Content Ordering in the Generation of Persuasive Discourse

Chris Reed
Department of Computer Science
University College London, Gower Street
London WCIE6BT England
C.Reed@cs.ucl.ac.uk

Derek Long
Department of Computer Science
University of Durham, South Road
Durham DH1 3LE England
D.P.Long@dur.ac.uk

Abstract

A framework is summarized which supports the planning of natural language argument structure. One key aspect of natural argument is the order in which components are presented. This is in part responsible for both the coherency and persuasive effect of an argument. One means of effecting such ordering is proposed, and an overview is provided of the various classes of ordering heuristics. These heuristics are based upon insights offered by rhetoric texts, psychological research, and a corpus study. Finally, it is demonstrated that such an approach can also contribute to the generation of surface textual features including formatting, punctuation and clue words.

1 Introduction

Generating arguments in natural language is of key importance in producing cogent justification in expert systems [Paris, 1991], critiquing user decisions [Rankin, 1993], explaining reasoning in logic-based systems [Fox and Das, 19961, tailoring medical advice [Marcu, 1996], etc. Argument is often more structured than other forms of natural language, and this fact can be exploited in designing a system for the generation of persuasive discourse. This paper is based upon the work of Reed *et al* [1996], in which a hierarchical framework is proposed which reflects the natural structuring of argument.

At the highest level of abstraction, the Argument Structure (AS) level is responsible for the logical form of the argument, i.e. the deductive, inductive and fallacial relationships holding between premises and conclusions. The operators which are available at the AS level are thus conceptually similar to a subset of those presented in [Maybury, 1993]. Below the AS level, the Eloquence Generation (EG) level effects stylistic and rhetorical refinements, based upon contextual and hearer-dependent parameters. The EG level does not alter the plan produced by the AS level; rather, it introduces structurally redundant rhetorical statements, marks stylistic preference (voice, mood, formality, etc.), inserts lexical information for subsequent realization, and so on.

Crucial to the interaction between the AS and EG levels is the notion *of salience.* The structure planned by the AS level does not explicitly employ linguistic goals: rather, it uses goals expressing the salience of facts to the hearer. It is the task of the EG level to determine how information should be made salient to the hearer - often, items need not be uttered for them to be salient, resulting in enthymematic argument.

Below the AS and EG levels, the inter-clause structure is resolved (employing relations similar to those of Rhetorical Structure Theory [Mann and Thompson, 1987]), and finally the grammatical, syntactic and morphological realization is performed through interface with the LOLITA system [Smith *et al.,* 1994].

The current work employs a concept of *argument* based upon the study of a corpus of natural arguments drawn from scientific papers, advertisements, editorial commentary and 'letters to the editor'. It is similar to that referred to as the "standard approach" in argumentation theory [Freeman, 1991], in which an argument is seen as comprising one or more premises contributing to one or more conclusions. Each premise can in turn be supported by a subargument, with the conclusion of the subargument acting as premise to the superargument. Finally, groups of premises contribute to a conclusion either independently (disjunct support) or in combination (conjunct support). This distinction is of crucial importance when attempting to determine the validity of an argument, [Freeman, 1991], and it is therefore surprising to find the distinction eschewed in many generation systems (e.g. [McConachy and Zukerman, 1996]).

