The Role of Saliency in Generating Natural Language Arguments

Chris Reed
Department of Applied Computing
University of Dundee
Dundee DD1 4HN
Scotland
Chris.Reed@computing.dundee.ac.uk

Abstract

Generating expressions which communicate information already known to the hearer, building enthymematic arguments, and characterising refutations all pose significant problems to traditional natural language generation techniques. After exploring these problems, an approach is proposed which through its employment of a notion of saliency handles them cleanly, and offers support for further features including clue word generation. It is argued that propositional salience and its interaction with intentional, attentional, epistemic and structural components of a text generation system have a key role to play in the design and realisation of persuasive text.

1 Motivation

It might be argued that the notion of saliency can be conflated with goals of discourse structure, or, in theories which handle intentionality separately, with intentional goals which are responsible for structure. In the former case, it could be maintained, including a proposition - or a relation between propositions - in a discourse plan carries with it an assumption that the proposition or relation will be realised directly into text and thereby made salient to the hearer. In the latter case, an intention such as BEL(H, P) (that the hearer H believe some proposition P) might be assumed to lead to the generation of textual structure which can be guaranteed to make P salient to H.

Such conflation of saliency with other components of a discourse loses a powerful distinction which can be demonstrated to be crucial in a number of concrete situations. Three generation problems are presented which rest upon an explicit handling of saliency.

The first is the enthymeme - a syllogistic argument with one component left implicit (i.e. with a premise or conclusion not realised into text). One of the most common forms of the enthymeme in natural language is the omission of the major premise from a Modus Ponens, as in, *You were born in Bermuda so you're a British citizen*, wherein the major premise, *if you were born in Bermuda then you're a British citizen* is left implicit. The arrangement is so common that it has led authors such as Sadock [1977] and Cohen [1987] to treat it as a separate argument structure altogether, labelling it the *Modus Brevis*. The reason for its frequency is clear from the foregoing example - to include the major premise would render a text hopelessly cumbersome and repetitive. In many cases it is impossible to utter the antecedent and consequent of a Modus Ponens without the implication between them being immediately obvious - or salient - to a hearer. In some cases, however, a speaker will decide that it is necessary to include the major premise: this may happen if some other component of the syllogism - say the conclusion - is being left implicit, as in *It's been raining, and rain makes the roads dangerously greasy*. More importantly, there are some situations where it is appropriate to make every component of a syllogism explicit. A good example is one which also demonstrates the importance of contextual features such as the hearer's bias. If a speaker has reason to believe that a hearer is particularly sceptical, she may decide to lay out each component of an argument explicitly - thus sacrificing brevity for the sake of clarity. Consider, for example, during a heated encounter between a Darwinian and a Creationist: *We see around us an intricate interplay of complex systems, and such intricacy requires a designer. So logically there must have been a designer.* In Modus Tollens, the default is reversed: normally, all three components of the syllogism are expressed (as in *If matter had always existed, there should be no radioactive elements left. The presence of uranium, etc. is scientific proof that matter has not always existed.*). Occasionally, however, the speaker may judge that one component can safely be omitted, e.g. *If it had been raining, the roads would be wet, but I assure you that it hasn't been raining for days.* There is a range of parameters and contextual features which impinge on the decision of whether or not to realise a proposition in order that it be salient (see, e.g., [Reed, 1998]), but the key point is that there are some cases where textual realisation is appropriate, and others where it is not. One approach to handling this decision is to include within an NLG system a module responsible for post-processing a given text plan and determining which components should be realised into text and which should not. Apart from problems of determining exactly how such a module should exert its influence, it introduces additional complexity unnecessarily. The planning machinery already available within many NLG systems is perfectly capable of performing the function, if saliency is treated as a goal. A given goal of saliency can then be fulfilled either by textual realisation, or by particular configurations of linguistic and extra-linguistic context.

