An Architecture for Argumentative Dialogue Planning*

Chris Reed[1], Derek Long[2] and Maria Fox[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
M.Fox@dur.ac.uk

Topics: Argumentation theory; Fallacies and their role in practical reasoning

Abstract. Argument represents an opportunity for a system to convince a possibly sceptical or resistant audience of the veracity of its own beliefs. This ability is a vital component of rich communication, facilitating explanation, instruction, cooperation and conflict resolution. In this paper, a proposal is presented for the architecture of a system capable of constructing arguments. The design of the architecture has made use of the wealth of naturally occurring argument, which, unlike much natural language, is particularly suited to analysis due to its clear aims and structure. The proposed framework is based upon a core hierarchical planner conceptually split into four levels of processing, the highest being responsible for abstract, intentional and pragmatic guidance, and the lowest handling realisation into natural language. The higher levels will have control over not just the logical form of the argument, but also over matters of style and rhetoric, in order to produce as cogent and convincing an argument as possible.

1. Introduction

In this paper an architecture is presented for the autonomous construction of argument, outlining the components required for persuasive communication. Argument plays a crucial role for a communicative system as a means of persuading the human interlocutor to adopt some belief which is of interest to the system. This might be simply to improve the coherence between their beliefs or to obtain a specific effect such as an adoption of some new goal by the hearer, elicitation of particular information, or instruction upon some particular topic.

It is assumed throughout that the argument will ultimately be rendered in natural language, but this assumption does not form a cornerstone of the theory: the framework employs intentional structures which are sufficiently abstract to remain unchanged by a shift to some more simple and restrictive artificial language.

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

After a brief introduction to argument as a form of natural language, the overview of the proposed architecture in §3 is followed by more detailed analyses of the five major components.

2. Argument

In order to model argument and design a rubric for its automatic synthesis, it is necessary to have a clear understanding of its nature. Unlike much natural discourse, arguments always have clearly defined goals, and in particular, to persuade a particular audience of a particular proposition. This perhaps oversimplifies the issue, since rarely will an argument be instigated between two parties who believe thesis and antithesis, and even more rarely does argument terminate with both parties believing the same proposition. This fuzziness concerning the beliefs of the participants is much more an intrinsic property of belief, rather than one of argument, since the actual process of argumentation can more accurately be perceived as the facilitation to the adoption of a belief on the part of the hearer. The process itself is the same, regardless of the myriad situations in which it can occur - public oration, debating house discourse, legal argument, newspaper commentary, scientific papers, etc.

3. Components of an Argumentative System

The architecture proposed in this paper rests in a hierarchical framework, reflecting the distinct though inter-related levels of structure within arguments identified through their analysis. There is a part of argument synthesis which is concerned with the resolution of syntax, expression and morphology, as with any natural language synthesis. This comprises the lowest level. Above this, there sits a level responsible for the coordination of sentence and intersentence relations, imposing structure at a more abstract level. This intermediate level is also responsible for handling aspects of communication such as focusing (Grosz and Sidner 1990). Finally, at the higher pragmatic levels, new machinery and techniques are required to produce cogent argument. This generation can be accomplished through two complimentary levels, one of which handles the structural aspects of argument, and the other effecting complex modifications to that structure and controlling stylistic devices. The former, the Argument Structure (AS) level, is at the highest level of abstraction, since it produces the logical form of the argument employing purely intentional data structures, by using logical operators, fallacy[1] operators and inductive operators. This form is then augmented and modified by the subordinate Eloquence Generation (EG) level, which employs heuristics based upon rhetoric and contextual parameters, and is concerned with such properties of a speech as its length, detail, meter, ordering of subarguments, grouping of subarguments,

[1]Lists of informal fallacies are common, though rarely identical: see, for example, Whately (1855) or (Johnson 1992).

enthymeme contraction, use of repetition, alliteration and so on. This architecture is summarised below in Fig.1.

Note that the 'intermediate level' has, in the figure, been labelled 'RST level': the functionality proposed for this level is in line with much of the work carried out on Rhetorical Structure Theory (Mann and Thompson 1986). (The inclusion of RST has also led to the Eloquence Generation level being so named: the more appropriate 'Rhetoric level' is open to confusion with RST).