2 Planning

The task of generating persuasive discourse is seen in this work, as in much -of NLG, as one of planning. NOAH [Sacerdoti, 1977] is frequently employed in text generation [Hovy, 1993], but is hampered by the fact that abstract operator bodies are primarily composed of operators. These operators are applied immediately the abstract operator is selected: the latter are thus acting as schema-like [4]recipes', and systems are liable to suffer from reduced flexibility as a consequence. In contrast, the current work makes use of AbNLP [Fox and Long, 1995], a hierarchical planner based upon the concept of *encapsulation,* whereby the body of an abstract operator contains goals rather than operators, and further, that the body of an operator is not opened up until an entire abstract plan has been completed (i.e. there are no goals left unfulfilled at that level of abstraction). On completion of an abstract plan (which can be seen, in discourse planning, as a skeletal outline of what is to be communicated), the *refinement* operation opens up all the abstract operator bodies, such that the structure and constraints determined at one level of abstraction are propagated to the next level down. As a consequence, many choices which might have been considered during planning of an argument at the detailed level can be pruned as they become inconsistent with the abstract plan. Such an approach has the potential to considerably improve upon the performance of a classical planner, [Bacchus and Yang, 1992]. The use of AbNLP in a framework for argumentative discourse planning is discussed in more detail in [Reed *et al,* 1996J.

The initial state in planning a persuasive text is described by two goals. Firstly, BEL (H, p), which represents the intention for the hearer, H, to believe some proposition, P. Thus it carries intentional rather than informational content [Young and Moore, 1994] and is classified as a communicative goal [Moore and Paris, 1994]. Secondly, the initial state also includes the goal IS_SALIENT (H, P, _), also a communicative goal. This represents the fact that the hearer must not just believe P, but that P must be salient to him (in the general discourse context, '_'- the use of this context parameter is discussed in more detail in |Reed *et al.,* 1997b]). In normal communication (i.e. communication which obeys Grice's maxims of cooperation, [Grice, 1975], and is therefore truthful and honest), every BEL goal will have a corresponding IS_SALIENT goal. It is the BEL goals which are planned for at the AS level: ultimately, the argumentation will have to reach beliefs which the speaker assumes the hearer either believes or will accept without further support. The IS_SALIENT goals, however, are usually satisfied by the MAKE_SALIENT operator which represents a primitive action at the AS level, and expresses the goal of making a proposition salient to the hearer (how this is achieved is determined by the EG level: some pieces of information in an argument are obvious and can be left implicit; others must be realized explicitly, in varying degrees of detail). The BEL goals are satisfied by the deductive, inductive and fallacial operators available to the AS level. Argument components are frequently linked by support relations which can be characterized by a Modus Ponens operator, in which the inference between antecedent and consequent is weaker than that of strict deduction. The MP operator is given below, in Figure 1 (note that similar operators are available for situations in which the hearer believes ~P, as opposed to those in which the hearer has no opinion on P).

In addition to the BEL and IS_SALIENT goals corresponding to two premises of Modus Ponens, the body of the operator also contains the topic manipulation goals, PUSH_TOPIC and POPJTOPIC. These are typically satisfied by primitive operators (of the same name, for convenience), which represent actions to be performed on a topic stack similar that proposed by [Grosz and Sidner, 1986]. The explicit operationalization of topic manipulation (rather than seeing focus as a constraint active over the planning process) offers a number of advantages in generating persuasive discourse. There is a close relationship between the position of topic manipulators and the occurrence of clue-words, which have been shown to be vital in ensuring the coherency of an argument [Cohen, 1987]. Similarly, at other levels of abstraction, topic manipulators co-occur with various punctuation and formatting (for example, three or four successive POPJTOPIC actions may warrant a paragraph break).

```
MP (H, P)
Shell: Precond:   ∃X: (X → P)
                  ~BEL (H, X)
       Add:       BEL (H, X)
Body:  Goals: 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)
```

Figure 1. Modus Ponens operator

Seeing PUSH_TOPIC and POPJTOPIC as goals suggests that it may also be appropriate to subdivide the class of communicative goals proposed by [Moore and Paris, 1994] into those which are intentional and those which are attentional (such an approach would also seem to be consistent with Moser and Moore's [1996] attempt to reconcile [Grosz and Sidner, 1986] with RST, [Mann and Thompson, 1987]).

