closely related and might be characterised as a single premise; CQ2 is clearly distinct and forms a second premise; but although P2 and CQ3 are similar, it may be argued that the critical question is subtley different, aiming for the very words that *a* spoke, as opposed to a paraphrasing or interpretation. It is not the aim here to resolve such potential disputes, but rather, to show that any particular interpretation of a list of premises and critical questions can be adequately characterised as a set of linked premises, many of which may be left implicit in an actual text.

Argumentation schemes in natural language generation

Previous work (Reed, 1998; 1999) has shown how deductive and non–deductive reasoning patterns can be operationalised in an AI planner. This implementation can then be used to generate the textual structure of a monologic argument. Thus, an argument structured around Modus Ponens (*a*, *a* ® *b*, *b*) is characterised as a single planning *operator*, with a postcondition that the hearer believes *b*, and preconditions that the hearer believe both minor (*a*) and major (*a* ® *b*) premise. Similar operators ae constructed for a range of deductive and non–deductive forms. The effect of such a characterisation is that arguments can be built up by a process of means–ends analysis: planning to bring the hearer to believe a conclusion by bringing him to believe the premises, which in turn can be achieved by the application of operators, which again have preconditions, and so on. There are a number of shortcomings with this work, including its reliance upon belief (rather than upon the more accesible notions of commitment or acceptability), and requirements of the knowledge base from which arguments are constructed (the implementation focuses upon the classical 'divisio' rather than 'inventio'). Nevertheless, the model provides a basis from which to develop an implementation of argumentation schemes.

Each argumentation scheme has a goal. That is, an argument's author is using a particular scheme to some specific effect. In many cases, the effect is one of lending support to a claim in the hope that some audience will accept the claim, though in some it is the determination of a course of action or the apportionment of resources that is the aim (cf. Walton and Krabbe's (1995) analysis of dialogue types). The claim−supporting (or, perhaps, standpoint−supporting) type of goal can be crudely captured for the purposes of automatically generating an argument in natural language through the mechanism described above, *viz*. as BEL(h, p), where BEL is a predicate capturing belief, h refers to the audience (or an individual stereotype of that audience), and p is the claim supported by the argumentation scheme. A goal of this form then constitutes the postcondition of an argumentation scheme operator. The preconditions of the operator include, in the first place, the premises of a particular scheme, similarly construed. So, for example, in *Argument from Consequences*, the scheme is constructed from a premise and a conclusion thus: (Walton, 1996:76)

(P1) If A is brought about, then good (bad) consequences will (may plausibly) occur

(C) Therefore, A should (not) be brought about.

In addition, there are the critical questions:

(CQ1) How strong is the likelihood that these cited consequences will (may, must, etc.) occur?

(CQ2) If A is brought about, will (or might) these consequences occur, and what evidence
supports this claim?

(CQ3) Are there other consequences of the opposite value that should be taken into account?

An implementation of this scheme, following the structure provided by the *Rhetorica* system in (Reed, 1999), would be constructed from the postcondition

BEL(h, *do(A)*)

and preconditions

BEL(h, *leads_to(do(A), good_consequences)*)
BEL(h, *not(leads_to(do(A), bad_consequences))*)

Like many other schemes described in (Walton, 1996), this characterisation of argument from consequences involves critical questions which probe for evidence. Rather than requiring additional, extra machinery, and explicit mention within a scheme operator, this evidential function is handled naturally as part of the system's processing. In order to gather evidence for any statement in an argument – including statements which play the role of premises in argumentation schemes – all that is required is another round of processing. Thus the first precondition goal, BEL(*h*, *leads_to(do(A), good_consequences)*), can in its turn be considered as a goal to be fulfilled by the application of further operators. Consideration of the evidence for this claim (and its linked sibling premise too) is thus a natural part of the functioning of the system.

Although it is clear from the brief example above that this approach places heavy demands on the knowledge representation system, it does suffice to demonstrate that defeasible argumentation schemes – replete with their critical questions – can be characterised quite formally in an existing implemented system. With a sufficiently rich knowledge base, textual arguments which involve schemes can thus be generated automatically, significantly extending the range of the arguments which can be produced by the implementation of (Reed, 1998).