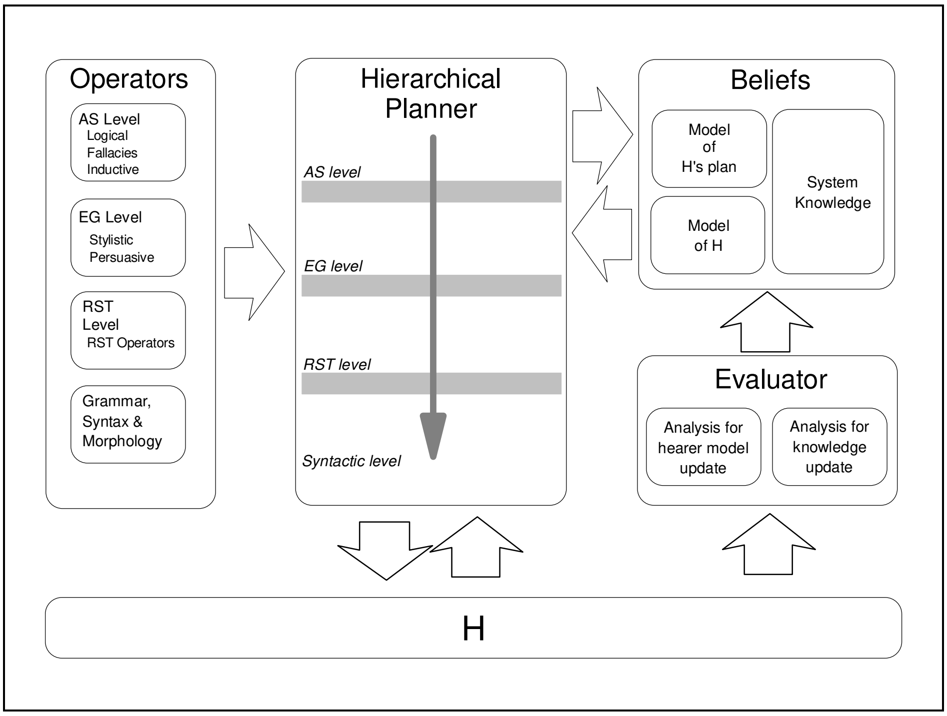

Fig. 1. System architecture overview

4. Supporting Components

4.1. Belief modelling

It has become clear that both AS and EG levels require access to the belief model of the audience, in addition, of course, to the beliefs of the system. The AS level needs to be able to assess in advance how different structures are likely to be received, and to take account of the beliefs of the audience in producing structure. The EG level employs the belief model to pitch an argument at the right level of detail, and to track the saliency of beliefs used during argumentation. A more detailed examination of the relationship between the higher levels and belief modelling is described below in §4.3 and §4.4.

There are a number of issues involved in analysing belief. There is more than one kind of belief - factual beliefs, which are either testable (at least in theory) or are definitional; opinions, which are based on morals and are thus ultimately personal and unprovable (since there is - and can be - no universally accepted and provably correct moral framework); and cultural beliefs, which are based upon sociocultural maxims (such as that which states that living into old age is desirable). This tripartite distinction has been based upon the views expressed implicitly by Blair (1838), when talking about general types of argument, but there are also numerous other divisions which could be detailed (see Ackermann (1972) for a number of examples). The way in which beliefs become manifest in argumentative dialogue - both their individual expression, and their interrelation - is chiefly dependent upon their class, so competent representation and recognition of these distinct belief types represents an important problem.

Another major problem is how to resolve two seemingly contradictory views generated by introspection, that of whether beliefs are best represented as dichotomous or scalar. A pragmatic resolution to this issue is in itself crucial to competent belief modelling, but it will also affect the approach taken to other problems (factual beliefs are usually dichotomous, whereas opinions and cultural beliefs are generally scalar, for example). The scalar nature of belief in particular is crucial to pin down precisely, for it is implicit in argumentation itself - arguments rely upon concepts such as 'persuasiveness' and 'strength of argument', and the very fact that an argument will often involve several separate and identifiably distinct component sub-arguments, suggests that the process involves some scalar value of 'mounting evidence', 'increasing conviction', or 'cogent context'. These and other problems associated with the concept of 'strength of belief' are discussed more fully in Galliers (1992).

