Argumentation Schemes: From Informal Logic to Computational Models

Trevor Bench-Capon and Katie Atkinson
Department of Computer Science,
University of Liverpool, UK
{tbc, katie}@liverpool.ac.uk

Abstract

Walton's argumentation schemes with associated characteristic critical questions have been the inspiration for a number of approaches to computational modelling of argumentation. Walton's schemes were originally intended for use in the analysis of natural language argument: once the scheme had been identified the critical questions were able to identify ways in which the presumptive conclusion of the argument could be challenged. For these purposes, the formalisation need not be too rigorous, and a number of schemes exist in a variety of similar but subtly different formulations. To be used as the basis of a computational model intended to generate arguments, however, it is necessary to make the formulation of the schemes precise and to provide a grounding in a concrete knowledge model. As yet, few schemes have undergone this process. In this paper we will illustrate the process of developing an argumentation scheme for computational use by considering the development of Walton's sufficient condition scheme for practical reasoning into a fully formalised basis for computational practical reasoning. We will then discuss some other schemes, related to the Argument from Position to Know, and how these might be used, or adapted for use, in computational contexts.

1 Introduction

Two of Douglas Walton's ideas have had a particular influence on modelling argumentation in Artificial Intelligence. One is the notion of dialogue types, developed with Eric Krabbe [14], and the other, which will be the topic of this paper, is the notion of argumentation schemes as presumptive justifications, subject to critique using a number of questions characteristic of the scheme [13] [15].

Walton's notion of argumentation schemes developed out of his long standing interest in fallacies. In particular there is a need to account for the fact that many of the well known fallacies often seem to be used quite properly to support positions in everyday argumentation. Thus, although fallacies such as the Argument from Ignorance, Argument from Expert Opinion and various forms of abductive argument, are strictly speaking logical fallacies, they also seem, in the right circumstances, to be accepted as justifying their conclusions. Thus the fallacy can be seen not so much in the form of the argument, but rather in the improper use of the argument. So, the notion of argumentation schemes was developed in order to explain the proper use of such arguments: argumentation schemes represent stereotypical patterns of reasoning which can presumptively support conclusions when used properly, but which have also the possibility of being fallacious when improperly used.

For proper use, first it has to be recognised that these arguments justify their conclusions only presumptively: anyone using such an argument must, when challenged, be prepared to offer further justification or else withdraw the conclusion. Second it has to be accepted that the conclusion depends on a number of assumptions, characteristic of the scheme. While these assumptions can legitimately be made, they need to be justified if questioned in the context in which the scheme is deployed. For example, if an argument from expert opinion is used, it is assumed that the expert was making a sincere, unbiased pronouncement on a topic within his field of expertise. The conclusion must therefore be withdrawn if reasons to think that the expert is biased can be produced: for example that his research was sponsored by a manufacturer with a vested interest in his conclusions.

The primary use for these schemes was for the analysis of naturally occurring arguments. Given a piece of prose setting out an argument, the text could be broken up into individual arguments. Where these represented instantiations of argumentation schemes, they could then be considered, using the critical questions characteristic of the particular schemes, to see whether they were acceptable, or whether they had failed to address the critical questions sufficiently to discharge the burden of proof. In this form they proved a useful tool for informal logic. With regard to AI, the first use of argumentation schemes was in argument diagramming systems. Argument diagrams had originally used Toulmin's scheme [12], as in e.g. [8], but the richer range of schemes proposed by Walton gave a greater scope which could be exploited in more general argument visualisation tools such as Araucaria [9]. Computational modelling of argument is not, however, limited to analysis of argument – there is also a desire to generate arguments from some underlying knowledge model. Since natural argumentation is not restricted to logical deduction, the use of argumentation schemes gave promise that a richer repertoire of arguments based on these schemes could be generated, extending the scope of computational argument. For example, the sufficient condition scheme for practical reasoning was used in [3] to provide a computational model of justifications for actions: this particular line of work will be further discussed in section 3. More generally, the Carneades system [7] explicitly uses argumentation schemes as the main driver for its inference mechanism.

