Multiple subarguments in logic, argumentation, rhetoric and text generation *

Chris Reed[1] and Derek Long[2]

[1] Department of Computer Science,
University College London,
Gower St.,
London, WC1E 6BT
C.Reed@cs.ucl.ac.uk
http://www.cs.ucl.ac.uk/staff/C.Reed

[2] Department of Computer Science,
Durham University,
South Road,
Durham
D.P.Long@dur.ac.uk

Abstract. A summary is provided of the problems of representing, determining generating and arranging disjunct multiple subarguments in several fields, including formal systems in uncertain domains, informal logic accounts of argument structure, rhetorical systems for maximising persuasive effect, and the automatic generation of persuasive discourse. Drawing upon the insights, problems, and partial solutions of these fields, a theory of subargument construction and organisation is presented, and is set in a framework for generating natural language argument.

Keywords: aggregating arguments, argumentation theory, rhetoric, defeasible reasoning, natural language generation, planning, rhetoric.

1. Introduction

One key task to be performed during the production of argument for some purpose (interagent negotiation, expert system explanations, decision support systems and other nontrivial human-computer discourse) is that of aggregating subarguments. In particular, it is necessary to make appropriate use of *multiple subarguments*, a term, which for the purposes of this paper, refers not simply to several lines of reasoning contributing to some conclusion, but specifically to lines of reasoning which offer *disjunctive* support. Thus in Fig. 1, only 1b contains multiple subarguments (1a is composed of a single subargument composed of two conjuncts).

(a)	(b)
Fred told me Socrates is human	Fred told me Socrates is mortal
<u>Fred's a reliable source of information</u>	<u>George told me Socrates is mortal</u>
∴Socrates is mortal	∴Socrates is mortal.

Fig. 1. Conjunctive and disjunctive support

*This work has been partly funded by EPSRC grant no. 94313824.

The question of how to deal with such support is one which is relevant not only in text generation, but also in formal approaches to argumentation (including defeasible reasoning and bespoke logics of argument), informal approaches (where the distinction between disjunctive and conjunctive support is particularly problematic) and in rhetoric (where the appropriate use of various forms of support is prescribed). By examining the approaches taken in these areas, it is possible to assail a number of key problems in the generation of natural language argument.

2. Formal approaches

In classical logic, the concept of multiple subarguments is of no use. If it is possible to prove P, say through X and X → P, it is unnecessary to then prove X once again with Y and Y → P: using both proofs rather than just one would in no way produce a 'better' argument. This accords with superficial intuitions that if a speaker, S, wishes to bring a hearer, H, to believe some proposition, all that is necessary is for S to show that the proposition follows from her beliefs. However, multiple subarguments occur with great frequency, both in natural language (Cohen, 1987), (Reed and Long, 1997) and increasingly in complex argumentation systems such as those of (Fox and Das, 1996). In both cases they are used to create arguments that are in some way 'better' than singly supported alternatives.

There are two (related) key points which necessitate the use of multiple subarguments in these circumstances. In the first place, reasoning about some domain in the real world (such as the medical applications of Fox and Das) can rarely use strict deductive inference; rather, it becomes necessary to use some weaker notion of support - and often then to express the degree of that support (either qualitatively - eg. (Parsons, 1996), or quantitatively - eg. (Sillince and Minors 1992)). If a system no longer relies solely upon strict inference then it can clearly no longer employ classical logic, and as a result may benefit from the use of multiple subarguments. These separate lines of support may then be aggregated under some flattening function, such as those discussed in (Das *et al*, 1996). In the second place, multiple subarguments offer a means of tackling the uncertainty and incompleteness inherent in an environment where agents must communicate. One agent cannot hope to maintain an accurate model of the beliefs of another, and yet it must be able to construct arguments which actually effect changes in other agents' beliefs. By planning several disjunctive supports for a proposition, an agent can increase the likelihood that the conclusion will be accepted. For although ideally an agent would want its model of the hearer's beliefs to be accurate, should there be a discrepancy, the alternative supports may turn out to be of crucial importance in persuading the hearer to concur (rather than those supports playing an ancillary role, as originally planned).

One particularly successful formal approach to dealing with uncertain and incomplete information is defeasible reasoning. Defeasible logics are also inherently able to accommodate multiple subarguments. For example, in Dung's (1995) theory of the acceptability of argumentation, the basic framework, *AF*, comprises two parts, a set of arguments and a set of attacks holding between arguments. Implicit in the former is the possibility that several arguments will contribute to a single conclusion (and it is anticipated that in natural communication, these multiple subarguments would form a *conflict free* set[1]). In other systems, the concept is more explicitly represented: Vreeswijk (1992), for example, introduces the *sub* feature of an argument, which defines the set of all subarguments (the model differs somewhat from that of Dung's, as it aims to capture structural information in addition to the content). Pollock's (1995) theory of defeasible reasoning also makes the concept of subargument explicit and more recently, also distinguishes conjunctive and disjunctive supports (Pollock, 1996).