The planning conducted at the AS level thus results in a highly parsimonious plan, in that it employs just three operators representing primitive actions at this level of abstraction: IS_SALIENT, PUSHJTOPIC and POPJTOPIC. For example, the goals in Figure 2 represents the abstract aim of persuading the hearer of *a.*

```
BEL (h, a)
IS_SALIENT (h, a, _)
```

Figure 2. Initial goals

If *a* could be supported by a single Modus Ponens argument, the BEL goal would be fulfilled by the MP operator, and the IS_SALIENT goal by MAKE_SALIENT, producing the abstract plan in Figure 3.

```
MP (h, a)    X = b
MAKE_SALIENT (h, a, _)
```

Figure 3. First complete abstract plan

During refinement, the MP body is opened up, from where the BEL goals are satisfied by matching those in the hearer model, the IS_SALIENT goals by MAKE_SALIENT operators and the topic manipulators by corresponding primitives. Thus the final plan of the AS level is as shown in Figure 4.

```
PUSH_TOPIC (a)
MAKE_SALIENT (h, b, a)
MAKE_SALIENT (h, b → a, a)
POP_TOPIC (a)
MAKE_SALIENT (h, a, _)
```

Figure 4. Plan of primitives for single step MP argument

However, although such a plan might convince a quite dispassionate audience (and may therefore suffice for inter-agent persuasion in a non-cooperative, multi-agent domain, [Reed *et* al., 1997a]), it may very well fail to be persuasive - or even coherent - simply because the order in which the information is presented is inappropriate.

3 Generating Order

Even within the coherency constraints imposed by the topic stack, ordering can still make the difference between coherent and incoherent—and at the very least, can improve the coherency of an argument (given that coherency seems to be a scalar attribute). Through analysis of arguments in the corpus, it was possible to determine structure, and then to modify that structure through permissible re-orderings, and finally to retranslate the structure back into text, and compare with the original. The results of re-ordering were often incoherent (see [Cohen, 1987] for examples of such incoherency) and almost always less persuasive than the original.

The fact that content ordering has an impact on the persuasive effect of an argument has also been demonstrated in psychology. McGuire [1968], for example, discusses how subarguments should be presented in *climax order,* i.e. with better supports coming later. Marcu [1996] provides a survey of how ordering and other techniques such as lexical choice affect the persuasive effect of arguments in the medical domain.

However, the most significant evidence of the importance of ordering comes from rhetoric, where heuristics are suggested for improving arguments so as to maximize their persuasive impact. Typical of such texts is [Blair, 1838], wherein are presented rules governing all aspects of public speaking, and in particular, a list of heuristics for determining an optimal ordering of subarguments. For these various heuristics to be implemented, it is first necessary to identify how the substrate of ordering is determined, i.e. given the structure of plans produced by the framework proposed in §2, over which components should any given ordering have scope.

Ordering occurs in two distinct phases in the planning process, corresponding to the two subtly different types of ordering necessary.

The first occurs with the fulfillment of BEL goals. The communicative goal BEL represents a problem for the classical planning framework, since it is inappropriate to consider it simply as an *achievement* goal (one which can be satisfied by a single operator). For a BEL is often best satisfied by several operators, i.e. by multiple subarguments (which have been shown to occur with great frequency in natural language, [Freeman, 1991]). However, it is also inappropriate to consider BEL a *maintenance* goal, such as the stylistic goals in Hovy's [1990] system, PAULINE. For it is not the case that BEL goals remain unsatisfied; rather, they are satisfied a number of times and are then considered fulfilled in the classical sense. This *iteration problem* in planning has generally been approached through the use of some 'for-all' operation. Maybury [1993], for example, makes explicit use of V in his operator descriptions, and Moore and Paris [1994] introduce a FORALL clause. In both papers, however, it is noted that the approach requires explicit, unprincipled modification of the plan language. In the absence of a principled solution, the current work employs a similarly pragmatic notion of universal quantification, to produce a maximal set of support by which to fulfill a BEL goal. This set of subarguments, in addition to the IS_SALIENT goal associated with the BEL (such as that in Figure 2), forms the substrate for a possible re-ordering. For example, *pre-order* (conclusion preceding premises) is generated by shifting the IS_JSALIENT ahead of all the subarguments. Similarly, *post-order* (all premises preceding the conclusion) and *hybrid* order are also effected at this phase.

