Context and Focusing in Argumentative Discourse Planning*

Chris Reed[1], Derek Long[2] and Maria Fox[2]

[1] Department of Computer Science,
University College London,
Gower St.,
London, WC1E 6BT
Tel. (+44) (0) 171 380 7214
Fax (+44) (0) 171 387 1397
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
Tel. (+44) (0) 191 374 2000
Fax (+44) (0) 191 374 2560
D.P.Long@dur.ac.uk
M.Fox@dur.ac.uk

Abstract. One of the key issues in natural language generation (NLG) is the question of how to ensure that text is *coherent*. In this paper, the problem is viewed as one of setting up and dynamically updating an appropriate context for each utterance, by maintaining a *topic stack* indicating the current focus of conversation. The research centres around the problem of producing argumentative discourse, where there are a number of close ties between discourse context, focusing and linguistic structure. The proposed model employs several conceptual tiers in a hierarchical planner, the highest of which is responsible for the structure (including focusing and coherency) of the argument, and which uses purely intentional data structures to achieve this structure. The notion of the *saliency* of a proposition to the hearer plays a key role in the planning process and the resultant focusing constraints.

Keywords: argumentation, context, focusing, natural language generation, planning

1. Introduction

Generation of argument is attracting increasing research effort due to its wide ranging applicability: to furnish an expert systems with a means to producing cogent justification (Paris, 1991); to improve agent negotiation (Parsons and Jennings, 1996), (Reed *et al*, 1996a); to assist in medical decision making (Fox and Das, 1996); to critique user decisions (Rankin, 1993), and so on. In cases where these applications demand interaction with a human user, the argumentation will be most effective when realised in natural language.

The generation of natural language argument can be approached by using a hierarchical planner applied in several stages: firstly to produce the logical structure of the argument, then to convert that structure into a plan for linguistic expression, and then to realise that plan into expressed utterances (Reed *et al*, 1996b). Similar 'pipeline' structures are frequently adopted in NLG to distinguish the text planning from the sentence realisation tasks, (eg. (Paris, 1991), (Maybury, 1991)) and, more recently, from the sentence planning task (eg. (Hovy and Wanner, 1996), (Panaget, 1994)). The pipeline architecture has been criticised due to a number of inherent problems[1], such as the possibility of lower level decisions affecting those at a higher level, and the difficulty of dealing with non-satisfiable maintenance goals (Hovy, 1990). Despite such problems, the generation of paragraph- and multi-paragraph- length text is usually assumed to rely, at least in part, on a phase of gross text planning which would only be hindered by consideration of lower level problems which receive attention as a subsequent, distinct planning phase.

This paper, then, concentrates on the first planning task, that of deciding which information to include, how the components interrelate, and in what order to present those components. This is the responsibility of the Argument Structure (AS) level. At this level of abstraction, the planning process

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

[1] Though a number of these problems seem to be inherited from restrictions under particular planning paradigms.

relies entirely upon *intentional* data structures. It has been noted in the NLG literature that representation of intentions, over and above the informational content of the text, is an important facet of generation, and in particular, is essential in order to accommodate a number of phenomena, including justification, explanation, failure recovery and question answering (Young and Moore, 1994), (Moore and Paris, 1994).

The process by which argument is planned is captured much more accurately in a hierarchical framework than in a sub-goaling system. At the most abstract levels of the plan, the operators are structural and thematic rather than explicitly detailed. As the plan is refined, the structure of the argument is developed until a level is reached at which planning can interact with speech acts. In dialogue, execution of parts of the argument plan can provide input into as yet unexpanded components of the remaining argument, allowing the development of the argument to respond to feedback from the other interlocutor.

In classical planning, operators can interact in a way that violates commitments made earlier in the plan development process. Typically this happens in the process of addressing new goals introduced as subgoals when the preconditions, or tasks in Hierarchical Task Network (HTN) planning (Erol *et al*, 1994), of operators are addressed. At the information-conveyance level of argument construction (that is, the level at which propositional beliefs are brought to the attention of the hearer in preparation for inferences on the part of the hearer that depend on these beliefs) the steps of the argument add monotonically to the beliefs of the interlocutors. For example, if a speaker wishes to persuade a hearer to accept a proposition Q, by demonstrating that the hearer believes both P and $(P \rightarrow Q)$, it will not be appropriate for the persuade operator to destroy the hearer's beliefs in P and $(P \rightarrow Q)$ when the desired effect has been achieved (quite apart from the speaker losing credibility, such a scenario is likely to violate rules of *commitment* (Walton and Krabbe, 1995)). This is traditional in classical planning in which the preconditions of an operator (in this case, the salience of P and $(P \rightarrow Q)$ for the hearer would be given as preconditions of the persuade operator) are often deleted by that operator. Thus, seen at the level of information conveyance, operators do not seem to have complex delete lists and therefore initiate few conflicts. This gives the impression that argument-planning is trivial, indeed that planning is not really a necessary part of argument construction.