Clearly, a system of defeasible reasoning such as that of (Pollock, 1995) needs to be extended to cope with the distributed aspects of argumentation occurring between complex, rational, *autonomous* agents (ie. agents that function without human intervention, that can interact with other agents, that can take proactive decisions on the basis of their own selfish goals and intentions, that cannot be guaranteed to have the same information as other agents, cannot be sure of the knowledge they have about other agents beliefs, and cannot be relied upon to be benevolent and cooperative - these issues are discussed in, for example, (Wooldridge and Jennings, 1995)). The notion of agency is touched upon in the work of both Vreeswijk and Dung, *inter alia*. However, it is only relatively recently that argumentation has been seen as a direct means of negotiation (Parsons and Jennings, 1996) and persuasion (Reed *et al*, 1996a) between agents (for after all, these are the primary ends served by argumentative communication between humans - though see (Walton, 1995) for a more detailed coverage).

Parsons and Jennings (1996) extended the defeasible argumentation system of Krause *et al* (1995) such that one agent could propose an argument (for a specific course of action) to another, and the recipient would then evaluate the argument on the basis of its own beliefs, intentions and plans. The evaluation is performed by searching for rebutting and undercutting counter-arguments (a rebuttal occurring when the conclusion of the argument is defeated, undercutting when one of the steps used in reaching the conclusion is defeated). Depending on the evaluation, the second agent may then agree, or communicate a counter-proposal. Thus the defeasibility of the system lies not only within a single agent's reasoning, but also in the subsequent processing in another agent with potentially quite different beliefs and goals. Importantly, the prioritization of the defeasible inferences is localised to each agent: in Parsons' system, this is due to the fact that the ordering relation is based upon the acceptability classes which in turn are dependent upon the beliefs of the agent involved. However, as he points out, it is equally possible to consider a system which performs the ordering based upon some valuation of the inferences (such as that of (Parsons, 1996)), which may offer some advantages in uncertain domains.

---

[1] Though multiple subarguments are unlikely to form an *admissible* set. For this to hold, the set, *S*, of multiple subarguments would have to attack any argument which itself attacks a component of *S* (this is Dung's definition of *acceptability*: *S* is then admissible iff each component of *S* is acceptable). In practice, this would mean that a speaker would have to create a 'water-tight' argument, anticipating all hearer counter-arguments (and this in turn would reduce the dialectic method to a one-shot monologue).

There are, however, problems with such an operationalization of distributed defeasible reasoning. The system proposed in (Parsons and Jennings, 1996) does not seem to be able to communicate multiple subarguments - a single proposal (albeit a conjunction of separate sub-proposals) is supported by a set of grounds which must *all* be taken together to infer the conclusion. As a result, the arguments which can be generated will be less sophisticated (for although it may be possible for one agent to enumerate all the disjunct subarguments for a particular course of action on successive turns in the dialogue, this is far from an optimal strategy for a resource-bounded agent[2]).

More importantly, defeasible reasoning across rational, autonomous agents seems to miss the intuitions of how the agents are functioning. Consider the scenario in Fig. 2, in which the speaker has two disjunct reasons for inferring[3] p, and the hearer may have a number of reasons for inferring ¬p. Let us assume that S knows that H believes ¬p. Following an account such as that offered by Dung (1995) or Vreeswijk (1992) would lead to S communicating exactly those subarguments which together either defeat, or are undefeated by, all the X_i which S presumes H to believe.

<u>S</u>	<u>H</u>
$a \rightarrow p$	$X_i \rightarrow \neg p$
$b \rightarrow p$	$\neg p$
p	

Fig. 2. Sample situation between speaker, S and hearer, H

However, in the real world, S may or may not know (or even be aware of the existence of) the X_i by which H believes ¬p. S is certainly unlikely to know the valuation that H places upon the various inferences from the X_i. She is even less likely to be able to anticipate how H will value her own inferences. It is not the case, therefore, that S constructs her argument through anticipating H's possible counter-arguments - rather, she is simply 'building a case' for her conclusion. Clearly, this process is going to involve consideration of what she thinks her hearer believes (in addition to other audience-specific information, such as possible bias and technical competence - see (Reed *et al*, 1996b) for details). But it does not require S to perform 'H-reasoning' to produce the arguments which she must ensure are defeated by her own. Although such reasoning may have a role to play in generating parts of a complex argument, the primary means of generating argument is, of necessity, a process of showing evidence to support a conclusion, resulting in an argument which stands independent of the possible subsequent reasoning by the hearing agent. It is this intuition which escapes a standard defeasible account of why one agent might utter a particular set of subarguments to another.

---

[2] The strategy would also suffer badly if used for communicating with a human: studies of rhetoric such as (Blair, 1838) devote great attention to issues of combining multiple subarguments: an overview is given in §4.

[3] The inference indicated by '→' is intended to be defeasible rather than strict.