The second problem which benefits from employing saliency goals is the generation of utterances which remind a hearer – i.e. which communicate something that the speaker believes the hearer already knows. Marcu [1996] explains that in the genre of persuasive text, “contrary to NLG wisdom, a system capable of generating persuasive text will also have to generate information that is known to the audience”. In the increasingly popular approach to NLG in which speaker’s intentions are fundamental (building on [Grosz and Sidner, 1986]), the only intention which can be fulfilled by the realisation of known information is one of saliency. If the speaker knows that the hearer believes P, then a typical intention of the form BEL(H, P) is trivially satisfied without the need for textual realisation.

Related to this is a third generation problem: refutation. It is often suggested in works on rhetoric and oration (e.g. [Blair, 1838]) that an argument which refutes a particular claim may, in certain situations, benefit from stating the claim which is to be refuted. There are two specialisations of this general rule which serve to illustrate the mechanism: first, the *reductio ad absurdum* argument strategy, whereby the reverse of the intended conclusion is assumed and then shown to lead to a contradiction; and second, the use of the *null hypothesis* in scientific experimentation, whereby the reverse of the anticipated conclusion is proven to be (significantly) wrong. In both these cases, it is necessary to realise a proposition that the speaker wishes to bring the hearer to **dis**believe. Again, within a planning framework based on intentions, there do not appear to be any alternatives to taking an approach employing goals of saliency, as there are no other perlocutionary effects of such discourse actions.

2 Modelling Saliency

The *Rhetorica* system implements an abstraction based nonlinear planner, based on AbNLP [Fox and Long, 1995]. The planning operators in the current version of the system represent forms of argument, including standard deductive and pseudo-deductive schemes (such as Modus Ponens and Modus Tollens amongst the former and undercutting amongst the latter), inductive schemes (currently only Inductive Generalisation), and fallacious schemes (including the *argumentum ad populum* and the *ignoratio elenchii*). Each operator has two parts: the shell, which resembles a conventional STRIPS operator with pre and post condition lists; and a body which comprises a set of goals. Planning proceeds by constructing a complete abstract plan to fulfil a given goal, and then refining that plan by opening up the bodies of the selected operators, and posting the goals found therein. Further details of the planning process and the means by which it can support sophisticated ordering heuristics can be found in [Reed and Long, 1997; Reed, 1998]. The central issue addressed here is the makeup of the operator bodies, and how saliency is operationalised in conjunction with focus of attention.

The bodies of the Modus Ponens (MP) and Modus Tollens (MT) operators are given below in Figure 1:

MP	***MT***
PUSH_TOPIC(arg(X, P))	**PUSH_TOPIC(arg(¬X, P))**
BEL(H, X)	**BEL(H, ¬X)**
IS_SALIENT(H, X, arg(X, P))	**IS_SALIENT(H,¬X, arg(X, P))**
BEL(H, X → P)	**BEL(H, ¬P → X)**
IS_SALIENT(H, X → P, arg(X,P))	**IS_SALIENT**
	(H, ¬P → X, arg(¬X, P))
POP_TOPIC(arg(X, P))	**POP_TOPIC(arg(¬X, P))**

Figure 1. Modus Ponens and Modus Tollens bodies.

Thus in MP and MT, as in all the other operators, goals of belief (representing the speaker’s intentions to engender belief by a particular agent in a particular proposition) are treated quite separately from goals of saliency (representing the speaker’s intentions to make a proposition salient to an agent). It is goals of belief which result in the adducement of further argument operators at each refinement. To generate an argument in which a supports b (i.e. a, $a \rightarrow b$, b) and b in turn supports c, *Rhetorica* first satisfies the goal BEL(h, c) by using an MP operator with b as the minor premise. At refinement, this leads to the posting, amongst others, of the goal BEL(h, b), which is fulfilled by the second instantiation of an MP operator with a as the minor premise. At this point, the process terminates: without further information available, all BEL goals are satisfied immediately (so long as certain conditions hold, such as the hearer not believing the converse of any of the propositions involved).