In fact, the role of planning in argument construction is not at the level of information-conveyance but at the level of maintaining the coherence and focus of the developing argument. An argument becomes incoherent if attending to a new goal diverts the focus of the argument before the objectives of the argument, introduced in that focus, have been achieved.

A lecture (or technical presentation) is an example of a sustained (and possibly complex and subtle) argument which requires significant planning. It is essential to the success of the argument that its early phases provide the context required for the audience to make sense of the premises and conclusions of the argument. This might require careful planning to ensure that the nesting of sub-arguments does not exceed a reasonable depth and to ensure that the focus of the lecture does not appear to veer dramatically between apparently unrelated subjects. This will affect the order in which ideas are presented and the amount of detail it is considered appropriate to include. At the level at which these considerations are planned for there is little concern with the actual technical content of the lecture, although it may be necessary to plan some parts in more detail to anticipate correctly how the phases of the lecture will fit together.

Planning is thus necessary to ensure that a complex, subtle or sustained argument is structured in such a way that earlier phases of the argument provide context for the later phases and that all of the objectives of the argument (and any sub-arguments) are achieved.

2. Context

There exists certain ambiguity over the meaning of the term *context* - in NLG alone, there seem to be at least two senses in which it is applied. Firstly, to describe the physical and mental environment shared by the interlocutors: this is thoroughly investigated as *mutual cognitive environment* in (Sperber and Wilson, 1986). In argumentation, the context would also include such features as the type of audience (specialised/general, heterogeneous/homogenous, biased/unbiased, etc.), the relationship existing between speakers (pedagogical, managerial, peer, etc.), the participants' investment in the outcome of the argument, and physical aspects of the argument's expression (brief/extended, structured/free, oral/written, etc.) in addition to the beliefs held by the participants making up the mutual cognitive environment. These contextual parameters and their impact on argumentation are explored in (Reed *et al*, 1996b). The second use of the term context expresses aspects of the local setting for a particular utterance (word, phrase, sentence, and possibly even paragraph). In this sense, it refers to the preceding discourse[2], the local and global aims and intentions (as discussed in (Hobbs, 1982)), and the local structure, and it is with this second interpretation that this paper is concerned.

The *coherency* of a text is a key notion which is clearly closely related to context: an utterance contributes to a coherent structure only if it is appropriate in the surrounding context. The generation of coherent text is therefore a problem of setting up an appropriate context for each utterance in the discourse. One of the most effective approaches to ensuring coherency is *focusing*, whereby the current focus of attention is recorded, and restrictions are placed upon how that focus can change. The role of focus in generation of argumentative discourse is particularly important, due to the fact that it is usually more structured than other forms of 'free' natural language. Indeed, this structuring has been pursued extensively in argumentation theory, and some of the results can be seen to shed light on the problems of generating coherency. Normative accounts of argument validity (such as those of Walton, (Walton and Krabbe, 1995), (Walton, 1992), *inter alia*) suggest constraints on the determination and presentation of information. Such constraints are of key importance to the current work, in which fallacies are being employed as a vital component of the generation process, so ensuring that the resultant text is not incoherent is therefore that much harder. Argumentation theory has also produced a number of accounts detailing the way in which natural argument is structured (Toulmin, 1958), (Freeman, 1991), (Walton and Krabbe, 1995), (Van Eemeren and Grootendorst, 1992): a number of these accounts have formed the basis for AI applications, both logical (eg. the defeasible logic of Fox *et al*, (Das, Fox and Krause, 1996), (Fox and Das, 1996), (Parsons and Jennings, 1996), which expresses the essence of Toulmin's schemas) and to a lesser extent, linguistic (eg. (Maybury, 1993), (McConachy and Zukerman, 1996)). Such accounts of structure help to explain what makes one argument coherent and another incoherent, such as the examples given in (Cohen, 1987).