3. Informal Approaches

In contrast to the 'bottom up' approach of the formal accounts of argumentation, argumentation theory is fundamentally empirical in its approach, driven by real examples of argumentation expressed in natural language. In naturally occurring argument, disjunct multiple subarguments occur with great frequency, and as a consequence are a feature of almost all theoretical accounts. However, many of these accounts do not explicitly discuss the ways in which subarguments can be combined, and the functional roles that they fulfil as a result. Fogelin (1991) is typical of such informal logic texts, explaining how argument structure can be determined and represented, through analysis of supporting relations holding between premises and conclusions. No distinction is made, however, between arguments supported by a single subargument, and those supported by several disjunctive subarguments - the latter is simply regarded as an extension of the former. Similarly, Fisher (1988) bases his theory of analysis on the *assertibility question*: "What argument or evidence would justify me in asserting the conclusion? (What would I have to know or believe to be justified in accepting it?)". Implicit is the notion that the argument or evidence may be composed of one or more subarguments.

Although there are minor differences in notation and analysis procedures, these and other accounts follow substantially the same approach to determining and describing the structure of argument. In sharp contrast is the theory proposed by Toulmin (1958) which expands upon the conventional premise/conclusion distinction to detail a six-fold division of utterances in argumentation (a *claim* supported by a *datum* to the degree specified by a *qualifier* which may include exceptional conditions which would otherwise cause *rebuttal*; the support between claim and datum is licensed by a *warrant* which may have a *backing*). Toulmin too regards the use of multiple subarguments (presumably multiple D--WB--QR links to the C) as a trivial extension to the theory.

Freeman (1991) discusses at length that the way in which premises can combine to support a conclusion (after discussion of Toulmin's theory, he eschews the six-fold division in favour of the 'standard' premise/conclusion approach). In particular he summarises (p2) the four main types of argument component structure, *viz. divergent* (whereby one premise can support several conclusions[4]), *serial* (whereby a single premise contributes to a single conclusion, which may then act as the single premise to another conclusion, and so on), *convergent* (whereby two or more premises contribute independently to a single conclusion) and *linked* (whereby two or more premises together contribute to a single conclusion). These forms are summarised in Fig. 3, below.

[4] He later points out (p93) that under some circumstances, it is more appropriate to view divergent structure as serial, with the premise repeated for each conclusion. This is the approach taken in (Reed *et al*, 1996b).

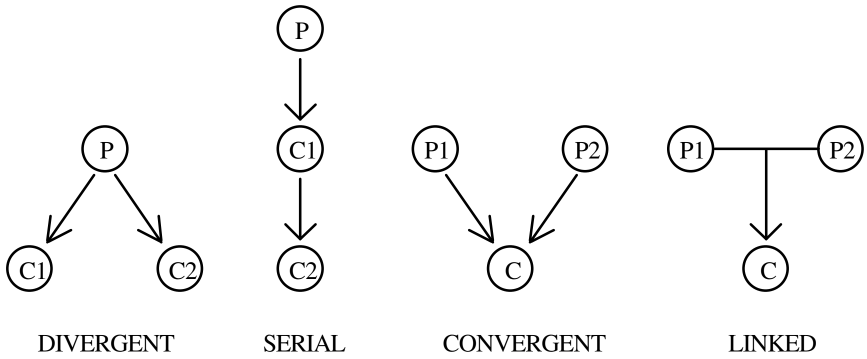

Fig. 3. The four basic argument structures, after (Freeman, 1991, p2).

Clearly, a convergent argument is one in which a conclusion is supported by disjunctive multiple subarguments. However, as Freeman explains, distinguishing between arguments which have a linked structure (such as that in Fig. 1a) from those which have convergent structure (such as that in Fig. 1b) is a particularly difficult task. This is as a direct result of the imprecise terminology used to define the classes[5]. There are two important classes of argumentation which pose a particular problem for classification: inductive generalisation and concluding a conjunction.

Inductive generalisation consists of generalising from a feature known to belong to a sample of a population, to claiming that feature of all members of the population. Thus it can be seen as presenting a number of examples (each member of the sample) from which the conclusion (the population as a whole) is inductively drawn. None of the examples presents particularly convincing evidence that the conclusion is true, but taken together the inference may be judged to be valid. On the one hand, the structure seems to be convergent: each datum is individually contributing to the conclusion, and does not depend on other data for its contribution to be valid. However, on the other hand, this fails to capture the nature of an inductive generalisation, which is that none of the data on their own would license the drawing of the conclusion - only when they are all taken together is the inference valid.

Arguments which involve two premises contributing to the conclusion which is their conjunction pose a similar problem. That is, if $(A \wedge B)$ is to be drawn from the premises A and B, is it the case that A and B have to be taken together (ie. linked structure) or do both A and B independently contribute to the conclusion as multiple subarguments (ie. convergent structure)?

Freeman's solution rests upon Toulmin's notions of argument as a dialectical construction, and potential questions that a challenger may put to the speaker. Consideration of a challenger asking the *relevance* question, "*Why is that reason relevant to the claim?*", gives rise to additional premises in a linked structure.

---

[5] For example, in the summarised definitions given in parentheses above, it is unclear exactly what the terms *contribute* and *independently* actually mean.