It is thus the BEL goals which create the structure of intentions forming the scaffold of an argument. The role of the saliency goals, in contrast, is to set out what needs to be available to the hearer as the argument proceeds. Saliency goals do not simply list what needs to be said: although some IS_SALIENT goals will ultimately be fulfilled by a MAKE_SALIENT operator (which in turn is then used as the basis for generating the textual components of an argument), many other IS_SALIENT goals will be discharged in other ways. For example, in the two step argument a, $a \rightarrow b$, b, $b \rightarrow c$, c, the two saliency goals associated with the major premises will (other things being equal) be trivially fulfilled through invocation of a heuristic stating that a hearer will usually accept an implication without further support. This is the heuristic which leads to the predominance of the enthymematic Modus Brevis form of MP. There are, of course, situations in which the heuristic is inhibited (high levels of hearer bias, known disagreement with the implication) – these ensure that when appropriate, an MP will be realised with all three components stated explicitly.

This distinction between goals of belief and goals of saliency appears superficially similar to previous distinctions between the intentional and informational facets [Young and Moore, 1994], but there is an important difference. Although the IS_SALIENT goals are responsible for the generation of informational content, they are not themselves informational; rather, they represent system intentions which may be realised textually in any number of ways – or may not be realised at all (if the heuristics demonstrate that saliency can be achieved by virtue of the contextual setting). As with Young and Moore's work on DPOCL, *Rhetorica* moves away from a traditional 'pipeline' view of the link between intention and information by allowing the two to be intermingled at the same level of abstraction, and having intentions give rise to further intentions. Unlike DPOCL, however, *Rhetorica* lists only goals in operator bodies, allowing saliency to be expressed as a goal which can then be fulfilled by context-dependent heuristics, or by further planning. Thus both belief and saliency goals should be classed as intentional, with informational content being associated only with the operators which fulfil remaining saliency goals.

The third and final type of goal employed in defining operator bodies are the topic manipulation goals. These are closely related to the saliency goals in several respects. First, they represent intentions of the speaker to manipulate the attention of the hearer. Second, they can be fulfilled either through explicit textual marking (clue words, formatting, etc.) or as a result of contextual information already available to the hearer (typically resulting from the explicit fulfilment of another topic manipulation goal in close proximity). This explicit operationalisation of topic manipulation represents a departure from the more conventional approach wherein focus of attention is an implicit feature of a discourse, and affords a number of benefits concerned with the reordering of argument components for persuasive effect. One key advantage lies in the interplay between propositional salience and focus of attention.

3 Saliency and the Focus of Attention

Saliency and the focus of attention are not equivalent concepts. In the first place, as Grosz and Sidner [1986] describe, "The attentional state is a property of the discourse itself, not of the discourse participants". This contrasts with propositional salience, which in the current work, is relativised to a particular agent. Thus *Rhetorica*'s topic manipulation goals represent intentions to alter the attentional state – so that, for example, if the hearer should interject, the attentional state for that interjection will be the one introduced by *Rhetorica*.[1] Saliency goals, on the other hand, represent system intentions to make a proposition salient to a given agent: what is salient to one agent in a discussion need not, from the system's point of view, necessarily be salient to all the interlocutors. One way of viewing this situation is to see the *focusing structure* (the model of the attentional state) as a (virtual) shared resource between the participants. Every participant in the dialogue has the ability both to change and inspect this structure, though inspection is subject to a 'mask' which restricts how much of the focusing structure can be seen by that agent - what the agent is attending to at a given moment. The successful fulfilment of a PUSH_TOPIC goal thus makes some change to that shared resource (namely, adding a new focus space). The execution of a MAKE_SALIENT action (i.e. an utterance) performs two related actions: first, it adds the proposition to the focus space; and second, it modifies the mask of the indicated agent such that that agent is able to view the proposition in that focus space.

The last feature of topic manipulation in the *Rhetorica* system is that focus spaces are named using the format arg(X,Y), read as *an argument from X to Y*. This represents an abstraction from any particular argument scheme (MP, MT, IG, etc.) and allows goals of saliency to be tied, at plan time, to particular focus spaces. The goal IS_SALIENT(h, x, arg(x, y)) thus specifies that the system intends to make proposition x salient to agent h in the context of arg(x, y), i.e. while the focus space arg(x, y) is on the top of the focus stack.