[2] It is in this sense that McCarthy has examined context in discourse (McCarthy and Buvac, 1995).

3. Focusing

The problem of implementing some mechanism for maintaining focus during text planning has been addressed in a number of different ways. Grosz and Sidner set out a basic framework (Grosz and Sidner, 1986), but as pointed out by McCoy and Cheng (1991), this was primarily concerned with discourse analysis rather than with issues of generation. The intuitions behind Grosz and Sidner's work, however, have pervaded much of the NLG research, and the current work is no exception.

The focusing mechanism relies upon a core notion of a *topic stack* onto which new topics are pushed and old topics popped. A topic is simply a proposition, to which others are related (thus in argumentation, a topic can be seen as a conclusion, to which premises are related[3]).

In this work, argument structure is simplified (but, not oversimplified - compare, for example, (Wilson, 1986)) to multiple premises followed by a single conclusion[4]; each premise can in turn be a sub-argument, with its conclusion acting as premise to the super-argument. Thus on entering a new sub-argument (ie. at the start of a discourse, or when creating a subargument in support of a premise) the conclusion is pushed on to the stack as the current topic. On concluding a sub-argument, the current topic on the stack is popped off.

Focusing in argumentation mirrors the hierarchical structure of the argument itself - this is also generally true in other forms of natural language (as pointed out in (Grosz and Sidner, 1986)), but the structure of an argument is more clearly visible at the textual level. As a consequence, the generation process can make more direct use of the focusing machinery for textual realisation. One clear example of this phenomenon is that push and pop actions immediately lend themselves to the generation of *clues*, words or phrases which explicitly indicate structural components in the argument (eg. "In addition", "But then", etc.). Clues have been found (from an analysis viewpoint) to be crucial components of argumentative discourse (Cohen, 1987). Clues become vital during complex nesting of subarguments (though as nesting becomes deeper, the processing costs of comprehension increase dramatically, and have a negative effect on the persuasive force of an argument; thus shallower nesting is implicitly preferred by the heuristics in the planning process which aim to maximise the persuasive effect of an argument).

Focusing is carried out entirely at the AS level. As a result, the substrate to the focusing process is exclusively intentional. Following the classification of goals by (Moore and Pollack, 1992) into *communicative* for those that are intention-laden, and *linguistic* for those that are not, the actions `PUSH_TOPIC` and `POP_TOPIC` are communicative goals. This implies that the speaker intends the hearer to appreciate the structure of the argument. The means by which such structural information is conveyed are decided by subsequent processing. Structure may be realised as explicit, conspicuous structural marking (such as "But returning to my point about X ..."), by the use of clues which may be more, or less, subtle, or alternatively, it may be appropriate to have no surface realisation whatsoever (for example, a chain of argumentation would be hindered by repeated clues indicating commencement of new subarguments, and would look ridiculous if it included multiple `POP_TOPIC` clues one after another at the close of the argument chain).

There is a strong relationship between the control of focus affecting local context through a discourse, and the use of context in the planning process itself. Unlike much work in NLG (surveyed, for example, in (Hovy, 1993)), the core component is based not on NOAH (Sacerdoti, 1977) (which by modern standards of planning technology is rather dated), but on AbNLP (Fox and Long, 1995). AbNLP is a hierarchical planner in which abstract operators encapsulate partially ordered and

---

[3] Of course, there are frequent cases in which either premise(s) or conclusion(s) are not realised into text at all (these are enthymemes), which is why it would be erroneous to view the terms conclusion and topic as synonymous.

[4] It has been pointed out that restricting an analysis to single conclusion arguments is not necessary for analytic purposes (Walton, personal communication), however, for synthesis, the simplification greatly reduces the complexity of the required machinery (and in particular, that responsible for focusing), without drastically compromising performance.

constrained goals appropriate at a refined level of detail. The encapsulated goals are revealed when abstract operators are refined, not at their moment of application as in HTN planners such as NOAH and SIPE (Wilkins, 1988). The goals associated with different abstract operators can be freely interleaved , whereas this is not possible in sound HTN planning (Erol *et al*, 1994). This provides greater flexibility in the plan development process and leads to more fluent argument plans. Further, abstract plans in AbNLP can be interpreted meaningfully as plans, without the need to edit plans to remove redundant steps as is required in HTN planning. The ability to interpret abstract plans without actually refining them is particularly important for argument planning, where the speaker's intentions may need to be communicated to the hearer as part of the development of the argument.