Another property of belief which is manifest in argument is that of saliency. Enthymeme contraction (where an unnecessary premise or conclusion is omitted) relies heavily upon saliency and so too does the process of focusing: keeping the argument to the point (Cf. the lower level focusing of Grosz and Sidner (1990)). The belief model used must therefore be able to competently handle the concept of saliency.

Argument modelling also relies on a treatment of the complex phenomenon of mutual belief: an argument is based upon common ground - a set of mutual beliefs (ie. which both parties hold, and which both parties also know the other to hold). Mutual belief is defined in terms of an infinite regress of nested beliefs: the problem is to pragmatically choose a level of nesting beyond which 'mutual' belief is to be assumed. In making this choice, it is understood that no matter how many levels a system can cope with, it is always possible to construct a (highly convoluted) example which exceeds the capabilities of that system. From a psychological (and intuitive) point of view, choosing some arbitrary level of nesting by which to define mutuality seems rather implausible. In humans, it would appear that belief nesting is a resource bounded operation with no known limit, and it is possible to construct deeply nested examples which present remarkably little difficulty (such as the example in Fig. 2, below), though handling further complexity rapidly becomes extremely difficult and time consuming. It may be possible to utilise this evidence and allow for a similar process in an implementation, such that some fundamental operator (say, BMB, following Cohen and Levesque (1990)) is assumed at a naively shallow level of nesting, and then to allow the operator to be quite corrigible in the light of new evidence.

In the classic Hitchcock film "North by Northwest", Cary Grant is the central character in a case of mistaken identity - he is mistaken for a goodguy agent by the evil James Mason. Towards the end, a scene occurs in which CG masquerades as the goodguy agent he is thought to be. As part of a goodguy ploy, he informs JM of his wish to defect, at which he is shot by EveMarie Saint, a goodguy agent working undercover as accomplice to JM. In her role as a villain, she needs to stop the defection, believing him to be a traitor to the goodguys. Thus with the proposition P that 'ES is a goodguy',

`BEL(ES, P)`	She knows she's a goodguy
`BEL(CG, P)`	He knows she's a goodguy ...
`BEL(ES, BEL(CG, ~P))`	... but she doesn't realise that ...
`BEL(CG, BEL(ES, BEL(CG, ~P)))`	... and he knows she doesn't realise ...
`BEL(Audience, BEL(CG, BEL(ES, BEL(CG, ~P))))`	... or at least that's what the audience thinks!

After JM has left however, CG gets up - the shooting had been faked. CG and ES must have been in league with each other after all. At this point the belief held must be changed to

`BEL(Audience, BMB(CG, ES, P))))`	They both know she's a goodguy, and both know that they both know it.

Hitchcock envisaged the whole situation and realised that it was unusual and interesting. There were in total, five levels of nested beliefs.

Fig. 2. Deep nesting of beliefs

4.2. Hierarchical planning

The role of planning in the construction of arguments is significant: arguments are constructed to achieve goals and are built out of primitive argument structures. this construction process includes consideration of the effects of the various components on the hearer's beliefs, the flow of the argument, its focus and force and the extent to which each part is supported by what precedes it. These issues are all interpretations, in the argument synthesis domain, of traditional planning problems, so it is no surprise that traditional planning has a role to play in this domain.

A considerable quantity of work already exists on the use of planners in the construction of monologue (Hovy (1993) contains a thorough overview). This work rests almost exclusively on planners employing a form of abstraction-based planning based on NOAH, a planner developed by Sacerdoti (1977) which is, by the standards of modern planning technology, rather dated. It is proposed that a more advanced style of abstraction-based hierarchical planning technology be used to support the architecture discussed in this paper, employing encapsulation-based abstraction (Fox and Long 1995) to exploit an intuitive structuring of the argument-planning domain. This approach involves the use of abstract operators which encapsulate bodies containing goals, the achievement of which sets the context for and then brings about the effects at a detailed level which are described at an abstract level in the abstract operator shells. The refinement of abstract operators into their more detailed bodies is carried out as an explicit and controlled stage in the planning process, rather than (as with the looser and less-powerful abstraction mechanism of NOAH) as the operators are applied. This allows a genuine hiding of detail until it is appropriate for the planner to consider it.