In this paper we will consider the use of argumentation schemes for the computational generation of arguments. In section 2 we will step back a little and consider some practices in AI developed independently of the informal logic work on argumentation schemes, but which share some of the motivations and characteristics. In section 3 we will discuss the adaptation of a particular scheme for computational use, and indicate some of the clarifications and decisions that are required if this is to be done successfully. In section 4 we will look at some other schemes, those related to Position to Know, and see how far they meet, or can be made to meet, the requirements for computational use. In section 5 we offer some concluding remarks.

2 AI Before Argumentation Schemes

Just as informal logic recognised that sound deductive reasoning was insufficient to analyse the whole range of legitimate arguments that can be used, so too AI found that it was necessary to go beyond deduction and make use of similar reasoning patterns in order to be able to derive the kind of conclusions it desired.

In classical Logic Programming, the need to use Horn Clause logic to make the computation tractable posed some difficulties for representing negation, since the heads of such clauses were positive literals, meaning that we could never conclude that something was *not* the case. The solution was to extend Horn Clauses with negation as failure, so that if a particular fact could not be demonstrated from the program, it was taken to be false. This is, of course, simply a form of the Argument from Ignorance:

I do not know that X is the case, so X is not the case

and is in general fallacious. In order to be a legitimate justification for X, the assumption that *if X were the case, then X would be known*, has to be satisfied. In Logic Programming this was addressed through the Closed World Assumption, which effectively assumes that everything, or at least everything of concern to the program, is known to the program. While this assumption is not generally true, the assumption can be made under circumstances formalised in [6]. What Clark proposes is the *completion of the database*, arguing that negation as failure may be used soundly for some relation R if the procedure for R is complete (a procedure for a relation is the set of clauses with the relation as head, and completeness means that no other clauses for that relation exist: i.e. the clauses constitute individually sufficient and collectively necessary conditions). If the procedure is complete, then the clauses of the procedure can be seen as individually sufficient *and collectively necessary* conditions, so that if none of the clauses are satisfied, the negation of the head is logically justified.

This aspect of logic programming bears striking similarities to Walton's notion of argumentation schemes. A style of reasoning which is potentially fallacious is used, but conditions under which it can be regarded as legitimate are provided. We could see the Closed World Assumption as a critical question: *can we be sure that the clauses of this procedure collectively supply a necessary condition for the truth of the predicate?*

Another common form of system dating from the 1980s is the expert system which operates by using backward chaining on a set of observations to find an explanation. This explanation is only *the* solution if the system has all possible explanations in the knowledge base. This can be seen as using an argumentation scheme, Argument from the Best Explanation, but relies for its legitimacy on there being no other explanations. Again therefore, this requires the Closed World Assumption to hold, and so is again subject to the critical question proposed above. Note here that the role of the Closed World Assumption is different in that it is required even if the system makes no use of negation.

For a third argumentation scheme relating to the logic programs and expert systems of the 1980s, consider the use of Query the User [11]. Often an expert system will gather facts from the user, and accept these facts unquestioningly. The justification of these facts, therefore, is simply that the user stated them. This might be seen as use of an argumentation scheme akin to Argument from Position to Know, since the user is assumed to be able to answer the questions correctly. It does not, however, satisfy the critical questions for such a scheme: there is no reason to believe that the assumptions of that scheme are met, and no assurance that the user will be sincere. It is perhaps better to regard the justification as being an Argument from Prior Commitment, *you should believe this because you told it to me*, which at least should have persuasive force for the user who supplied the information.

We can see therefore that these early programs were effectively making use of argumentation schemes – exploiting potentially fallacious patterns of reasoning in order to produce conclusions which could be justified even though not logically inevitable, subject to certain conditions being satisfied. Note, however, that these conditions – the analogue of critical questions – are couched in terms which are very specific to the knowledge model being used. This is essential if they are to be given a precise and formal statement suitable for use within a computational model. The informally expressed critical questions associated with the schemes are not, however, completely avoided. For example, with respect to the Argument from Ignorance, we can pose critical questions such as *Has an appropriate effort to find the information been made?* and *Is this the kind of knowledge which can be said to be complete?* These questions are properly directed at the designer of the program: the program is acceptable only if the claim that the procedure is complete can be made, and this requires proper care on the part of the designer, and that the Closed World Assumption be applied only to predicates for which complete knowledge is possible. Thus a user will only accept the conclusions using these argumentation schemes if there is confidence that the critical questions have been considered and resolved in the design of the system.