Similarly, consideration of a challenger posing the first *ground adequacy* question, "*Can you give me another reason?*", gives rise to additional premises in a convergent structure. Classifying an argument becomes a case of identifying which question could have been asked to elicit the additional premises. Crucial to Freeman's account is a distinction between *relevance* and *modality* between a premise and the conclusion it supports (the two aspects giving rise to the relevance and ground adequacy questions). The modality attached to a support arc (ie. the strength of that support, ranging from strict deductive inference to weak inductive inference) can affect intuitions for how a particular example should be classified[6]. In particular, where the strength of support between each premise and the conclusion is high, the temptation is to classify the argument as convergent, since each premise can be seen as providing an independent subargument. Conversely, when strength of support is low (such as from the individual examples of an inductive generalisation), the temptation is to assume a linked structure, since each premise on its own is 'too weak' to contribute to the conclusion by itself.

Using the relevance and ground adequacy questions, Freeman concludes that both inductive generalisation and concluding the conjunction are in fact examples of convergent structure (that is, each premise answers a phantom ground adequacy question). These results are controversial, but if valid may have important ramifications for logical models of natural argumentation.

The fact that concluding a conjunction in natural argument is the product of a convergent structure is particularly surprising, though Freeman's terminology rather hides the fact. By convergent structure, he is referring to premises which provide *disjunctive* support for a conclusion. Thus to claim that a conjunction is supported by premises in a convergent structure, is to claim that individual conjuncts support their conjunction *disjunctively*. It is unclear how such a notion is best represented formally, though in a system such as that of (Das *et al*, 1996), it would presumably mean modifying the aggregation function (- each premise of a conjunction is independently relevant to the conclusion, so it is the modalities which remain to be combined).

4. Rhetorical Approaches

Like argumentation theories, systems of rhetoric are concerned with real examples of argument, and as a consequence, must detail the use of subarguments: their structure, arrangement, presentation and so on. Unlike argumentation theory, however, texts on rhetoric (and in particular, those of the eighteenth and nineteenth centuries) aim to guide the synthesis of good argument - rather than guiding analysis to determine whether or not an argument is good.

The first consideration offered in most texts is the dissection of discourse

---

[6] Indeed, modality formed the basis of a theory of classification espoused by Yanal (1984). There were, however, a number of problems with the approach, as discussed in (Freeman, 1991, pp105-108). In particular, determining the modality of a link involves evaluating an argument: a process which is separate from, and intuitively should follow, the process of argument analysis.

into successive phases. Blair (1838), for example, suggests six stages: introduction, division of subject, narration of facts, argumentative part, pathetic part (essentially constituting an appeal to the emotions of the audience), and conclusion. In addition, he suggests that the argumentative part (seemingly the only part under investigation by either formal or informal logicians[7]) is then further subdivided into invention, arrangement and expression. This subdivision of argumentation follows the classical scheme of the branches of rhetoric, supported by Aristotle, Quintilian and Cicero (Billig, 1996, p81). As Blair notes, however, the task of 'inventing' the arguments is beyond the scope of rhetoric:

> "Art cannot go so far, as to supply a speaker with arguments on every cause, and every subject; though it may be of considerable use in assisting him to arrange and express those, which his knowledge of the subject has discovered. For it is one thing to discover the reasons that are most proper to convince men, and another, to manage these reasons with the most advantage." (Lect. XXXII, p427)

Once the arguments are available, there is the choice of choosing which to use. Blair, again summarises:

> "... one of the first things to be attended to is, among the various arguments which may occur upon a cause, to make a proper selection of such as appear to one's self the most solid; and to employ these as the chief means of persuasion." (Lect. XXXII, p429)

The notion of selecting and employing a number of 'best' subarguments poses several serious problems for computational systems. In a system such as that of (Das *et al*, 1996), it would mean setting some arbitrary threshold; in (Reed *et al*, 1996a), the process of deriving all possible arguments and then selecting the best would be arbitrarily restricted due to resource bounding.

After subarguments are chosen, they must be arranged in such a way as to maximise their persuasive impact. The crucial importance of this stage is noted by many writers (Blair, 1838), (Billig, 1996), but is eloquently explained by Whately:

> "... Arrangement is a more important point than is generally supposed; indeed it is not perhaps of less consequence in Rhetoric than in the Military Art; in which it is well known, that with an equality of forces, in numbers, courage, and every other point, the manner in which they are drawn up, so as either to afford mutual support, or on the other hand, even to impede and annoy each other, may make the difference of victory of defeat."(Whately, 1855, Ch. I, §3, p35)

Various systems of rhetoric provide rules for arranging subarguments, varying as much in specificity as they do in number. Here, a brief summary will be presented of the rules of Blair and Whately, which, between them, are typical of those which are available elsewhere.