4.3. Argument Structure

Through an analysis of a corpus of arguments drawn from a number of the sources mentioned in §2, the following structure has been identified:

(1) An argument consists of one or more premises and exactly one conclusion

(2) A premise can be a subargument (which itself consists of one or more premises and exactly one conclusion: the conclusion then stands as the premise in the superargument)

(3) A subargument is an integral unit whose components cannot be referred to from elsewhere, nor can the conclusion of a subargument rest upon premises extraneous to that subargument

(4) The only exception to (3) is where a conclusion in a distant subargument is restated locally as a premise

Analyses based on similar theories are performed in many texts: see for example, Johnson (1992) and Wilson (1980). It is this structure which the AS level constructs, linking premises to conclusions through the use of three groups of operators: standard logical relations, rhetorical fallacies, and inductively reasoned implications.

The first group comprises Modus Ponens, Modus Tollens, Conjunction, Disjunctive Syllogism, Hypothetical Syllogism and Constructive Dilemma (MP being by far the most common). The second group comprises around two dozen fallacies which, though illogical, are very common in natural language. Some are, without a doubt, tricks of the rhetorician, designed to mislead the audience (Red Herring and Straw Man, for example), whilst most are used quite inadvertently, but nevertheless assist in strengthening an argument (at least in the perceptions of some audiences) without recourse to factual support. To implement the use of fallacies could be seen as a violation of Grice's first maxim for successful communication, the maxim of quality, which demands sincerity and truthfulness (if the system employs a rhetorical trick, it is in some ways 'cheating'), but since humans generally utilise fallacies unwittingly, there would appear to be no lack of sincerity unless there is devious intent behind their application. Their use would be heavily constrained by tight preconditions specifying contextual factors in addition to both system knowledge and assumed audience beliefs. (Incidentally, if the preconditions were to include some analysis of the hearer's susceptibility to the fallacy, and its consequent success, it could then be argued that there is some measure of devious intent).

Finally, there are the inductive operators (inductive strength being defined as "...more probable than not that the conclusion follows [from the premises]."

(Johnson 1992)). The inductive operators are of three types: inductive generalisation, causal (based on Mill's methods) and analogical (eg. Long and Garigliano (1994)). All three will have preconditions that are again tightly specified for context and belief, but with the additional constraint of the 'criteria of inductive strength', ie. the requirements for their application to produce an inductively strong argument.

4.4. Eloquence Generation

Although the EG level also employs a number of operators, the bulk of its functionality is based on the application of heuristics. The first task of the EG level is to control stylistic presentation, comprising a number of relatively low level 'tweaks' such as the vocabulary range, syntactic construction (eg. the ratio of active to passive constructions) and the frequency of specific devices such as alliteration, repetition, metaphor and analogy. There are a large number of such factors detailed in the stylistic literature - see Sandell (1977) for an overview. To date, such factors have been controlled by quite artificial and simplistic means, typically, the user setting a number of variables to particular values (eg. Vreeswijk (1995)). For the instantiations to be made and altered 'on-the-fly', the EG level must refer to a body of parameters, in addition to the belief model of the audience and context, all of which are modified dynamically.

One vitally important parameter affecting the argument is the relationship which the speaker wishes to create or maintain with the hearer. This relationship is established through stylistic rather than structural means, and is not necessarily divorced from other aims: if a hearer accepts the speaker's authoritative stance, for example, the speaker may be able to use the relationship to reinforce his statements. Attempts at instigating different relationships between speaker and hearer account (at least in part) for a number of complex phenomena - humour, for example, is frequently used to establish the speaker as a friend (and therefore reliable and truthful). Such phenomena are not intended to fall inside the scope of this work.