However, despite the fact that most re-ordering can be accomplished during goal fulfillment, a separate phase is required to resolve ordering between operator body components. For example, the Modus Ponens operator presented in Figure 1 has two contributory parts in its body: X and X -> P. The order between these (i.e. between the two pairs tl-t2 and t3-t4) can only occur with refinement, and not with goal fulfillment.

These two forms of ordering are demonstrated in the example below, in which an argument has the structure shown in Figure 5, with a conclusion *a* supported by three premises *(b, c* and *d),* two of which are supported by further premises *(e* and *f)* (this example has been taken, with minor modification, from the corpus).

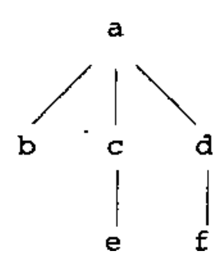

Figure 5. Sample argument structure

From the initial goals in Figure 2, the BEL (h, a) is fulfilled by three Modus Ponens arguments, as shown in Figure 6a. This structure then undergoes ordering, to produce Figure 6b (as discussed below, the ordering is the result of a persuasion-oriented heuristic due to the weakness of the conclusion, *a).*

(a)	(b)
MP (h, a) *X=b*	MP (h, a) *X=b*
MP (h, a) *X=c*	MAKE_SALIENT (h, a,)
MP (h, a) *X=d*	MP (h, a) *X=c*
MAKE_SALIENT (h,a,_)	MP (h, a) *X=d*

Figure 6. Result of planning (a) before and (b) after ordering

The three Modus Ponens subarguments are then opened up during refinement. The position of the primitives in the abstract plan (in this case, only one - the MAKE_SALIENT) remains the same at the next level of abstraction, and similarly, the arrangement of the subargument units is also preserved. This produces the situation in Figure 7a. However, the coherency of the plan in Figure 7a is compromised, due to the (relatively) large subargument which intervenes between the expression of the conclusion, a, and the support leant to it by the third subargument, *d.* One means of repairing the coherency is to reverse the order of the components in the final subargument such that *(d --> a)* is expressed before *d* itself, thus indicating to the hearer the relevance of *d* to the conclusion. This reordering is effected in Figure 7b.

(a)	(b)
PUSH_TOPIC (a)	PUSH_TOPIC (a)
BEL (h, b)	BEL (h, b)
IS_SALIENT (h, b, a)	IS_SALIENT (h, b, a)
BEL (h, b→a, a)	BEL (h, b→a, a)
IS_SALIENT (h,b→a,a)	IS_SALIENT (h,b→a,a)
POP_TOPIC (a)	POP_TOPIC (a)
MAKE_SALIENT (h,a,_)	MAKE_SALIENT (h,a,_)
PUSH_TOPIC (a)	PUSH_TOPIC (a)
BEL (h, c)	BEL (h, c)
IS_SALIENT (h, c, a)	IS_SALIENT (h, c, a)
BEL (h, c→a, a)	BEL (h, c→a, a)
IS_SALIENT (h,c→a,a)	IS_SALIENT (h,c→a,a)
POP_TOPIC (a)	POP_TOPIC (a)
PUSH_TOPIC (a)	PUSH_TOPIC (a)
BEL (h, d)	BEL (h, d→a)
IS_SALIENT (h, d, a)	IS_SALIENT (h,d→a,a)
BEL (h, d→a, a)	BEL (h, d, a)
IS_SALIENT (h,d→a,a)	IS_SALIENT (h, d ,a)
POP_TOPIC (a)	POP_TOPIC (a)

Figure 7. Result of refinement (a) before, and *(b)* after ordering

The argument is then completed with another round of planning (to plan the supports *e* and /), followed by refinement (to open up the bodies of the *e* and / subarguments), and finally, planning again (to resolve the remaining IS_SALIENT and topic manipulation goals). Clearly, however, the ordering depicted in Figure 7 is achieved through knowledge of the fact that the subargument for c: is relatively large, inasmuch as it has a supporting subsubargument, *e.* This knowledge is not directly available at this level, since subargument *c* has not yet been planned. In order to account for this phenomenon, and others based upon information regarding planning at lower levels of abstraction, AbNLP offers the concept of resource estimation. This process heuristically allocates resources to particular parts of an abstract plan, based upon contextual and hearer-model knowledge.