Blair suggests four rules. In the first place, he claims, arguments should be grouped together according to their type (where *type* is defined as being one of truth, morality or profitability). Clearly, for such a heuristic to be implemented, it would be necessary to have some means of determining the 'type' of a subargument, something

[7] Also, possibly, including the narration: the facts established during the narrative part could be seen as basic grounds for the subsequent argumentation.

which is beyond even a tailored logical argumentation schemes such as LA (Krause *et al*, 1995). The second and third rules concern the strength of individual subarguments: that they should increase in strength (except in particular situations, as discussed in (Reed *et al*, 1996b)) and that any subargument which is particularly strong should be "brought out by itself, placed in its full light, amplified, and rested upon" (Blair, 1838, p431). These heuristics become amenable to implementation as soon as a means of rigorously determining argument strength is devised: in systems such as (Das *et al*, 1996) and (Parsons and Jennings, 1996), strength is indicated explicitly; §5 details how a text generation system might also make use of the distinction, proposed by Freeman (1991), between inferential and persuasive force. Lastly, that subarguments should not be extended too far, nor be too great in number. This is a particularly important fact in human communication, and one that is particularly difficult to deal with formally, where the aim is usually to employ precisely all the subarguments available to lend maximum support to a conclusion[8]. For this last rule builds upon the intuition that, in persuasive communication at least, more can in fact be less: additional subarguments may *weaken* a position rather than bolstering it[9].

Whately's list (1855) is a little longer, and encompasses a slightly different range of variables. Firstly, the ordering is crucially affected by whether the argument is setting out to convince an audience which is already aware of the subject matter, or whether to inform an audience which is not: this is called the *primary aim* in (Reed *et al*, 1996b), and its effects on a computational system are discussed there. Secondly, arguments from cause to effect should precede others (Whately (p34) provides an example: if an honest man were accused of corruption, the evidence is far more likely to be listened to if it is first shown that he may be greedy). The computational implications of this rule are similar to those of Blair's first rule, in that some means of tagging argumentation may be necessary, though in this case, identification of a subargument from cause to effect may be possible on purely structural grounds (if, for example, causality is a distinguished form of support). Thirdly, Whately claims that "Refutations of Objections should generally be placed in the midst of other Arguments, but nearer the beginning than the end" (p38). For, as mentioned above, although an argument is not based entirely on defeating anticipated counter-arguments, the technique should not be eschewed altogether. The same sentiments are echoed by Blair[10] as well as by more recent texts on informal logic[11]. Whately goes on to give two further rules which present a particular problem for a system which is to generate natural language arguments: that well known propositions

---

[8] The system of Parsons and Jennings (1996) system clearly does not suffer from this problem - but only because arguments are exchanged one at a time, at each dialogical turn.

[9] For example, an additional weak subargument may give room for an opponent to build a 'thin-end-of-the-wedge' or 'straw man' counter-argument.

[10] "Every speaker should place himself in the situation of a hearer, and think how he would be affected by those reasons, which he purposes to employ for persuading others" (Blair, 1838, pp429-430)

[11] Eg. Fogelin (1991, p41) discusses the use of "*Discounting*: anticipating criticisms and dismissing them".

should be stated at once, and that a recapitulation of the main points of an argument should occur in reverse order. Both of these techniques entail telling the hearer *something he already knows*, which runs contrary to the intuitions implemented in many natural language generation systems - a point made by Marcu (1996) in his psycholinguistic analysis of what makes a text persuasive. Blair, too, relies upon the fact that repetition is required, when he states (p440) that the conclusion should contain no new subject matter.

This last point makes particularly salient a common criticism of such rhetoric texts - Richards (quoted in (Billing, 1996, p90)) complains that they offer nothing more than

> "prudential Rules about the best sorts of things to say in various argumentative situations ... we get the usual postcard's worth of crude commonsense :- be clear, yet don't be dry: be vivacious, use metaphors when they will be understood not otherwise ..."

However, the careful, exhaustive enumeration of all this commonsense knowledge is a requisite precursor for implementing a system which might be able to produce natural language persuasive discourse.

5. Approaches in Natural Language Generation

In natural language generation (NLG), the problems of multiple subarguments fall naturally into three categories: generation, inclusion and ordering, ie. how subarguments are invented (to borrow the term used by Blair (1838)), how the decisions are made over which to use, and finally, how they are arranged as a persuasive whole.

Ideally, an agent should be able to generate all possible arguments supporting a cause, and then select those which best suit the situation at hand. There are, however, several problems with such an approach. Firstly, intuitions suggest that this is psychologically implausible (despite the implications of Blair's dictum to "make a proper selection of [arguments] as appear to one's self the most solid", as quoted in §4). In addition, the idea of expending significant computational energy on generating arguments which are never used is unappealing. More importantly, for a resource-bounded agent (see, for example, (Bratman *et al*, 1988)), this exhaustive procedure may simply be too costly, and would certainly not make optimal use of the limited resources available.