The lessons from this are first that it is very natural to use argumentation schemes in computer systems: similar problems were being encountered and similar solutions proposed independently in both AI and informal logic. The contribution of Walton's notion of argumentation schemes is that it gives a better rationale for the common AI practices. The second lesson is that informal analysis of argument is not required – and because it is designed to operate on natural language cannot reasonably be required – to give very precise and rigorous definitions of the schemes and the questions. In contrast, in computational contexts, because the schemes are operating on a well defined knowledge model, it is both possible and necessary (if they are to be applied by the program) to provide precise definitions in terms of the underlying logical model. This is clearly shown by a comparison of the various formulations of Argument from Ignorance and the specification of the Closed World Assumption for logic programs. In the next section we will give an extended case study of the adaptation of an informal argumentation scheme for computational use.

3 Development of the Sufficient Condition Scheme for Practical Reasoning

Argumentation schemes can be used both for reasoning about beliefs and for reasoning about actions, known as *practical reasoning*. Philosophical work on practical reasoning has centered on the practical syllogism introduced by Aristotle, but is regarded as somewhat problematic because it is essentially abductive and thus potentially fallacious. Walton's notion of argumentation schemes offers a solution to these difficulties, and in [13] he offers two argumentation schemes for practical reasoning: the sufficient and the necessary condition schemes for reasoning about action. These schemes treat the practical syllogism not as deduction, but as a presumptive argumentation scheme. Our discussions concentrate on the sufficient condition scheme which is as follows:

W1: G is a goal for agent *a*
Doing action A is sufficient for agent *a* to carry out goal G
Therefore agent *a* ought to do action A.

Walton associates with this scheme four critical questions:

CQ1: Are there alternative ways of realising goal G?
CQ2: Is it possible to do action A?
CQ3: Does agent *a* have goals other than G which should be taken into account?
CQ4: Are there other consequences?

This scheme is perhaps sufficiently precise for analysing informal arguments, where the goal can be interpreted as the context demands. But, as we have previously argued in [3], the notion of a goal here is ambiguous, potentially referring indifferently to any direct results of the action, the consequences of those results, and the reasons why those consequences are desired.

Consider someone who wishes to travel to London to meet a friend, John, who will shortly be leaving the country. Using W1 any of the following three arguments could justify his action:

W1a: I wish to be in London
Going to London is sufficient for me to be in London
Therefore I ought to go to London.

W1b: I wish to see John
Going to London is sufficient for me to see John
Therefore I ought to go to London.

W1c: I wish to maintain my friendship with John
Going to London and meeting John is sufficient for me to maintain my friendship with John
Therefore I ought to go to London (and meet with John).

That the distinction is important can be seen by the objections that can be made against the different arguments. 'John is in Manchester' is an objection to W1b and W1c but not W1a. Similarly telephoning John is an alternative for W1c but not for the other two. Moreover we need to establish which of these we mean for the purposes of mapping into our knowledge model: if we want W1c, then we must explicitly represent meeting John in our state, whereas for W1a it is sufficient to allow it to be a consequence of the propositions true in the state. The representation of the goal of W1b also has implications for the knowledge model.

We therefore see uses of W1a-c as enthymemes for a compete argument along the lines of *I want to be in London to meet John in order to maintain our friendship: going to London is sufficient to achieve this, so I ought to go to London*. Accordingly, in [3] Walton's scheme is developed into the more elaborated scheme:

AS1 In the circumstances R
we should perform action A
to achieve new circumstances S
which will realise some goal G
which will promote some value V.