Estimates of subargument size are not the only determinants of ordering. Most other features are directly available at the current level of abstraction: structural features (such as inductive operators having precedence over fallacial operators); the surrounding structure (such as that causing the reordering in Figure 7); the source of the support (speaker beliefs, beliefs the speaker knows the hearer to hold, or beliefs grounded in an authority - *grounding* is discussed in [Reed *et al,* 1997a]). Perhaps the single most important feature to determine ordering is argument *strength.* Following [Freeman, 1991], strength is subdivided into *inferential force* and *persuasive force:* the former is purely a matter of determining the validity of the inference, whilst the latter is determined by assessing how the inference would be accepted by the hearer based on the speaker's models of his beliefs, competence, bias, scepticism, etc.

Persuasion-oriented ordering heuristics, in particular, make frequent use of this hearer-dependent facet of argument strength. Blair's [1838] dicta suggest that arguments should increase in strength (which accords with the psychological results of [McGuire, 1968], mentioned above), unless the hearer is particularly sceptical (in which case, a strong argument should come first, so that the hearer should attend to what follows), or if there is one weak argument amongst many strong (in which case it may be appropriate to hide it in the midst of the others). Furthermore, that if one argument is particularly strong it should be dwelt upon—and conversely, that many weak arguments should be grouped together (this is resolved as an issue of resource allocation).

Often at odds with those of persuasion, the coherency-oriented heuristics aim to keep the argument as simple as possible, minimizing the amount and length of information employed. When persuasion heuristics introduce or order information which could produce poorly coherent results, re-orderings are available to repair the structure, such as the reversal of antecedent and inference in an operator body (as in Figure 7). Persuasion ordering heuristics are not limited to planning phase ordering, nor are coherency ordering heuristics limited to refinement phase ordering. For example, the conjunction operator CONJ, has two conjuncts in its body, and these are subject to the same persuasion heuristics (such as climax ordering) as disjunct multiple subarguments. Equally, if the persuasion heuristics had not suggested the hybrid order in Figure 7, coherency heuristics would have restricted the ordering such that the conclusion was immediately adjacent to the final subargument.

The AS and EG levels can thus be seen as mediating between heuristics of persuasion and coherency. However, both coherency and persuasive effect can be improved at subsequent levels in the framework, employing the results of the ordering process. In particular, the positioning of the topic manipulators, PUSH_TOPIC and POP_TOPIC, can be used in the generation of a number of surface features, including punctuation, clue words and formatting.

4 Effects of Ordering

As suggested in [Cohen, 1987], *clue words* can be used to repair text which is weakly coherent, or even in some cases, completely incoherent. The EG level can in this way effect repair on a plan produced by the AS level which is less than optimally coherent (due, for example, to persuasion-oriented ordering heuristics). The original plan is not altered during this process; rather, clues are introduced into the plan to explicate its structure, and thus improve its overall coherence.

Cohen [1987] distinguished two groups of clues: redirective and connective. Clues of re-direction refer back to earlier pieces of information. For example, in a complex argument where a deeply nested subargument marks the end of several of its 'parent' superarguments, a clue phrase is often required to indicate which superargument is being returned to (e.g. *To return to the question of).* When the result of such an ordering is encountered by the EG level, the multiple consecutive POP_TOPIC operations which characterize the scenario cause the insertion of a re-direction clue. In natural language, however, clues of re-direction are often avoided, since they seem to indicate a lack of clear underlying structure. Correspondingly, many orderings are prohibited by the AS level - Cohen's *parallel evidence ordering*[1], for example, simply cannot be generated (due to incontrovertible coherency constraints). As a result, redirective clues are required infrequently.