In practice, though, the reasoning by which arguments are produced is not properly a problem of NLG (in much the same way that Blair claims that the invention of arguments does not come within the purview of rhetoric). Although the boundary is a little unclear, it is important to distinguish the means by which the content of an argument is determined (eg. through the use of *reasoning agents*, (McConachy and Zukerman, 1996)) from the way it is subsequently structured and expressed. Part of the reason that the boundary is unclear is that the processes of invention and expression appear to be interleaved: it is unreasonable (and counter-intuitive) to assume that a full set of complete arguments is available to the NLG system. Rather, for any given belief there are supports which represent final key steps in chains of reasoning, and as the argument progresses, particular chains of support are pursued at greater length, depending upon the strength of the links, and on the hearer's knowledge and sympathy in the areas involved.

Given, then, the existence of argument components in some knowledge base (stored as 'potential', rather than complete, argument units), the first task is to decide which of them to include. As a planning task, this represents a particular problem, due to the nature of the goal to be fulfilled. In standard hierarchical planning, such as that implemented in NOAH (Sacerdoti, 1977), and widely used in discourse planning, (Hovy, 1993), the task is expressed in terms of a number of goals which are then met by applying appropriate operators. In (Reed *et al*, 1996a), for example, the goal BEL(H, P) (that the hearer believe some proposition, P) might be fulfilled by application of the Modus Ponens operator which lists BEL(H, P) on its postcondition list. However, once such an *achievement* goal is fulfilled, it is considered 'finished', such that no further planning is required to support it.

There have been some attempts at provision for a notion of *maintenance* goals. Hovy's (1990) system PAULINE, for example, could express stylistic goals which had continuing effect throughout the planning process. Goals which might give rise to multiple subarguments, however, are neither achievement nor maintenance in nature: they are planned for several times and then considered fulfilled in the classical sense. Several systems have implemented goals of this sort, through the use of some 'for all' function in operator definitions, such as Moore and Paris' (1994) FORALL clause, and Maybury's (1993) explicit use of the $\forall$ symbol: in both these cases, the authors have noted that the clauses have required explicit, unprincipled modification of the plan language.

Even if a rigorous planning foundation were available for such goals, NLG would continue to suffer from a related problem, that with a rich knowledge base available, a NLG system must be able to determine an appropriate level of detail. In argumentation, this becomes even more difficult since, as mentioned above, it is often useful to tell the hearer things he already knows (Marcu, 1996). One partial solution to this problem lies in the notion of coherency (discussed in this regard in more detail in (Reed and Long, 1997)). Including complex or large subarguments is detrimental to the resultant coherency of the argument as a whole. Similarly, every additional subargument employed decreases the overall coherency. The means by which this technique may contribute to a solution is discussed in the next section.

One of the most important factors determining whether or not a subargument is included in an argument is also the primary means of effecting an appropriate ordering between subarguments, namely, the strength of a subargument. Although many formal (and informal) approaches recognise the concept of argument 'strength', few make the distinction between *inferential force* and *persuasive force*, due to Freeman (1991):

> "The persuasive force of an argument is its ability to move an (intended) audience. Inferential strength is a completely normative issue. How well does this inference satisfy the canons of deductive or inductive logic?", p243

Formal systems - (Das *et al*, 1996), (Parsons, 1996), etc. - whether qualitative or quantitative, express only the inferential force. Whereas rhetorical maxims are based almost exclusively on persuasive force. As a consequence, NLG systems such as that of Maybury (1993) have generally been unable to implement persuasion related heuristics. One notable exception was Sycara's (1989) system, PERSUADER, which had a limited notion of persuasive strength, fixed as a rigid, domain-specific hierarchy between her nine argumentation techniques. As a result, PERSUADER was able to perform some degree of choice over the inclusion and ordering of subarguments. However, in order to flexibly plan persuasive arguments, the distinction needs to be made and implemented explicitly.

6. Towards a Solution

The task of generating natural language argument is being approached through the use of the hierarchical framework proposed in (Reed *et al*, 1996b), in which the planning process is conceptually divided into four layers of abstraction. At the highest, *Argument Structure* (AS) level, the structural form of the argument is produced; below this, the *Eloquence Generation* (EG) level performs rhetorical and stylistic refinement; at the next level, the interclausal structure is refined through the use of Rhetorical Structure Theory (Mann and Thompson, 1986); finally, linguistic form is produced at the lowest levels of syntactic and morphological realisation. The problem of the 'generation gap' (Meteer, 1993) is minimised through the underlying use of LOLITA, a large scale, domain-independent, natural language system (Smith *et al*, 1994). The distinction between AS and EG levels mirrors the rhetorical distinction between arrangement and expression (and as mentioned above, both these phases succeed that of invention).

The planning task is carried out by AbNLP (Fox and Long, 1995), which makes use of encapsulation (such that an operator body is completely hidden until the abstract plan is completed) and refinement (which opens up all the operator bodies on completion of the abstract plan). These techniques both correspond closely to the structure of natural argument, and also lead to significant computational savings (Baccus and Yang, 1992). Investigation is under way of an optimal means of implementing the 'forall' problem discussed in the previous section (though it is anticipated that the solution will be similar to that of (Moore and Paris, 1994)).