What this scheme does in particular is to distinguish three aspects which are conflated into the notion of goal in Walton's scheme. These aspects are: the *state of affairs* which will result from the action; the *goal*, which is those aspects of the new state of affairs for the sake of which the action is performed; and the *value*, which is the reason why the agent desires the goal. As indicated by the example above, making these distinctions opens up several distinct types of alternative to the recommended action. We may perform a different action to realise the same state of affairs; we may act so as to bring a different state of affairs which realises the same goal; or we may realise a different goal which promotes the same value. Alternatively, since the state of affairs potentially realises several goals, we can justify the action in terms of promoting a different value. In coming to agreement this last possibility may be of particular importance: we may want to promote different values, and so agree to perform the action on the basis of different arguments. Furthermore, different agents may have their own different values that they want to promote leading each to propose different actions to achieve a goal.

For an example suppose that Trevor and Katie need to travel to Paris for a conference. Trevor offers the argument "we should travel by plane because it is quickest". Katie replies with the argument "we should travel by train because it is much pleasanter". Trevor and Katie may continue to disagree as to how to travel, but they cannot deny each other's arguments. The conclusion will be something like "we should travel by train because it is much pleasanter, even though travelling by plane is quicker". Because two people may have different preferences, values, interests and aspirations, people may rationally choose different options: if Katie prefers comfort to speed she will rationally choose the train, but this does not mean that Trevor cannot rationally choose the plane if he prefers speed to comfort. To relate this back to the AS1 scheme, we can through this example see the distinction between a goal and a value. The goal is to be in Paris for the conference, and this is not in dispute: the dispute is how that goal should be realised and turns on the values promoted by the different methods of travel. What is important is not the state reached, but the way in which the transition is made from the initial state to the goal state.

Starting from the definition of the AS1 scheme, we can make a systematic effort to identify critical questions. We can question the various elements of the scheme, for example, whether the action is possible, or whether the value is something worthy of promotion. Also we can question the connections between elements: does the action bring about the state; does the the state realise the goal, etc. Finally we can consider alternatives in a more articulated way: we can realise the goal by reaching a different state, or promote the value by realising a different goal. In this way we can hope to determine a complete list of critical questions. As given in [3], AS1 has associated with it sixteen critical questions:

CQ1: Are the believed circumstances true?
CQ2: Assuming the circumstances, does the action have the stated consequences?
CQ3: Assuming the circumstances and that the action has the stated consequences, will the action bring about the desired goal?
CQ4: Does the goal realise the value stated?
CQ5: Are there alternative ways of realising the same consequences?
CQ6: Are there alternative ways of realising the same goal?
CQ7: Are there alternative ways of promoting the same value?
CQ8: Does doing the action have a side effect which demotes the value?
CQ9: Does doing the action have a side effect which demotes some other value?
CQ10: Does doing the action promote some other value?
CQ11: Does doing the action preclude some other action which would promote some other value?
CQ12: Are the circumstances as described possible?
CQ13: Is the action possible?
CQ14: Are the consequences as described possible?
CQ15: Can the desired goal be realised?
CQ16: Is the value indeed a legitimate value?

Thus far we have elaborated the scheme so as to give a specification of what we require. This is, however, not yet enough for computational use: we need to relate the scheme and the critical questions to an underlying knowledge model. One way of doing this would be to use a BDI model, as in [2], but a more natural approach is to use a state transition system. In [1], AS1 and CQ1-16 were defined in terms of an Action-based Alternating Transition System (AATS), a structure originally defined in [16].

For an AATS we begin with a finite set Q of possible *states*, with $q_0 \in Q$ designated as the *initial state*. Systems are populated by a set Ag of *agents*. Each agent $i \in Ag$ is associated with a set Ac_i of possible actions, and it is assumed that these sets of actions are pairwise disjoint (i.e., actions are unique to agents). The set of actions associated with the set of agents Ag is denoted by Ac_{Ag}, so $Ac_{Ag} = \bigcup_{i \in Ag} Ac_i$.