Connective clues, in contrast, are used frequently in natural language to make explicit the relationship holding between adjacent components of an argument. Cohen identifies six categories of connective clues—a full exposition of how each category can be accounted for within the proposed framework is beyond the scope of this paper, but by way of example, the generation of inference, detail and parallel clues is considered below.

Inference clues occur between a premise and the conclusion to which it lends support—these are frequently employed in analyses in argumentation theory (e.g. the *conclusion locators* of [Wilson, 1986]). Inference clues *(so, therefore,* etc.) can be generated on locating a POP_TOPIC action immediately preceding the conclusion in an argument exhibiting pre-order.

Similarly, *detail* clues *(since, which follows from,* etc.) occur between a conclusion and its first supporting premise. These are the result of a PUSH_TOPIC action which immediately follows a conclusion in hybrid or post-order.

The relationship between subarguments at the same level of abstraction contributing to a single conclusion is termed *parallel* by Cohen. This is characterized in the plan produced by the AS level by repeated occurrences of POP_TOPIC (X) followed immediately by PUSH__TOPIC (X) (see Figure 7). In the simplest case, the co-occurrence can lead to the insertion of the lexeme *and.* However, the EG level has available information regarding scope, due to the PUSH_TOPIC and POP_TOPIC 'end-stops', from which the hierarchical structure can be determined. This enables more complex lexical realization such as numbering of subarguments.

Indeed, parallel clues are not just expressed lexically. Depending upon the level of abstraction, parallel structure can also be indicated through the use of formatting conventions (such as paragraph breaks and section numbering), and through the use of punctuation (particularly semi-colons between weak and briefly expressed subarguments).

The generation of clues is thus dependent upon the ordering determined at the AS level, and is also a means of repairing coherency which has been sacrificed in favour of persuasive effect.

5 Conclusion

The problem of generating a persuasive argument in natural language depends upon determining a good balance between coherency on the one hand and purely persuasive aspects on the other. Both these factors rely on a process of resolving appropriate ordering of content. Heuristics from a number of sources can be brought to bear upon this process, including those prescribed in rhetoric texts, those derived from empirical c bservations in psychology, and those inferred through the study of a corpus.

A hierarchical planning framework employing a parsimonious set of abstract primitives has been shown to be capable of generating the coarse-grained structure of an argument in an intuitive way. Furthermore, the approach affords the opportunity for clear, principled application of ordering rules. Finally, it has been demonstrated that this representation facilitates the subsequent realization of structure into text, and in particular, the generation of clue words which are known to be crucially important in argument comprehension.

[1] Parallel evidence ordering, Cohen claims, would, in the example in Figure 7, permit expressing items in the order *c, d, e,f* Such ordering requires significant clue word repair, and in the current work is simply considered incoherent. This can be justified on the grounds that the apparent occurrence of parallel evidence in natural language may in fact be the result of a summary preceding an argument proper—in Cohen's [1987] example, the conclusions are then left implicit in the argument for reasons of proximity to the summary.

References

[Bacchus and Yang, 1992] Fahiem Bacchus and Qiang Yang. The expected value of hierarchical problem-solving. In *Proceedings of the National Conference on*

Artificial Intelligence, pages 369-374,1992.

[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.

[Fox and Das, 1996] John Fox and Subrata Das. A Unified Framework for Hypothetical and Practical Reasoning (2) - Lessons from Medical Applications, In Dov Gabbay and Hans-Jiirgen Ohlbach, editors, *Practical Reasoning,* Springer Verlag, Berlin, pages 73-92, 1996.

[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.