Pedagogy

One of the advantages of adopting argumentation schemes as a component of argument analysis is in teaching critical thinking skills. In the first place, the structure provided by a scheme narrows the options, and can serve as an aid to the student in identifying missing premises. Thus if a *circumstantial argument against the person* (Walton, 1996:58) is a strong candidate for explaining a particular stretch of text, the analyst might look in more detail for the appropriate premises – that the speaker claims that everyone should act in a particular way, and so on. The analyst is then led to identify premises specified by the scheme but not present in the text – such as the premise usually left implicit in circumstantial ad hominems, that if any speaker claims that something should hold for everyone, then it should hold for that speaker too.

The critical questions associated with each scheme have several roles to play: first in the initial identification – if critical questions are inappropriate in a given case, the chances are that the argument scheme too is inappropriate. Of course, the critical questions are also crucial in guiding an analyst towards an evaluation of the argument, supporting in some detail a critical approach to the argumentation presented. In more formal approaches, the only tools that a student has available are soundness and validity; critical questions offer rich contextual prompts to support the analysis process. The use of argumentation schemes, and their close relation to fallacies, also introduces another key advantage: flexibility. By accepting the fact that there may be a scale from good to bad argument, it becomes possible to equip learners with the flexibility necessary for handling real – rather than logic–textbook – arguments. The scale thus encompasses deductively strong arguments and downright fallacies (in the broadest sense, including examples such as affirming the consequent, which is not only 'bad' but also fails by most standards to appear 'good') – and also examples lying in between where support might be 'substantial' or 'slight'.

Two further features of argumentation schemes in the context of pedagogy are the ease with which they can be integrated into traditional diagramming techniques, and the possibility of supporting such diagramming with software tools. Current collaborative research is under way at the University of Dundee to build a software tool that integrates traditional argument reconstruction and diagramming with the specification of argument schemes. A key feature of this software is that selecting schemes allows students to view critical questions and supply appropriate missing premises. In practical terms, the software is also designed for portability; it runs on Windows, Mac and UNIX, and saves arguments using XML, a common interlingua which is easily converted into web pages. A prototype of this software is available for demonstration purposes at OSSA'2001.

Inter−agent communication

Research into multi−agent systems (MAS) is a rapidly expanding area of AI and computer science, and as such has not yet achieved universally accepted definitions or delimitations. The concept of an agent employed here follows in the trend typified by the work of Wooldridge and Jennings (1995), in which an agent has several defining properties. First, an agent is autonomous and persistent. When an agent is started, it typically persists for an extended period, and, crucially, has no further direct manipulation from either a user or another computer process. It may interact with either a human or another software agent, but its internal state is not under the direct influence of either. Second, an agent is proactive. An agent will typically be characterised through a mentalistic definition, employing logics of beliefs and goals. Endowed with more or less sophisticated reasoning capabilities, an agent will determine a course of action which furthers its own goals. The demands of autonomy and proactivity typically produce agents which are ultimately 'selfish'. Finally, an agent is communicative. The power of multi−agent systems lies not in the capabilities of individual agents, but rather in the abilities of heterogeneous agents to form teams and joint plans which exploit individual capabilities. The crux of designing agent systems which work effectively is in building mechanisms by which agents can reach agreements. Examples of such mechanisms include auctions, rules of information exchange, and contract specification.

Many approaches to the design of inter−agent communication protocols have enjoyed considerable success through application of traditional deductive logics in supporting the distributed reasoning which is required for reaching agreements. The work of Cohen and Levesque (1990) provides a good example of this approach. By defining complex communicative acts in terms of the mental states of the agents involved in a communicative encounter, they have developed mechanisms for the formation of closely−knit teams and the design of plans distributed across members of those teams.

This approach, however, has important limitations. In particular, a number of works (McBurney & Parsons, 2000; Parsons & Jennings, 1996; Reed, 1998; Parsons *et* al., 1998) argue that the inherent incompleteness and uncertainty of data available to individual agents can effectively stymie monotonic, deductive reasoning, and that a more flexible alternative is required. Argumentation has been proposed and implemented as a means of providing this flexibility, in two distinct ways.