The technical and general competence of the hearer are also important parameters to be considered at the outset. General competence determines the hearer's ability to understand complex argumentation (and to some extent, complex grammar); technical competence enables the argument to pitched at the right level, and affects the choice of appropriate vocabulary.

The initial goals of the argument are also best viewed as a parameter: as mentioned above, the primary goal in different situations (debating house, soap box, law courts, etc.) is subtly different: convincing an opponent is probably the most usual goal, but arguments are also used to impress, to instruct, to confound, to counter, to deceive and to provoke. Importantly, there are often several concurrent goals, such that a number of goal-dependent heuristics could be active at any one stage. Many of the EG level's suggestions are affected - vocabulary, grammar, gross structure and content all depend, to some extent, upon the goals of the system.

A number of other parameters have also been identified, including the speaker's potential gain from 'winning' the argument (in addition to the speaker's gain as perceived by the hearer), the speaker's investment, or potential losses, suffered in losing the argument, the medium in which the argument is expressed and aspects of the speaker's model of the hearer's beliefs such as possible scepticism or bias. For example, advertisements usually have the aim of convincing the hearer to buy the product. In addition though, there are aims connected with the product identity, the particular advertising campaign, and often the speaker-hearer relationship. An attempt is frequently made to hide or play down the speaker's involvement and gain, whilst very careful analysis is made of the audience at which the advertisement is aimed, and the consequent weaknesses and beliefs of that audience.

Some of the EG heuristics are much less dependent upon these parameters: those aiding in the premise-conclusion ordering, for example, rely almost entirely upon general principles (though do also make use of factors such as the general tone of the argument). There are three possible statement orderings in an argument: Conclusion-First (C P*), Conclusion-Last (P* C), and Conclusion-Sandwich (P* C P*). The first is usually used where the P* are examples (in the case of factual conclusions: for opinions, the P* would often be analogies), especially when the hearer is being led to a hasty generalisation from the P* to the C. Conclusion-First is also used when the initial conclusion is deliberately provocative - the construction being used to draw attention to the argument (and consequently it is unusual to find a weak argument structured with the Conclusion-First ordering). It can also be forced should a premise not be accepted by the hearer, and require a sub-argument in its support. Conclusion-Last is the usual choice for longer, more complex or less convincing arguments (indeed some arguments are less convincing precisely because they are longer or more complex). It is also used for 'thin end of the wedge' argument and for grouping together premises which individually lend only very weak support to the conclusion. The Conclusion-Sandwich construction is rarer, usually occurring when the speaker has completed an Conclusion-Last argument which has not been accepted and therefore requires further support.

Finally, there are a number of features controlled by the EG level which lie somewhere between the stylistic and the rhetoric: well known artifices such as repetition and alliteration which can prove extremely effective in constructing eloquent and compelling argument. It is interesting to note that almost all of EG heuristics are devices listed by classical texts as figures and tropes. Indeed this is a uniquely interesting property of the EG level: the emphasis it puts in utilisation of ideas posited in precomputational treatises, especially the classical texts of Cicero and Aristotle (Freese (1926), for example) and the ideas developed in the late middle ages and renaissance (Whately (1855) and Blair (1838) have been used as source material for much of the analysis). This fact poses unique challenges as well as affording unique advantages: on the one hand, the ideas are clear and unbiased, whilst on the other, may be difficult to transcribe to implementation.

One of the key functions of both the AS and EG levels is to maintain a representation of the intention behind utterances. RST has been criticised for its restrictive handling of intention (Moore and Pollack 1992), (Young and Moore

1994), an undeniably vital ingredient in discourse. Thus the AS and EG levels have to manipulate rather different data structures from those employed at lower levels (which would appear quite reasonable, given the difference in data structures at the RST level compared with lower levels). Due to these important differences, the current work is concentrating specifically on the higher, more abstract functionality. The remaining, lower level functionality required to realise the abstract intentional structures into natural language utterances will be provided by a system such as LOLITA (Smith et al. 1994), a large scale, domain independent natural language system, in which a natural language generation algorithm is already implemented which subsumes responsibility for solving certain low-level text generation planning problems.