[Freeman, 1991] James B. Freeman. *Dialectics and the Macrostructure of Arguments.* Foris, Berlin, 1991.

[Grice, 1975] H.P. Grice. Logic and Conversation. In P.Cole and J.L. Morgan, editors, *Syntax and Semantics III: Speech Acts,* pp 41-58, Academic Press, NY, 1975.

[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.

[Hovy, 1990] Eduard Hovy. Pragmatics and Natural Language Generation. *Artificial Intelligence* 43:153-197, 1990.

[Hovy, 1993] Eduard Hovy. Automated Discourse Generation Using Discourse Structure Relations. *Artificial Intelligence* 63:341-385, 1993.

[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.

[Mann and Thompson, 1987] William C. Mann and Sandra A. Thompson. Rhetorical structure theory: description and construction of text structures. In G. Kempen, editor, *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics.* Martinus Nijhoff, pages 85-95, 1987.

[Maybury, 1993] Mark T. Maybury. Communicative Acts for Generating Natural Language Arguments. *Proceedings of the 11th National Conference on Artificial Intelligence,* pages 357-364, 1993.

[McConachy and Zukerman, 1996] Richard McConachy and Ingrid Zukerman. Using Argument Graphs to Generate Arguments. *Proceedings of the 12th European Conference on Artificial Intelligence,* pages 592-596, Budapest, August 1996. John Wiley.

[McGuire, 1968] W.J. McGuire. The nature of attitudes and attitude change. In G. Lindzey and E. Aronson, editors, *The Handbook of Social Psychology Volume 3,* pages 136-314, Addison Wesley, 1968.

[Moser and Moore, 1996] Megan Moser and Johanna D. Moore. Towards a Synthesis of Two Accounts of Discourse Structure. *Computational Linguistics* 22(3):409-419, 1996.

[Moore and Paris, 1994] Johanna D. Moore, and Cecile L. Paris. Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information. *Computational Linguistics* 20(4):651-694, 1994.

[Paris, 1991] Cecile L. Paris. Generation and Explanation: Building an Explanation Facility for the Explainable Expert System Framework. Cecile L. Paris, W.R. Swartout, and William C. Mann, editors, *Natural Language Generation in Artificial Intelligence and Computational Linguistics.* Kluwer, pages 49-82, 1991.

[Rankin, 1993] Ivan Rankin. Natural language generation in critiquing. *Knowledge Engineering Review* 8(4):329-347, 1993.

[Reed *et al.,* 1996] Chris Reed, Derek Long, and Maria Fox. An Architecture for Argumentative Discourse Planning. In Dov Gabbay and Hans-Jiirgen Ohlbach, editors, *Practical Reasoning,* Springer Verlag, Berlin, pages 555-566, 1996.

|Reed *et al.,* 1997a] Chris Reed, Derek Long, Maria Fox, and Max Garagnani. Persuasion as a Form of Inter-Agent Negotiation. *Proceedings of the 2nd Australian Workshop on DAI,* Springer Verlag, Berlin, 1997.

[Reed *et al.,* 1997b] Chris Reed, Derek Long, and Maria Fox. Context and Focusing in Argumentative Discourse Planning. *Proceedings of the International Conference on Context,* Rio de Janeiro, Brazil, February 1997.

[Sacerdoti, 1977] E.D. Sacerdoti. *A Structure for Plans and Behaviour,* Elsevier, Amsterdam, 1977.

[Smith *et al.,* 1994] Mark H. Smith, Roberto Garigliano, and Richard G. Morgan. Generation in the LOLITA System: An Engineering Approach. *Proceedings of the 7th International Workshop on NLG,* Kennebunkport, Maine, 1994.

[Wilson, 1986] Barrie A. Wilson. The Anatomy of Argument (Revised Edition). *University Press of America.* Lanham,MD, 1986.

[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.

NATURAL LANGUAGE PROCESSING AND GRAPHICAL PRESENTATION

Natural-Language Processing 5: Dialogue and Discourse