To deal with the inherent uncertainty of the communicative situation (both in the speaker's own beliefs, and in her beliefs of what her audience believes), a standard approach is taken, such as that proposed by Parsons (1996). However, two important caveats are noted. Firstly, there is no assumption of 'distributed' defeasibility (a concept introduced in §2). That is, agents do not construct their arguments (solely) on the basis of anticipated counter-arguments. This approach to uncertainty is thus similar to that of (Parsons and Jennings, 1996), except that an agent is capable of constructing complex arguments, composed of multiple subarguments. In the second place, it is recognised that the representation is limited to detailing only inferential force, and that the notion of persuasive force needs to be handled explicitly and separately.

Freeman (1991) suggests that persuasive and inferential force are quite distinct and unrelated facets of an argument. However, the claim seems to be rather too strong, certainly for NLG, where the assessment of persuasive force will need to include reference to the inferential force (an utterly illogical argument, for example, is often far less likely to convince an audience). Indeed, in a situation in which the speaker knows nothing about a hearer other than that he is rational, the persuasive force of an argument may very well be the *same* as the inferential force. Clearly, though, persuasive force is normally determined by more than just the inferential strength due only to the structure of an argument. As discussed above, an evaluation of persuasive force is required in order to make various rhetorical decisions over which subarguments to include, and how then to determine an effective ordering between them. Such an evaluation can be conducted by considering both the reception an argument would receive, on the basis of the model of the hearer's beliefs, in addition to various factors which would affect that reception, including hearer bias, scepticism, competence, etc.[12]. Importantly, this process is not equivalent to the counter-intuitive distributed defeasibility which builds arguments on the basis of anticipated hearer counter-arguments. Rather, it gives rise to precisely that functionality suggested by intuition, namely, that arguments can involve refutations of rebuttals, but that those refutations are included (or not) and ordered on the basis of their persuasive force. In many cases, (as expected), such refutations will be an effective means of persuasion, and will play a primary role in an argument as a consequence. But, there will also be instances when other factors (such as high levels of hearer scepticism, or particularly low confidence in aspects of the hearer model) drastically reduce effectiveness of counter-counter-arguments, resulting in their marginalization.

The AS level can thus be seen as mediating between conflicting pressures. On the one hand, coherency constraints (Reed and Long, 1997) aim to minimise complexity, reduce the number of subarguments and avoid deeply nested argumentation. At the same time, structural planning is producing additional complexity; this planning includes persuasion related aspects such as informing the hearer of propositions he already believes and employing fallacies (a technique which is similar to McConachy and Zukerman's (1996) *licentious* mode, and which extends Sycara's (1988) use of fallacies, by granting them the same status as the deductive and inductive operators, thus avoiding the necessity of placing them in a fixed preference hierarchy). This conceptual process of mediation is the primary means of tackling the NLG detail problem, and implementing the associated rhetorical rules.

7. Conclusion

The phenomenon of disjunct multiple subarguments, occurring in many systems of argumentation, present a number of problems: how to reason defeasibly in a multi-agent society without eschewing the intuitions of argument construction; how to distinguish between linked and independent subargument support; how to deal with rhetorical maxims which depend on a notion of persuasive force which is primarily empirical; how to determine an appropriate amount of detail whilst generating an argument in natural language.

This paper has presented a synthesis of these ideas, which has been used as

---

[12] These factors are crucial to the system, at both AS and EG levels, as discussed in (Reed *et al*, 1996b)

the basis for proposing an approach to the generation of natural language argument. As a consequence, it is hoped that the approach benefits from the findings of several disparate fields, whilst avoiding trivialising the problems in any one. The solution put forward is not complete. The conjunction operator used to combine linked subarguments, for example, is currently rather naive and does not fully capture the intuitions discussed by Freeman. The problem of a principled approach to the 'forall' operator required to plan for multiple subarguments remains unresolved. However, a framework has been described which includes a number of important and novel aspects: a clear distinction between disjunctive and conjunctive support; a means of implementing rhetorical heuristics through a definition of 'strength'; a distinction between persuasive and inferential strength; a means of determining an appropriate level of detail. But more importantly, this paper has attempted to clearly identify the problems which have a bearing on the issue of natural language argument generation, and to motivate the continuing research on an interdisciplinary basis.

References

Baccus F. & Yang Q. "The expected value of hierarchical problem-solving", in *Proceedings of the National Conference on AI (AAAI'92)* (1992)

Billig, M., *Arguing and thinking* (Second edition), Cambridge University Press (1996)

Blair, H. *Lectures on Rhetoric and Belles Lettres*, Charles Daly, London (1838)

Bratman, M.E., Israel, D.J., Pollack, M.E., "Plans and resource-bounded practical reasoning", *Computational Intelligence* **4** (1988) 349-355

Cohen, R., "Analyzing the Structure of Argumentative Discourse", *Computational Linguistics* **13** (1) (1987) 11-24

Das, S., Fox, J. & Krause, P., "A Unified Framework for Hypothetical and Practical Reasoning (1): Theoretical Foundations", in Gabbay, D. & Ohlbach, H.J. *Practical Reasoning*, Springer Verlag (1996) 58-72

Dung, P.M., "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and *n*-person games", *Artificial Intelligence* **77** (1995) 321-357