This tying of utterances to focus spaces is necessary in providing the freedom to order argument supports in any arrangement selected by the rhetoric component of the system, whilst maintaining the level of argument coherence discussed by Cohen [1987]. The key constraint is that components of separate arguments should not be interleaved. So for example, if it has been determined that two MP arguments should both be rendered complete with both major and minor premise, then each minor should be adjacent to its related major – even though the order within each of the pairs, and between the two pairs should be left unconstrained. It is not at all clear that these disjunctive constraints – crucial for achieving the flexibility required to produce persuasive component orderings – are supported by related systems such as DPOCL [Young and Moore, 1994] and NAG [Zukerman *et al.*, 1998].

4 Characterising Saliency

It has been demonstrated that propositional salience is related to both the focusing structure of a discourse and the belief structure of an agent, but it remains to define

[1] The current implementation of *Rhetorica* focuses on monologue, and does not support dialogic interaction. The point is made here to illustrate the independence of the attentional state from either participant.

rather more carefully what these relationships entail. As discussed above, the belief component of saliency cannot be BEL(H, P) since propositions may be made salient by a speaker who neither believes a hearer holds *p* to be true, nor wants to persuade a hearer that *p* is true. There are several possible alternatives. The first is BEL(H, BEL(S, P)) – i.e. that if P is salient to H, then H is led to believe that S believes P. Related to this is a slightly stronger possibility BEL(H, INTENDS(S, BEL(H, X)). Both of these seem flawed, as demonstrated by the case of refutation, where the speaker neither believes nor intends the hearer to believe the claim to be refuted. To maintain that by making the claim salient to the hearer, he is brought to believe that the hearer intends or believes something which she does not, is to imply that every case of refutation involves deception, which runs counter to intuition.

Another possible, weaker, interpretation of the belief component of saliency rephrases the foregoing in terms of absence of contradiction, i.e. BEL(H, BEL(S, ¬BEL(H, ¬P)) or, perhaps, BEL(H, INTENDS(S, ¬BEL(H, ¬P)). Again, however, although these seem reasonable in many cases, stating the claim in a refutation violates these definitions. Perhaps more importantly, all these candidates violate an intuition that propositional salience should not involve sense – if P is salient to an agent it should not necessitate that agent holding a belief that either P or its negation is true.

This intuition is formalised explicitly in the belief model employed in *Rhetorica*, which distinguishes between conventional belief and 'awareness', introducing further possible epistemic states – namely, being unaware of a proposition, and being aware but 'undecided'. Agnosticism, for example, might be characterised using the 'aware but undecided' state. Although the primary aim of this distinction is to capture the differences in argumentation style adopted in situations of varying hearer belief, it permits an alternative definition of the epistemic component of saliency. Thus BEL(H, ?P) – representing the hearer being aware of, but undecided with respect to P – forms a part of saliency such that after the speaker utters some proposition, the hearer is then, at the very least, aware of that proposition. Of course, it may be the case that in uttering something the speaker brings the hearer to believe, or, in certain situations (such as irony), disbelieve the proposition. The speaker is guaranteed, however, that if nothing else, the hearer will be undecided with respect to the proposition: he will certainly not be unaware (assuming the utterance to be successful).

There must, however, be a second, attentional component, for even if BEL(H, ?P) is (guaranteed to be) in the hearer's model, the hearer may not be attending to it. The attentional model adopted in the *Rhetorica* system is based upon the hierarchical, stack model proposed by Grosz and Sidner [1986], rather than the cache model sketched out in [Walker, 1996]. There are several factors contributing to the conclusion that Walker's model is inappropriate in the current context. The first is that argumentation represents an unusually structured and strict form of natural language, in which focusing constraints seem to be carefully adhered to. It is presumably for this reason that it has been possible to derive theories of coherent argument over and above basic textual coherence [Cohen, 1987]. Walker cites evidence to suggest that even in argument, the cache model is more appropriate because "referring expressions in argumentative texts treat the conclusions of popped sisters as salient" (p261) – consistent with propositions remaining in the cache until being dislodged, rather than being actively removed at a pop. The proposed cross-boundary reference would seem to violate the rules of argument integrity proposed in [Cohen, 1987] inter alia, and although it may be that these rules are too constraining, a simpler solution is that the conclusion of a syllogism is active in a different focus space to that in which the premises are found. In this way, focus spaces concentrate on the supporting function of statements, with conclusions belonging to the focus spaces in which they act as premises to some higher conclusion. This approach, employed in *Rhetorica*, is crucial in arranging multiple supports for a given proposition – if, for example, *a* and *b* both support some conclusion *c*, then *c* must lie in a focus space lower down the stack than *a* or *b* – otherwise it would be stated twice (once as the conclusion of the *a* argument and once as the conclusion of the *b* argument).

Another key area of evidence that Walker cites in favour of the cache model is the occurrence of *informationally redundant utterances*, IRUs. Of the three main classes of IRU identified in [Walker, 1996b], it is those of *attention* and *consequence* which are of importance here (since the function of *attitude* IRUs are exclusively dialogic). Consequence IRUs "make inferences explicit" (p187), a definition which relies upon the LOGICAL OMNISCIENCE ASSUMPTION (LOA), that interlocutor agents are logically omniscient – the existence of consequence IRUs, Walker then takes to be evidence of the failure of the LOA. Walker's view, however, is that such consequence IRUs have to be introduced to a discourse plan, since the initial plan has been produced under the LOA[2], and therefore would not involve components which are logically redundant. Walker thus considers enthymemes to be the starting point, and fully specified arguments to be extensions of those enthymemes involving IRUs. Apart from the problem that Walker's account only enables enthymemes in which the conclusion is omitted (surprisingly, it doesn't allow for the Modus Brevis form, which would presumably require IRUs of a different class) the proposal seems, in the context of argumenta-

[2] In fact, Walker's account is rather more subtle and involves assuming deductive capabilities across what is salient, rather than what is believed. There is still, however, an assumption in the initial plan that if the minor and major premise are available then the hearer will deduce the conclusion without it being made explicit.

tion, to be somehow convoluted, making use of the LOA which it goes on to show is unreasonable. An alternative approach is to eschew the LOA altogether, and start from the assumption that an audience needs to be led from one proposition to the next – that a presented argument should be the record of an informal proof procedure in natural language. From this starting point, enthymemes and implicit components are generated by deleting goals, rather than by failing to add IRUs. As with consequence IRUs, attention IRUs (which are employed to make a proposition salient) are also already present in the plan at a deep level, and then may or may not be realised at a surface level, depending on contextual pressures.

Apart from allowing greater flexibility in the types of enthymemes which can be produced, this approach is also closer to Aristotelian and argumentation theoretic notions of argument presentation, and for these reasons, *Rhetorica* employs neither IRUs nor a cache based model of focusing. Walker's criticism [1996] that stack based models have "no constraints related to length, depth, or amount of processing required [for a given focus space]" is circumvented because many such constraints are explicitly represented in the heuristics drawn from studies of rhetoric, psychology and argumentation theory.

5 Worked Example

To demonstrate the interplay between the components of saliency employed in the *Rhetorica* system, a short example is here presented and analysed.

The example is based on a real argument from source cited in [Fisher, 1988]: the original has been analysed to produce the propositions in Figure 2.

a: *radioactive elements disintegrate and eventually turn into lead*
b: *matter has not always existed*
c: *there are no radioactive elements left*
d: *uranium etc. is still present*

Figure 2. Initial propositions

The input to *Rhetorica* in Figure 3 is a representation of the relationships between these propositions, expressed as a set of system beliefs (other input, including a broad parameterisation of the intended hearer is omitted for clarity). The hearer is assumed initially to be unaware of all propositions, hence the absence of any hearer beliefs.

bel(s, a)	**bel(s, not(b))**
bel(s, supports(not(b), c))	**bel(s, d)**
bel(s, not(c))	**bel(s, supports(d, not(c)))**
bel(s, supports(a, supports(not(b), c)))	

Figure 3. System input

The initial goal of the system is BEL(*ag0*, *b*) – i.e. to persuade the hearer that *matter has not always existed*. Associated with this goal is a saliency goal for the same proposition. After one round of planning, *Rhetorica* identifies that a Modus Tollens argument can be employed to fulfil BEL(*ag0*, *b*). After refinement, the goals within the body of that MT are posted: these are shown in Figure 4 (notice that although not(*c*) is the conclusion of the major premise of the MT, it is *b* which is the conclusion of the MT as a whole).

PUSH_TOPIC(arg(not c, b))
BEL(ag0, not c)
IS_SALIENT(ag0, not c, arg(not c, b))
BEL(ag0, supports(not b, c))
IS_SALIENT(ag0, supports(not b, c), arg(not c, b))
POP_TOPIC(arg(not c, b))

Figure 4. System processing after refinement

Both the major and the minor premise of the MT can be supported by further argumentation: the major with an argument from *d*, and the minor with an argument from *a*. Both these supports take the form of MP arguments. Thus Figure 4 leads to the abstract plan in Figure 5.

push_topic(arg(not c, b))
argModusPonens(ag0, not c, d)
make_salient(ag0, not c, arg(not c, b))
argModusPonens(ag0, supports(not b, c), a)
make_salient(ag0, supports(not b, c), arg(not c, b))
pop_topic(arg(not c, b))

Figure 5. An abstract plan during the planning process

The topic manipulation goals are fulfilled by operators of the same name which are regarded as primitive in *Rhetorica*. Similarly, both IS_SALIENT goals are fulfilled by MAKE_SALIENT operators which are also primitive at this abstract level. As discussed in section 2, it is rarely appropriate to perform enthymematic contraction within MT, so neither IS_SALIENT goal gets fulfilled by heuristic manipulation.

At the next stage, the goals in the bodies of the two MPs are posted. The goals of the first are shown below in Figure 6:

PUSH_TOPIC(arg(d, not c))
BEL(ag0, d)
IS_SALIENT(ag0, d, arg(d, not c))
BEL(ag0, supports(d, not c))
IS_SALIENT(ag0, supports(d, not c), arg(d, not c))
POP_TOPIC(arg(d, not c))

Figure 6. Body goals from the first MP

The two belief goals are trivially satisfied because no further information is available, and because the hearer is not known to disbelieve them. Of the two saliency goals, only the first (expressing *d*) is fulfilled by a MAKE_SALIENT operator. The second is discharged by one of the enthymematic contraction heuristics. The other MP operator is planned for in a similar way.

In the interests of brevity, this summary of the planning process has ignored details of rhetorical ordering heuristics and the clue word introduction mechanisms, and has omitted the specification of the partial order; details of this functionality can be found in [Reed and Long, 1997; Reed, 1998]. The final step is to produce a fully ordered list of primitives: this system output is shown in Figure 7.

1. push_topic(arg(not c, b))
2. push_topic(arg(d, not c))
3. make_salient(ag0, d, arg(d, not c))
4. pop_topic(arg(d, not c))
5. push_topic(arg(a, supports(not b, c)))
6. make_salient(ag0, a, arg(a, supports(not b, c)))
7. pop_topic(arg(a,supports(not b,c)))
8. make_salient(ag0, clue(inference), arg(not c ,b))
9. make_salient(ag0, supports(not b, c), arg(not c, b))
10.make_salient(ag0,clue(polarity,negative), arg(not c ,b))
11.make_salient(ag0, not c, arg(not c, b))
12.pop_topic(arg(not c, b))
13.make_salient(ag0, clue(inference), null_topic)
14.make_salient(ag0, b, null_topic)