A joint action j_C for a set of agents Ag is a tuple $\langle \alpha_1,...,\alpha_k \rangle$, where for each α_j (where $j \leq k$) there is some $i \in Ag$ such that $\alpha_j \in Ac_i$. Moreover, there are no two different actions α_j and $\alpha_{j'}$ in j_{Ag} that belong to the same Ac_i. The set of all joint actions for a set of agents Ag is denoted by J_{Ag}, so $J_{Ag} = \prod_{i \in Ag} Ac_i$. Given an element j of J_{Ag} and an agent $i \in Ag$, i's action in j is denoted by j^i.

As given in [1], this allows an AATS to be defined as follows:

An *Action-based Alternating Transition System* (AATS) is an $(n + 7)$-tuple $S = \langle Q, q_0, Ag, Ac_1, \ldots, Ac_n, \rho, \tau, \Phi, \pi \rangle$, where:

- Q is a finite, non-empty set of *states*;
- $q_0 \in Q$ is the *initial state*;
- $Ag = \{1,...,n\}$ is a finite, non-empty set of *agents*;
- Ac_i is a finite, non-empty set of actions, for each $i \in Ag$ where $Ac_i \cap Ac_j = \emptyset$ for all $i \neq j \in Ag$;
- $\rho : Ac_{Ag} \to 2^Q$ is an *action precondition function*, which for each action $\alpha \in Ac_{Ag}$ defines the set of states $\rho(\alpha)$ from which α may be executed;
- $\tau : Q \times J_{Ag} \to Q$ is a partial *system transition function*, which defines the state $\tau(q, j)$ that would result by the performance of j from state q. Note that, as this function is partial, not all joint actions are possible in all states (cf. the precondition function above);
- Φ is a finite, non-empty set of *atomic propositions*; and
- $\pi : Q \to 2^\Phi$ is an *interpretation function*, which gives the set of primitive propositions satisfied in each state: if $p \in \pi(q)$, then this means that the propositional variable p is satisfied (equivalently, true) in state q.

In order to express preferences between states, the AATS was extended in [1] to include a notion of values. The idea is that a value may be promoted or demoted (or neither) by a transition between two states.

- Av_i is a finite, non-empty set of values $Av_i \subseteq V$, for each $i \in Ag$. The set of all values for a set of agents Ag is denoted by Av_{Ag}.
- δ: $Q \times Q \times Av_{Ag} \to \{+, -, =\}$ is a *valuation function* which defines the status (respectively, promoted, demoted or neutral) of a value $v \in Av_{Ag}$ ascribed to the transition between two states: $\delta(q_i, q_j, v)$ labels the transition between q_i and q_j with one of $\{+, -, =\}$ with respect to the value $v \in Av_{Ag}$.

Given this knowledge model we can define AS1 and its critical questions in terms of it.

AS2 The initial state $q_0 = q_x \in Q$,
Agent $i \in Ag$ should participate in joint action $j_n \in J_{Ag}$ where ${j_n}^i = \alpha_i$,
Such that $\tau(q_x, j_n)$ is q_y,
Such that $p_a \in \pi(q_y)$ and $p_a \notin \pi(q_x)$, or $p_a \notin \pi(q_y)$ and $p_a \in \pi(q_x)$,
Such that for some $v_u \in Av_i$, $\delta(q_x, q_y, v_u)$ is +.

CQ1: $q_0 \neq q_x$ and $q_0 \notin \rho(\alpha_i)$.

CQ2: $\tau(q_x, j_n)$ is not q_y.

CQ3: $p_a \notin \pi(q_y)$.

CQ4: $\delta(q_x, q_y, v_u)$ is not +.

CQ5: Agent $i \in Ag$ can participate in joint action $j_m \in J_{Ag}$, where $j_n \neq j_m$, such that $\tau(q_x, j_m)$ is q_y.

CQ6: Agent $i \in Ag$ can participate in joint action $j_m \in J_{Ag}$, where $j_n \neq j_m$, such that $\tau(q_x, j_m)$ is q_z, such that $p_a \in \pi(q_z)$ and $p_a \notin \pi(q_x)$ or $p_a \notin \pi(q_z)$ and $p_a \in \pi(q_x)$.

CQ7: Agent $i \in Ag$ can participate in joint action $j_m \in J_{Ag}$, where $j_n \neq j_m$, such that $\tau(q_x, j_m)$ is q_z, such that $\delta(q_x, q_z, v_u)$ is +.