Abstract operators encapsulate all of the planning effort involved in achieving their effects. In this sense abstract operators can be seen as placeholders for the sub-arguments which will eventually form the completed spoken argument. They will coordinate, at the abstract levels of the plan, the various shifts in focus involved in interfacing the sub-arguments. Each abstract operator refines into a local context in which the objectives of the operator can be achieved in more detail. These local contexts are brought into the global context of the plan after refinement, in a way that is consistent with the structural properties of the abstract plan.

The use of encapsulation means that some of the goals in the body set up a context which, in conjunction with contextual information provided by the surrounding abstract plan, facilitate fulfilment of the remaining body goals. The body of the Modus Ponens operator, for example, includes topic manipulation goals (corresponding to PUSH_TOPIC and POP_TOPIC), the achievement of which sets local context for the subsequent planning for goals expressing desired hearer beliefs. The body of the MP operator is given below in Figure 1; the topic manipulators surround the BEL goals (which may give rise to further planning at the next level of abstraction) and the IS_SALIENT goals (which mark the possible need for a communicative act).

```
Modus Ponens (H, P)

Body:    t0:  PUSH_TOPIC (P)
         t1:  BEL (H, X)
         t2:  IS_SALIENT (H, X, P)
         t3:  BEL (H, X ⊃ P)
         t4:  IS_SALIENT (H, X ⊃ P, P)
         t5:  POP_TOPIC (P)
```

Fig 1. Modus Ponens Operator Body

The result of the planning at the AS level is a partially ordered plan of primitive operators[5] which correspond to intentions. These intentions are of just two types: the PUSH_TOPIC and POP_TOPIC actions, used for effecting control over the topic stack, and the MAKE_SALIENT action, used to indicate information to be communicated. The topic stack manipulators take a single parameter: the topic to be pushed or popped[6]. The MAKE_SALIENT intention involves three parameters: the agent to whom the intention is addressed, the proposition to be conveyed, and the topic to which the proposition pertains. The example in Figure 2, below, shows how the result of the AS level is related to the argument structure, by giving a brief argument, its standard analysis and structure, and then also a plan which might be generated by the AS to give rise to such an argument.

[5] Or at least, primitive at this planning level

[6] Indicating the topic to be popped is simply for readability: clearly a popping action can only refer to the item at the top of the stack.

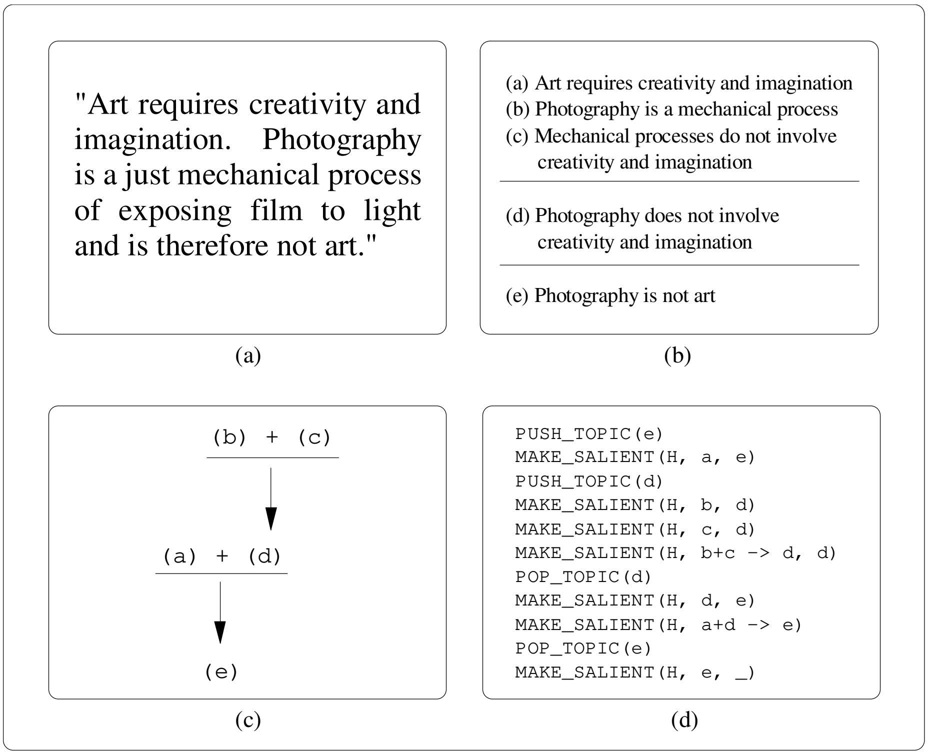

Fig. 2. Sample argument[7] (a) text, (b) analysis, (c) structure and (d) intention plan