In the first place, traditional defeasible logics have been extended (e.g. by Dung (1995)) to capture formal abstractions of arguments. These systems typically borrow terminology from Toulmin (1958), referring to undercutters and rebutters, and then distinguish sets of propositions and their closures on the basis of 'acceptability classes'. (Note that the use of the term acceptability here is only very distantly related to its meaning in argumentation theory). An argument belongs to the highest acceptability class if its closure is consistent with current knowledge, or to some lower class depending upon the availability of undercutting propositions, rebutting propositions, and counters to those propositions. This work then served as the foundation for a multi–agent interpretation where sets of these formal structures are exchanged between agents and then evaluated (Parsons and Jennings, 1996).

The second use of argumentation focuses upon the dialogical aspects, exploiting work such as Walton and Krabbe's dialogue typology (1995), and Mackenzie's (1990) dialogue games. In these situations it is the structure of the communication language itself which avails itself of argumentation theoretic concepts. This is typified by the work on dialogue frames described in (Reed, 1998) and on the *ludens* language in (McBurney and Parsons, 2000).

Though successful and the focus of ongoing research, these two classes of applications of argumentation to multi–agent communication both suffer from a further limitation. Consider again the argument from expert opinion. If such an argument were to be employed by an agent capable of producing defeasible arguments, then the example situation in Ex2 might be rendered between two agents as something along the lines of Ex3:

(Ex2) *Bob and Elaine are working on a computer science assignment and come across a question asking whether or not a particular machine possesses an accumulator. They discuss it for a moment, and it turns out that Elaine thinks the answer is 'No'. "But wait," says Bob, remembering his expert friend, "Mary said that this computer did have an accumulator."*

(Ex3) *agent1 sends agent2 the message*
(C) has_accumulator(computer1) ∧
(P1) said(mary,has_accumulator(computer1)) ∧
(P2) defeasibly_implies(P1, C)

There are two problems that occur with a conventional approach to handling Ex3. The first is that sending the message in this form misses a crucial piece of information. This is an example of an argument from expert opinion. Such arguments are characteristically different from other types of argument, and are to be evaluated in a characteristically different way. By reducing all argument types to plain, 'vanilla' defeasible logic, this valuable distinguishing feature is lost. One key difficulty arising from this loss manifests itself as the second problem. The agent receiving the message has to perform processing and evaluation before updating its beliefs. The predominant approach would focus upon propositions which deductively or defeasibly rebut or undercut any of C, P1 or P2. It is quite possible that the receiving agent has stored *not(expert_in(mary, computer_science))* but there is no mechansim by which this can be brought to bear in evaluating the message received. The agent has no means of determining that this information is relevant, because there is no recognition that the message fits the pattern of an argument from expert opinion.

By explicitly handling argumentation schemes – either in the message itself, or as a component of agent reasoning – it becomes possible for agents to at once broaden the scope of relevant information, and at the same time, narrow down selection on the basis of the argument scheme detected. Argumentation schemes may thus offer at least one part of an attack against the thorny problem of relevance – which becomes critical as the size of agents' knowledge bases grow past the trivial.

Conclusions

The importance and potential utility of argumentation schemes within AI and related areas has been briefly explored, but perhaps the starkest conclusion is the dearth of research upon which AI can build. Even after the work of Hastings, Kienpointner, Walton and others, argumentation schemes are still ill understood, with many issues remaining to be addressed. Researchers in AI are eager to get to work building upon the results of argumentation theory, so as a contribution towards stimulating work in the area, this paper concludes with a four–point list of desiderata of issues that AI needs to see addressed. The list is not intended to be exhaustive, but rather, focuses on key issues motivated by the areas discussed above.

First, although work here, and the research upon which it builds, has attempted to elucidate somewhat the internal structure of argumentation schemes, a fuller and more principled solution is still required. Should all critical questions be viewed as implicit premises? Are all the premises of a scheme always linked? Or are critical questions nothing more than rephrasings of premises, or annotations for the student?

Second, the problem of relevance is as substantial and pressing in philosophy as it is in AI; to what extent do argumentation schemes structure or contribute to a definition of relevance that is computationally tractable?

Third, how are schemes organised? Many relationships between schemes seem to suggest a taxonomic hierarchy, but how rigid is this hierarchy, and how are properties of schemes lower down inherited from those higher up? Along how many dimensions does the hierarchy run?