CQ8: In the initial state $q_x \in Q$, if agent $i \in Ag$ participates in joint action $j_n \in J_{Ag}$, then $\tau(q_x, j_n)$ is q_y, such that $p_b \in \pi(q_y)$, where $p_a \neq p_b$, such that $\delta(q_x, q_y, v_u)$ is –.

CQ9: In the initial state $q_x \in Q$, if agent $i \in Ag$ participates in joint action $j_n \in J_{Ag}$, then $\tau(q_x, j_n)$ is q_y, such that $\delta(q_x, q_y, v_w)$ is –, where $v_u \neq v_w$.

CQ10: In the initial state $q_x \in Q$, if agent $i \in Ag$ participates in joint action $j_n \in J_{Ag}$, then $\tau(q_x, j_n)$ is q_y, such that $\delta(q_x, q_y, v_w)$ is +, where $v_u \neq v_w$.

CQ11: In the initial state $q_x \in Q$, if agent $i \in Ag$ participates in joint action $j_n \in J_{Ag}$, then $\tau(q_x, j_n)$ is q_y and $\delta(q_x, q_y, v_u)$ is +. There is some other joint action $j_m \in J_{Ag}$, where $j_n \neq j_m$, such that $\tau(q_x, j_m)$ is q_z, such that $\delta(q_x, q_z, v_w)$ is +, where $v_u \neq v_w$.

CQ12: $q_x \notin Q$.

CQ13: $j_n \notin J_{Ag}$.

CQ14: $\tau(q_x, j_n) \notin Q$.

CQ15: $p_a \notin \pi(q)$ for any $q \in Q$.

CQ16: $v_u \notin V$.

Only now – having come a very long way from W1 – do we have an argumentation scheme capable of use for the computational generation of arguments. Of course, in producing these definitions, some interpretation of the original questions in terms of the target knowledge model was necessary[1]. Also, of course, it would have been possible

[1] In [1] a seventeenth critical question was added to the list to distinguish cases where an action failed due to the intervention of another agent, as permitted through the use of joint actions in the underlying AATS structure. This question could be covered by CQ2, but the use of the AATS allows us to make the useful distinction between actions which fail because of the actions of another agent and those that fail from natural

to use some other underlying model. What is essential, however, is that we have a well specified model of the knowledge we will use to generate the arguments, and a well specified description of the scheme and the critical questions in terms of that model. Whereas in the analysis of informal argumentation the scheme and the questions can be interpreted (and re-interpreted) in terms appropriate to the particular context, for computational use all this needs to be fixed in advance.

Thus we can see that there are a number of steps that need to be undertaken when developing an argumentation scheme for computational use:

1. Any ambiguous terms in the argumentation scheme must be made precise. This may involve replacing one term by several, since the scope of the term appropriate to the context cannot be determined at run time;

2. Given the elaborated scheme, all critical questions must be identified in a systematic manner;

3. A suitable knowledge model must be selected;

4. The scheme and the critical questions must be restated in terms of the model; this may involve some further interpretation.

So far this full process has been performed for very few schemes. In the next section we will look at the use of a family of related schemes in a computational context.

4 Position to Know in a Computational Context

In this section we will consider a family of argumentation schemes which can be termed Position to Know schemes. Essentially all of them have the form *X is in a position to know whether P: X says that P, so P*. The different schemes arise from what puts X in a position to know that P. Some examples are:

- a period of study which has made X an expert (Argument from Expert Opinion);
- being present at a particular place and time (Argument from Witness Testimony);
- being a particular person, as when we say *I should know whether that hurts or not* (Argument from Privileged Access);
- some kind of public standing or community acceptance (Argument from Authority).

Some of these, particularly the first two, have received attention in the AI and argumentation literature (see e.g. [7] for Expert Opinion and [4] for Witness Testimony). However, these accounts do not necessarily have the precise analysis necessary for computer generation of arguments. In [4], no specific critical questions are given; instead there is a discussion:

causes, and so the additional question better exploits the knowledge model we are using.