Through the standard analysis treatment, this argument can be shown to have two missing components: the premise (c) in the sub-argument and the conclusion (d) of the sub-argument (which stands as a premise in the main argument), and that ‘reconstructing’ the argument in its complete form yields the structure in 1c.

The generation of this argument would involve the AS creating[8] a plan like that in Figure 1d. The details of the planning process are beyond the scope of this paper, but it follows the same hierarchical structure as is present in Figure 1c, ie., the goal of persuading the hearer of (e) is refined to one of showing that (a) and (d) are true, and that they support belief in (e); similarly, the goal of persuading the hearer of (d) is refined to showing that (b) and (c) are true and support (d).

The ordered list of primitive actions in Figure 1(c) expresses intentions to

– push the overall aim of the argument, (e), onto the topic stack
– make the hearer aware of the first premise, (a) (under the topic (e))
– push the aim of the sub-argument, (d), onto the topic stack
– make the hearer aware of the second premise (b) (under the topic (d))
– make the hearer aware of the third premise (c) (under the topic (d))
– make the hearer aware of the support given to (d) by (b) and (c) (under the topic (d))
– pop the aim of the sub-argument, (d), from the topic stack
– make the hearer aware of the conclusion of the sub-argument (d) (under the topic (e))
– make the hearer aware of the support given to (e) by (a) and (d) (under the topic (e))
– pop the aim of the argument , (e), from the topic stack

[7] This example is a paraphrase of Snowden’s argument, as cited in (Johnson, 1992). The example exhibits a number of phenomena which clearly demonstrate aspects of the planning process, but in practice, generating such trivially short argumentation through AbNLP-style planning is probably unnecessary.
[8] Details of the planning process which produces such a final plan of primitive operators are beyond the scope of this paper.

– make the hearer aware of the conclusion of the argument (e) (under some outer, unspecified topic[9])

It is worth noting that the topic stack fulfils a dual role: it exists at plan-time to impose coherency constraints (such as those mentioned in (Reed *et al*, 1996b) p560), and also exists at run-time as a real world object. Clearly, though, these two roles are very closely linked: the state of the plan-time topic stack at any particular stage in the planning process should be identical to the state of the run-time stack at the same stage in expression of the plan.

Expression of the component intentions of the plan does not follow a one-to-one correspondence with surface textual features: as mentioned above, the topic manipulators may or may not have textual realisation. Similarly, the `MAKE_SALIENT` intentions may not need to be uttered explicitly. From the example there would appear to be two types of such inexpression: enthymeme contraction, wherein a premise or conclusion is omitted (eg. `MAKE_SALIENT(H, c, d)`), and the absence of explicit support relations (eg. `MAKE_SALIENT(H, b+c -> d, d)`). In deciding not to express particular intentions, the post-AS processing is not simply discarding the intentions. Rather, the decision is taken that the information is already salient to the hearer[10], and thus no more work is required to fulfil the intention (which after all, simply demands no more than that the proposition be salient to the hearer). Thus the two forms of inexpression are both examples of the same phenomenon.

Given this brief explanation of how intentions need not be expressed, the relationship between 1d and 1a becomes clearer. `MAKE_SALIENT(H, a, e)` is realised into “Art requires creativity and imagination”. `MAKE_SALIENT(H, b, d)` is similarly realised into “Photography is a mechanical process of exposing film to light”, with the introduction of the loaded word “just” (through some mechanism for lexical realisation such as that discussed in (Marcu, 1996)). The phrase “and therefore” can be seen as fulfilling both the structural intention `MAKE_SALIENT(H, a+d -> e, e)`, and the focus intention `POP_TOPIC(e)` (indeed this co-occurrence of structural intention and focus shift with a common textual component is frequent - though there are times when it is necessary to permit individual expression, such as the termination of a sub-argument which does not close with the conclusion). Finally, “... is not art” expresses the overall aim of the argument, `MAKE_SALIENT(H, e, _)`. It is interesting to note that many of the decisions over which intentions can be fulfilled without realisation are, in this example, quite arbitrary - the final conclusion for example, could easily have been omitted without detriment to the hearer’s understanding of the argument. It is for this reason that this processing stage relies heavily upon theories of rhetoric and eloquence (this is examined in detail in (Reed *et al*, 1996b)).