4.5. Interaction

It is intended that the system should be fully interactive, and enter in to dynamic natural language dialogue with a human interlocutor. There are a number of issues specific to this goal, most obviously, the differences between monologue and dialogue. Both monologue and dialogue require belief modelling of the audience (at the very least for argument, and quite possibly for all natural language); during production of a monologue however, no opportunity is presented for checking, revising and refining its audience model, which makes extended arguments simpler to construct (with no need to modify the plan according to new information regarding the beliefs of the hearer), but as a consequence, rather less likely to succeed at convincing a hearer of the proposition in hand. That is not to say that monologue is of no interest: quite to the contrary, since Debating House argument could be seen as the monologues of the two opposing interlocutors as opposed to true dialogue. Work on monologue-formed argument would probably be easier to tackle in the short term, but it ignores many critical aspects of modelling which are of longer term interest, in particular, in monologue the classical planning assumptions of perfect knowledge of effects and environment state can be employed, while with dialogue it is necessary to consider the problems of uncertainty, imperfect knowledge (and the need to acquire knowledge) and plan failure.

During true dialogue, then, competent belief revision is essential. This revision process is nontrivial and forms an important part of the functioning of the AS and EG levels, drawing in particular from the work of Gardenfors (1992) and Galliers (1992) for implementation.

Dialogue also affords an opportunity for the application of argument to the areas in which it is most useful: negotiation and conflict resolution (where the interlocutors may not already be disposed to helping one another). The literature details much work on negotiatory communication (eg. Zlotkin and Rosenschein (1990), emphasise conflict resolution in non-cooperative domains, ie. those in which the negotiation set may be empty). Much of this work has assumed agent-agent communication, whereas in this paper it is proposed that a similar approach could be adopted in modelling natural language human-computer interaction. There are situations in which the user has differing goals to that of the system, especially where the system has predefined aims. For example, a system which provides an interface to an online information provider might have goals of persuading the user that it has the required information, that it is the cheapest or most reliable source, and, therefore, that the user should buy from it immediately.

Another benefit (or complexity) of dialogue is that communicative failure can be detected and recovery instigated (a vital ability if misunderstandings are to be avoided: see for example, Moore and Swartout (1991)). It would appear that failure can occur at any of the levels of processing, each level producing characteristic failures: at the lowest level, lexical ambiguities such as homonyms can induce failure. At the RST level too, ambiguities in the hearer's RST analysis of an utterance can lead to serious failure. The EG level's enthymemes could misfire (a conclusion may be obvious and omitted by the speaker, but may be non-obvious to the hearer who then misunderstands the argument). Finally, the structural operators of the AS level could lead to misunderstanding in one of two ways: either the hearer fails to see the applicability of an induction or fallacy in trying to prove the speaker's point (a naive failure), or else the hearer identifies and draws attention to the rhetoric trick (an informed failure).

There would seem to be two possible methods by which failure could be recognised and dealt with. Firstly, there is the intuitively appealing idea of having the recognition and recovery occur at the level at which the error belongs. For example, if the failure was syntactic, it would become evident at the syntactic level that communication had failed, and there would then be immediate recovery. This may at first sight seem appealing, but it does not appear to be either computationally or psychologically very plausible: computationally, it is very difficult to see how the syntactic level could possibly recognise that a communication failure had occurred, and psychologically, from analysing argument, it is clear that even something as trivial as homonymity can rupture an argument to the point of stepping outside the entire argument and performing a little meta-conversation to rectify the disparity in the definitions employed. So it would appear that regardless of the type of the error, recovery is carried out at the highest level of planning. This would seem to suggest that perhaps recognition only occurs here too. This is supported by the fact that the major means of discovering failure (checking for disparity between the intentional structure behind the hearer's utterances and the system's belief model of the hearer) is only available to the higher planning levels. It is also only the higher level which can check the intentional structure to decide whether or not a communicative failure needs to be replanned because it is an intentional, core part of the argument, rather than some extraneous side effect (this is an important distinction, recognised in the fields of both communication recovery and of planning - see, for example, McCoy (1986) and Young and Moore (1994), respectively).

The reactivity inherent to such an approach to dialogue forces a style of planning in which plan construction is interleaved with plan execution, following the work of, for example, Ambros-Ingerson and Steel (1990).