> The role of such critical questions has been discussed extensively in the legal literature on witness testimony and examination. Schum ([10], p. 325) has identified three requirements of the credibility of the testimony of a witness that can be questioned: (1) veracity, or whether the witness believes what she said, (2) objectivity, or whether what was reported corresponds to the event believed, and (3) observational sensitivity, or observations of linkages between events. Bromby and Hall [5] devised a system to advise on the credibility of witness testimony by citing factors of (1) competency, (2) compellability, including the connection between the witness and the accused and any immunity the witness may have, and (3) reliability, which includes position to know factors. There remain many fine points to be clarified.

The kind of considerations mentioned here are quite appropriate for human use on an informal instance (and for the kind of analysis undertaken in [4]), but would be less suitable for computer realisation. This is perhaps unsurprising, since determining whether someone is telling the truth is an advanced social skill, and one in which there are few very reliable performers.

Expert Opinion may be more tractable. In [7], Walton's definition is first cited:

> *Major Premise.* Source E is an expert in the subject domain S containing proposition A.
>
> *Minor Premise.* E asserts that proposition A in domain S is true.
>
> *Conclusion.* A may plausibly be taken as true.

Six critical questions are given:

1. How credible is E as an expert source?
2. Is E an expert in the field that A is in?
3. Does E's testimony imply A?
4. Is E reliable?
5. Is A consistent with the testimony of other experts?
6. Is A supported by evidence?

Later in [7], when considering the representation of this scheme for use in the Carneades system, the critical questions are recast as:

- Premise. E is an expert in the subject domain S containing the proposition A.
- Premise. E asserts A.
- Assumption. E is a credible expert.
- Exception. E is not reliable.

- Exception. A is not consistent with the testimony of other experts.
- Assumption. A is based on evidence. Conclusion. A.

The authors of [7] do not specify any form of the knowledge base, but one must assume that either there is a knowledge base which allows the required propositions to be derived, or there is some means of querying the user to determine, for example, whether E is reliable. To generate arguments in a specific application, however, we believe that there is rather more work to be done, so that notions such as reliability can be specified with sufficient precision in terms of the knowledge model to be used by the application.

Considered as a general problem, addressing the above issue would be a daunting task. In particular applications, however, things may in fact be easier. Firstly a computer application is likely to use specific resources, rather than trawling a large number of unspecified resources to find support. Suppose I want to argue that Peter Lever played more cricket matches for England than Paul Allot. I can do so on the basis that Lever played 17 times for England and Allot only 13. To justify these claims I can cite the expert opinion of `www.cricinfo.com`, which is endorsed by Wisden and considered completely decisive for cricket matters. This site is credible, reliable, based on evidence, and if anyone disagrees with it, *tant pis* for them. Effectively I can use the Argument from Expert Opinion here because I can be confident that the critical questions are not problematic. Similar definitive resources exist in other areas, such as `www.imdb.com` for facts about films and actors. If, however, the application is to use some less authoritative source, such as Wikipedia, the same confidence cannot be assumed. Here we may well wish to make some effort to answer at least the question of consistency with other sources, and find some corroborative information, or make some effort to discover conflicting opinion.

Suppose then we wish to use specific Internet resources to solve questions posed to us. In order to do so we will need to rely on some analogue of Argument from Expert Opinion, call it Argument from Internet Resource. Use of such an argument is sound only if the chosen source has the right credentials, that is, if the designer who has selected this resource is satisfied with respect to a set of critical questions similar to those posed against the Argument from Expert Opinion. It is essential that the chosen source be reliable, credible, based on evidence, and be believed preferentially in cases of disagreement. The justification, in terms of the diligence of the designer, is thus very similar to the justification of the use of the Argument from Ignorance in logic programs as discussed in section 2.

In Multi-Agent Systems, where an agent will be told things by other agents, we would need to consider justifying acceptance of information received in this way using a scheme akin to that of Witness Testimony. In many current proposals, the problem is circumvented by the assumption of benevolence, in which agents can be assumed to be sincere and reliable. This is relatively plausible for closed systems, but less so for open systems, where it is dangerous to assume anything about other agents. This is a real problem for open multi-agent systems, and one which would doubtless benefit from hard thinking about how the insights from the Witness Testimony scheme and its critical questions can be adapted for use in this context.