Finally, perhaps the single most important task – which has as prerequisites at least the three mentioned above – is to pin down definitions of argumentation schemes. The characterisation needs to be formal enough to support computational implementation, whilst retaining the unique advantage possessed by schemes of having the flexibility requisite for untidy, unsanitised real world argumentation.

References

Carbogim, D.V., Robertson, D.S., Lee, J.R. "Argument–based applications to knowledge engineering" *The Knowledge Engineering Review,* 2000.

Cohen, P.R. and Levesque, H.J. "Intention is Choice with Commitment", *Artificial Intelligence* **42**, pp213–261, 1990.

Dung, P.M. "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n–person games", *Artificial Intelligence* **77**, pp321–357, 1995.

van Eemeren, F.H. and Grootendorst, R. *Argumentation, Communication and Fallacies*, Hillsdale, N. J. Erlbaum, 1992.

Fox, J. and Das, S. *Safe and Sound:Artificial Intelligence in Hazardous Applications*, MIT Press, 2000.

Freeman, J.B. *Dialectics and the Macrostructure of Argument*, Foris, Dordrecht, 1991.

Grasso, F., Cawsey, A. and Jones, R., "Dialectical Argumentation to Solve Conflicts in Advice Giving: a Case Study in the Promotion of Healthy Nutrition", *International Journal of Human–Computer Studies*, **53 (6)**, p. 1077–1115, 2000.

Hastings, A.C. *A Reformulation of the Modes of Reasoning in Argumentation*, Evanston, Illinois, Ph.D. Dissertation, 1963.

Josephson, J.R. and Josephson, S.G. *Abductive Inference : Computation, Philosophy, Technology*, New York, Cambridge University Press, 1994.

Kienpointner, M. *Alltagslogik: Struktur und Funktion von Argumentationsmustern*, Stuttgart, Fromman–Holzboog, 1992.

Mackenzie, J.D. "Four Dialogue Systems", *Studia Logica* **49**, pp567–583, 1990.

McBurney, P. and Parsons, S., "Games that agents play: A formal framework for dialogues between autonomous agents." *Journal of Logic, Language and Information. To appear.*

Parsons, S. and Jennings, N.R. "Negotiation through Argumentation – a preliminary report", in *Proceedings of the International Conference on Multi Agent Systems, ICMAS'96*, 1996.

Parsons, S. Sierra, C. & Jennings, N.R. "Agents that reason and negotiate by arguing" *Journal of Logic and Computation* **8 (3)**, 1998.

Perelman, C. and Olbrechts–Tyteca, L. *The New Rhetoric,* Notre Dame, University of Notre Dame Press, 1969.

Reed, C.A. "Representing and applying knowledge for argumentation in a social context", *AI & Society* **11 (3–4)**, pp138–154, 1997.

Reed, C.A. *Generating Arguments in Natural Language*, PhD Thesis, University College London, 1998.

Reed, C.A. "Dialogue Frames in Agent Communication" in Demazeau, Y. (ed) *Proceedings of the 3rd International Conference on Multi Agent Systems (ICMAS'98)*, 1998.

Reed, C.A. "The Role of Saliency in Generating Natural Language Arguments", in *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI'99),* 1999.

Reed, C.A., Norman, T.J. & Gabbay, D.(eds) *Handbook of Argument and Computation* (*to appear*).

Sadock, J.M. "Modus Brevis: The truncated argument" in *Papers from the 13th Regional Meeting of the Chicago Linguistics Society*, pp545–554, 1977.

Toulmin, S.E. *The Uses of Argument,* CUP*,* 1958.

Walton, D.N. *Argumentation Schemes for Presumptive Reasoning,* Mahwah, N. J., Erlbaum, 1996.

Walton, D.N. *Appeal to Expert Opinion*, University Park, Penn State Press, 1997.

Walton, D.N. *The New Dialectic: Conversational Contexts of Argument*, Toronto, University of Toronto Press, 1998.

Walton, D.N. "The place of dialogue theory in logic, computer science and communication studies"*, Synthese* **123**, pp327–346, 2000.

Walton, D.N. & Krabbe, E.C.W. *Commitment in Dialoge*, SUNY Press, 1995.

Wooldridge, .M. & Jennings, N.R. "Intelligent Agents: Theory and Practice" *The Knowledge Engineering Review* **10 (2)**, pp115–152, 1995.