5. Conclusion

This paper has presented an outline of the architecture required for constructing extended arguments from the highest level of pragmatic, intention-rich goals to a string of utterances, using a core hierarchical planner employing structural and rhetorical operators and heuristics. The proposed architecture is the first clear framework which presents a computational view of the rhetorical means used by humans to create cogent argument.

Work is currently underway to characterise the parameters and operators involved in the early stages of argument synthesis, and to relate these to aspects of argument unique to dialogue, such as turn taking, checking moves, etc., building upon the work of Levin and Moore (1977).

The architecture proposed will facilitate implementation of an HCI system capable of dealing with situations in which its own goals conflict with those of the user, or in which persuading the user to adopt new beliefs is a nontrivial task, whose success cannot be guaranteed.

References

Ackermann, R.J. "Belief and Knowledge", Anchor, New York (1972)

Ambros-Ingerson, J.A., Steel, S. "Integrated Planning, Execution and Monitoring", in Readings in Planning, (1990) 735-740.

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

Cohen, P.R., Levesque, H.J. "Rational Interaction as the Basis for Communication" in Cohen, P.R., Morgan, J., Pollack, M.E., (eds), Intentions in Communication, MIT Press, Boston (1990) 221-255

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

Freese, J.H. (trans), Aristotle "The Art of Rhetoric", Heinmann, London (1926)

Galliers, J.R. "Autonomous belief revision and communication" in Gardenfors, P., (ed), Belief Revision, Cambridge University Press, Cambridge (1992) 220-246

Gardenfors, P. "Belief Revision: An Introduction" in Gardenfors, P., (eds), Belief Revision, Cambridge University Press, Cambridge (1992) 1-28

Grosz, B., Sidner, C.L "Plans for Discourse" in Cohen, P.R., Morgan, J. & Pollack, M.E., (eds), Intentions in Communication, MIT Press, Boston (1990) 418-444

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

Johnson, R.M. "A Logic Book", Wadsworth, Belmont CA (1992)

Levin, J.A & Moore, J.A. "Dialogue Games: Metacommunication Structures for Natural Language Interaction", Cognitive Science **1**, (1977) 395-420

Long, D., Garigliano, R. "Reasoning by Analogy and Causality: A Model and Application", Ellis Horwood (1994)

McCoy, K.F. "Contextual Effects on Responses to Misconceptions", in Kempen, G., (ed), Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics, Kluwer (1986) 43-54

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

Moore, J.D., Pollack, M.E. "A Problem for RST: The Need for Multi-Level Discourse Analysis", Computational Linguistics **18** (4), (1992) 537-544

Moore, J.D., Swartout, W.R. "A Reactive Approach to Explanation: Taking the User's Feedback into Account", in Paris, C.L., Swartout, W.R. & Mann, W.C. (eds), Natural Language Generation in AI & Computational Linguistics, Kluwer, (1991) 3-48

Sacerdoti, E.D. "A Structure for Plans and Behaviour", Elsevier, Amsterdam (1977)

Sandell, R. "Linguistic Style and Persuasion", Academic Press, London (1977)

Smith, M.H., Garigliano, R., Morgan, R.C. "Generation in the LOLITA system: An engineering approach" in Proceedings of the 7th International Workshop on Natural Language Generation (INLG'94), Kennebunkport, Maine (1994)

Vreeswijk, G. "IACAS: an implementation of Chisolm's Principles of Knowledge", Proceedings of the 2nd Dutch/German Workshop on Nonmonotonic Reasoning, Utrecht, Witteveen, C. et al. (eds) (1995) 225-234

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

Wilson, B.A. "The Anatomy of Argument", University Press of America, Washington (1980)

Young, R.M., Moore, J.D. "DPOCL: A principled approach to discourse planning" in Proceedings of the 7th International Workshop on Natural Language Generation, Kennebunkport, Maine, (1994) 13-20

Zlotkin, G. & Rosenchein, J.S. "Negotiation and Conflict Resolution in Non-Cooperative Domains", in Proceedings of the National Conference on AI (AAAI'90), (1990) 100-105