Fisher, A., *The Logic of Real Arguments*, Cambridge University Press (1988)

Fogelin, R.J. & Sinnott-Armstrong, W., *Understanding Arguments* (Fourth Edition), Harcourt Brace Jovanovich College Publishers (1991)

Fox, J. & Das, S., "A Unified Framework for Hypothetical and Practical Reasoning (2): Lessons from Medical Applications", in Gabbay, D. & Ohlbach, H.J. *Practical Reasoning*, Springer Verlag (1996) 73-92

Fox, M. & Long, D.P. "Hierarchical Planning using Abstraction", *IEE Proceedings on Control Theory and Applications* **142** (3) (1995)

Freeman, J.B., *Dialectics and the Macrostructure of Arguments*, Foris (1991)

Hovy, E.H., "Pragmatics and Natural Language Generation", *Artificial Intelligence* **43** (1991) 153-197

Hovy, E. H. "Automated Discourse Generation Using Discourse Structure Relations", *Artificial Intelligence* **63** (1993) 341-385

Krause, P., Ambler, S., Elvang-Gøransson, M. & Fox, J., "A Logic of Argumentation for Reasoning under Uncertainty", *Computational Intelligence* **11** (1995) 113-131

Mann, W.C. & Thompson, S.A. "Rhetorical structure theory: description and construction of text structures" in Kempen, G., *Natural Language Generation: New Results in AI, Psychology and Linguistics*, Kluwer (1986) 279-300

Marcu, D., “The Conceptual and Linguistic Facets of Persuasive Arguments”, in *Working Notes of the ECAI'96 Workshop on Planning and NLG* (1996) 43-46

Maybury, M.T., “Communicative Acts for Generating Natural Language Arguments”, in *Proceedings of the 11th National Conference on AI (AAAI’93)* (1993) 357-364

McConachy, R. & Zukerman, I., “Using Argument Graphs to Generate Arguments”, in *Proceedings of the 12th European Conference on AI (ECAI’96)* (1996) 592-596

Meteer, M., *Expressibility and the Problem of Efficient Text Planning*, Francis Pinter (1993)

Moore, J.D., Paris, C.L. “Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information”, *Computational Linguistics* **19** (4) (1994) 651-694

Parsons, S., “Defining Normative Systems for Qualitative Argumentation”, in Gabbay, D. & Ohlbach, H.J. *Practical Reasoning*, Springer Verlag (1996) 449-463

Parsons, S. & Jennings, N.R., “Negotiation Through Argumentation - a Preliminary Report”, in *Proceedings of the International Conf on Multi-Agent Systems (ICMAS’96)* (1996)

Pollock, J.L., *Cognitive Carpentry*, MIT Press (1995)

Pollock, J. L., “Implementing Defeasible Reasoning”, in *Working Notes of the FAPR'96 Workshop on Computational Dialectics* (1996)

Reed, C.A., Long, D.P., Fox, M. & Garagnani, M., “Persuasion as a Form of Inter-Agent Negotiation”, in *Proceedings of the 2nd Australian Workshop on DAI* (1996a, to appear)

Reed, C.A., Long, D.P. & Fox, M. “An Architecture for Argumentative Dialogue Planning”, in Gabbay, D. & Ohlbach, H.J. *Practical Reasoning*, Springer Verlag (1996b) 555-566

Reed, C.A. & Long, D.P., “Ordering and Focusing in an Architecture for Persuasive Discourse”, under review for *European Workshop on Natural Language Generation (EWNLG’97)* (1997)

Sacerdoti, E.D. *A structure for plans and behaviour*, Elsevier, North Holland (1977)

Sillince, J.A.A & Minors, R.H., “Argumentation, Self-Inconsistency, and Multidimensional Argument Strength”, *Communication and Cognition* **25** (4) (1992) 325-338

Smith, M.H., Garigliano, R., Morgan, R.C. "Generation in the LOLITA system: An engineering approach" in *Proceedings of the 7th International Workshop on NLG*, Kennebunkport, Maine (1994)

Sycara, K. “Argumentation: Planning Other Agent’s Plans” in *Proceedings of the International Joint Conference on AI (IJCAI'89)* (1989) 517-523

Toulmin, S.E., *The Uses of Argument*, Cambridge University Press (1958)

Vreeswijk, G., “Reasoning with Defeasible Arguments”, in Wagner, G. & Pearce, D. (eds) *Proc. of the European Workshop on Logics in AI (JELIA’92)*, Springer Verlag (1992) 189-211

Walton, D.N. & Krabbe, E.C.W., *Commitment in Dialogue*, State University of New York Press (1995)

Whately, R. *Logic*, Richard Griffin, London (1855)

Wooldridge, M. & Jennings, N.R., “Intelligent Agents: Theory and Practice”, *The Knowledge Engineering Review* **10** (2) (1995) 115-152

Yanal, R.J., “‘Convergent’ and ‘Linked’ Reasons.”, *APA Newsletter on Teaching Philosophy*, **4** (5), (1984)1-3