5 Concluding Remarks

In this paper we have considered one of the strands of Walton's work which has had a significant impact in AI, namely his notion of argumentation schemes as providing presumptive justification of a conclusion subject to critical questioning. The underlying insight, that several forms of reasoning traditionally classified as fallacious are in fact desirable and necessary in informal argumentation, and can be legitimately used if certain considerations are observed, has a mirror in the practice of AI in logic programs and expert systems. They too made use of possibly fallacious forms of reasoning, subject to constraints, such as the Close World Assumption, which were able to legitimise their use in particular contexts. Argumentation schemes provide an interesting rationale for these practices. If we wish to use the schemes for purposes other than human analysis of pre-existing arguments, for example to generate arguments, the schemes as stated in informal logic cannot be used immediately. Rather it is necessary to rethink them in terms of their intended computational use and supporting knowledge model, so that all need for contextual interpretation is removed. A case study of the process was given for one particular scheme, and some considerations relating to other schemes were advanced.

Douglas Walton's influential argumentation schemes remain an important insight, one which needs to be embraced by computational modelling of argument. His work comes from informal logic and was designed for manual analysis of argumentation. For computational purposes much detailed work remains to be done before the use of argumentation schemes can be as standard and well understood as logical deduction.

References

[1] K. Atkinson and T. J. M. Bench-Capon. Practical reasoning as presumptive argumentation using action based alternating transition systems. *Artificial Intelligence*, 171(10–15):855–874, 2007.

[2] K. Atkinson, T. J. M. Bench-Capon, and P. McBurney. Arguing about cases as practical reasoning. In *Proceedings of the Tenth International Conference on Artificial Intelligence and Law*, pages 35–44. ACM Press, 2005.

[3] K. Atkinson, T. J. M. Bench-Capon, and P. McBurney. Computational representation of practical argument. *Synthese*, 152(2):157–206, 2006.

[4] F. Bex, H. Prakken, C. Reed, and D. Walton. Towards a formal account of reasoning about evidence: Argumentation schemes and generalisations. *Artificial Intelligence and Law*, 11(2-3):125–165, 2003.

[5] M. C. Bromby and M. J. J. Hall. The development and rapid evaluation of the knowledge model of advokate: An advisory system to assess the credibility of eyewitness testimony. In *Proceedings of Jurix 2002*, pages 143–52, 2002.

[6] K. L. Clark. Negation as failure. In *Logic and Data Bases*, pages 293–322, 1977.

[7] T. Gordon, H. Prakken, and D. Walton. The Carneades model of argument and burden of proof. *Artificial Intelligence*, 171(10-15):875–896, 2007.

[8] C. C. Marshall. Representing the structure of a legal argument. In *The Second International Conference on Artificial Intelligence and Law*, pages 121–27. ACM Press, 1989.

[9] C. Reed and G. Rowe. Araucaria: Software for argument analysis, diagramming and representation. *International Journal on Artificial Intelligence Tools*, 13(4):983–, 2004.

[10] D. A. Schum. *The Evidential Foundations of Probabilistic Reasoning.* John Wiley and Sons, New York, 1994.

[11] M. Sergot. A query the user facitiy for logic programs. In M. Yazdani, editor, *New Horizons in Educational Computing*, pages 145–163. Ellis Horwood, Chichester, 1984.

[12] S. Toulmin. *The Use of Argument*. Cambridge University Press, Cambridge, UK, 1958.

[13] D. N. Walton. *Argumentation Schemes for Presumptive Reasoning*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1996.

[14] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Press, Albany, NY, USA, 1995.

[15] D. N. Walton, C. Reed, and F. Macagno. *Argumentation Schemes*. Cambridge University Press, Cambridge, 2008.

[16] M. Wooldridge and W. van der Hoek. On obligations and normative ability: Towards a logical analysis of the social contract. *Journal of Applied Logic*, 3:396–420, 2005.