This arbitrary nature of decisions made during the generation process also indicates one key problem: evaluation. Quite apart from the inherent difficulties in assessing the performance of NLG systems, simply answering the question ‘Is this output coherent?’ is far from trivial. It seems that coherency is a scalar notion: it is quite possible to give examples of text which lie at either end of the scale, and then to classify them, but it is equally possible to give examples which depend for classification upon essentially subjective judgements (this point is touched upon by Cohen (1987)). Given this scalar coherency, the task of a natural language generator becomes one of *maximising*, rather than ensuring, coherency. As such, other factors must then also be taken into account: a maximally coherent choice may also be a less persuasive choice (or an inappropriately long-winded choice, or a choice that adversely affects the speaker-hearer relationship, etc.) and may therefore not be the optimum choice.

---

[9] Appropriate handling of the external topic, referred to in the last step, has yet to be resolved, but relies upon an account of common ground, such as the discussion of *mutual initial discourse context*, given in (Thomason and Moore, 1995).

[10] Of course, this conclusion may be mistaken, in which case misunderstanding may ensue.

4. Conclusion

A model of focusing in the generation of argumentative discourse has been put forward that can assist in the production of coherent text. This model therefore also describes a method for dynamically managing context through a discourse, since coherency is dependent upon appropriate context. Focusing in argument is of particular interest as the effects are more immediately manifest (with linguistic features such as clues and structure markers).

The proposed model employs a number of hierarchical planning stages. The first develops the structure of the argument, producing an abstract plan of intentions. These intentions are of just two types: topic manipulation and propositional saliency. The next planning stage then refines the plan by deciding how to express these intentions, if at all.

Work is currently under way to implement a system to generate simple, monologic, natural language arguments based on knowledge and low level linguistic realisation offered by LOLITA, (Smith *et al*, 1994), a large scale, domain independent natural language system.

References

(Cohen, 1987) Cohen, R., "Analyzing the Structure of Argumentative Discourse", *Computational Linguistics*, **13** (1), pp11-24 (1987)

(Das, Fox and Krause, 1996) Das, S., Fox, J, & Krause, P. "A Unified Framework for Hypothetical and Practical Reasoning (1): Theoretical Foundations", in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning*, Springer Verlag, Berlin, pp58-72 (1996)

(Erol *et al*, 1994) Erol K, Hendler J and Nau D., *Semantics for Hierarchical Task-Network Planning*, Technical Report CS-TR-3239, University of Maryland, USA (1994)

(Fox and Das, 1996) Fox, J. & Das, S., "A Unified Framework for Hypothetical and Practical Reasoning (2) - Lessons from Medical Applications", in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning*, Springer Verlag, Berlin, pp73-92 (1996)

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

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

(Grosz and Sidner, 1986) Grosz, B.J., Sidner, C.L., "Attention, Intentions and the Structure of Discourse", *Computational Linguistics* **12** (3), pp175-204 (1986)

(Hobbs, 1982) Hobbs, J.R., "Towards an Understanding of Coherence in Discourse" in Lehnert, W.G., Ringle, M.H., (eds), *Strategies for Natural Language Processing*, Lawrence Erlbaum Associates, pp223-243 (1982)

(Hovy, 1990) Hovy, E.H., "Pragmatics and Natural Language Generation", *Artificial Intelligence* **43**, pp153-197 (1990)

(Hovy, 1993) Hovy, E.H., "Automated Discourse Generation Using Discourse Structure Relations", *Artificial Intelligence* **63**, pp341-385 (1993)

(Hovy and Wanner, 1996) Hovy, E.H. & Wanner, L., "Managing Sentence Planning Requirements", in *Working Notes of the ECAI'96 Workshop on Planning and Natural Language Generation*, pp53-58 (1996)

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

(Marcu, 1996) Marcu, D., "The Conceptual and Linguistic Facets of Persuasive Arguments", in *Working Notes of the ECAI'96 Workshop on Planning and Natural Language Generation*, pp43-46 (1996)