Figure 7. Final plan of primitives

Although *Rhetorica*, like comparable systems such as NAG [Zukerman *et al.*, 1998], does not yet produce textual output, the text associated with the propositions in the original argument can be linked with only very minor syntactic modification, and clue words introduced by assuming a one to one mapping (such that inference clues become *so*, negative polarity *but*, and so on). The plan in Figure 7 can thus be seen as responsible for the following version of the argument (numerical tags indicate responsible steps in the plan): [*Elements such as uranium are still present.*]$_3$ [*Radioactive elements disintegrate and eventually turn into lead,*]$_6$ [*so*]$_8$ [*if matter had always existed, there should be no radioactive elements left.*]$_9$ [*But*]$_{10}$ [*there are such elements left,*]$_{11}$ [*so*]$_{13}$ [*matter cannot always have existed.*]$_{14}$

6 Conclusions

It is difficult to evaluate output of an NLG system, particularly where evaluation is to be based upon a highly subjective notion like persuasiveness. The argument produced by *Rhetorica* in the previous section is certainly coherent at both textual and Cohen-like argument levels, despite differing from the original in terms of the selection and ordering of components, and the number and type of clue words employed. Preliminary investigation has been undertaken to evaluate (rather longer) arguments generated by the system, by setting them against the original version and asking web users to rate the two versions. Although the results are very encouraging, more rigorous experimentation is required to confirm that automatically generated arguments are as persuasive as (or better than) their natural counterparts. Entirely automatic generation of textual output from a discourse plan is also left to future work, but it is clear even from the current *Rhetorica* implementation that using salience as a way of coordinating which parts of an argument are expressed, and through interaction with explicit control of focus, how arguments are arranged, is a powerful approach to the generation of persuasive text.

References

[Blair, 1838] Hugh Blair. *Lectures on Rhetoric and Belles Lettres*. Charles Daly, London, 1838.

[Cohen, 1987] Robin Cohen. Analyzing the Structure of Argumentative Discourse. *Computational Linguistics* 13(1):11–24, 1987.

[Fisher, 1988] Alec Fisher. *The logic of real arguments*. Cambridge University Press, 1988.

[Fox and Long, 1995] Maria Fox and Derek Long. Hierarchical Planning using Abstraction. *IEE Proc on Control Theory and Applications* 142(3):197–210, 1995.

[Grosz and Sidner, 1986] Barbara J. Grosz and Candace L. Sidner. Attention, Intentions and the Structure of Discourse. *Computational Linguistics* 12(3):175–204, 1986.

[Marcu, 1996] Daniel Marcu. The Conceptual and Linguistic Facets of Persuasive Arguments. *Proceedings of the ECAI'96 Workshop on Planning and Natural Language Generation*, pages 43–46, Budapest, Aug. 1996.

[Reed and Long, 1997] Chris Reed and Derek Long. Content Ordering in the Generation of Persuasive Text. *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, Nagoya, Japan, pages 1022-1027, 1997.

[Reed, 1998] Chris Reed. Generating Arguments in Natural Language. PhD Dissertation, Department of Computer Science, University College London, 1998.

[Sadock, 1977] J. Sadock. Modus Brevis: The Truncated Argument. *Papers from the 13th Regional Meeting of the Chicago Linguistics Society*, pages 545-554, 1977.

[Walker, 1996] Marilyn A. Walker. Limited Attention and Discourse Structure. *Computational Linguistics* 22(2): 255-264, 1996.

[Walker, 1996b] Marilyn A. Walker. The effect of resource limits and task complexity on collaborative planning in dialogue. *Artificial Intelligence* 85: 181-243, 1996.

[Young and Moore, 1994] Richard M. Young and Johanna D. Moore. DPOCL: A principled approach to discourse planning. *Proceedings of the 7th International Workshop on NLG*, Kennebunkport, Maine, pages 13–20, 1994.

[Zukerman *et al.*, 1998] Ingrid Zukerman, Richard McConachy and Kevin B Korb. Attention during argument generation and presentation. *Proceedings of the 9th International Workshop on Natural Language Generation* Niagara-on-the-Lake, Canada, pages 148-157, 1998.