(Maybury, 1991) Maybury, M.T. *Planning Multisentential English Text Using Communicative Acts*, PhD. dissertation, Cambridge University (1991)

(Maybury, 1993) Maybury, M.T., "Communicative Acts for Generative Natural Language Arguments" in *Proceedings of the National Conference on Artificial Intelligence (AAAI-93)*, AAAI, pp357-364 (1993)

(McCarthy and Buvac, 1995) McCarthy, J. & Buvac, S., "Formalizing Context (Expanded Notes)", in *Working Notes of the AAAI'95 Workshop on Formalizing Context*, (1995)

(McConachy and Zukerman, 1996) McConachy, R. and Zukerman, I., "Using Argument Graphs to Generate Arguments", in *Proceedings of the Twelfth European Conference on Artificial Intelligence (ECAI'96)*, pp592-596 (1996)

(McCoy and Cheng, 1991) McCoy, K.F., Cheng, J. "Focus of attention: Constraining what can be said next" in Paris, C.L., Swartout, W.R. & Mann, W.C., (eds), *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, Kluwer, pp103-124 (1991)

(Moore and Paris, 1994) Moore, J.D., Paris, C.L., "Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information", *Computational Linguistics* **20** (4), pp651-694 (1994)

(Moore and Pollack, 1992) Moore, J.D., Pollack, M.E., "A Problem for RST: The Need for Multi-Level Discourse Analysis", *Computational Linguistics* **18** (4), pp537-544 (1992)

(Panaget, 1994) Panaget, F., "'Using a textual representational level component in the context of discourse or dialogue generation" in *Proceedings of the 7th International Workshop on Natural Language Generation*, Kennebunkport, Maine, pp127-136 (1994)

(Paris, 1991) Paris, C.L., "Generation and Explanation: Building an Explanation Facility for the Explainable Expert System Framework" in Paris, C.L., Swartout, W.R. & Mann, W.C., (eds), *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, Kluwer, pp49-82 (1991)

(Parsons and Jennings, 1996) Parsons, S. & Jennings, N.R., "Negotiation through argumentation - a preliminary report", in *Proceedings of ICMAS'96* (1996, to appear)

(Rankin, 1993) Rankin, I., "Natural language generation in critiquing", *Knowledge Engineering Review* **8** (4), pp329-347 (1993)

(Reed *et al*, 1996a) Reed, C.A., Long, D.P., Fox, M. & Garagnani, M., "Persuasion as a Form of Inter-Agent Negotiation", in *Working Notes of the 2nd Australian Workshop on Distributed AI*, pp13-27 (1996)

(Reed *et al*, 1996b) Reed, C.A., Long, D.P. & Fox, M., "An Architecture for Argumentative Discourse Planning" in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning*, Springer Verlag, Berlin, pp555-566 (1996)

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

(Sperber and Wilson, 1986) Sperber, D. & Wilson, D. *Relevance: Communication and Cognition*, Basil Blackwell, Oxford (1986)

(Smith *et al*, 1994) 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*, Kennebunkport, Maine (1994)

(Thomason and Moore, 1995) Thomason, R.H. & Moore, J.D., "Discourse Context", in *Working Notes of the AAAI'95 Workshop on Formalizing Context* (1995)

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

(Van Eemeren and Grootendorst, 1992) Van Eemeren, F.H. & Grootendorst, R., *Argumentation, Communication, and Fallacies: A Pragma-Dialectical Perspective*, Lawrence Erlbaum, Hillsdale, N.J. (1992)

(Walton, 1992) Walton, D.N., *The Place of Emotion in Argument*, Pennsylvania State University Press (1992)

(Walton and Krabbe, 1995) Walton, D.N. & Krabbe, E.C.W., *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning* (1995)

(Wilkins, 1988) Wilkins D, *Practical Planning - Extending the Classical AI Planning Paradigm*, Morgan-Kaufmann, CA (1988)

(Wilson, 1986) Wilson, B.A., *The Anatomy of Argument*, 2nd edition, University Press of America, Washington, D.C. (1986)

(Young and Moore, 1994) Young, R.M., Moore, J.D. "DPOCL: A principled approach to discourse planning" in *Proceedings of the 7th International Workshop on Natural Language Gene*

ERROR: stackunderflow
OFFENDING COMMAND: rlineto

STACK: