
Generating Arguments in Natural Language

Chris Reed

**Submitted in fulfilment of the requirements
for the award of the degree of Doctor of Philosophy**

University College London

1998


ProQuest Number: U643524

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest U643524

Published by ProQuest LLC(2016). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

ABSTRACT

Automated generation of persuasive arguments has a wide range of potential applications, but represents a major challenge to existing natural language generation (NLG) techniques. In this thesis, it is argued that existing approaches fall short in several fundamental ways, and that handling argumentation demands major extensions to the conventional NLG model. Five key extensions are discussed. First, a distinction between the logical and rhetorical components of a text is advocated which is reflected in a similar modularisation of the planning task. Second, the adoption of an advanced style of hierarchical planning is proposed which is shown to mirror the hierarchical structure of argument, to increase generative flexibility, and to reduce computational cost. Third, the insufficiencies of a coherence-relation account are enumerated, and employed to motivate a more abstract representation layer drawing on the structural theories developed in argumentation theory. Fourth, conventional models in NLG have represented informational content; more recently, the role of intentional content has been emphasised; here, the importance of the attentional state and its explicit manipulation is also incorporated in a uniform way. Fifth, it is demonstrated that the generation of cue phrases between argument components relies not upon relations holding between clauses, but upon relations between more abstract units of text, and that those cues must necessarily therefore be introduced at an earlier stage of the planning process.

An architecture is proposed which integrates these extensions and formalises components of accounts offered in argumentation theory. This formalisation is carried out through a characterisation of deductive, inductive and 'fallacious' argument operators, including Modus Ponens, Modus Tollens, Inductive Generalisation and *Ignoratio Elenchi*. These argument forms are operationalised (in much the same way as Rhetorical Structure Theory relations have been) as planning operators which employ basic notions not only of belief, but also of saliency. Through a careful analysis of this distinction, argument forms such as the enthymeme, and rhetorical devices such as informing the hearer of facts which he is known already to believe, are shown to be easily accounted for. The architecture is implemented in the *Rhetorica* system, which encompasses layers of processing responsible for argument structuring and eloquence generation. *Rhetorica* also employs a body of thirty heuristics, which uniformly represent a variety of the most common guidelines listed in rhetoric and oratory texts of classical, renaissance, Victorian and contemporary authors.

The output of the *Rhetorica* system is a partially ordered plan of primitives which can be refined to lower levels of representation – and in particular, to coherence relation structures. This plan is expressed in a highly parsimonious language involving goals of attention manipulation and saliency, where the latter make reference to the attentional state through a context mechanism. Instances of potentially affect-laden cue phrases of an appropriate class are indicated by saliency goals introduced at the same level of abstraction as the textual units which the cues serve to link. The final plan represents the structure of an argument which, given the available information pertaining to the hearer and other situational factors, maximises both coherency and persuasive effect.

Table of Contents

EXORDIUM	8
THE PROBLEM SPACE	12
2.1 LANGUAGE	12
2.1.1 THE DIALOGIC SITUATION	16
2.1.2 AGENT COMMUNICATION	19
2.1.3 PLANNING	21
2.1.4 NATURAL LANGUAGE GENERATION	24
2.2 ARGUMENTATION	27
2.2.1 ARGUMENTATION THEORY: ANALYSIS TECHNIQUES AND STRUCTURAL MODELS	28
2.2.2 ARGUMENT AS A SOCIAL PHENOMENON	31
2.2.3 APPLIED ARGUMENTATION IN FORMAL REASONING	33
2.2.4 COMPUTATIONAL APPROACHES TO ARGUMENTATION	34
2.2.5 ARGUMENTATION IN MULTI-AGENT SYSTEMS	37
FRAMING A SOLUTION	41
3.1 DESIGN DECISIONS	41
3.1.1 MONOLOGUE OR DIALOGUE?	41
3.1.2 BLACKBOARD OR PIPELINE?	51
3.1.3 WHERE TO STOP?	54
3.1.4 RST OR ... SOMETHING ELSE?	55
3.1.5 FUNCTIONALIST OR FORMALIST?	59
3.1.6 BELIEF AND SALIENCY	60
3.2 ARCHITECTURE	63
3.3 WORKED EXAMPLE	70
FOCUS AND ORDER	75
4.1 MAINTAINING FOCUS	75
4.1.1 AS LEVEL PLANNING OPERATORS	79
4.1.2 AS LEVEL PLANNING	89
4.1.3 INTENTION AND ATTENTION	91
4.1.4 THE TOPIC STACK	93
4.2 ORDERING FOR COHERENCY	97
4.2.1 HEURISTICS FROM RHETORIC	97
4.2.2 HEURISTICS FROM INTUITION	99
4.2.3 HARD AND SOFT CONSTRAINTS	100
4.3 ORDERING FOR PERSUASIVE EFFECT	100
4.3.1 PERSUASIVE AND INFERENTIAL STRENGTH	101
4.3.2 HEURISTICS FROM RHETORIC	102
4.3.3 HEURISTICS FROM PSYCHOLOGY	104
4.4 PULLING IT TOGETHER	105
SURFACE FEATURES	111
5.1 CLUES	111
5.1.1 INTERCLAUSE CUE PHRASE GENERATION	112
5.1.2 INTERARGUMENT CLUE GENERATION	118
5.1.3 PUNCTUATION AND FORMATTING AS CLUES	122
5.2 EXPLICIT AND IMPLICIT SALIENCY	126
5.3 RHETORIC	129
5.4 TAGGING	132
5.5 PULLING IT TOGETHER	135
SYSTEM OUTPUT	141
PERORATION	149
7.1 APPRAISAL	149
7.2 FUTURE WORK	154
7.3 CONTRIBUTIONS	160
APPENDICES	166

Table of Figures

FIGURE 2.1 THE FOUR BASIC ARGUMENT STRUCTURES, AFTER (FREEMAN, 1991, p2).....	29
FIGURE 2.2 TOULMIN SCHEMA AND EXAMPLE.....	30
FIGURE 2.3 SAMPLE SITUATION BETWEEN SPEAKER, S AND HEARER, H.....	39
FIGURE 2.4 TYPES OF DIALOGUE, ADAPTED FROM (WALTON AND KRABBE, 1995, p66).....	40
FIGURE 3.1 SAMPLE PROCESS, (A), PRODUCT, (B), AND ANALYSIS, (C) OF MONOLOGUE.....	50
FIGURE 3.2 PIPELINE MODELS AND TERMINOLOGY.....	53
FIGURE 3.3 DEEP NESTING OF BELIEFS.....	61
FIGURE 3.4 ARCHITECTURE OF THE <i>RHETORICA</i> SYSTEM.....	63
FIGURE 3.5 (A) SAMPLE RST TREE AND (B) EQUIVALENT SEMNET REPRESENTATION.....	64
FIGURE 3.6 ABNLP OPERATOR CHARACTERISING MODUS PONENS.....	65
FIGURE 3.7 SAMPLE ARGUMENT STRUCTURE.....	70
FIGURE 3.8 FIRST AS PLAN WITH PARTIAL ORDER.....	71
FIGURE 3.9 RESULT OF REFINEMENT (A) BEFORE, AND (B) AFTER ORDERING.....	72
FIGURE 3.10 HIGH LEVEL PROCESSING CYCLE.....	73
FIGURE 4.1 MAJOR THEORIES OF FOCUS.....	78
FIGURE 4.2 THE NINE DEDUCTIVE OPERATORS AT THE AS LEVEL.....	82
FIGURE 4.3 SAMPLE SCENARIO FOR REBUTTING AND UNDERCUTTING.....	83
FIGURE 4.4 THE OPERATORS UCP AND UCI.....	84
FIGURE 4.5 THE HIERARCHICAL STRUCTURE OF AN INDUCTIVE GENERALISATION.....	85
FIGURE 4.6 THE OPERATORS IG AND ISUP.....	86
FIGURE 4.7 THE AP OPERATOR.....	88
FIGURE 4.8 THE IE OPERATOR.....	88
FIGURE 4.9 THE MAKE_SALIENT OPERATOR.....	89
FIGURE 4.10 THE PUSH_TOPIC AND POP_TOPIC OPERATORS.....	93
FIGURE 4.11 THE LIMITS OF PARTIAL ORDER.....	94
FIGURE 4.12 THE RUNTIME SPECIFICATION FOR MAKE_SALIENT.....	95
FIGURE 4.13 THE RUNTIME SPECIFICATION FOR PUSH_TOPIC AND POP_TOPIC.....	96
FIGURE 4.14 THE DEFINITION OF HCR1: GROUPING TOPICS.....	98
FIGURE 4.15 THE DEFINITION OF HCR2: REDUCE BREADTH.....	99
FIGURE 4.16 THE DEFINITION OF HCI1: REMINDING OF IMPLICATION.....	100
FIGURE 4.17 THE DEFINITION OF HPR1: CLIMAX ORDERING.....	102
FIGURE 4.18 THE DEFINITION OF HPR2: CONCLUSION FIRST.....	103
FIGURE 4.19 THE DEFINITION OF HPR3: WELL KNOWN FIRST.....	103
FIGURE 4.20 THE DEFINITION OF HPR4: REFUTATIONS FIRST.....	104
FIGURE 4.21 THE DEFINITION OF HPP1: NO WEAK REFUTATIONS.....	104
FIGURE 4.22 THE PREFERENCE ORDER AMONGST HEURISTICS FROM MOST TO LEAST PREFERRED.....	105
FIGURE 4.23 THE ARGUMENT STRUCTURE FOR THE VEGETARIANISM ARGUMENT.....	107
FIGURE 4.24 THE INITIAL KNOWLEDGE BASE FOR THE VEGETARIANISM ARGUMENT.....	108
FIGURE 4.25 CONTEXT PARAMETERS FOR THE VEGETARIANISM ARGUMENT.....	108
FIGURE 4.26 MIDWAY STATUS OF THE PLANNING PROCESS FOR THE VEGETARIANISM ARGUMENT.....	109
FIGURE 4.27 FINAL PLAN OF PRIMITIVES FOR VEGETARIANISM ARGUMENT.....	110
FIGURE 5.1 THE CLUE-MT HEURISTIC.....	115
FIGURE 5.2 EXAMPLES 6.10 AND 6.11 FROM (KNOTT, 1996, p103).....	115
FIGURE 5.3 THE CLUE-UCP AND CLUE-UCI HEURISTICS.....	118
FIGURE 5.4 CLUE CATEGORIES, AFTER (COHEN, 1987, p15).....	119
FIGURE 5.5 THE CLUE-PARALLEL HEURISTIC.....	119
FIGURE 5.6 THE CLUE-INFERENCE HEURISTIC.....	120
FIGURE 5.7 THE CLUE-DETAIL HEURISTIC.....	121
FIGURE 5.8 THE PUNC-BREAK HEURISTIC.....	124
FIGURE 5.9 THE PUNC-FN1 HEURISTIC.....	125
FIGURE 5.10 THE PUNC-FN2 HEURISTIC.....	126
FIGURE 5.11 THE EC-EPISTEMIC2 HEURISTIC.....	128
FIGURE 5.12 THE EC-LENGTH1 HEURISTIC.....	129
FIGURE 5.13 THE RHETORIC-PE HEURISTIC.....	131
FIGURE 5.14 THE EC-QUALIFIER1 AND EC-QUALIFIER2 HEURISTICS.....	132
FIGURE 5.15 THE AFFECT-UC HEURISTIC.....	134
FIGURE 5.16 THE AFFECT-GOOD HEURISTICS.....	134
FIGURE 5.17 THE CLUE-DS HEURISTIC.....	135

FIGURE 5.18 THE CLUE-DS INSTANTIATION IN THE VEGETARIANISM ARGUMENT.....	136
FIGURE 5.19 THE PUNC-BREAK HEURISTIC INSTANTIATED IN THE VEGETARIANISM ARGUMENT	137
FIGURE 6.1 STRUCTURE OF THE 'TOURIST FACILITY SIGNS' ARGUMENT	142
FIGURE 6.2 SYSTEM BELIEFS AND PARAMETERS FOR THE 'TOURIST FACILITY SIGNS' ARGUMENT	142
FIGURE 6.3 FINAL <i>RHETORICA</i> OUTPUT FOR THE 'TOURIST FACILITY SIGNS' ARGUMENT.....	143
FIGURE 6.4 STRUCTURE OF THE 'CLARE SHORT' ARGUMENT	145
FIGURE 6.5 SYSTEM BELIEFS FOR THE 'CLARE SHORT' ARGUMENT	145
FIGURE 6.6 FINAL <i>RHETORICA</i> OUTPUT FOR THE 'CLARE SHORT' ARGUMENT.....	146
FIGURE 6.7 STRUCTURE OF THE 'IRISH CENSUS' ARGUMENT	147
FIGURE 6.8 SYSTEM BELIEFS IN THE 'IRISH CENSUS' ARGUMENT.....	147
FIGURE 6.9 FINAL <i>RHETORICA</i> OUTPUT FOR THE 'IRISH CENSUS' ARGUMENT.....	148
FIGURE 7.1 SUMMARY OF RESULTS FROM PRES AS RAW DATA AND RANKINGS	152
FIGURE 7.2 SUMMARY OF RAW RESULTS FROM PRES.....	153
FIGURE 7.3 SUMMARY OF RESULTS FROM PRES ANALYSED AS RANKINGS.....	153
FIGURE 7.4 SUMMARY OF RESULTS ON PRIMACY/RECENCY EFFECTS IN PRES	155
FIGURE C.1 <i>RHETORICA</i> PROCESSING FOR THE 'TOURIST FACILITY SIGNS' ARGUMENT.....	173
FIGURE C.2 <i>RHETORICA</i> -MINUS-EG PROCESSING FOR THE 'TOURIST FACILITY SIGNS' ARGUMENT	176
FIGURE C.3 <i>RHETORICA</i> PROCESSING FOR THE 'CLARE SHORT' ARGUMENT.....	177
FIGURE C.4 <i>RHETORICA</i> -MINUS-EG PROCESSING FOR THE 'CLARE SHORT' ARGUMENT	179
FIGURE C.5 <i>RHETORICA</i> PROCESSING FOR THE 'IRISH CENSUS' ARGUMENT	180
FIGURE C.6 <i>RHETORICA</i> -MINUS-EG PROCESSING FOR THE 'IRISH CENSUS' ARGUMENT.....	181
FIGURE C.7 THE RESULTS OF PRES	183

Acknowledgements

I gratefully acknowledge the financial support provided by EPSRC, and the additional contributions of the following bodies: AISB, Brunel University Department of IS & Computing, Durham University Department of Computer Science, IJCAI, SSHRC, UCL Department of Computer Science, and the UCL Graduate School.

During the course of my doctoral studies, I have had the privilege to discuss preliminary (and often incoherent) ideas with eclectic, exciting, stimulating and exceedingly patient researchers. I would like to thank the following for ideas and comments which have contributed to the development of my work: Tony Blair, Robert Dale, Robin Fawcett, John Fox, Maria Fox, James Freeman, Roberto Garigliano, Michael Gilbert, Tom Gordon, David Green, Stephen Green, Leo Groarke, Mandy Haggith, Hans Hansen, Jaakko Hintikka, Graeme Hirst, Helmut Horacek, Ed Hovy, Xiaorong Huang, Kristiina Jokinen, Alistair Knott, Eric Krabbe, Daniel Marcu, Tim Norman, Franck Panaget, Cecile Paris, the late Elsa Pascual, Larry Powers, Ivan Rankin, Ehud Reiter, Francisca Snoeck Henkemans, Bart Verheij, Doug Walton, and John Woods. I would also like to thank my examiners, Donia Scott and Simon Parsons, for their diligent reading of my work and their useful comments.

I have also had the benefit of the support and encouragement of friends and family who put up with me going on and on about argumentation. Thanks to all, and particular thanks to Mum, Dad, John, Aspasia, Nancy and Jon.

For teaching me how to write, for long discussions, for giving me the freedom just to get on with it, for answering the phone and discussing papers in the wee small hours of a deadline day, for an excuse to visit Durham regularly, for occasional motivation boosts, for a scrupulous approach to research, for pointing out my mistakes and helping me towards solutions, for lending me good books that he will now finally get back, for the best vegetarian cooking, and for unbelievable patience, I would like to thank my supervisor, Derek Long. Without Derek, it would not have been nearly so much fun.

Finally, I owe a huge debt of gratitude to Cathy, who has been a steadfast support from the very start. Through the good times of drinking rum and discussing monologue, and through the bad times of seemingly interminable thesis writing, Cathy has always been there for me. Thank you for everything.

I

Exordium

Persuasive text has a variety and abundance of entomological proportion. Advertisements, editorials, academic papers, letters to the editor, parliamentary speeches, materials for education, religious pamphlets and more, all offer examples of text which aims to alter the beliefs of some audience, and might therefore be considered persuasive. Yet despite the enormously important role played by persuasion in natural communication, only a handful of models have been built to investigate the process by which such text is created. Furthermore, there is a strong trend amongst this work to consider only the logical structure of an argument, and to denigrate textual argument from the status of an elegant, complex interplay between linguistic, psychological and interpersonal factors, to that of little more than a set of propositions. The reasons for this approach are doubtless rooted in pragmatism, expediency, and a desire for simplicity, but the result is a reduction in flexibility and expressiveness to a level at which almost none of the phenomena explored by argumentation theorists from Aristotle onwards can be accounted for. To describe a praying mantis and a giant peacock moth alike, by enumerating legs, wings and antennae, is to miss the point rather.

The abundance of examples of argumentation in natural communication comprises only one of a much wider set of features motivating such a specific focus in the current work. Chief amongst these is the highly structured nature of argument. Although almost all text exhibits an intrinsic, functional structure, argument often occurs in situations in which (a) there is a large degree of interconnection between many textual units, and (b) it is important to ensure that the hearer is following. As a result, argument is often more highly structured than other forms of naturally occurring text, and also demonstrates more explicit marking of that structure. Although the analysis of argument is fraught with problems of equivocation, subjectivity and a lack of consensus on even basic terms and techniques, research in argumentation theory has developed a range of methods for determining the underlying structure of an argument, and these methods are significantly more advanced than comparable approaches to linguistic structure in general. Thus although structure is difficult to determine, the techniques of argumentation theory better equip the analyst to develop a richer, more informed

representation than alternatives available in the pragmatics community as a whole. In addition to the propensity for argument to employ and manifest a high degree of structuring, it is also clear that most argument involves, specifically, a hierarchical organisation of components. This hierarchical nature can be exploited in the employment of a hierarchical generation technique: such a technique is not only available, but also has a good pedigree in artificial intelligence, and has been demonstrated to be highly suited to the natural language generation (NLG) field. The use of hierarchical planning in NLG is an established approach to which the current work subscribes, but with one important departure with respect to the style of hierarchical planning employed.

One final aspect motivating investigation of argument in particular is the potential utility of a system capable of generating NL argument. Persuasive argument is primarily concerned with shifting belief in a given audience (the picture is actually rather more complex than this, as explained below). Thus wherever it is necessary to alter the beliefs of a user of a system, the ability to generate persuasive argument is critical. Many examples are currently under investigation in NLG – health education materials aimed at inducing a disinclination to smoking, expert systems justifying their conclusions, decision support systems offering critiques of user decisions, etc. Many more can be conceived of – a service provider demonstrating its superiority, generating adverts tailored to a particular audience, creating political campaign materials, etc.

The current work, then, focuses on the genre of persuasive text, and in so doing, cuts through a range of important issues in NLG. One recent concern in NLG has been the role played by intentions in the generation process, contrasting with the primarily informational approach predominant until the early 90's. The current work proposes an approach which not only integrates informational and intentional facets of generation, but also draws in the attentional component into a single framework, thus uniformly modelling each element of the triumvirate controlling discourse structure. Another current NLG issue addressed concerns the introduction and placement of cue phrases (or clues in the specific of argumentation). A major tenet of the current work is that clue introduction is often dependent upon very high level processing – specifically, that as a clue may function to relate two large segments of text (e.g. a section break), it is appropriate for a generation algorithm to introduce such a clue at the same level of abstraction as the units of text which it connects.

In exploring the central topic of argument generation, a number of novel claims are posited. One of the key claims forms a direct response to that research, mentioned in the first paragraph, which exhibits a tendency towards stripping argument of all but its logical composition. Although this logical aspect clearly has a role to play, the thesis returns frequently to the idea that 'logical' and 'rhetorical' components can and should be distinguished and separated in the generation model. The quoting around the terms emphasises that they are used to refer to a slightly different set of features than might be suggested by intuition: by 'logical' is meant the propositional content and the interrelations (deductive, inductive and fallacious) holding between propositions; by 'rhetorical' is meant structural features such as ordering and constraints on subargument length, stylistic features such as affect, and other extra-'logical' aspects such as propositionally void text. There is a close relationship between the logical/rhetorical distinction embodied in modular partitioning of functionality, and two key resultant properties of the text: coherency and persuasive effect. In the same way that previous research has

concentrated predominantly on the logical structure of argument, so there has been an implicit assumption in much relevant generation work that coherency will ensure persuasiveness. Yet psychological research has demonstrated that such a notion of idealised human rationality is mistaken: ordering of pro and con arguments, phrasing, and repetition have all been demonstrated to significantly impact the reception of a particular text. Work in argumentation theory and rhetoric contributes to an exploration of the distinction and its manifestation in text, and from there, its role in the generation process.

In addition to this dualistic underpinning explored primarily through the distinctions between logic and rhetoric, and coherency and persuasive effect, there are also a number of further claims of a more technical nature. These include an approach to the handling of the disjunctive constraints holding over the partial order of a plan, the development of a highly parsimonious representation language, the explicit computational characterisation of the relationship between linked and convergent argument support, and the design of a control structure for high-level processing which departs from the classical pipeline model.

One swathe of NLG issues which has been largely ignored in the current work is that of tactical generation – the problems of deciding which lexical, syntactic, and morphological constructions best convey the intended message. In this division of the generation task between tactical and strategic levels of processing, the work follows the prevailing view in NLG, and similarly has several precedents in its implementation of only the strategic level. Although following this route distances the work from real text, and as a result stores up problems for evaluation, there are a number of benefits, chief of which is the ability to focus – as discussed below, if the assumptions made about other, unimplemented, stages in the framework are realistic (e.g. by reference to existing research investigating those stages), then the decoupling is less worrisome, and provides an opportunity to concentrate solely on higher level functionality.

The theoretical claims that are laid out are further supported through their integration into the *Rhetorica*¹ system which implements the upper (i.e. more abstract) layers of the generation framework. Each of the various facets of argument generation are formalised and implemented, resulting in a coherent system which not only demonstrates the validity of the theoretical approach, but also offers a means of evaluating the components of that approach empirically.

The intended output of *Rhetorica* is a text plan; a (tagged) representation of the structure of an argument. To produce this output, the system relies upon operators and heuristics derived from seven sources. First, a core body of structuring techniques are adopted from classical logic, supporting the basic techniques of rational reasoning. Second, contributions from computational linguistics are adduced, particularly from computational accounts of the linguistic structure and form of argumentation. Third, argumentation theory itself lends further components which extend the classical paradigm. Fourth, maxims from rhetoric and oration are included – these transcend the rational core and provide hearer-centred heuristics which focus on techniques of persuasion, rather than of reasoning. Fifth, rules can be inferred from studies of attitude change in social psychology – again,

¹ The name is both an acronym – Realising Heuristic-based Eloquence and Textual Ordering Resulting In Coherent Arguments – and a foreshortening of its full title, *Rhetorica ab Computatis*, a reference to the mediaeval text book on rhetoric, *Rhetorica ad Herennium*.

these focus on persuasive rather than deductive aspects of argument. Sixth, some few additional rules are suggested merely by intuition – these are primarily features which rhetoricians and psychologists have taken for granted, but which need to be made explicit in an operationalisation. Finally, there are also contributions from the study of a small corpus. This corpus is composed of several dozen arguments taken from a variety of sources, including letters to the editor of a newspaper, editorials and leader comments in magazines, advertisements, and a precis of a monograph in paleontology. In most cases, generalisations of structuring rules observed in this corpus are already covered by components of one of the first six categories, but several additional rules were also extracted. It is the identification, formalisation and integration of rules and heuristics from these various sources which forms the core of the current work.

The thesis is broadly structured in such a way that the two key threads – NLG and argumentation – are first presented and reviewed separately, then introduced to one another and gradually woven together into a coherent whole, from which it is demonstrated that a variety of benefits and avenues for future investigation can be derived. Chapter two performs the initial task of reviewing first NLG research, and then argumentation theory. Though primarily compendious and exegetical, the chapter also identifies problems in argumentation theory, misunderstandings and misconstructions of argument in the computational domain, and current issues in NLG with particular consideration for the generation of argument. Chapter three then offers solutions to some of these issues in the process of dissecting and motivating the major design decisions taken in building the *Rhetorica* system. These decisions range perforce over argumentation, NLG, and the intersection of the two. Thus for example, the discussion positing a definition for persuasive monologue draws heavily upon argumentation theory; the investigation of the relative appropriateness of blackboard and pipeline models of generation is primarily concerned with NLG research; and the problems of how to deal with the interaction between belief and saliency, and of the role of Rhetorical Structure Theory in generation are tackled with both fields impinging on the discussion. The second role of Chapter three is to perform a preliminary integration of the various ideas into an overarching framework; this is achieved through the investigation of a small worked example. Chapter four then develops the architecture in more depth, fleshing out the functionality of *Rhetorica*, and discussing both the planning operators and coherency-oriented heuristics which impinge upon the high level generation process. To illustrate the various components, an extended example is presented in some detail. Chapter five further motivates the adoption of the architecture laid out in chapters three and four by exploring how it lends itself to the generation of surface features, and in particular, of clue phrases. The example of chapter four is continued to illustrate the machinery of clue introduction. Chapter six then offers a concrete demonstration of the system's activity through three examples of input and output. Finally, chapter seven summarises directions for future work which have been opened up by this study, and presents an evaluation of *Rhetorica* based in part upon definitions of the terms persuasion and coherency, and in part upon a small experimental study.

II

The Problem Space

The two threads from which the thesis is composed are here introduced and their history explored: firstly, the linguistic components, including a brief scene-setting summary on the theoretical background to many assumptions made in natural language generation (primarily on the nature of speech act theory), followed by a discussion of the dialogic situation and its components. Finally, the key milestones in the development of NLG theory are discussed, and a survey of more recent work presented. The second thread is then introduced with a survey of current work in argumentation theory (again motivated by a brief discussion of the roots of the work), summarising several major models of argument analysis including those founded on social and interpersonal considerations. This work in ‘informal logic’ is then contrasted with more conventional, formal approaches to argument. Such formalisation carries with it a range of problems which are then inherited by computational systems adopting the formal approach: the formalisations, systems and associated problems are explored in order to motivate the framework presented in the next chapter.

2.1 Language

Before the work of Austin, straightforward sentences had generally been regarded in the philosophy community as constative statements with which a truth value could be associated². In (Austin, 1976), however, it is demonstrated that sentences such as *I bet you sixpence it will rain tomorrow* cannot be constative – it makes little sense to claim that the statement is true or false. Rather, the sentence itself is an act:

“... to utter the sentence is not to *describe* my doing of what I should be said in so uttering to be doing or to state that I am doing it: it is to do it.” (Austin, 1976, p6)

Sentences of this form Austin terms *performatives*: utterances which are in themselves actions. The notion is then extended through consideration of a more general class of utterances: stating *I believe X*

is not simply a constative communication, it is also an act – one of stating belief (compare, for example the more explicit performative *I state that I believe X*).

Austin then goes on to explore the variety of ways in which such performatives can be ‘unhappy’ or infelicitous, i.e. the situations in which they are void or abused. He identifies six constraints on the successful execution of a performative, and each then forms the basis for a characterisation of a particular infelicity. A *misapplication* occurs in a situation where a conventional procedure is incorrectly instantiated by actors and objects (e.g. a layman declaring two people married); a related but rarer infelicity occurs where such a conventional procedure does not exist at all. Together, these two are termed *misinvocations*. In a situation where the conventional procedure is incorrectly or incompletely executed, the infelicity is termed *misexecution* (and, specifically, a *flaw* refers to incorrect and *hitch* to incomplete execution). Finally, an *abuse* results from the insincerity of one or more of the actors, either in not possessing the thoughts and feelings demanded by the conventional procedure (e.g. betting without intending to pay) or in failing to meet constraints on subsequent conduct (e.g. failing to pay a gambling debt).

For the purposes of the current discussion, however, this taxonomy of the infelicities of speech acts is of less direct concern than the subsequent complimentary development of a theory of the necessary and sufficient conditions for the successful performance of speech acts, proposed in (Searle, 1969). Searle builds on Austin’s distinction between *locutionary*, *illocutionary* and *perlocutionary* acts, all of which are present in any utterance. The locutionary component is simply the production of noise which accords with the phonetic, morphological and syntactic constraints of a language; the illocutionary component characterises the act performed *in* making the appropriate noise (as opposed to the performance of the act *of* making that noise); the perlocutionary component is the effect that that noise may have – the act performed *by* making the noise. Austin’s example (Austin, 1976, p102) distinguishes a locutionary act *He said ‘Shoot her!’* from a corresponding illocutionary act, *He urged me to shoot her*, and again from a perlocutionary act, *He persuaded me to shoot her*. Searle focuses upon the structure of illocutionary acts, and sets out precise prerequisites for their successful execution (Searle, 1969, pp66-7). These prerequisites make use of two types of constraint: those concerning the beliefs of the speaker, and those concerning the intentions of the speaker. For the latter, Searle draws upon Grice’s intention analysis, (Grice, 1957), where it is proposed that “a speaker *S* meant something by *X* where *S* intended the utterance of *X* to produce some effect in a hearer *H* by means of the recognition of this intention.” (paraphrased in (Searle, 1969, p43)). Searle develops this proposal in two respects: first by characterising the role played by conventions in determining meaning, and second by modifying the implicit assumption that the speaker intends some perlocutionary effect. This second development is in recognition of the subtle distinction between a hearer *understanding* a locution, and being persuaded/convinced/etc. The latter is the perlocutionary effect described in the Gricean analysis, the former the illocutionary effect employed in Searle’s definition (though of course the definition then requires some analysis of what it is to ‘understand’ a locution: Searle’s analysis employs both the

² As Austin points out, earlier philosophy had not failed to recognise that some sentences are nonsensical (in a non-grammatical way) or involve clauses which are not descriptive, hence the restriction on the generalisation to ‘straightforward’ sentences.

characterisation of conventions in determining meaning and the reflexive account of illocutionary effect).

The conditions which must be met for the successful carriage of various classes of illocutionary acts constrain four facets: the propositional content (e.g. in a promise, that the utterance predicates some future act *A* of *S*), the necessary preparatory conditions (e.g. in a promise, that *H* would prefer *S*'s doing *A* to her not doing *A*, *S* believes *H* has this preference, and *S* would not normally perform *A*), the requisite sincerity (e.g. in a promise, that *S* intends to do *A*), and the 'essential' condition, or what the act constitutes (e.g. in a promise, that *S* intends that the utterance will place her under an obligation to perform *A*). The means by which uttering a statement meeting the propositional content, preparatory and sincerity conditions constitutes the realisation of the essential condition are then defined in terms of the illocutionary effect mentioned above³.

The clarity of Searle's characterisation of the constraints by which illocutionary acts are bound, and the lack of a formal analysis of the adequacy of his account, led to an attempt by Cohen and Perrault to provide such an analysis through a computational model of speech act arrangement (Cohen and Perrault, 1988). The first premise of their model is that language can be characterised as a goal directed behaviour (and indeed, a similar assumption underpins most research in computational approaches to the pragmatic aspects of natural language). Although a similar assumption is present in the work of both Austin and Searle, it is important to recognise a potential misunderstanding which may be the result either of the translation from Searle to computational work, or from the development of the computational theories over time. In these theories, the goal invariably involves the state of some hearer; in Austin and Searle (and more explicitly in the latter), this need not be the case, particularly as it could lead to the conflation of illocutionary and perlocutionary effect – the disentanglement of which was one of Searle's major contributions. Searle gives an example: "I may make a statement without caring whether my audience believes it or not but simply because I feel it my duty to make it." (p46). Although this example could well be accounted for in terms of goal fulfilment, it is rather different to the speech act goals explicated in (Cohen and Perrault, 1988).

Cohen and Perrault's assumption of goal-directed behaviour permits their development of a plan-based approach, viewing individual illocutionary acts as plan operators with specific preconditions (i.e. those described by Searle under the heads 'propositional content', 'preparatory', etc.), and

³ For a promise, this is defined by Searle (p60) thus: *S intends (i-1) to produce in H the knowledge (K) that the utterance of T is to count as placing S under an obligation to do A. S intends to produce K by means of the recognition of i-1, and he intends i-1 to be recognised in virtue of (by means of) H's knowledge of the meaning of T.*

characteristic effects (i.e. the essential condition)⁴. The planning itself employed by Cohen and Perrault is unremarkable (its predecessors and successors are discussed in more detail in §2.1.3) but its use in organising the pragmatic structure of text represents a crucial development in natural language generation (surveyed in §2.1.4) – despite the primary aim of the paper being one of substantiating claims in the philosophy of language, rather than introducing a new framework for the construction of applied text generation systems. It is interesting to note that speaker intentions at the level of individual illocutionary acts were recognised to play a key role; at the level of larger pieces of text, addressed in subsequent NLG work, this role was heavily underestimated until quite recently. The development of intention-laden approaches is traced in §2.1.4, and the discussion returns to the issue in chapters three and four.

An important development of Grice's formulation is presented in (Sperber and Wilson, 1986), which adduces the aforementioned results from Searle, in addition to alternative reformulations by Strawson, Schiffer, etc. (placing particular emphasis on (Strawson, 1964)). Sperber and Wilson focus upon the situated act of communication and its cognitive components. Their starting point is that communication is *ostensive* behaviour: "behaviour which makes manifest an intention to make something manifest" (p49), where manifestness is defined as "[being] capable of representing [a fact] mentally and accepting its representation as true or probably true" (p39)⁵. Facts are more or less manifest within an agent's *cognitive environment*, a combination of his perceived environment and his cognitive abilities (i.e. of percepts and stored facts). Sperber and Wilson then define the situation of communication as one involving a *mutual cognitive environment*, the intersection of interlocutors' individual cognitive environments, in which it is also manifest which agents share it. They also demonstrate (p41) that the notion of a mutual cognitive environment, unlike mutual knowledge or mutual belief, is tractable (the discussion returns in more detail to the problems of mutual belief and the notion of saliency – akin to that of manifestness – in §3.1.6).

One of the key claims in (Sperber and Wilson, 1986) concerns the *informative intention* of a speaker, and the means by which this can lead to a definition of relevance. Given the definitions of manifestness and the ostention of communicative behaviour, Sperber and Wilson claim (p58) that a speaker has an informative intention in making an utterance to make (more) manifest a set of facts $\{I\}$ (though an explicit individuation of $\{I\}$ is not required). A speaker's intention is thus to modify the

⁴ Note that Searle distinguishes the essential condition (p60) "S *intends* that the utterance T will place him under an obligation to do A" (emphasis added), from the essential rule (p63) "T counts as the undertaking of an obligation to do A". The former constrains the application of the illocutionary act, the latter the linguistic devices used to convey the associated illocutionary force. This distinction seems not to have been recognised by Cohen and Perrault, with the result that they map from the essential *rule* (which they term the essential *condition*) to the effect of a planning operator (i.e. an illocutionary act). This seems to lose an important aspect of Searle's account, in which mapping from the essential condition (i.e. S's intention that an utterance should count, in the case of a promise, as an obligation to perform A) to the essential rule (i.e. the that the utterance actually does count as that obligation) is nontrivial. Indeed, the mapping requires invocation of the post-Gricean definition of illocutionary effect mentioned above: this invocation is termed by Cohen and Perrault the *force condition*, about which they state "we have chosen not to deal with the force condition until we have a better understanding of the plans for speech acts and how they can be recognized" (Cohen and Perrault, 1988, p175). It is unclear to what extent this omission damages the account and those drawing upon it, but it is worrying that the effects of the plan operators – i.e. the illocutionary acts – are characterised in terms of what Searle intended to be viewed as constraints on linguistic realisation.

⁵ Thus definition of ostensive communicative behaviour thus admits of the full range of non-verbal communication.

cognitive environment of the hearer (and thereby their mutual cognitive environment). On hearing an utterance, an interlocutor thus has a number of sets of information available: information that is already known (and manifest), information manifest in the perceptual environment, and information made manifest through ostensive communication. The first set, $\{C\}$, is 'old' information; the two others together present 'new' information, $\{P\}$. There are three potential relations between the two sets $\{C\}$ and $\{P\}$: (i) $\{P\}$ may be so different as to yield no *contextual implications* – no new deducible information based on the union of $\{C\}$ and $\{P\}$ (e.g. reading an advanced academic paper in a technical field which is not one's own); (ii) $\{P\}$ may be so small or so trivial (with respect to $\{C\}$) as to allow no contextual implications (e.g. reading a glossary which serves only to equate terms with which one is already familiar); (iii) $\{P\}$ may be sufficiently new and interesting – but not so unconnected with $\{C\}$ – as to allow a multiplicity of contextual implications (e.g. reading a technical paper in one's own field). In the third case only is the information in $\{P\}$ *relevant* – the more contextual implications which are derivable, the more relevant the information⁶. Finally, Sperber and Wilson demonstrate that "every act of ostensive communication communicates the presumption of its own optimal relevance" (p158), which they term the *principle of relevance* – the 'presumption of optimal relevance' draws on two parts: first, that the set $\{I\}$ which the speaker intends to make manifest is sufficiently relevant (using the above definition) to make it worth the while of the hearer to process the communication, and second, that the ostensive act is the most relevant that could have been used to communicate $\{I\}$ ⁷.

One of the key developments in (Sperber and Wilson, 1986) is the emphasis placed upon the whole dialogic situation, rather than taking a more restricted, heavily speaker-centred view, as adopted by Searle, and developed explicitly in (Cohen and Perrault, 1988) as the "point of view principle". The next section is devoted to characterising in more detail the various aspects from which this rich notion of the dialogic situation is comprised.

2.1.1 The Dialogic Situation

To engage in dialogue, a system needs to maintain a model of the user. This necessity arises from a number of demands placed upon such a system. In the first place, a basic understanding of user utterances relies upon more than just an ability to analyse the communicated message: deictic reference, anaphora, ellipsis and indirect speech acts all require some understanding of the cognitive state of the user (in Sperber and Wilson terms, of the user's cognitive environment). Secondly, most dialogue (task-oriented, information-giving, etc.) is cooperative in nature; to achieve this cooperation, a system must be able to infer the beliefs, goals, plans, and intentions of the user. A third requirement, extending a minimal level of cooperation, also relies upon competent user modelling: the ability to volunteer unsolicited information, as found in Jameson's IMP system (Jameson, 1989). Fourthly, the generation of system utterances needs to be tailored to the user: ellipsis, for example, can be tested against a model of the user to determine whether or not it would be understood (using an *anticipation feedback loop*), as employed in the HAM-ANS system (Jameson and Wahlster, 1982). The same process can also be used to determine appropriate *affect* (Hovy, 1986) – i.e. how to select appropriate

⁶ In fact, Sperber and Wilson's characterisation is rather more subtle than this, in that it considers the relative weights of information in $\{C\}$ and $\{P\}$: this refinement is not relevant to the current discussion.

⁷ The second part of the definition of presumption of optimal relevance thus seems to be a refinement of Searle's essential rule, though this analogy is not drawn by Sperber and Wilson themselves.

lexicalisation to pander to a user's particular bias, also employed in HAM-ANS (Morik, 1989). Fifthly, the ability of the user to understand various versions of a text should not be assumed: Paris' TAILOR system (Paris, 1993), for example, distinguished novice and expert users in tailoring explanation construction. Finally, another imperfection in the user's competence makes a further demand on a system's user model: the ability to cope with user misconceptions (McCoy, 1986), (Calistri-Yeh, 1991). Together, these demands suggest that "user models constitute an indispensable prerequisite for any flexible dialog [system]" (Wahlster and Kobsa, 1989, p5) (though Sparck Jones (1991) cautions against building models which are unnecessarily sophisticated).

Although the structure of user models exhibits great diversity, several features are common to a number of systems (Kass and Finin, 1988). The starting point for many user models is a set of default assumptions about a user. Such defaults are often arranged into sets associated with a particular user class, providing a stereotypical characterisation of that class of user. One of the earliest systems to adopt this approach was GRUNDY (Rich, 1989), which requested the user to give a short description of themselves, on the basis of which, GRUNDY would identify appropriate stereotypes. In extensions to the basic notion, individual users could belong to numerous stereotypical classes, with the latter arranged in specificity hierarchies to resolve conflicts (Finin, 1989). If appropriate, specific observations could then override the values determined by reference to the stereotype (thus demanding nonmonotonicity in the supporting reasoning mechanism).

The model itself is typically founded upon some notion of belief, frequently one derived from Hintikka's (Hintikka, 1962) modal characterisation (though popular alternatives are surveyed in (Wahlster and Kobsa, 1989)). Furthermore, the beliefs of the user are rarely captured adequately by a static model: in order to handle the necessary dynamic nature, a truth maintenance system (either justification based (Doyle, 1979), as used in TRUMP (Bonarini, 1987) or assumption based (de Kleer, 1986) as used in GUMS (Finin, 1989)) is required.

Dynamic update of the user model during a dialogue involves a number of challenging tasks: key amongst these are belief ascription and plan ascription. Though clearly related, the former concentrates on inference drawn at the level of the individual sentence (considering presupposition, affect, etc.) (Jameson, 1989), whilst the latter focuses on higher level recognition of the broader goals and associated plans of the user (Carberry and Pope, 1993), (Goodman and Litman, 1992). A range of approaches have been proposed to the latter problem, including those based on defeasible reasoning (Konolige and Pollack, 1989), and relatedly, abduction (Appelt and Pollack, 1992); those concentrating on coping with ill-formed user input (Eller and Carberry, 1992); and those that recognise the need to constrain the inference process (Mayfield, 1992).

The dialogic situation, however, comprises more than just the beliefs of the speaker and hearer (in bilateral communication). In addition, there are the norms which characterise the exchange that can occur between them. Levin and Moore (1977) propose that the characteristic patterns of exchanges which occur between interlocutors can be represented as stereotypical encounters which they term, after Wittgenstein, *dialogue games*. They identify a number of commonly occurring examples of these dialogue games (based upon a large corpus of pre-experimental teletype conversations), including

helping, action-seeking, information-seeking, information-probing, instructing and griping. Associated with each dialogue are a number of roles (to be adopted by the parties playing a particular game), a topic (the propositional content of the game) and specific subgoals (which the parties undertake to fulfil according to their roles). In a style similar to the ‘adjacency pair’ theory of Schegloff and Sacks (1973), they identify canonical openings and closings to the various dialogue games – typically, for example, a game is initiated with a ‘proposal’ utterance followed by an ‘acceptance’ utterance.

Identification of these dialogue games, Levin and Moore suggest, facilitates an extension to the speech act account summarised above whereby understanding of indirect speech acts (i.e. those whose linguistic structure alone is insufficient to enable identification of illocutionary force, such as “Can you pass the salt?”) can be achieved by using knowledge of the current dialogue game to ‘disambiguate’. A given party attempts to understand an utterance of the other party by adopting the *meta-goal of comprehension*: “To comprehend an utterance, find some previously known goal of the speaker which this utterance can be seen as furthering” (Levin and Moore, 1977, p415) (the approach is thus similar in spirit to the Gricean notion of conversational postulates (Grice, 1975)). In their conclusion, Levin and Moore note that the dialogue games characterisation thus rests upon a bilateral analysis of action, in contrast to the unilateral actions embodied in speech acts; through such an analysis it then becomes possible to reduce the number and complexity of the speech acts involved.

The internal structure of the dialogue games themselves is quite unconstrained – the rules of exchange and of interaction between utterances are only very loosely specified. A more formal approach is provided in the logical accounts of the microstructure of dialogue offered in, for example, (Lorenz, 1982), which sets out a formal characterisation of the rights and duties incurred by interlocutors on a turn-by-turn basis. Concentrating on the structure of reasoned dialogue, Lorenz specifies restrictions on which interlocutor may put forward attacks or defences of a position at particular points in a dialogue. In particular, he claims that attacks are a right (i.e. either party may table an attack at any point) and the defences are a duty (and, specifically, that a defence must be provided in immediate response to an attack unless the defender has a counterattack available). Lorenz employs a game-theoretic foundation in his model, but other authors opt for a variety of alternative techniques: Heidrich (1982) uses a Montague grammar, Leopold-Wildburger (Leopold-Wildburger, 1982) decision theory, and Apostel (1982) action theory, for example. Apostel also makes the point that in addition to dialogue games, there is also a need for meta-games, “having as their object other games or leading to the modification of the rules of the game” (Apostel, 1982, p109) – these are of particular importance in jurisprudential dialogue (Feteris, 1997).

The characterisation of the permissible moves by the interlocutors is also highly heterogeneous across the field: Hintikka (Hintikka and Hintikka, 1982), (Hintikka *et al.*, 1996), for example, advocates a primarily interrogative approach based upon question-and-answer exchanges, whereas Carlson’s model (Carlson, 1983), offers a rich model encompassing (amongst others) assertions, questions, acceptances, presuppositions and intonation.

Girle (Girle, 1996) comments on the diversity of these various dialogue logics, but notes that they generally possess a set of features in common: (i) a set of locutions, (ii) a representation of the ‘commitments’ of the interlocutors, (iii) rules of how the contents of those commitment stores are

updated during a dialogue, and (iv) rules of interaction. His system, DL3, is an attempt to generalise over these logics, drawing in particular on his earlier work on DL and DL2, and on the related BQD system of Mackenzie(1979). Mackenzie's work is of particular interest because it has been shown to be directly amenable to implementation: Pilkington *et al.* (1992) demonstrate how an implementation of BQD can form a core component in a computer aided learning system.

Indeed, both linguistic approaches to dialogue structure (such as those of (Schegloff and Sacks, 1973) and (Levin and Moore, 1977)), and dialogue logics in general are often specified to a sufficient level of detail to facilitate computational interpretation. Freeman and Farley (Farley and Freeman, 1995), (Freeman and Farley, 1994) offer a recent example of the latter; of the former, accounts of linguistic structure (particularly those based upon Rhetorical Structure Theory, discussed in more detail in §2.1.4 and §3.1.4) extended to integrate a model of exchange structure, have been demonstrated to be of use in natural language generation(Daradoumis, 1996), (Fawcett and Davies, 1992).

One of the most promising routes for computational application of research into dialogue structure, however, is the formalisation of dialogue logics – particularly the rich system developed in (Walton and Krabbe, 1995) – for the communication demanded by increasingly complex multi agent systems.

2.1.2 Agent Communication

Though there is continuing debate over the meaning and scope of the term *agent* (Nwana, 1996) (the discussion returns to this point in §2.2.5), it seems clear that in any non-trivial multi-agent system there will need to be a means of communication between component agents, and indeed the capability for an agent to be able to communicate with its peers is often taken as a defining feature of agenthood (Wooldridge and Jennings, 1995) (and in some cases, the *sole* defining feature (Genesereth and Ketchpel, 1994)).

Though pre-dating much of the multi agent terminology, Smith's Contract Net (Smith, 1980) represents one of the first examples of a system to employ a common communication language between (mostly) autonomous 'nodes' (in Smith's terminology, the *common internode language*). This language comprises around a dozen primitives, each with a specific structure, purpose and set of conditions: though not recognised at the time, the primitives thus bear a striking resemblance to both illocutionary speech acts and the locutions of dialogue logics mentioned by Girle (Girle, 1996). The primary content of these locutions is a simplistic contract, a notion which has recently undergone a revival and can be found in several contemporary systems – as the *service level agreement* in the ADEPT project (Jennings *et al.*, 1996), and explicitly as contracts in (Sierra *et al.*, 1997), (Sierra *et al.*, 1997a) (Reed, 1998), *inter alia*. It is these contracts which form the subject of negotiation, a concept examined in more detail with regard to its relation to argumentation in §2.2.5.

Despite these and other projects such as (Laasri *et al.*, 1992) employing idiosyncratic communication protocols (usually driven by a need for a richer means of communication), a de facto standard is emerging amongst current research. The Knowledge Query and Manipulation Language (KQML) is a component of the research of the Knowledge Sharing Effort consortium and represents a performative based language which is independent of both underlying transport mechanisms and,

importantly, of the language in which agents (and the content of their messages) are written. Finin *et al.* (1997) offer a seven point desiderata for an agent communication language: (i) declarative and simple locutions, (ii) distinction between the communicative 'layer' (at which the illocutionary acts are expressed) and the content layer (at which domain facts are expressed – in whichever language is indicated at the communicative layer), (iii) well defined semantics, (iv) efficient implementation, (v) compatibility with networking technology, (vi) ability to cope with dynamic, heterogeneity, (vii) support for reliability and security. KQML performs well under each head, though of particular importance is the recent formalisation of its semantics (Labrou and Finin, 1997).

One of the very few major alternatives to KQML is the speech act based communication component of Shoham's Agent Oriented Programming (AOP) language, Agent-0 (Shoham, 1993). Agent-0 is closer to the dialogue logic and speech act account of communication than KQML, both because of the explicit operationalisation of illocutionary acts, and the role played by commitment: along with beliefs and obligations, commitments represent a primitive modality in AOP (in contrast to, for example, the belief-desire-intention account of (Rao and Georgeff, 1992)). The notion of commitment, as mentioned above, plays a key role in the definition of dialogue logics (and this point is carefully explicated in (Walton and Krabbe, 1995)). A key disadvantage with AOP however, is that unlike KQML, the communication language is tightly coupled to the design of the agents themselves. It is for this reason that the attempt to support communication between AOP and KQML – Agent-K – is coded as an agent in AOP which can interface to a KQML world (Davies and Edwards, 1994).

As multi agent systems become increasingly complex, so the demands placed upon the communication protocol between agents become ever greater. This complexity has led some authors to propose a more flexible 'open' approach to protocol design, whereby the agents themselves can negotiate the protocol to be employed (Vreeswijk, 1995).

Perhaps the highest possible level of complexity of agent communication would be represented by the use of natural language. As an agent-agent communication language, the approach would clearly introduce far more problems than it would solve (although a move in this direction seems to have been taken in the TRAINS project (Allen *et al.*, 1995)), but as a human-computer interface, designing the computer agent with natural language capabilities is highly desirable (see, for example, (Blandford, 1993)).

Determining an appropriate utterance at a given moment requires some internal processing within the agent (as Turner (1994) points out, given resource bounding of the kind discussed in (Bratman *et al.*, 1988), even determining what to include in a communication for it to be cooperative is a difficult task). When the utterance is to be expressed in natural language, that processing assumes an even more important role, since there will no longer be a one-to-one relationship between message content and message form. As suggested by (Searle, 1969), deciding upon the appropriate surface form of an utterance is a goal directed activity – planning, along the lines suggested in (Cohen and Perrault, 1988), is therefore a strong candidate for the reasoning subsystem. Before examining in detail the

variety of approaches to plan-based communication, a brief survey of the wider planning literature is in order⁸.

2.1.3 Planning

The first three decades or so of planning research have a well rehearsed genealogy, a typical example of which is offered in (Chapman, 1987). The root is usually taken to lie with GPS (Newell and Simon, 1963) and the introduction of means-ends analysis, by which new steps are introduced to a plan to fulfil particular goals: this is step addition⁹. STRIPS (Fikes and Nilsson, 1971) characterises actions as plan operators with preconditions and postconditions – the latter are lists of what explicitly gets added to or deleted from the world as a result of executing the action. Sussman's HACKER (Sussman, 1975) introduces in heuristic form many of the ideas which become formalised in later nonlinear planners, particularly the idea of promotion (constraining one operator to precede another). Promotion was first formalised in INTERPLAN (Tate, 1975), but the first truly nonlinear planner was NOAH (Sacerdoti, 1975) which in employing promotion and separation (as well as step addition) offers a straightforward solution to the Sussman anomaly (a scenario which was insoluble by linear planners such as HACKER without recourse to a 'hack'). NONLIN (Tate, 1977) extends NOAH by adding backtracking to enable better coverage of the search space, and employing simple establishment to force codesignation between a variable and an atom. Stefik's MOLGEN (Stefik, 1981) formalises constraints (of codesignation and ordering), and in a similarly 'neat' spirit, Chapman provides a precise characterisation of nonlinear planning which illustrates the restrictions necessary to ensure soundness and completeness (Chapman, 1987). The characterisation is summarised in the modal truth criterion which is implemented in the TWEAK algorithm (though more recently, it has been argued that the modal truth criterion does not in fact represent a necessary condition (Fox and Long, 1993)).

One of the problems with this tranche of planners is the inflexibility of the underlying representation language which remains virtually unchanged from STRIPS to TWEAK. There are a number of key problems with STRIPS-based planning languages which have been addressed in more recent research. The first problem is of distinguishing primary effects from side effects. As discussed in (Fink and Yang, 1997), the distinction can lead to significant computational savings by pruning the search space without compromising either soundness or completeness of the resulting planning algorithm. One of the first planners to represent primary effects explicitly was Wilkins' SIPE (Wilkins, 1988), where the distinction was employed in simplifying conflict resolution. The notion has also been developed in explicitly handling the links between one operator's postconditions and another's preconditions – 'causal' links – which may then be threatened by the existence of other operators in the partial plan. Resolution of these threats (i.e. the formulation of *safety conditions* (McAllester and Rosenblitt, 1991)) then forms a key component of the functionality of the planning algorithm. Partial order, causal link (POCL) planners such as UCPOP (Penberthy and Weld, 1992) and SNLP (McAllester and Rosenblitt, 1991) have also been demonstrated to be sound and complete. Acknowledging the demands of recent natural language generation research, the notion of primary

⁸ The role of planning mentioned here is in producing an utterance; related work has examined the reverse relationship, i.e. the role of dialogic communication in 'distributed' planning, e.g. (Shadbolt, 1992).

⁹ Throughout this discussion of planning, the terminology of Chapman (1987) is adopted for convenience and consistency.

effect has been associated with that of intention (a key component of the generation process, as discussed in the next section) in the development of the DPOCL planner (Young and Moore, 1994) (which although inspired by NLG work is nevertheless a domain independent planner: Young and Moore emphasise the point that discourse planning does not demand a domain specific planner (Young and Moore, 1994a)).

Several other fundamental restrictions on the expressiveness of the STRIPS language are addressed in UCPOP (Penberthy and Weld, 1992) through its adoption of Pednault's ADL language (Pednault, 1989) which represents a STRIPS-like characterisation of actions in the situation calculus (and was originally formulated for linear planning, as demonstrated in PEDESTAL (McDermott, 1991)). Within operator descriptions, ADL – and therefore UCPOP – permits the representation of (i) conditional effects and (ii) universal quantification, and as Penberthy and Weld point out, although neither of these are individually novel, UCPOP represents the first planner to admit both into a nonlinear framework, and be demonstrably sound and complete. The use of universal quantification introduces another problem: what is the domain of that quantification? More specifically, is it reasonable to make the closed world assumption (Reiter, 1978), and take it that absent information is false? Etzioni *et al.* (1997) review literature making use of the *open-world assumption* (that information not explicitly represented is unknown), which has the undesirable effect that universal quantification is simply not possible (since a planner cannot guarantee it is aware of every element in the domain). They go on to propose an alternative approach based upon *local closed-world information*, whereby information in a limited, local domain is complete (e.g. *ls -a* can be guaranteed to return a complete list of files).

In order to cope with the complexity of real world planning domains (exacerbated by the additional overheads of representing the various constraints associated with the codesignation and partial order of nonlinear planning), a technique is required for limiting the combinatorial explosion of the search tree. Sacerdoti proposed a formal characterisation of the intuitive concept of abstraction, whereby a system constructs a plan in an abstract, simplified description of the domain, before moving to a less abstract, more detailed domain description – the details in the latter are not considered until a plan is generated in the former, thus drastically reducing the options available to the planner, without impinging upon soundness or completeness. ABSTRIPS (Sacerdoti, 1974) represents the first abstraction based planner (implementing abstraction within a STRIPS planning framework); the unification of abstraction-based planning and nonlinear planning is formalised in the ABTWEAK system (Yang *et al.*, 1996).

Employing abstraction in the planning process, however, introduces a new set of problems. Chief amongst these is the potentially expensive need to backtrack between abstraction levels (Bacchus and Yang, 1992), which can be obviated through an abstraction hierarchy possessing the *downward refinement property* (DRP), which states that if a problem is solvable, then any correct abstract solution must have a correct refinement (Bacchus and Yang, 1991), (Bacchus and Yang, 1994). Related to the DRP, but orthogonal to it, is Knoblock's *ordered monotonicity property* (OMP) which demands that literals established at an abstract level may not be altered by any subsequent refinements (Knoblock *et al.*, 1991), (Knoblock, 1994). As discussed in (Bacchus and Yang, 1994), the DRP is a strong demand

to make; the OMP, in contrast, is weaker, and though holding subproblems invariant, does not prevent backtracking through the abstraction hierarchy (Smith and Peot, 1992). (Weaker again is ABTWEAK's *monotonic protection* (Yang *et al.*, 1996) which ensures that preconditions on an abstract operator are protected at lower levels in the abstraction hierarchy, i.e., they are guaranteed to have at least one establishment which is not necessarily clobbered).

Fox (1997) offers a clear classification of the various approaches to abstraction in hierarchical planning, based upon a tripartite division: Hierarchical Task Network (HTN) decomposition (in which abstraction is achieved by representing compound tasks), model-reduction (in which abstraction is performed by assigning criticality levels to literals), and operator decomposition (in which compound tasks are represented as abstract operators with associated pre- and postconditions). The HTN decomposition approach discussed in (Erol *et al.*, 1994) includes NOAH, NONLIN, O-PLAN (Currie and Tate, 1991) and Erol's own UMCP (Erol *et al.*, 1994a); model-reduction includes ABSTRIPS, ABTWEAK and ALPINE (Knoblock, 1994); and operator decomposition, DPOCL, UCPOP, SIPE and SNLP. Fox goes on to demonstrate that a concept of abstraction would benefit from aspects of both HTN and operator decomposition approaches, whereby a complete abstract plan undergoes refinement on completion and has a clear semantics¹⁰ (HTN decomposition) and manipulates abstract operators representing compound tasks (operator decomposition). Furthermore, the use of goals rather than operators in the bodies of abstract operators is shown to increase the flexibility of the planner, by eschewing the rigid, prescriptive, 'recipe'-like abstract operators whose bodies are composed of other operators (such as are used in NONLIN, O-PLAN and SIPE, *inter alia*). Fox and Long (1995), (1996) formalise these desirable features in the AbNLP planner, which also provides a much richer characterisation of time (associating, for example, an operator with an inseparable pair of time points, the earlier representing the moment of application – at which point the preconditions must be true – and the latter, the moment at which effects become true).

Increasingly, planning research is also focusing on practical issues of efficiency. A good example is the work of Gerevini and Schubert (1996) on improving the search mechanism used in executing nondeterminism: the approach is a pragmatic one rather than, for example, the less well understood and computationally expensive approach of metaplanning (Stefik, 1981a). Another is the translation mechanism which can map from the rich but computationally expensive domain representation of UCPOP to the concise but impoverished language of an efficient language such as Graphplan (Gazen and Knoblock, 1997). Finally, the computational benefits afforded through adoption of a hierarchical planning approach are catalogued in (Knoblock *et al.*, 1991), (Bacchus and Yang, 1992), and of automatically generating abstraction hierarchies in (Knoblock, 1994).

The next section explores in more detail the demands made upon planning systems by the natural language generation domain, reviewing the reasons for the current affinity with POCL planners, and motivating the adoption of AbNLP in the current work.

¹⁰ Fox and Long (1995) advocate an approach whereby the semantics of an abstract operator forming part of an abstract plan are defined in terms of a transformation between sets of states, rather than the more conventional definition based on the set of primitive linearizations.

2.1.4 Natural Language Generation

This section does not aim to review the field of natural language generation (NLG) as a whole. Such an objective would not only be more suited to a textbook, but would also demand inclusion of a range of issues which are not directly relevant to the current work: lexical and syntactic generation, multilingual generation, spoken dialogue, text summarisation, machine translation, natural language understanding, plan recognition, etc. Wider reviews of NLG encompassing these topics, and relating them to research on natural language processing in general can be found in (Cole *et al.*, 1995), (Dale and Reiter, 1999). Instead, the survey presented here concentrates specifically upon the generation of discourse structure, and the trends which have emerged on that topic over the past twenty years, omitting only review of research into focus and cue phrases which is provided in chapters four and five, respectively.

The ‘prehistory’ of discourse generation was characterised by ad hoc approaches which were difficult to generalise – a survey encompassing many such systems is presented in (Hovy, 1993). One of the first attempts at domain independence was McKeown’s TEXT system (McKeown, 1985). Though applied as a component of an interface to a naval database, the approach she advocated used generic *schemas*, stereotypical supra-sentence textual structures. Schemas are defined in terms of *rhetorical predicates* drawn from stylistic analyses of relations between clauses of text (such as those of Grimes (Grimes, 1975) and (Williams, 1893)) – examples of these rhetorical predicates used by McKeown include Evidence, Amplification and Particular-illustration. Schema-based generation offers a straightforward, fast way of generating paragraph sized chunks of text, and as a result remains a popular method of generation – see for example, (Rambow and Korelsky, 1992) for an applied example of recent schema-based work, Mellish *et al.*’s IDAS system employing schemas for content selection (Mellish *et al.*, 1996), (Reiter *et al.*, 1992), (Reiter *et al.*, 1995), Paris’s TAILOR system (Paris, 1993), Jones’s discourse component of the large scale language engineering project LOLITA (Jones, 1994), and (Jonsson, 1996) for an argument advocating the utility of the schema-based approach in general.

One of the key problems with the schema approach, however, is its inherent inflexibility. To better cope with novel, unexpected generation requirements, and to build larger, coherent pieces of text, a more dynamic approach to generation is required; such is the approach offered by planning. Early plan-based approaches to the generation task were predominantly concerned with generation at the utterance level (seminally, by Appelt (Appelt, 1985) and Cohen (Cohen and Perrault, 1988)). Soon after, a number of systems adopted the planning paradigm for discourse level generation, foremost amongst which was Hovy’s PAULINE, which also integrated stylistic concerns and a sensitivity to the hearer (Hovy, 1986a), (Hovy, 1990).

The idiosyncratic nature of this work, however, drove further research into a more stable, generalisable, extensible and justifiable underpinning. Such an underpinning was provided by a direct operationalisation of the rhetorical relations in Rhetorical Structure Theory (RST) (Mann and Thompson, 1986), (Mann and Thompson, 1988), an analytical tool for describing the structure of text. These rhetorical relations are similar to those used in McKeown’s schemas, but are more precisely defined, and are founded upon several important claims which endow the theory with strong predictive power. Hovy’s operationalisation (Hovy, 1988), (Hovy, 1991), (Hovy, 1993) of these relations as planning operators relies on several of Mann and Thompson’s claims. First, the distinction between

nucleus and satellite: Mann and Thompson claim, on the basis of corpus examination, that relations generally hold between two text spans, one of which (the nucleus) is more important than the other (the satellite). Hovy's planning algorithm, making use of 'growth points' for generating the two spans, requires this assumption of binary relations. Secondly, that each relation has an intended effect – this forms an operator's postcondition in Hovy's formulation (expressed in terms of Cohen and Levesque's BEL and BMB modal logic). Finally, Hovy's characterisation of the planning process also makes use of the canonical orderings found by Mann and Thompson in their corpus – particular relations are associated with a typical ordering between satellite and nucleus (e.g. in a Circumstance relation, nucleus typically precedes satellite).

Planning with RST operators thus offers a clear, principled approach to generation motivated by extensive corpus analysis – and as a result, has enjoyed significant popularity (Hovy, 1993). There are however, several critical problems (even after the extensions proposed in, for example, (Marcu, 1997) and (Scott and deSouza, 1990) are taken into account): two of these are examined here, and those remaining are discussed in §3.1.4. The first problem concerns the relations themselves – Hovy summarises: "Which relations? How many?", (Hovy, 1993, p359). In an attempt to condense the multifarious profligate views on rhetorical relations, Hovy and Maier (1993) list a taxonomy including some fifty relations and a rough measure of confidence in the utility of each (calculated using nothing more than the number of researchers in their review who had mentioned each relation). Mann and Thompson themselves are at pains to emphasise that the list they present is not exhaustive: that RST could be easily extended to handle further relations as deemed appropriate by analysts was considered one of its advantages. This is the prevailing view: although occasionally compendia are proposed which claim exhaustiveness (most notably, (Hobbs, 1982)) it is much more common to leave such lists open ended. The approach begs not only the question posed by Hovy, but also a meta-theoretic question: how can a given set be justified (other than on an atheoretic, ad hoc basis)? Knott attempts to answer this question (Knott, 1993), (Knott and Dale, 1996) by analysing the cue words which signal particular relations, and in particular, comparing applicability of related cues in a range of rhetorical situations. The discussion returns to Knott's analysis in more detail in chapter five.

Moore *et al.* propose an RST-based planning mechanism for generating explanations interactively (Moore and Paris, 1989), (Moore and Swartout, 1991), (Moore and Paris, 1994), and explain that Hovy's operationalisation suffers from a key failing – a failing which constitutes the second major problem with RST. Without any representation of *why* a particular operator is being employed at a particular point, there is no way to answer follow-up questions, or recover from failure (such as the hearer failing to understand, or misunderstanding (McCoy, 1986)) – tasks required for competent dialogic explanation¹¹. That is, RST offers no way of recording, in Grosz and Sidner's terms (Grosz and Sidner, 1986), the *intentional structure* of the discourse as distinct, from the *structure of the utterances*. Moore *et al.* circumvent the problems inherent to Hovy's account by explicitly distinguishing intentional from rhetorical structure; for the former, they employ a characterisation of goals drawing on a rich predicate language (including, e.g., PERSUADED, KNOW, KNOW-ABOUT), rather

¹¹ In fact, as pointed out in (Moore and Swartout, 1991) the shortcomings of Hovy's approach, with respect to planning regime and operator descriptions, effectively reduce it to a schema-based generation system.

than restricting definitions to Cohen and Levesque's BEL/BMB. They demonstrate (Moore and Paris, 1994, p670) that the rhetorical means available for fulfilling various intentions are multifarious – there is a “many-to-many” relationship between intentions and rhetorical relations. Moore and Paris acknowledge that the precursor to their distinction was Mann and Thompson's original distinction between *presentational* and *subject matter* relations: the former aim to induce a state in the hearer, the latter, to induce an understanding of some relation holding in the content of the utterance. Crucially, the latter cannot record the intention which lead to their application. Moore and Paris' planner thus employs goals of two kinds; *communicative goals* which aim to “affect the beliefs or goals of the hearer” (Moore and Paris, 1994, p668) and *linguistic goals* which characterise rhetorical relations – and are thus similar to the goals in Hovy's planner. Fulfilment of communicative goals leads to the posting of linguistic goals, so the intentional and linguistic structures of an utterance are built in tandem.

More recently, work has started on a principled unification of the tripartite distinction proposed in (Grosz and Sidner, 1986) and Rhetorical Structure Theory, (Moser and Moore, 1996), whereby dominance is equated with nuclearity. One of the outstanding problems is how to deal with attentional states in this framework; this is one of the key issues addressed in this thesis and is explored in more depth below.

Unfortunately, there are deeper problems with an operationalisation of RST which arise not from the shortcomings of RST, but from those of the operationalisation itself – and in particular, from the style of planning employed. Almost all discourse planning (e.g. that reviewed in (Hovy, 1993)) draws upon the NOAH planning paradigm (Sacerdoti, 1975), yet there are well understood problems and a quarter of a century of developments – charted in the previous section – since Sacerdoti's work. Young *et al.* (1994) first explain that, as a result of the underlying planning algorithms, in systems such as (Hovy, 1991), (Moore and Paris, 1994), (Cawsey, 1993) and (Maybury, 1992), “there is nothing ... to prevent them from generating incorrect plans, generating plans with redundant steps, or failing to find plans in situations where they exist.” They go on to propose an alternative approach based on a planner which explicitly handles both decomposition (i.e. hierarchical) links and causal (i.e. postcondition-precondition) links (in addition to the standard freedoms of partial ordering and partial instantiation, both of which can be subject to specified constraints). Their DPOCL¹² planner (further detailed in (Young and Moore, 1994)) forms the core of the LONGBOW discourse planning system (Young, 1996) currently under development at Pittsburgh. The adoption of a POCL based planning mechanism not only facilitates explicit manipulation of the system's communicative intentions, but is also demonstrably sound and ‘primitive complete’ (i.e. every permissible order of primitives can be reached, but not necessarily all hierarchical structures) (Young *et al.*, 1994a).

DPOCL too, however, suffers from a range of shortcomings. Fox (1997) offers a careful analysis of the problems, from which four key points emerge: first, abstract plans in DPOCL do not have a clear interpretation; second, the DPOCL planning process may involve frequent backtracking between levels of refinement (which is very costly); third, DPOCL does not competently handle interactions between operators at different levels of abstraction (though this has not hampered DPOCL since discourse planning typically involves operators without delete lists (Reed *et al.*, 1997)); fourth,

¹² Decompositional, Partial Order, Causal Link

DPOCL is unable to handle interleaving of the subtasks of multiple operators. The AbNLP planner (Fox and Long, 1995) addresses each of these issues in a coherent framework which follows in the rigorous style of TWEAK (Chapman, 1987). AbNLP is discussed in §2.1.3 above, and its adaptation for the current work in chapters three and four.

Of course, the NOAH-RST approach to discourse planning, though dominant, has a wide range of variants and dissenters. One of the foremost alternatives to RST is Kamp's Discourse Representation Theory (DRT) (Kamp, 1981), used as the basis for implementation in the work of (Asher and Lascarides, 1994) and in the VERBMOBIL project (Wahlster, 1993). Other, less general work employs a number of approaches similar in spirit to RST operationalisation: Moore's explainable expert system, EES (Moore and Swartout, 1991), (Moore and Paris, 1989, (Paris, 1991), Dale's EPICURE (Dale, 1990), Horacek's OFFICE-PLAN (Horacek, 1992), (Horacek, 1994), Meteer's SAGE (Meteer, 1994) and the work of (Cawsey, 1993). A closely related trend employs Goldman's (Goldman, 1970) characterisation of the semantic relations GENERATION and ENABLEMENT, represented computationally in the work of (Sider and Burger, 1992) and (Scott and Paris, 1995), *inter alia*. There have also been reactions to the standard planning architectures: Haller *et al.* (Haller, 1994), (Haller and Shapiro, 1996), for example, suggest a more reactive approach, echoed by criticisms of the "top-down" approach to generation in general (Marcu, 1997a), (Mellish *et al.*, 1998). Others – such as (Meteer, 1991), (Inui *et al.*, 1992), (de Rosis *et al.*, 1997) and (Robin, 1994) – have proposed a revision based approach to generation whereby a complete draft plan is created and then subsequently improved (this is similar to the select-and-repair method of the HealthDoc system (DiMarco *et al.*, 1995), (Hirst *et al.*, 1997)). Finally, work founded on systemic functional grammars should also be included, though as Hovy points out in (Cole *et al.*, 1995), much of it is focused upon single sentence generation. Nevertheless, systems such as PENMAN (Mann, 1983), POPEL (Reithinger, 1991) and FUF (Elhadad, 1992) play a pivotal role in much NLG research.

This first thread thus has several key milestones in its development: the use of schemas, the use of non-linear planning, the adoption of RST as an operational model, and more recently, the criticisms of both RST and the planning paradigm. The next section proceeds to develop the second thread in a complimentary way – though its roots are somewhat more ancient.

2.2 Argumentation

The tradition of argumentation has its foundations in Ancient Greece, with the first rhetoric texts accompanying formal training by orators such as Gorgias in the fourth century BC. The techniques of rhetoric were (and still are) regarded with mistrust as being tools for confusing audiences and disguising truth (consider, for example, idioms such as *empty rhetoric*). However, such criticism is only justified when rhetoric is teased apart from argumentation and seen as nothing more than "at worst, the empty jingle-making of its critics, or, at best, a guide to tasteful figures of language" (Billig, 1996, p63). It is the richer sense of rhetoric, including all the structural and compositional facets of argumentation, which formed the basis of Aristotle's *Art of Rhetoric* (Aristotle, 1926), and of most studies of the field thereafter.

Rhetoric formed a core component of education from before the Middle Ages through to the eighteenth and nineteenth centuries, supported by prescriptive texts such as (Blair, 1838) and (Whately, 1855). These works are based heavily upon classical sources, with a corresponding emphasis on regurgitation and some few contemporary examples, rather than novel analytic work. The twentieth century, however, has seen a resurgence in the theoretical approach towards the analysis of argument. Despite the fact that logic itself is grounded in argument analysis, the formal, mathematical entities that constitute modern, post-Frege logical systems are generally *too* formal to accurately characterise naturally occurring arguments. Two research trends addressing this issue are identifiable. First, there have been a number of attempts to extend various aspects of conventional logics to incorporate more of the phenomena found in natural argumentation: these generally come under the heads of dialogue and interrogative logics discussed above. However, even with such extensions to formal systems, there are noticeable shortcomings in the general applicability of these theories. In contrast, ‘informal logic’ has concentrated on building models which can cope with the complexity and freedom of real dialogic situations, at the expense of rigid formality (and the computational, semantic and ontological benefits thereof). The models proposed under informal logic are motivated by a variety of specific problems: identification of fallacies (Hamblin, 1970), (Woods and Walton, 1989); characterisation of ‘burden of proof’ (Farley and Freeman, 1995), (Rescher, 1977), and the dynamics of legal reasoning (Alexy, 1989); (Feteris, 1996); the role of relationship between rhetoric and the social sciences and humanities (Perelman and Olbrechts-Tyteca, 1969); (Billig, 1996); etc. Each has its own method of analysis upon which the models are founded. Clearly, for any computational system which hopes to build complex argumentation comparable to that found in the real world, it is crucial to examine both the techniques of argument analysis (for they may be of direct use during synthesis) and the models of argument structure (for they may represent some form of goal state), offered by the research conducted in argumentation theory.

2.2.1 Argumentation Theory: Analysis Techniques and Structural Models

One of the key problems facing argument analysis is that there is no general way of determining whether or not a particular analysis is ‘right’ or ‘complete’. The analysis process itself is for the most part an intuitionistic, heuristic process guided by experience. Furthermore, there exists little consensus on the various features of analysis. The unit size of argument ‘chunks’, for example, ranges from pseudo-logical propositions to hugely complex paragraphs. In (Johnson, 1992, p328), the following (from a court extract) is analysed as a single premise: *Because Nancy Beth Cruzan did not have the foresight to preserve her constitutional right in a living will, or some comparable “clear and convincing” alternative, her right is gone forever and her fate is in the hands of the state Legislature instead of in those of her family, her independent neutral guardian ad litem, and an impartial judge – all of whom agree on the course of action that is in her best interests.* Compare this with the well-worn premise (*ibid.*, p316), *All humans are mortal.* No distinction is drawn between the different roles these premises play in their respective arguments; they are both simply analysed as discrete functional units. Similar extremes exist in another facet of argument analysis: the phase of *reconstruction*. It is claimed that the argument form communicated does not exactly correspond to the underlying argument structure: to analyse the structure of an argument therefore requires the reconstruction of underlying

form from surface form¹³. This seems to be supported by the frequent occurrence of enthymemes where one of the premises or the conclusion is left implicit (e.g. *All men are mortal, so Socrates is mortal*) – indeed, they occur so frequently that they have been seen as separate rules of inference (namely, the *Modus Brevis* form (Sadock, 1977)). However the degree to which such reconstruction is carried out varies wildly from the minimal reconstruction of the Woods-Walton approach (Woods and Walton, 1989), to the extensive reconstructive process of pragma-dialectics (Eemeren *et al.*, 1993).

Despite a lack of consistency in the process of analysis, the subsequent representation and manipulation of argument form is a principled, structured activity. The most common form into which arguments are analysed is an intuitive framework little changed since Aristotle. This *standard treatment* is explained in many text books of informal logic (e.g. (Johnson, 1992)), but one of the clearest expositions is to be found in (Freeman, 1991). The four main types of argument component structure are *divergent* (whereby one premise can support several conclusions¹⁴), *serial* (whereby a single premise contributes to a single conclusion, which may then act as the single premise to another conclusion, etc.), *convergent* (whereby two or more premises contribute independently to a single conclusion) and *linked* (whereby two or more premises together contribute to a single conclusion). These forms are summarised in Figure 2.1, below.




Figure 2.1 The four basic argument structures, after (Freeman, 1991, p2).

Complete arguments are composed of various combinations of these forms, as there is no distinction inherent in a textual clause between the role of premise and that of conclusion (i.e. a premise to a super-argument can be the conclusion of a subargument).

The main rival to this representation is the Toulmin schema, proposed in (Toulmin, 1958), in which the usual distinction between premise and conclusion is eschewed in favour of a more refined, six-fold classification. Under Toulmin schema, a claim, *C* is supported by a datum, *D*, qualified to a degree *Q*. The support relationship is licensed by a warrant, *W*, which is in turn founded upon a backing, *B*. Finally, an argument can include a caveat anticipating a rebuttal, *R*. This schema is shown in Figure 2.2, below, along with Toulmin's example which clarifies the various functional roles.

Toulmin's aim was to indicate shortcomings not just of existing techniques of argument

¹³ Note that this notion has important implications for a generation system: if surface form is the result of processing on an underlying form, then we may need a two-stage process in generation. The discussion returns to this point below.

¹⁴ He later points out (p93) that under some circumstances, it is more appropriate to view divergent structure as serial, with the premise repeated for each conclusion. This is the approach taken in (Reed *et al.*, 1996).

analysis, but of formal logic as a whole (a summary of Toulmin's agenda is offered in (Eemeren *et al.*, 1996, pp129-138&190)). As a result he was heavily criticised, with criticisms from the philosophical community attacking both the lack of rigour in his approach, and specific downfalls, such as the problem of distinguishing data from warrant, and the absence of a satisfactory analysis for arguments based upon probabilities (Cooley, 1959). More recently, (Freeman, 1991) has re-evaluated Toulmin's work and concluded that as it stands, it contains serious errors which become manifest in arguments based on inductive generalisation or probability. He proposes a solution based upon an amalgamation of the standard approach and Toulmin schema, in which the standard premise/conclusion distinction is enriched to include Toulmin-esque modality and rebuttals.




Figure 2.2 Toulmin schema and example

The predominant aim in argumentation theory is not, however, simply to diagram arguments (indeed, extended accounts of such diagrammation are rare – (Finocchiaro, 1980) offers one of the few examples). Rather, it is the *evaluation* of arguments – determining soundness, validity, or whether an argument is good or bad – which motivates most argumentation research. There are five broad categories into which this research falls (this taxonomy has been adapted from (Blair, 1996)). Firstly, it has been claimed that an argument is good if it is *logically* good, i.e. that its premises are true and the inferences are deductively valid. This represents the original aim of logic as construed by Aristotle – logic having been developed to determine whether or not an argument is valid. However, as mentioned above, this is generally too strong for natural argument which often (i) makes use of non-deductive inference and (ii) employs premises which are plausible – or acceptable to all parties – rather than true. Secondly, an argument can be evaluated through the use of *schemes*, specific argument forms to which empirical data can be compared (Kienpointner, 1992). Each scheme has associated with it a set of critical questions, by which validity can be determined. The Toulmin schema is seen in (Kienpointner, 1992) as a metascheme. There are a number of problems with this approach – quite apart from the inherent lack of a principled basis precluding the construction of a definitively exhaustive list, there are practical difficulties in accurately matching data with schemes. Thirdly, arguments might be assessed by reference to criteria of relevance, sufficiency and acceptability – the so-called RSA-triangle (Johnson and Blair, 1993). Premises must be relevant to the conclusion, must represent sufficient support for the conclusion, and must be acceptable to arguer and audience. Defining these three terms with sufficient rigour as to enable unequivocal analysis, however, has proved highly problematic (see

(Eemeren *et al.*, 1996, p178-181) for a review). Fourthly, argument validity is equated with absence of fallacy. Like the scheme-based approach, the ad hoc way in which lists of fallacies are constructed makes it impossible to claim either exhaustiveness or accuracy of those lists. There is a core component upon which most authors agree, comprising about eighteen fallacies – see, for example, (Johnson, 1992) or (Fogelin and Sinnott-Armstrong, 1991), but beyond this there are significant differences (compare, for example, those of (Wilson, 1980))¹⁵. Furthermore, the schemes of classification of fallacies also differ wildly, though recent research has looked at potential means of uniting various taxonomies (Hansen, 1996). The dissensus has prompted investigation into the development of a metatheory of fallacies, which might lead to a formalisation and justification of both an exhaustive list of fallacies, and a comprehensive system of classification. The pioneering work of Hamblin (1970) in this pursuit has recently been extended by Walton and Krabbe (1995), who sketch out a taxonomy of dialogue types and then propose that fallacies are generally associated with illicit shifts from one type of dialogue to another. The importance of the preliminary results offered by Walton and Krabbe, both for the current work, and for communication research in the multi-agent domain, is discussed at greater length below. Finally, the fifth general means of evaluating argument eschews the absolutist view of objective argument validity in favour of a relativist, context-dependent metric, whereby an argument is valid precisely if it is persuasive to a particular audience. This view has been developed by Perelman and Ohlbrechts-Tyteca (1969), where the notions of *persuasive* and *convincing* are distinguished: the former expresses an argument's successful effect on its hearers; the latter expresses the successful effect on the *universal audience*, an artificial construct representing the speaker's notion of an idealised rational judge. (This distinction is rooted in the divergence of rhetoric and logic identified by earlier writers such as Whately, (1855)). Clearly, this last evaluation technique is closely linked to issues of effective communication in real social situations. As natural argument is only ever found in real social situations, it is important to give consideration to the various factors affecting argumentation which are the direct result of the social psychological context.

2.2.2 Argument As A Social Phenomenon

The importance of viewing examples of argumentation in their natural environment is discussed in (Gilbert, 1995):

“Argumentation, first and last, is a subspecies of communication, and communication is a complex act that integrates cultural and sub-cultural symbolism, social actors and local context. This means that any given argument or part thereof may be acceptable or appropriate or useful or sensible when used by one set of persons in one place and time, and not acceptable, etc., when any or all of those variables are altered.”, p127

Clearly, this makes accurate analysis (such as that based upon identification of fallacies) impossible, since the analysis will inevitably be based upon some impoverished record of the original argument. (Gilbert goes on to argue that the paragraph-length extracts offered for analysis in critical reasoning undergraduate courses are therefore of limited value). But equally, this situated nature of argumentation is of crucial importance during the synthesis – either by human or computer – of extended argument.

¹⁵ A summary of the major fallacies is provided in Appendix A

One of the key factors which is generally lost in the process of abstracting from real, social argument to the idealised version subsequently analysed is the relationship between the interlocutors. For two or more parties to engage in discourse, some form of relationship must exist between them, or its creation be one of the aims of the discourse. And yet argumentation is inherently conflict-based, and that conflict is liable to threaten the existing or incipient relationship. This problem is discussed in terms of a division in the types of goals which impinge upon the argumentation process: broadly, there are *task goals* which specify a participant's direct aims of the discourse (e.g. to convince an opponent that a particular proposition is true), and *face goals* which specify the limits of appropriate behaviour, including maintaining 'face', and respecting 'face' of the interlocutor, (Gilbert, 1996), (Tracy, 1990), (Waldron *et al.*, 1990). Indeed, the role of *facework* is of crucial importance in argumentation, where conflict is almost unavoidable. As discussed in (O'Keefe, 1995), there are several means of managing face threats in conflict situations, including toning the threat down and offering redress: these various methods relate to levels of 'politeness' in discourse, an issue which has received increased attention after the seminal work of Brown and Levinson (1987). As Penman points out, (Penman, 1990, p17), the problems of maintaining politeness can usefully be seen as subsumed by those of facework, as face is both a prerequisite of politeness, and facework can be performed using other strategies. One such strategy is the propensity for equivocation in situations where truthfulness would be hurtful or embarrassing, but lying too risky (Bavelas *et al.*, 1990)¹⁶. Another technique, more closely related to issues of politeness, is *phatic* communication, which Laver (1976) claims is used during initial stages to avoid hostility, establish a working consensus and open interaction. This seems closely related to the concept of establishing common ground in argument, a phase classically seen as distinct from the 'argumentative' phase (Blair, 1838) (this distinction is elaborated upon below).

Real argument, then, involves fulfilling and mediating between multiple goals expressing various task and face aims (and given this fact, it should be briefly noted that the notion of *goal* is far from clear and unequivocal, (Craig, 1990)). This goal-based view of the argument process is characteristic of the speaker-oriented approach (c.f. (Eemeren *et al.*, 1996, p210)), and equates closely with O'Keefe's (1977) definition of argument₂: a process which "two or more persons have (or engage in)". This contrasts with the discourse-oriented, argument-as-product view, defined by O'Keefe as argument₁: "something one person makes (or gives or presents or utters)"¹⁷. It is this latter view of argument which predominates both in studies of argument structure and, perhaps somewhat surprisingly, in studies of arguments in cognitive psychology (e.g. (Kuhn, 1991)) – even the Toulmin approach to arguments₁ has been advocated as a plausible model of cognitive organisation of arguments (e.g. the mental-models approach of Green (1996)). The problem with such approaches, of course, is that they fail to capture the social, situated issues of facework (Willard, 1976). The same problem faces the psychological investigations of argument *generation* (rather than representation) – Rips (1994), for example, discusses the central role played by natural deduction in communication, but that role is only applicable to the generation of arguments₁. The generation of real arguments (i.e.

¹⁶ There are serious issues involved in resorting to lying, and numerous problems involved analysis of non-truthful discourse, e.g. (Smeltzer, 1996)

¹⁷ Note that, as pointed out in (Reed and Long, 1997), this implies that someone engaged in argument₂ is perforce presenting arguments₁.

participation in arguments₂) – whether as a psychological model of human behaviour or as an implementable computational system – relies upon adducing to a system of argument₁ structure an appropriate characterisation of social considerations. For a competent model, the interaction between the two facets will be complex and rich, to account for the interplay of goals, and the multifunctionality of individual utterances (Waldron *et al.*, 1990). The construction of such a model could be of benefit to the multifarious fields in which aspects of argumentation are applied.

2.2.3 *Applied Argumentation In Formal Reasoning*

It might seem unlikely that argumentation would be a useful technique in formal systems – for two complimentary reasons. Firstly, it has been claimed that where an argument's conclusion is logically and undeniably entailed by its premises, the result doesn't qualify as argumentation at all, (Brockriede, 1975). Such a scenario offers no room for dialogue, discussion, concession, retraction or any of the other characteristic aspects of dialectic. Similarly, one of the key features of argument is its use of convergent multiple subarguments, each contributing further to the claim (see Figure 2.1). In classical logic, however, the concept of multiple subarguments is redundant. If it is possible to prove P , say through X and $X \rightarrow P$, it is unnecessary to then prove P once again with Y and $Y \rightarrow P$: using both proofs rather than just one would in no way produce a 'better' argument (Reed and Long, 1997a).

There is increasing interest in using argumentation for systems based upon formal logic which need to reason about the real world, and in particular, which must be able to cope with uncertain and incomplete information. Reasoning about such domains can rarely employ strict deductive inference; rather, it becomes necessary to use some weaker notion of support – and often then to express the degree of that support (either qualitatively – e.g. (Parsons, 1996), or quantitatively – e.g. (Sillince and Minors, 1992)). If a system no longer relies solely upon strict inference then it may as a result benefit from the use of multiple subarguments. These separate lines of support may then be aggregated under some flattening function, such as those discussed in (Das *et al.*, 1996). Furthermore the set of arguments contributing to a claim can itself be evaluated as a first class data object to determine the acceptability of the argument as a whole. This is the approach adopted in the argumentation logic LA (Krause *et al.*, 1995), which uses a labelled deductive system (Gabbay, 1992) to record sets of supports and determine acceptability. This approach is motivated by the need to reason under uncertainty, (Elvang-Goransson *et al.*, 1993), and has successfully been applied in a number of medical domains (Fox and Das, 1996).

LA represents a highly specific logic, based, in the first instance, upon intuitions of evaluating a claim on the basis of 'pro' and 'con' arguments. A more generic approach to uncertainty and incompleteness which has been at least as successful is defeasible reasoning, e.g. (Pollock, 1987), in which there are two closely related trends. Firstly, argumentation is used as a technique for implementing systems of defeasible reasoning, both those involving priorities between defeasible rules such as (Prakken and Sartor, 1996), (Antoniou, 1996), (Vreeswijk, 1992), and those based upon a probabilistic underpinning, (Geffner, 1996). The second trend is closely allied with the first in terms of its results, but its motivations differ in that the aim is to *model* argument *using* defeasible reasoning, (Loui, 1994) for example, and Prakken's (Prakken, 1996) 'dialectical proof theory'. This last work is characteristic of the field in that it does not attempt to model generic, free argument, but instead

concentrates on legal reasoning. Argumentation in jurisprudence benefits not only from the practical advantage of a plentiful existence of transcribed source, but also from the theoretical advantage of possessing clearer rules of exchange and dialectical progression – legal argument is more ‘rigorous’, demanding greater adherence to logical consistency, and admitting little retraction (Kowalski and Toni, 1994), (Prakken and Sartor, 1996); (Verheij, 1996), etc.

Legal argument is also one of the foremost domains for intelligent tutoring systems (ITS) teaching argumentation techniques – for example, the CATO system (Aleven and Ashley, 1994), based on their HYPO model of case-based legal argumentation (Aleven and Ashley, 1992), aimed to develop student skills of analogical reasoning, through identification of relevant features with past cases. The HYPO model – and ITS in general – represents one of a number of areas in which computational models of argumentation are required.

2.2.4 Computational Approaches to Argumentation

Broadly, there are two fields of computational research investigating argumentation as an end in itself (rather than as a technique for dealing with uncertain domains, for example). Firstly, there are a range of projects concentrating on representing and occasionally analysing natural argument. In contrast is the smaller research trend focusing on the automatic generation of arguments from some knowledge base.

Under the first research trend, there are several distinct identifiable aims. A number of systems support humans in constructing or following argumentation in order to reach decisions: a variety of medical decision support systems based on Krause's (Krause *et al.*, 1995) LA are discussed in (Fox and Das, 1996), and similarly, the Negoplan system of (Matwin *et al.*, 1989) supports human negotiation using an expert system to represent the structure of the arguments. Distinct from those which offer support are a group of systems which offer the medium for argument. In particular, there have been attempts to integrate argumentation frameworks with the world-wide web (WWW), which though an attractive arena for debate, suffers, with its current modes of interaction, from a number of inherent problems. In particular, (Jackson, 1997) points out that newsgroup postings – a prime example of rich, wide-ranging and unregulated debate – are exceedingly difficult to employ successfully for constructive argument: it is difficult to see the thread of an argument, and see which points have been addressed, which are contentious, and which need resolving. She suggests a solution whereby structure is imposed upon the debate; postings become hypertext documents which are arranged hierarchically according to their functional role in the argument (rather than ordered sequentially and chronologically). A similar approach has been suggested by Gordon *et al.* (Gordon, 1994), (Gordon and Karacapilidis, 1996) in their *Zeno Framework*, which intends to offer mediation on the WWW and uses for its underlying argument structuring the IBIS system of (Rittel and Webber, 1973) as extended by (Conklin and Begeman, 1988). Distinct again from support- and medium-oriented research, computational representation of argument is also employed to pedagogical ends (in teaching skills of both argument production, e.g. (Cavalli-Sforza *et al.*, 1992), (Pilkington *et al.*, 1992), and argument criticism, (Cavalli-Sforza *et al.*, 1993)), and finally, as a means of abstracting from particular data sources and thus facilitating representation of arguments drawing on disparate and possibly conflicting sources (implemented in Haggith's FORA system (Haggith, 1995), (Haggith, 1996)).

The second research trend, aiming to create – rather than represent – argument can also be subdivided into work focusing primarily on the structure of argument, and that focusing primarily on the language of argument. Clearly these two tasks are not entirely separable, but nevertheless, systems such as Zukerman's NAG (McConachy and Zukerman, 1996), (Zukerman and McConachy, 1995), (Zukerman *et al.*, 1996) are chiefly concerned with the generation of the structure of an argument: using 'reasoning agents' to select information from a range of knowledge bases (in this it bears resemblance to Haggith's FORA, though the latter offers a more principled approach to resolving inter-KB conflict). The premises supporting a conclusion form nodes in an argument graph, producing a structure similar to that arrived at by analysis in informal logic (although, as pointed out in (Reed and Long, 1997a), NAG does not distinguish between linked and convergent structures – see Figure 2.1 – and its expressive power is thus restricted). Further processing then determines an optimum 'path' through the argument graph, based on parameters concerning the user's abilities and the system's honesty (e.g. whether or not it is appropriate to exploit misconceptions held by the hearer). Finally, NAG determines an appropriate presentation strategy, though the system described in (Zukerman *et al.*, 1996) is limited to a very narrow range of options, with a naive approach to the problems of component ordering and linguistic style, issues examined in (Reed *et al.*, 1996), (Reed and Long, 1997b) and in more detail below. The earlier work of Birnbaum *et al.* (Birnbaum, 1982), (Flowers *et al.*, 1982) also concentrates on determining the content of argument, and in particular, on identifying and implementing schema-like argument exchanges termed argument *molecules*.

Standing in contrast is the work in natural language generation (NLG), where the goal is to produce the text of an argument, rather than to restrict effort to generating the underlying structure. Elhadad (1992a), (1995) concentrates on generating arguments comprising a single paragraph, in order to investigate the impact of argumentative 'orientation' on lexical choice. In particular, he builds on the distinction between evaluation functions and topoi proposed by (Anscombe and Ducrot, 1983) – the former determine the force of particular propositions in a particular context; the latter then link these evaluations through a generic form "the more/less X is P, the more/less Y is Q". These topoi relations as construed in Elhadad's work are similar to the links in Sycara's belief graphs (Sycara, 1989), (Sycara, 1990) (which form the basis of her PERSUADER argumentation system discussed below) and also to the arcs in qualitative probabilistic networks (QPNs) (Wellman, 1990), used, for example, in Parsons's argumentation reasoning system, (Parsons, 1996), (Parsons, 1997). These relations, however, seem unable to express the full range of argumentation moves (it is difficult to see, for example, how a categorical syllogism could be expressed accurately in these terms). A different approach to the linguistic realisation of arguments has been adopted by Maybury (Maybury, 1993), which builds on plan-based models of communication (rather than Elhadad's unification based approach). Maybury proposes abstract plan operators which encode argument strategies (similar to the molecules proposed by (Birnbaum, 1982)), such as *convince-by-cause-and-evidence*. Although the general approach appears promising, there are a number of specific problems with Maybury's system from the viewpoints of both argumentation and NLG. Firstly, Maybury offers an abstract taxonomy of communicative acts, which at its highest level divides *argue* into *deduce*, *induce*, and *persuade*. In light of informal logic research this seems highly implausible – deduction and induction (along with various

other forms of non-standard reasoning, some of which Maybury notes) are utilised during the process of persuasion (and also during other, similar processes such as negotiation, information-seeking, and so on). Secondly, Maybury implies that his system employs NOAH, which suffers inherently from inflexibility and inability to cope with uncertain domains (as discussed above in §2.1.4). Furthermore, the components of a plan are not subject to any reordering – the partial order imposed by the planning process represents the final result, thus precluding arrangement of premises and conclusions to optimise persuasive effect (despite the fact that Maybury mentions the rhetoric and stylistic literature which deals with this problem). In addition, Maybury's system has no notion of focus or context, and it is not at all clear whether intention is expressly represented, and if so how it is achieved (cf. (Moore and Pollack, 1992)). Finally, it also fails to be sensitive to the audience (compare Perelman: "For since argumentation aims at securing the adherence of those to whom it is addressed, it is, in its entirety, relative to the audience to be influenced", (Perelman and Olbrechts-Tyteca, 1969, p19)). Despite these shortcomings, (Maybury, 1993) represents the first attempt at plan-based generation of natural language arguments.

Maybury's work was not solely concentrated upon generating argument: another area in which his *communicative acts* have been applied is in explanation generation (Maybury, 1992). This overlap is important and in a sense, unsurprising, for argumentation clearly shares many features with explanation. In Paris's (1991) Explainable Expert System (EES) framework, a number of the aims (e.g. dealing with user questions such as "why is this action being recommended?"), plan operators (e.g. the 'persuade' operator using the 'evidence' relation), and resultant intentional structures, are very close to those discussed in (Maybury, 1993) and the current work. Again, however, EES is based upon NOAH and as a consequence, suffers from the same problems as (Maybury, 1993), mentioned above. Unlike Maybury, however, Paris discusses the key importance of user modelling, with EES tailoring the kind of information, the level of detail, and lexical choice to the current user: such careful dependence upon a user model is clearly in line with the audience-based analysis of Perelman (Perelman and Olbrechts-Tyteca, 1969), and the prescriptive manuals of, for example, (Blair, 1838). Argumentation, in the fields of both informal logic and computational modelling, has been claimed to be fundamentally dialectical – the same claim has also been made for explanations (Moore and Swartout, 1991), (CavalliSforza *et al.*, 1993). Further similarities between argument and explanation can be found in Moore's (1994) discussion of the importance of *discourse markers* – a similar analysis has been performed for argument by Cohen, (1987) (aiming to investigate the general dictum found in informal logic texts such as (Wilson, 1980) that these discourse markers are useful in identifying structural components of an argument). Interestingly, these markers (or cues or clue words) have also been investigated in order to *distinguish* argument from explanation, though so far with only limited success (Snoeck Henkemans, 1995), (Snoeck Henkemans, 1997). Finally, it has been argued that Rhetorical Structure Theory (RST) (Mann and Thompson, 1988) is fundamentally unable to account for the high level structure present in explanation (Mooney *et al.*, 1990): a similar case is put forward below that RST is unable to model or determine coherence of argument.

Explanation represents a single close relative of persuasive argumentation: as Walton and Krabbe (Walton and Krabbe, 1995) point out, however, there are a number of others (negotiation,

deliberation, etc.), and though these are attracting limited interest in NLG, they have been core problems in distributed AI for some time.

2.2.5 Argumentation in Multi-Agent Systems

One of the most common problems in recent work on multi-agent systems has been a lack of consensus on an operational definition of *agent*. Most researchers seem to acknowledge the origin as Hewitt (1977), and then proceed to offer their own definition. The diversity has been recognised and has given rise to a plethora of review (e.g. (Wooldridge and Jennings, 1995)), classification (e.g. (Nwana, 1996)) and taxonomic papers (e.g. (Franklin and Graesser, 1996)). Despite the wide range of opinions and definitions, there are a number of key characteristics which feature in most definitions of agency. First is the notion of *autonomy*, explored in detail by a number of authors ((Castelfranchi, 1995) and (d’Inverno and Luck, 1996), for example), but summed up adequately by Wooldridge and Jennings:

“agents operate without the direct intervention of humans or others, and have some kind of control over their actions and internal state”, (Wooldridge and Jennings, 1995, p116).

Next, the fact that agents are *proactive* – they are not simply passive, stimulus-driven units; rather they can plan and take action to achieve their goals. Finally, it is often assumed that agents may be *socially situated* in an environment where there are other agents (i.e. a scenario with scope for inter-agent communication) – and that that multi-agent environment may well be highly heterogeneous and unpredictable.

Together, these three aspects give rise to the need for one crucial skill in an agent’s repertoire. One agent, A, needs to be able to communicate with another agent, B, and persuade B to assist in the fulfilment of A’s goals, to their mutual advantage: this is negotiation¹⁸. The field of research into negotiation is extremely wide (see, for example, (Kemper and Kemper, 1994)): the current discussion will focus quite precisely on recent work into models of inter-agent negotiation, and some problems thereof.

One common model for negotiation is grounded in game-theoretic approaches such as (Nash, 1950), (Nash, 1953). The foremost example of this trend is the work of Zlotkin and Rosenschein (1994), which is based on the Harsanyi’s (Harsanyi, 1977) concept of a *negotiation set*, which, roughly, represents the set of all joint plans which agents in a domain are prepared to employ. Agents in this framework follow a number of limiting assumptions including: (i) every agent is rational and maximises personal utility (and calculation of that utility is absolute and unchanging, and can be computed for any satisfaction or partial satisfaction of a set of goals); (ii) every agent has complete knowledge of the environment (including other agents plans, beliefs and valuations) – though this condition is relaxed in certain circumstances in (Zlotkin and Rosenschein, 1994) (iii) each negotiation is a distinct event, unrelated to any other – there is “no history” (agents can’t make promises about future action, for example); (iv) between them, agents can compute and agree on a common evaluation of agent-wide utility for a particular course of action; (v) all agents have common abilities and identical

¹⁸ Note that the term *negotiation* is the one selected in DAI; in Walton and Krabbe’s (1995) more refined analysis, the term negotiation conveys a rather different meaning – typically, relating to discussion of scarce resources – and the process described here would be termed by them *deliberation*.

costs associated with those abilities; (vi) all agents are bound to their commitments. Earlier formalisations included restrictions that the negotiation set between agents never be empty (i.e. there was always a solution which fulfilled participant's goals), and that agents' goals remain unchanged; these were addressed formally in later work, (Zlotkin and Rosenschein, 1990) and (Zlotkin and Rosenschein, 1996), respectively. However, even with the absence of these last two restrictions, the assumptions in the above list – and in particular, the second and fifth – leave the model highly unrealistic for real domains (indeed, they even appear to impinge upon agent autonomy).

An approach which makes far fewer assumptions about agents' rationality and environment is one based upon defeasible reasoning, such as that discussed above in §2.2.3, and in particular on the argument-based defeasible reasoning of Dung (1995). A good example of such *distributed defeasible reasoning*, as it is termed in (Reed and Long, 1997a, p499), is offered in (Parsons and Jennings, 1996). In this framework agents negotiate by communicating proposed joint plans which may involve defeasible inferences – the plans are arguments in the logic of (Elvang-Gøransson *et al.*, 1993). Agents then evaluate incoming proposals on the basis of the calculated *acceptability class* of the plan. The highest class comprises arguments which are tautological (i.e. which require no support) – any such argument would be immediately accepted. The next class comprise arguments for which the recipient agent cannot build an *undercutting* argument (i.e. an argument attacking one of the supports of the proposal) – these too would be accepted, unless steps in the proposal would preclude the recipient fulfilling one of its own objectives. Below this is the class of arguments for which no *rebutting* arguments can be constructed (i.e. arguments which counter the conclusion of the proposal) – but for which undercutting arguments may be available. Proposals which fall into this class would not be rejected out of hand: rather, the recipient agent would convey the undercutting argument to the proposer, who would then be obliged to find an alternative scheme of support for the proposal, in order that the recipient agree to its execution. The penultimate class comprises simply those arguments which are consistent, but for which the recipient agent has rebuttals available. This would lead the recipient agent to communicate both the rebuttal and a counterproposal back to the proposer. The final class of arguments are those which are inconsistent, which the recipient would reject absolutely. Thus the defeasibility of the system lies not only within a single agent's reasoning, but also in the subsequent processing in another agent with potentially quite different beliefs and goals.

Although such a scenario is much less restricted and consequently much more realistic and applicable than the work of Zlotkin and Rosenschein, it does suffer from a key problem¹⁹. This problem arises not from any idiosyncratic implementation issue, but rather, from this kind of use of distributed defeasible reasoning in general. Defeasible reasoning across rational, autonomous agents seems to miss the intuitions of how the agents are functioning. Consider the scenario in Figure 2.3, in which the speaker has two disjunct reasons for (defeasibly) inferring p , and the hearer may have a number of reasons for (defeasibly) inferring $\neg p$. Let us assume that S knows that H believes $\neg p$. Following an account such as that offered by Dung (1995) or Vreeswijk (1992) would lead to S communicating

¹⁹ In fact, there are several problems with the system proposed in (Parsons and Jennings, 1996): they are covered in full in (Reed and Long, 1997a); only the most important problem is discussed here.

exactly those subarguments which together either defeat, or are undefeated by, all the X_i which S presumes H to believe.

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

Figure 2.3 Sample situation between speaker, S and hearer, H

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

The few systems which are suitable for application to real domains (those which can deal with uncertain information, which preserve agent autonomy, etc.) and yet do not rely on distributed defeasible reasoning with its associated counter-intuitive result, are generally founded on empirical observation of human techniques of argumentation. A prime example of such work is Sycara's PERSUADER system (Sycara, 1989), (Sycara, 1990). The remit of this seminal work is concisely offered in (Sycara, 1990):

"Researchers in decision sciences have recognized the need for behavioral mechanisms that underlie the process of changing perceptions, goals and constraints which in turn underlie the process of concession making. Our work presents a model of the process, *persuasive argumentation*, and mechanisms that are used by the parties to effect these changes. Our model represents the belief structure of the parties in a symbolic form that is capable of being dynamically updated as a result of arguments." (p204)

She goes on (p207) to emphasise that PERSUADER concentrates on modelling the process – rather than the participants – of negotiation, and that PERSUADER ...

"...can be an actively participating third party that proposes settlements, produces evaluations and justification of the proposed settlements and modifies proposals to increase their acceptability." (Sycara, 1990, p207).

The PERSUADER system, then, aims to participate in negotiation by creating persuasive arguments

which (i) are tailored to the beliefs of the recipient agent; (ii) are constructed and presented in an appropriate format for the current situation (e.g. determining the appropriateness of the enthymematic form, the use of non-deductive or fallacious reasoning, etc.); (iii) take account of the overall aim of the argument – whether to change the recipient agent’s beliefs or behaviour. In this way, the system is not only very ambitious, but it is also drawing upon psychological theories of persuasion between humans, such as (Abelson, 1959). In a similar way, the proposal of (Reed *et al.*, 1997a) discusses the use of persuasive argument in multi-agent systems (though not exclusively for negotiation).

It is illuminating, however, to compare this work to the recent advances in argumentation theory of Walton and Krabbe (1995), and in particular, their typology of dialogue forms. The title of (Sycara, 1990) is *Persuasive Argumentation in Negotiation*. Walton and Krabbe propose precisely such a technique – after distinguishing persuasion from negotiation (amongst a total of six basic argument-oriented dialogue types, listed in Figure 2.4, below), they describe the use of *functional embedding*, whereby one characteristic dialogue type can be embedded within another (pp72-77). The persuasion type is often found embedded within the negotiation type. Importantly, Walton and Krabbe also touch upon other examples of embedding (though unfortunately, they do not go into any great detail). For example, the information-seeking dialogue can almost always be embedded in any other type of dialogue when one agent is in possession of more information relevant to the current issue. A preliminary investigation of the dialogue types proposed in (Walton and Krabbe, 1995), the potential combinations of embedding, and the relationship between the dialogue types and the notions of agent *cooperation* and *collaboration*, all within a multi-agent framework, is presented in (Reed and Long, 1997c), (Reed, 1998).

Dialogue Type	Initial Situation	Main Goal	Participant’s Aims
Persuasion	Conflict of views	Conflict resolution	Persuade the other
Negotiation	Conflict of interests & need for cooperation	Making a deal	Get the best out of it for oneself
Inquiry	General ignorance	Growth of knowledge & agreement	Find a proof
Deliberation	Need for action	Reach decision	Influence outcome
Information-seeking	Personal ignorance	Spreading knowledge	Gain/convey knowledge

Figure 2.4 Types of Dialogue, adapted from (Walton and Krabbe, 1995, p66)

Identifying the various forms of argument is a crucial step in defining the functional remit of a computational system to generate argument. In the next chapter, this remit is characterised in terms of Walton and Krabbe’s conception of persuasive argument, and the various design decisions which impinge upon the construction of a system so oriented are explored.

III

Framing a Solution

This chapter offers an overview of the *Rhetorica* system, covering the distinction between the two abstraction layers responsible for argument structure and eloquence, and exploring their interaction and inter-dependencies. The system sits in a wider framework first discussed in (Reed *et al.*, 1996), and also investigated here. Before examining either the modules of the framework, or those implemented in *Rhetorica*, a number of design decisions are discussed, some of which explain the relationship between *Rhetorica* and the framework as a whole.

3.1 Design Decisions

3.1.1 Monologue or Dialogue?

The majority of NLG systems are not designed to engage in extended dialogue with a user – consider for example, the tasks of machine translation, generation of software documentation, and automated provision of health education materials. Between them, these constitute a significant proportion of all recent text generation work, and in every case the process of producing the text is completely decoupled from the user's subsequent inspection and assimilation of its content. This decoupling of the NLG process from both user interaction and often, underlying reasoning, has been criticised on the grounds that it involves simply too many unrealistic assumptions. In contrast to research concentrating only on generation, there is increasing interest in integrated 'end-to-end dialogue systems' – those which take user input in natural language, parse, process and perform appropriate reasoning before designing appropriate discourse goals, and producing natural language text in response. A prime example is the TRAINS project of Allen *et al.* (1995), in which the functionality of the NLG process is heavily dependent upon feedback from the hearer, and upon the planning subsystem; these dependencies represent a very tight coupling. A similar coupling is to be found in the argumentation system of Zukerman *et al.* (1996), where the processes of argument generation and analysis are separate, but contribute to a unified 'Argument Strategist'.

How then might the decision to focus on 'non-engaged' (Blair, 1997) discourse – i.e. discourse produced without the hearer playing an interactive role – be justified? The first answer to this question is a practical one: eschewing issues of dialogic interaction from a study of NLG offers an opportunity to focus quite precisely on very specific areas of NLG. In a research endeavour with strictly bounded resources, this focus is critical. Areas of work which are unique to investigation of dialogue can be excised. Several processing modules can be ignored: discourse comprehension and parsing, user model belief revision, and user plan recognition. Indeed the entire architecture can be simplified, since it becomes unnecessary (and impossible) to monitor effects: there is no need to interleave planning with execution, no need to detect failure in execution of the plan, or its understanding, and effect subsequent repair (e.g. dealing with misconceptions), and no need to be able to account for and deal with interruptions.

By itself, however, expediency would not justify a research approach. There are two other important considerations. The first is that such non-engaged discourse research has real, useful applications – the broad areas mentioned above constituting good examples. Indeed, applications where the aim is to produce written text are generally going to use these techniques. Producing (written) health education materials tailored to particular clients is one major area of research activity: examples include the HealthDoc (DiMarco *et al.*, 1995), (Hirst *et al.*, 1997), Smoking Letters (Reiter *et al.*, 1997), and Goldfish and Piglit projects (Binstead *et al.*, 1995), (Grasso, 1997), which aim to influence patients' decisions on smoking cessation, and the RAGs project (Cooper *et al.*, 1996) which assists in genetic counselling. Another area focuses on technical documentation generation, DRAFTER (Paris and VanderLinden, 1996), (Hartley and Paris, 1996), TECHDOC (Rosner and Stede, 1992) and IDAS (Reiter *et al.*, 1995), for example. In all these cases, the written textual output is produced without intervention by a user – there is often a knowledge elicitation or acquisition stage, but this entirely precedes the generation phase, which is itself irrevocably completed before the user sees the result. Even where written textual reports are not the aim – even where a user is sat at a terminal 'interacting' with some NL interface, the generation component is very often designed as a non-engaged monologic discourse system. A good example is the use of NL interfaces for database querying, such as (Maybury, 1992), where there is rarely any history constructed between consecutive user requests.

The other consideration in support of a monologic investigation is that such text is created naturally. This point is relevant to natural language in general, but is of particular pertinence to argumentation. For *persuasive monologues* (Reed and Long, 1997) are very common – advertisements, editorials, political addresses, theses and academic papers can all offer examples of persuasive monologues. Yet despite their widespread use, argumentation theory does not have a competent theory for accurately characterising monologues and distinguishing them from the more usual substrate of investigation, dialogue. For this reason, in conjunction with the fact that the generation of persuasive monologue is one of the key objectives of the *Rhetorica* system, an extended digression is warranted, in which a definition for persuasive monologue is put forward and defended. The discussion is based upon ideas presented in (Reed and Long, 1997), and those offered in reply by Vorobej (1997), and is divided into four sections: the first three characterise persuasive monologue on the basis of its aims, physical situation and internal structure; the fourth then points out some common misconceptions of what

comprises monologue, which are then rejected on the basis of the three preceding sections.

The aims of persuasive monologue (and indeed persuasive dialogue as well) fall into three groups. Firstly, to alter the beliefs of either the hearer (e.g. a letter from one academic to another discussing some matter upon which they disagree), a particular audience (e.g. an academic paper presented at a small, focused workshop), or a general audience (e.g. an article in *Scientific American*). As discussed in (Perelman and Ohlbrechts-Tyteca, 1969), the difference between the constructs *particular audience* and *general audience* is crucial – as mentioned above in §2.2.1, it is used in defining the distinction between persuasive and convincing argument. The further distinction between particular audience and single hearer (which in the work of Perelman and Ohlbrechts-Tyteca are conflated) is important for determining an appropriate level of confidence in the model of the hearer (broadly, that a model of a single hearer is likely to be more reliable than a less specific model which abstracts the beliefs of an entire audience). It is often not transparently obvious who the intended audience is in any given situation – in the debating chamber, for example, the speaker has one or more opponents to whom she is supposed to be addressing herself – the primary aim of her discourse, however, is to change the beliefs of the nonparticipatory audience. This form of ‘misdirection’ is very common in monologue, especially in those examples where a particular position is being attacked. Other permutations are rarer, but one could imagine a scenario in which a monologue was addressed to a general audience and yet the speaker hope only to influence the beliefs of some particular subset of that audience. It should also be noted that Perelman’s terminology is a little misleading, for under the heading of ‘altering belief’ is included more than just persuading and convincing, *viz.* shedding doubt, confusing, confounding and dissuading. Often, a speaker’s ‘best hope’ may be to persuade, but would settle for simply reducing the audience’s certainty in their belief.

Changing the beliefs of an audience is not the only – or even the most common – aim of persuasive discourse. For although most such discourse is constructed in such a way that it *appears* that the speaker’s aim is to influence belief, in point of fact, orators frequently

“aim principally to alter behaviour, generate enthusiasm, or create feelings of various sorts (guilt, pleasure, solidarity), rather than alter beliefs.” (Vorobej, 1997, p2)

The second type of monologue aim, then, involves changing hearer behaviour. As with discourse aimed at altering belief, that concentrating on changing behaviour can be aimed at an individual, a particular audience or a general audience, and has similar scope for ‘misdirection’. Indeed the similarities between epistemic and behavioural change are very great, since commitment to action can be defined as propositional belief – (Walton and Krabbe, 1995) (though as Walton and Krabbe point out, p15, such a relationship may break down if commitment is incurred by an unstructured, heterogeneous audience). It is useful to class these behavioural aims distinctly, because the arguments which service them often involve characteristic reasoning patterns and stylistic constructions.

The third and final group of aims of persuasion are emotive in nature, engendering particular feelings in the audience (- notice that Vorobej’s ‘generating enthusiasm’ can be classed either under this head if it is undirected, or as a behavioural aim if it is directed towards a particular action). This sort of manipulation is unlikely to meet with acquiescence from the audience were it blatant, hence the common technique of building a façade that a monologue’s aim is to alter belief. There is a wide

variety of emotive aims which can be fulfilled through persuasive monologue, which, in addition to Vorobej's list, include impressing the audience, inducing fear or shock, and causing amusement through humour or wit, (and of course, these are far from mutually exclusive). Despite this wide range of characteristic aims – both epistemic, behavioural and emotive – together they distinguish between persuasive discourse and the other argument forms listed in (Walton and Krabbe, 1995), and summarised above in Figure 2.4 (but note that the use of persuasion monologue to alter behaviour clouds the distinction slightly between persuasion and deliberation – this situation can be remedied in part by consideration of the action-oriented nature of deliberation and its typical use of means-ends reasoning). The aims alone, however, fail to distinguish between persuasive monologue and persuasive dialogue. This distinction rests in part on the physical situation in which the argument is conducted.

Brockriede's reply (Brockriede, 1977) to O'Keefe's (1977) proposal that the term *argument* should be divided into the argument₁ – "something one person makes (or gives or presents or utters)" – and argument₂ – "something two or more persons have (or engage in)" is enlightening in that it highlights the distinction between seeing an argument as a process on the one hand and as a product on the other. From an NLG perspective, this is a particularly important distinction to recognise since although monologue is generally viewed (e.g. implicitly by O'Keefe) from the argument-as-product stance, the creation of a monologue from a set of beliefs and goals is necessarily a process. And, crucially, the process of creating a persuasive monologue is assumed to be complete before it is uttered to an audience. Vorobej voices concerns that although a persuasive monologue may not admit linguistic response from the hearer, there may nevertheless be nonverbal indication of a monologue's reception. He thus distinguishes *veiled persuasive monologue* – "where there is no possibility of any physical, verbal, or symbolic contact between the audience and the speaker" – from *face-to-face persuasive monologue* – "where the audience is verbally silenced, but may symbolically interact with the speaker in other ways." (Vorobej, 1997, p3). In a computational setting, such 'face-to-face' persuasive monologue is difficult to envisage, since the channels for non-linguistic communication would have to be expressly designed and built, but it is important nevertheless to emphasise that the *Rhetorica* system assumes that no modification to the monologue plan occurs after realisation of that plan has commenced. To permit such run-time modification would be to re-introduce almost all of the problems of a full dialogue system – indeed it could well be argued that the scenario represents an – albeit impoverished – dialogue. (Furthermore, eschewing the generation of face-to-face monologue also side-steps Vorobej's criticisms concerning the claim in (Reed and Long, 1997) that the potential for true retraction – a defining feature of persuasive dialogue – is absent in persuasive monologue).

The physical situation and involvement of the hearer also forms one facet of the distinction proposed by Blair (1997) between *fully-engaged dialogue* and *non-engaged dialogue*. In examples of the former,

"what is supplied by each participant at each turn is a direct response to what was stated or asked in the previous turn", p5

In contrast, the interlocutors in a non-engaged dialogue

“take up the same topic, defending (apparently) incompatible positions on it, but they do not interact directly with one another ... Even where they interact, each side chooses which of the views of the other side it wants to attempt to refute and which of its own claims it wants to support, and is not forced by questions or challenges from the other side to address the issues that other side deems important.”, p8

Clearly, Blair too conflates into his second category the limited interaction available in Vorobej’s face-to-face monologue with the absolute absence of interaction in veiled monologue. However, the key distinction between fully-engaged and non-engaged dialogues, Blair maintains, is not the physical situation, but the permitted complexity of each turn in the dialogue. He identifies thirteen levels of complexity: at the level of greatest simplicity are question and answer dialogues in which the questions are designed to elicit yes/no answers, and the respondent may only answer yes or no. At the next level of complexity, questions may elicit single propositions. The third level allows an admixture of these two (and is characteristic of Plato’s *Dialogues*). The next level, Blair proposes, is in a separate class, whereby the proponent can offer simple arguments, and the opponent can question the propositions or inferences employed in those arguments. At the next level of complexity, more than one simple argument is permitted. At level six, the opponent is allowed to offer arguments for his doubts. At level seven, the roles of proponent and opponent are allowed to fluctuate dynamically. Level eight again represents a new class, in which arguments can be chained (with supports for support). At the next level, the length of these chains is unrestricted. At level ten, more than one line of argument can be put forward at each turn, and at the next level, multiple lines of argument each of arbitrary length are permitted. Level twelve again enters a new class, where refutations of opposing arguments may be offered. Level thirteen, the most complex, represents the combination of twelve and eleven.

It seems, however, that such an approach is characterised on the basis of the *result of the process* rather than on the *process itself*. Blair’s ‘level-thirteen complexity’ is characteristic of non-engaged dialogue precisely because it comprises the most appropriate forms of reasoning for the *process* of such dialogue to employ.

Persuasive monologue is composed of two forms of reasoning. Firstly, the intuitive ‘case-building’ of presenting arguments in support of the thesis. Premises are supported by subarguments, which themselves are further supported, and so on until basic premises are reached which fulfil one of three conditions: (i) the speaker believes them and has no further information available with which to support them; (ii) the speaker believes the hearer believes them (irrespective of whether the speaker herself believes them); (iii) the speaker believes the hearer will accept them without further argumentation (even though, as far as the speaker’s model of the hearer goes, he doesn’t currently believe them). Without opportunity for the speaker to defer supporting argumentation until prompted by her audience, this case-building is clearly essential. Furthermore, the speaker will often employ multiple chains of support – not because she believes that one particular line of support is insufficiently strong, nor even because she assumes that the hearer will find one line of support weak. Rather, she is ‘hedging her bets’ – given the fact that the hearer model is assumed to be imperfect, it may turn out that a premise assumed to be acceptable to the hearer is in fact rejected, and in such a situation, auxiliary arguments may become vital.

Secondly, there is the more complex technique of presenting counterarguments to the thesis propounded, and then offering arguments which defeat those counterarguments. One example of accomplished use of the technique is Turing's (1950) *Computing Machinery and Intelligence* in which he proposes that human intelligence is theoretically and fundamentally reproducible in a computer, and goes on to counter nine common objections from various philosophical, theological and intuitionistic viewpoints. Each counterargument is aimed at a different hearer, the theological to the theologian, etc., and is constructed precisely for that hearer. Thus the theological objection is countered from theological premises, which Turing indicates he considers dubious at best. Turing also explicitly identifies the two components of monologue which appear in his paper (the counter-counterarguments and the case-building):

"The reader will have anticipated that I have no very convincing arguments of a positive nature to support my views. If I had I should not have taken such pains to point out the fallacies in contrary views. Such evidence as I have I shall now give ..."
(p454)

Turing thus claims that the counter-counterarguments he has presented would not be required if he could offer unassailable arguments for his thesis, and indeed this seems to be generally the case: counter-counterarguments play an ancillary role to the more central case-building argumentation (Reed and Long, 1997a). Again, however, counter-counterargument represents an appropriate strategy for the process of creating non-engaged dialogue: without the opportunity to deal with counterarguments if and when an opponent tables them, a speaker runs the risk of losing the hearer. If the hearer believes he has a valid counterargument for some claim in the speaker's monologue, he may conclude that – regardless of the content of the remainder of the monologue – the speaker's argument is flawed (and therefore not worthy of any further attention). By anticipating and countering as many counterarguments as possible, a speaker improves the likelihood that a hearer will remain unbiased to the end. This claim is supported by noting that in the Turing example, which argued on a very emotive and contentious issue, his own arguments came *after* his long list of the various counter-counterarguments. This type of ordering turns out to play an important role in argument construction, as evinced by results in cognitive psychology, and examined in depth in chapter four.

Thus rather than defining monologue from a product-oriented stance (as Blair does), a more incisive approach is to offer a definition from a process-oriented stance. Using multiple lines of reasoning, for example, is not simply the defining feature of 'level-five complexity' – rather, it is a technique employed in response to situations in which the speaker is aware of her imperfection in modelling the hearer and wants therefore to maximise the likelihood of her thesis being accepted through utilisation of a whole battery of support. Considering only the product of argument leaves any definition susceptible to weakness since no such product can be a true record of the argument – the context will have been lost, and with it, the information necessary to perform classification. The importance of context (a process attribute) can be demonstrated by considering the problems with Blair's scale of complexity. Employing counter-counterarguments, he claims, is at the highest level of complexity (i.e. at the furthest 'solo argument' end of the scale). Somewhat less complex is the use of multiple chains of support; less complex again, single lines of support; and much less complex again,

single argument units. However it is perfectly possible to envisage a persuasive monologue (i.e. a non-engaged, solo argument such as a letter-to-the-editor) which employs nothing more complex than a single argument unit. Equally, it is possible to imagine a debate – involving true engaged argument – in which the first question from the floor involves counter-counterarguments and multiple lines of support. Thus the scale of complexity does not seem to coincide well with a scale ranging from monologue to dialogue. Indeed, the same text of either of the previous examples could be found in situations characterised as either unequivocally solo or unequivocally duet argumentation. In order to distinguish monologue from dialogue, then, it is essential to examine the physical and cognitive context in which the process of argument occurs.

Blair's complexity hierarchy also suffers from another problem in the way in which it implicitly characterises monologue as subordinate to dialogue. The hierarchy discusses the complexity of an individual turn; when that complexity reaches a sufficiently high level, the result can be termed a monologue. However, it seems inappropriate to class a monologue as an extended turn in dialogue, and the reason again turns upon consideration of the process of creating the argument. For that process is not constrained by what the opponent has previously uttered, it has no (external) concept of 'local thesis' or 'current topic', and is not in any way constructed from rules of some super-system. It also makes many more assumptions about the beliefs of the hearer, as monologue is not afforded the opportunity for maieutic elicitation of those beliefs. The speaker is obviously aware that these assumptions concerning hearer beliefs (and attitudes – scepticism, bias, etc.) are not verifiable, and as a result, makes rather more careful use of them, perhaps placing less reliance (or less obvious reliance) upon them than she might in a dialogue, where oversights or carelessness can be addressed at subsequent turns. A speaker recognises that a monologue is a one-shot deal, and that no extra explanation or backtracking can be performed if she misjudges the hearer in some respect. Monologue, then, is constructed with rather more diligence and with greater consideration given to its reception by the intended audience than is a turn in dialogue which is generally more forgiving due to the inherently dynamic nature of its environment. This distinction clearly relies upon examining the process of monologue, and taking into consideration the various contextual factors. For the resulting product could then not only be analysed as a dialogue turn, but could in fact function as a turn in dialogue – a good example is that offered in both (Reed and Long, 1997) and (Blair, 1997) of an academic paper followed by a published criticism: each is constructed as a monologue but can be retrodictively analysed as a turn in dialogue (and indeed this is the thrust of the second half of Blair's paper). The fact that the monologue product is functioning as a turn in dialogue in no way alters the fact that the process was one of monologue (with the various contextual expectations mentioned above) rather than one of constructing a turn in dialogue (which would not have had those expectations). Again, the same piece of text could be the result of the process of monologue in one situation and the process of creating a turn in dialogue in another. So again, identification of monologue relies upon an analysis of the process by which the text was created and the contextual factors thereof.

The assumption that monologue is akin to a turn in dialogue is one of the most common misconceptions regarding its nature. This is demonstrated by the fact that it is held not just in argumentation theory, but also in other areas, including computational research (e.g. (Fawcett and

Davies, 1992)). It is not the only such misconception, however, and mention of several others will bring this digression into a definition of persuasive monologue to a close.

Monologue is not simply a record of a line of reasoning entertained by the speaker to reach some conclusion for her own benefit. For a persuasive monologue has an aim – to alter the beliefs, behaviour or emotions of an audience, and to this end, makes careful use of the hearer model. In contrast, the reasoning processes of the speaker are neither hearer sensitive nor directed towards affecting the beliefs of anyone but the speaker. Similarly, the vital role played by consideration of the hearer's beliefs means that monologue is not soliloquy. The fact that persuasive monologue is constructed around the aim of affecting the hearer is termed by Vorobej the 'intention condition'.

Monologue is not an account of an internalised dialogue between the speaker and the speaker's model of the hearer – or between the speaker and some other conflicting model maintained by the speaker (such as a devil's advocate position). This is a particularly strong claim to make, since many authors agree that any argumentative text – whether monologic or dialogic – can be analysed as an 'implicit dialogue'. The point is made by van Eemeren and Grootendorst:

"Argumentative discourse can, in principle, always be dialectically analysed, even if it concerns a discursive text that, at first sight, appears to be a monologue.... A speaker or writer who is intent on resolving a dispute will have to take just as much account of implicit doubt about his standpoint as of doubt that has been expressed explicitly. His argumentative discourse is ... part of a real or imagined *implicit discussion*" (Eemeren and Grootendorst, 1992) pp42-3.

Similarly, Freeman, extending original ideas discussed by Toulmin (1958), suggests that precise implicit questions give rise to the various types of argument structure (*viz.* divergent, serial, convergent, linked – see §2.2.1) – the relevance question, 'why is that relevant?', causing the further premises to be adduced in a linked structure, and the ground adequacy question, 'can you give me another reason?', causing convergent structure, etc. (Freeman, 1991, pp38-9).

There is, however, a crucial difference between the process of dialogue and the process of creating a monologue, an explanation of which requires the identification of two subsets of a speaker's beliefs. Firstly, the set, S , of beliefs pertaining to propositions the speaker herself holds to be true. And secondly, the set, H_m , of beliefs the speaker believes the hearer to hold. There are two relevant facts about these sets: (1) $S \cup H_m$ can be either consistent or inconsistent; (2) H_m can be either a perfectly accurate model of the hearer's true beliefs (in the current arena of discussion) or can be flawed.

In situations where $S \cup H_m$ is inconsistent (regardless of the accuracy of H_m), the speaker may potentially commit errors by selecting arguments which the hearer can effectively counter (*i.e.* counter to the satisfaction of the speaker) – this might be characterised as the speaker not having 'thought it through'. In other words, she is aware of hearer beliefs which contradict her own, but has not yet had the opportunity to deal with them (either by creating or selecting arguments which defeat those beliefs, or even by retracting some of her own beliefs). This seems to be a common situation given the fact that significant cognitive resources may be required to assimilate a hearer's complex belief set – especially as the model is continually changing throughout a dialogue.

In a dialogue in which the hearer model is imperfect (regardless of the consistency of $S \cup H_m$), the speaker will need to detect the success or failure of her actions and perhaps replan subsequent parts of her argument if appropriate. She will also have the opportunity to dynamically update H_m at each turn. In the real world, beset with uncertainty and limited processing capabilities, an inconsistent $S \cup H_m$, and imperfect H_m are the norm. This leaves three permutations, each of which display interesting characteristics.

Situations where the speaker has diligently prepared for a discussion, and has given careful consideration to the various counter-arguments and beliefs of the hearer, might be characterised as the H_m imperfect, $S \cup H_m$ consistent permutation. Clearly, the preparation would leave the speaker in a strong position – and would constitute good advice to any orator or debator – yet the result of the dialogue is far from a foregone conclusion, since the speaker's model of the hearer is still flawed.

Such is not the case in the situation where the hearer model is perfect and $S \cup H_m$ is consistent. From this, a bizarre dialogue may ensue, in which the speaker will (a) be able to completely predict each hearer response (except perhaps the order in which they are given) and (b) be able to predict with absolute certainty the effect of her utterances. Any dynamic aspect is lost, and it is thus extremely difficult to imagine any real world dialogue in which this could happen. Given the complete absence of any dynamic flow, it would be perfectly possible for the hearer to offer her entire argument in a single turn. Or, to put it another way, the dialogue could be recorded and every utterance of the hearer discarded, leaving only the speaker's utterances. If such a dialogue were to be internalised and conducted between the voice of the set S and the voice of the set H_m , then we have the process of monologue. H_m is obviously perfect in this process, since H_m is acting as a model of itself – the dialogue at this stage is being conducted between H_m and S . This process can indeed be seen as dialogic, but with the caveat that such a dialogic characterisation is one which differs importantly from real world dialogue, since H_m is perfect. Notice that it is not claimed that a real world dialogue simply *couldn't* be held between a speaker and a hearer of whom the speaker has a perfect model. Rather, such a dialogue (a) is very strange and (b) could be used to create a monologue to convince the same hearer.

There is also one further permutation for consideration: a perfect H_m but inconsistencies between H_m and S . Such a scenario is very similar to the real process of creating an extended monologue – one in which the speaker changes their mind part way through and changes what she already intended to say because she realises that the hearer could offer a counterargument (for example). This permutation seems, therefore, to be a component of the process of generating a complex monologue. Importantly, however, it is not a phase which can be inferred from the final structure of that monologue. For the final monologue product will not involve *any* retraction on the part of the speaker. Similarly, a dialogic analysis of the creation of the monologue will also not involve any retraction – it will appear as though it was constructed using a perfect H_m and consistent set $S \cup H_m$. By way of example, consider the simple example in Figure 3.1, below. Figure 3.1 (a) shows the process employed to create the monologue – the inconsistency between the sets S and H_m manifests itself as a retraction by S at S_5 . The final monologue might run as in (b). An analysis of the monologue in (b), however, would run something like (c), in which there is no retraction on the part of S because the sets S and H_m are consistent.




Figure 3.1 Sample process, (a), product, (b), and analysis, (c) of monologue.

The ‘pure’ process of monologue (i.e. the process determined through analysis of the product, such as Figure 3.1(c)) can thus involve no retraction – that is, a speaker cannot directly assert a proposition and its negation within a single monologue. This fact further underscores the difference between the dialogic process involved in creating monologue and that occurring in real world persuasive colloquy, for the latter is usually characterised by the presence of the potential for retraction – without this potential, there would be no hope of one party successfully changing the beliefs of another (Walton and Krabbe, 1995, p10). (It is noted however, that there are situations in which it would be possible to have a dialogue, with one party – even the speaker – refusing any retraction: Vorobej offers an example of discussing Catholicism with the Pope (Vorobej, 1997, p6). It is clear however, that such dialogues represent rather unusual examples of persuasive discourse).

This absence of retraction in monologue is also true in instances where the monologue actually voices some of the Hm counterarguments generated during the internalised dialogue between S and Hm. This generally occurs where the speaker wishes to offer counter-counterarguments (as discussed above), and needs to make the counterarguments clear in the first instance. At no stage in the ‘pure’ process does the speaker perform retraction – to do so would render the monologue incoherent and irrational.

In summary then, a definition of persuasive monologue requires first to distinguish the process of monologue from the resulting product, since the latter has no intrinsic indicator of whether it is monologue or a turn in dialogue. The distinction rests entirely on the various factors which form the context of the process, such as the speaker’s expectations concerning potential for recovery from various communication failures, the precise aims of the discourse, the amount of time allotted for preparation and of space for presentation, the possibility for and frequency of hearer model update, etc. The intrinsic structure of the argument is unable to determine absolutely, but can contribute to the distinction since certain forms (in particular, those that are highly complex) are characteristic of monologue, whilst others (those that are less complex) are characteristic of dialogue turns – due to contextual pressures. Furthermore, any monologue or turn in dialogue can be analysed dialogically. The dialogic process involved in creating monologue, however, differs importantly from usual real-

world dialogue in that the speaker's model of the hearer position is perfect, and as a result, the speaker is never led to retraction.

The aim of the *Rhetorica* system therefore is to generate persuasive monologue. This fact has a number of computational ramifications. Firstly, as mentioned above, restricting investigation to monologic discourse enables a tight focus to be maintained on the core issue – how to generate argument structures and realise them linguistically. Secondly, the process of generating a monologue operates in a certain, predictable environment. The speaker plans the monologue by considering the simulated effects of the actions on a simulated model of the hearer's beliefs within the speaker herself. Within this internal environment of the speaker's beliefs and simulated hearer's beliefs, the planned utterances forming the monologue have predictable effects (even if those effects model the expected variation in responses of a hearer, the model will rest on a representation of the specific range of variation). By exploiting an internal environment the speaker avoids the need to interact during the planning process and therefore is not bound by the constraints of social verbal interaction at that time. Thus, the resources available during the planning process are far less constrained than during dialogue. Often the plans themselves are less rigorously bound by resource constraints during execution. Finally, focus is entirely under the control of the speaker and plans which direct it very carefully between successive elements of a monologue are typical. The machinery which holds this tight rein over focus shift is discussed below in chapter four.

3.1.2 Blackboard or Pipeline?

One core assumption adopted in the vast majority of NLG research is that the generation task can be divided into 'earlier' and 'later' tasks, with chronological order of processing corresponding to decreasing levels of abstraction. The traditional broad division of this *pipeline* model has been into the two tasks of *text planning* (also termed the 'strategic' level by McKeown (1985)) and *sentence realisation* (or 'tactical' level generation). The former is concerned with the large-scale selection and organisation of content, whereas the latter concentrates on determining appropriate syntactic and lexical choice. Both tasks are traditionally approached using planning-based formalisms.

A consensus is emerging, however, that an intermediate level of processing is required, to handle the tasks which neither fit properly into the domain of lexical selection, nor into that of content selection. Such an intermediate level is not only demanded by consideration of the kind of domain information required by a planning process – in addition, it has been recognised that the output of a traditional text planning process is in several ways incomplete with respect to the realisation task which is supposed to immediately follow. In the first place, that message is underspecified, with too little information to narrow the vast choice of realisation options, (Panaget, 1994), (Panaget, 1997). Relatedly, the text plan is most usually realised on a clause-by-clause basis, leading to massive under-utilisation of natural language, which can employ great richness and subtlety to pack information into single clauses and noun phrases (Meteer, 1991) (this fact, of course, is the very reason that speech act analysis is so complex). Finally, giving the text planner unlimited freedom to dictate structure means that it may create strings which are inexpressible in the target language:

“Allowing a generation system to select concrete resources directly ... makes available many more degrees of freedom than the language actually permits.”
(Meteer, 1991, p298)

The fact that generation systems, as Hovy puts it, “run out of steam” at this intermediate level, despite the existence of excellent linguistic realisation tools, is termed by Meteer the *generation gap*.

Claims in NLG of the necessity of an intermediate processing level are supported by psycholinguistic research into human speech: Levelt (Levelt, 1989) discusses three phases in a pipeline model which correspond closely to those defined in NLG. In his speech act inspired account, illocutionary intentions form the basis of the *macroplanning* level, at which content selection and organisation is performed (and a significant body of evidence exists to support this hypothesis through analysis of pause timing in extended monologue – pauses being associated with the cognitively expensive task of macroplanning). Following this, *microplanning* performs various tasks to ensure the hearer will understand the speakers intentions, including generation of appropriate referring expressions. Finally the *formulator* produces the utterance from the preverbal message.

An number of approaches have been adopted in the design of the ‘microplanning’ phase. These are discussed below, and for reference a summary is provided in Figure 3.2. This figure is intended to serve simply to offer an overview of the various approaches and their individual idiosyncratic terminology: no claims are implied as to the precise equivalence of processes to which the terms refer.

Meteer’s own solution to bridging the generation gap is the notion of *text structure*, built from component *abstract linguistic resources* (Meteer, 1991). In this latter term lies the crucial difference between the text structure level and the realisation level: the latter is concerned with *concrete resources*, real lexemes and morphological manipulation. In contrast, the former abstracts from and unifies these concrete resources, in effect creating mini-schemas with which the text structure can be planned.

Panaget extends and refines Meteer’s approach in a number of ways. In (Panaget, 1994), he points out that Meteer’s claim that the text structure level can be used in place of Rhetorical Structure Theory, RST (Mann and Thompson, 1988), i.e. at the text planning level, is detrimental to her theory as it clouds the boundary between the macro- and micro-planning levels (Panaget adopts Levelt’s terminology here). He advocates microplanning for *textual structure* which he implicitly defines as being a separate, subordinate process to the more coarse grained macroplanning. In (Panaget, 1997) he carefully elaborates the notion of abstract linguistic resource honing it to include only lexical, grammatical and morphological resources (rather than Meteer’s rather more broad definition). In doing this, he points out that the traditional lexical/grammatical distinction is of limited use because the same intention can be fulfilled using either sort of device, and that there is no optimal processing order between them. Panaget’s approach is to eschew the distinction altogether, unifying them into the abstract linguistic resource which is composed of various attributes including in particular ideational and textual components, borrowed from systemic linguistic theory.




Figure 3.2 Pipeline models and terminology

Hovy and Wanner offer a significantly different approach, at the level of what they term *sentence planning*. Under this head, they include a diverse range of phenomena which belong neither to text planning nor to realisation, viz. sentence delimitation, sentence content organisation, reference, aggregation, and so on (Hovy and Wanner, 1996). Each of these form a separate processing module in their model.

The pipeline model (or close variants such as Reithinger's (1991) POPEL) is an extremely popular NLG architecture. However, as Levelt points out,

"This .. view may turn out to be too restrictive; there may be situations in which macroplanning is affected by microplanning. Until then, however, the ... view is to be preferred because it is more restrictive" (Levelt, 1989, p110).

And more recently, there have also been criticisms in the NLG community – Hovy and Wanner, for example, pointing out that the pipeline model is founded upon the assumption

"... that 'later' planning tasks such as lexical and clause internal matters do not affect 'earlier' ones such as content selection and text structure organization." (Hovy and Wanner, 1996, p53).

They claim that the sentence planning task in particular is ill suited to the pipeline approach. Instead, they propose that each of the sentence planning tasks should be afforded equal importance, by functioning as a module with access to a common blackboard, from which partial plans are taken, modifications made, and the results then reposted. In their conclusion, they state:

"The phenomena [of NLG] are too interwoven; the problems are too complex ... [and] the more tightly interwoven subtasks of sentence planning remain to be dealt with ... It is important to study them in concert ... precisely in order to determine their interdependencies and their interactions" . (Hovy and Wanner, 1996, p57).

Despite criticism of the pipeline model, Hovy and Wanner nevertheless sit the sentence planning conglomeration (albeit one involving a blackboard) in an overarching pipeline flow from text planning through sentence planning to realisation, and as such, still rely upon the foundational assumption of monodirectional influence.

The framework in which the *Rhetorica* system sits follows the received wisdom in that it too is constructed around a core conceptual pipeline. However, *Rhetorica* itself deviates from the model in an important respect. For reasons discussed in §3.1.3 and §3.1.4, the system implements text planning at a more abstract level than conventional RST-based generation systems – that is, the entire functionality of *Rhetorica* is conceived as preceding RST-based text planning in the pipeline of the framework as a whole. The processing carried out at this level of abstraction is divided into two parts, whose functionality is examined in detail in §3.2, but whose tasks are, roughly speaking, to control the propositional content and organisation of an argument on the one hand and the persuasive effect of that argument on the other. These two tasks are clearly very closely related (a poorly structured argument is likely to be less persuasive, for example), and as such are similar to the sentence planning tasks discussed by Hovy and Wanner (1996). Unlike these latter, however, there is a clear precedence between argument structuring and eloquence generating – namely that the former precedes the latter. For an argument must first and foremost be coherent and understandable to the hearer; processing to enhance persuasive effect can only work within such constraints.

The two abstract planning tasks implemented in *Rhetorica* are thus neither rightly seen as pipelined (because there are many complex interactions violating the assumption of monodirectional influence), nor as modules of a blackboard system (since there is an important precedence between them). The solution adopted is thus to cycle between them before passing control to lower level functionality. The means by which this is accomplished, and examples of the functionality are offered in §3.2 and §3.3, respectively.

3.1.3 Where to Stop?

Given the predominance of the pipeline model in the design of generation systems, the focus of the current work (and consequently, the functional remit of the *Rhetorica* system) can be defined in terms of a segment of the pipeline. Quite simply, the focus is squarely upon the very earliest (i.e. most abstract) phases of argument construction – the organisation of premises, conclusions, and subarguments. Clearly, the information concerning such units and their interrelation is only available at these early stages, and as a result, any features traditionally seen as ‘lower level’ issues but that are a direct result of ‘higher level’ relations must be controlled from the higher levels at which they are demanded. A good example is the use of ‘clue words’, discussed in much greater detail in chapter five. These lexical items relate argument units together, and need, therefore, to be introduced at the same level as the argument units themselves are planned. However, making lexical choice (or, at least, heavily influencing lexical choice) must then affect all subsequent planning – choosing a lexeme to stand between two clauses will obviously affect the composition and presentation of those clauses. The *Rhetorica* system is thus responsible for planning the structure of an argument and all other aspects of the text directly influenced by that structure.

There are several reasons for restricting the scope of the implementation in this way. In the first place, as with restricting the system to monologue, concentrating on the early stages of argument construction facilitates a detailed study of the issues affecting those stages – rather than being obliged to spread research effort along the entire generation pipeline. But as with the dialogue-to-monologue

focus, expediency alone is not a sufficient justification.

The first factor supporting such an approach is the existence of a large field of research concentrating on lower level issues. Indeed, the field has reached a stage of maturity at which off-the-shelf realisation packages are increasingly being selected for use in larger systems. Two prime examples are the FUF/SURGE package of Elhadad and Robin (1999), and the PENMAN and NIGEL work of (Matthiessen and Bateman, 1991), both widely used in NLG research. The *Rhetorica* system itself sits in a framework of which the responsibility for lower level realisation is assumed to be subsumed by processing such as that of the LOLITA system (Smith *et al.*, 1994), (Smith *et al.*, 1994a). (The role of LOLITA and its relationship to *Rhetorica* is discussed in §3.2).

The decision to focus upon a well defined phase of the generation pipeline is a common one, and can be highly rewarding if the assumptions made in regard to the other phases are realistic. As mentioned in the previous section, however, criticism has been levelled at the approach for a number of different reasons, including the possibility of complex interaction (Hovy and Wanner, 1996) and bidirectional influence (Levelt, 1989). Furthermore, the decision to concentrate only on the earlier levels of generation may fall foul of the general criticism of the generation gap (Meteer, 1991) – that a planner will “run out of steam” at the level of sentence planning (Hovy and Wanner, 1996). However, the current work in no way eschews the importance of building bridges across the generation gap nor a wariness of wholehearted acceptance of the pipeline model. Indeed, the sentence planning tasks required in argument generation play a crucial role in the effectiveness of the final text – Huang and Fiedler (1996), for example, examine the control of aggregation in the presentation of proofs in natural language under Huang’s (Huang, 1994) PROVERB system. Rather, it has been found that persuasive text represents an example of NLG in which conventional approaches fall short because they omit higher level, more abstract control over paragraph and supra-paragraph sized texts. This claim may seem surprising, chiefly because of the hierarchical nature of the RST-based planning which sits at the top of most pipeline models. Rhetorical Structure Theory by itself, however, seems to be insufficient for generating the text of an argument, and it is to this issue that the next section is addressed.

3.1.4 RST or ... something else?

The adoption of Rhetorical Structure Theory as a core component of discourse planning has been one of the most striking features of NLG research over the last decade. Tentative claims about its utility in this respect were presented by Mann and Thompson themselves, (Mann and Thompson, 1986), with a diverse range of research subsequently implementing operationalisations of RST (Moore and Paris, 1989), (Hovy, 1991), *inter alia*. The various applications of RST in the field are examined above in §2.1.4, as is a key problem – the lack of representation within the RST framework of communicative intent (first discussed by Moore and Pollack (Moore and Pollack, 1992)). Here, two further problems will be explored: the unsatisfactory relationship holding between RST and the notion of coherency, and the apparent breakdown of the descriptive and generative powers of RST at high levels of abstraction.

Coherence is a notoriously imprecise notion, with a variety of loose definitions in AI and linguistics. To a greater or lesser extent, these definitions attempt to capture the fact that a coherent text seems to ‘hang together’. There seem to be two broad categories into which the various definitions fall: in the first place, are those which refer to extralinguistic referents of discourse, and in the second, those

which refer to the internal structure of the discourse (and then, in some cases, the appreciation of that structure by the interlocutors). Dahlgreen (Dahlgreen, 1996) offers a prime example of the first:

“What coheres is a mental model of the situation a discourse describes, and that the way the mental model coheres is explained by naïve theories of causal and other structure in the actual world.” pp113-4

This claim entails giving up any notion of inherent coherence to a text, and would seem to be susceptible to traditional problems of reference (e.g. (Quine, 1960)). The alternative, positing the existence of internal structure which can be classified absolutely as coherent or incoherent is more popular – from the very loose terms of Ellis (1983): “Coherence implies the presence of some organisational structure of relational principles”, p222, to a slightly more exact suggestion from Hovy (Hovy, 1993): “the speaker’s and hearer’s beliefs agree about how each segment relates to its neighbours (and thus to the whole”, p353. To characterise the general principles by which a coherency judgement should be made, however, is extremely hard. Mann and Thompson’s RST (Mann and Thompson, 1988) has presented an attractive route for investigating these principles, for although they do not explicitly suggest that RST can be used normatively to characterise coherence, there are strong suggestions that they saw this as a potential development, discussing the crucial role of relational propositions in coherency and the difference in coherency of artificial texts created of solely of satellites on the one hand, and solely of nuclei on the other. Furthermore, computational linguistics – and in particular, NLG research – has latched on to RST as precisely a means of ensuring coherency. By way of example, the quote from (Hovy, 1993) continues: “Coherence is enforced by the constraints of intersegment discourse structure relations” (and these relations represent Hovy’s operationalisation of RST).

However, there are a number of problems in trying to define coherency in terms of the presence or absence of an RST analysis (and conversely, of trying to generate only coherent text through an operationalisation of RST). In the first place, there is the notion of *canonical orderings* – that the nucleus usually precedes the satellite in an elaboration, for example. With such ordering requirements unspecified by RST, and yet determinants of resultant coherency, the desired one to one mapping is clearly absent. Marcu (1996), (1997a) expands on Mann and Thompson’s original comments and after examining the results of a corpus study, integrates a statistically founded complete account of canonical orderings into an operationalisation of RST. Although this represents a pragmatic approach to improving coherent generation, Marcu discusses the more general problem that RST analyses can be determined for blatantly incoherent text. This indicates the presence of an underlying mismatch between RST analysis and resulting coherency, which clearly presents a major problem for any NLG system founded upon RST. Finally, and most importantly, there are some coherent texts to which RST analyses cannot be assigned: a generation system may thus not be able to fulfil its goals in an optimal or even appropriate way, simply because RST is insufficiently expressive. Mann and Thompson themselves identify a number of classes of text which are not amenable to RST analysis (and therefore, presumably, would be inappropriate domains for RST based generation), and although Marcu goes on to point out some general shortcomings, Mann and Thompson’s original list includes laws and contracts: two prime examples of argumentation. Clearly for a generation system focusing on

argument, it would be highly undesirable to excise two large classes (and furthermore, Mann and Thompson do not offer a justification for why RST fails to capture these genres, so it is quite possible that they would admit to the list numerous other examples of argumentation). However, RST also suffers from a more fundamental problem, which becomes manifest in argument analysis, for despite their opening claim that “it is insensitive to text size”, RST seems to be unable to adequately represent the high level abstract structure found in larger texts and, crucially, argumentation.

Mann and Thompson discuss the key role played by the notion of *nuclearity* – that relations hold between one nucleus and one satellite. They do, however, concede (p269) that there are a few cases in which nuclearity breaks down – and these they regard as rather unusual. The two types of multi-nuclear constructs they identify are *enveloping structures* – “texts with conventional openings and closings” and *parallel structures* – “texts in which parallelism is the dominant organizing pattern”. Both of these are not just common in argument, but form key components. Enveloping structures are precisely what are described by, for example, Blair (1838), when presenting the dissection of argument into introduction, proposition, division, narration, argumentative, pathetic and conclusion (these are by no means obligatory in every argument, nor is there any great consensus over this particular characterisation; most authors, however, would agree that some such gross structure, usually involving introduction and conclusion, is appropriate). These structures are found with great frequency in natural argument, and cannot, therefore, be ignored. Parallel structures form the very basis of argument, since only the most trivial will involve lines of reasoning in which a single premise supports a single conclusion. Multiple subarguments conjoined to support a conclusion are the norm (see for example, (Cohen, 1987), (Reed and Long, 1997a)), and these, necessarily form parallel structures.

Another point of dissonance between RST and argument analysis is that it is accepted that a text may be amenable to multiple RST analyses – not just as a result of ambiguity, but because there are, at a fundamental level, “multiple compatible analyses”. Mann and Thompson comment:

“Multiplicity of RST analyses is normal, consistent with linguistic experience as a whole, and is one of the kinds of pattern by which the analyses are informative.”,
(Mann and Thompson, 1988, p265).

This contrasts with the view in argumentation theory, where one argument has a single, unequivocal structure. There may, of course, be practical problems in identifying this structure, and two analysts may disagree on the most appropriate analysis (and indeed this latter has a close parallel in RST, since different analysts are at liberty to make different ‘plausibility judgements’ as to the aims of the speaker). The presence of these problems, however, is not equivalent to claiming that arguments may simply have more than one structure, a claim which would pose insurmountable problems to the evaluation process (argumentation theory aims to determine a means of classifying an argument as either good or bad, and the presence of inherent structural multiplicity would present the possibility of an argument being simultaneously good and bad).

Finally, there is a more intuitive problem with RST, highlighted by analysing argument structure. Although there is much debate over the number and range of rhetorical relations (e.g. (Nicholas, 1995), (Hovy, 1993), (Hovy and Maier, 1993)) there seems to be no way of dealing with the idea of argumentative support. In the first place, as Snoeck Henkemans (1997) points out, Motivation,

Evidence, Justification, Cause, Solutionhood and other relations could all be used argumentatively (as well, of course, as being applicable in non-argumentative situations). Thus it is impossible to identify an argumentative relation on the basis of RST alone. Secondly, RST offers no way of capturing higher level organisational units, such as Modus Ponens, Modus Tollens, and so on. For although their structure (or at least the structure of any one instance) can be represented in RST – and, given Marcu's (Marcu, 1996) elegant extensions, even their hierarchical use in larger units – adopting this approach necessitates a lower level view. It becomes no longer possible to represent and employ an MT subargument supporting the antecedent of an MP – rather, the situation can only be characterised as P supporting through one of the potentially argumentative RST relations Q, and showing that $\neg Q$, so $\neg P$, and $\neg P$ then supporting through one of the potentially argumentative RST relations R, therefore R. Apart from being obviously cumbersome, the representation has lost the abstract structure of the argument altogether, and is not generalisable and comparable to other similar argument structures. (It could perhaps be maintained that such structures could be represented as RST schemas, but there are several problems with such an approach: in the first place, schemas cannot abstract from individual relations, so there would need to be a separate 'MP' schema for each possible argumentative support relation; furthermore, the optionality and repetition rules of schema application (p248) are not suited to argument, as they license the creation of incoherent argument structure).

It is for these reasons, and particularly, the last, that although RST plays an important role in the framework of *Rhetorica*, it is subsumed by a layer which explicitly represents argumentative constructs. These constructs can be mapped on to the most appropriate set of RST relations (thus, for example, the implicature in an MP may be realised into any one of the potentially argumentative relations mentioned above). The approach thus maintains the generative capabilities of RST (particularly when extended along the lines of (Marcu, 1996) to ensure coherency through adducement of canonical ordering constraints), whilst embracing the intuitive argumentative relationships at a more abstract level. It is these latter relationships which characterise the structure of the argument (i.e. the structure which argumentation theory strives to determine). The relationships are also unambiguous: a single argument has exactly one structure at this level of abstraction (though multiplicity is not thereby prevented at the RST level). Further, parallelism occurs only at the higher level of abstraction (multiple subarguments contribute to a conclusion, but each subargument is mononucleaic), and similarly, enveloping structures are also characterised only at the higher level (thus RST is restricted to a predominantly mononucleaic structure). Additionally, complete argument texts are not obliged to have complete RST trees. For although most parts of a text are likely to have unifying RST analyses, and although there must be a single overarching structure at the highest level of abstraction, the refinement to RST need not enforce the introduction of rhetorical relations between parts. This expands the flexibility and generative capacity of the system encompassing a greater proportion of coherent arguments (including, for example, those found in laws and contracts). Finally, by separating argument representation from RST, the system is raised to an independence from RST – or any other particular coherence relation theory. Thus it is in principle possible to reformulate the framework in terms of, for example, Hobbs' more parsimonious theory of coherence relations (Hobbs, 1982) without altering the functionality of the more abstract layers which have responsibility for representing argument structure.

This subsumption of RST (or something similar) by additional explicit structural planning leaves open an important question of classification: RST fits well into the traditional ‘functionalist’ framework, but how should a system which introduces further abstracted constructs be classified, and how does that classification affect the design of the system’s other components?

3.1.5 *Functionalist or Formalist?*

Hovy distinguishes between two different approaches to the analysis of text (Hovy, 1993, pp344-5): the formalist camp concentrate on the framework which structures the semantic information, while the functionalist camp focus instead on the structure of the speaker’s goals and the manifestation of that communicative intent in the text. His citations for the former include Kamp’s Discourse Representation Theory, DRT (Kamp, 1981), and, importantly, Cohen’s argument analysis (Cohen, 1987); for the latter, he cites amongst others, (Hobbs, 1982) and of course, RST (Mann and Thompson, 1988) (Though this latter has been criticised for being only ‘partly functional’ as it fails to account for the interactional goals of the speaker, (Snoeck Henkemans, 1997)).

Identifying which of these two approaches to employ – or how to integrate them – in a text generation system clearly has a major impact on the design process, not only in terms of structural representation, but also of the type of goals the system maintains – intentional or semantic.

The current work draws heavily from both camps. Cohen’s work represents an important antecedent for the arrangement of premises and conclusions and their lexical marking; Reichman’s work (1985) similarly contributes to the appropriate use of various clue words. Thus the parts of *Rhetorica* responsible for ordering of premises and conclusions and for introducing clues might be viewed as formalist. Yet the planning of argument structure is clearly descended from the functionalist tradition, being based exclusively upon communicative intentions, and interfacing to a stage of planning at the level of RST relations.

Furthermore, the formalist decisions can at least in part be rationalised with a functionalist approach: a particular ordering of premise and conclusion is not simply an arbitrary decision, but one which is based on surrounding structure – both semantic and intentional (for example, changing a belief deeply entrenched in the hearer’s knowledge may require the generation of many supporting subarguments which are then ordered with innocuous premises preceding the conclusion which can then be further bolstered by stronger, but more antagonistic supports).

This alloy of functionalist and formalist approaches follows Grosz and Sidner (Grosz and Sidner, 1986), in which the linguistic and intentional structures are seen as distinct, complementary components of discourse structure. Unlike much of the research surveyed in (Hovy, 1993), however, the current work also explicitly integrates the attentional state of the discourse, which forms the third component of the theory proposed by Grosz and Sidner. For although most NLG systems have been obliged to implement some form of focusing control, few have attempted to employ the complete model presented in Sidner (Grosz and Sidner, 1986). As discussed in the previous section, many discourse planners have employed RST at their core, and as a result have had to import a notion of focus (a good example of such importation is given in (Hovy and McCoy, 1989)). The subsequent identification of discrepancies between RST-based generation and the informational-intentional distinction in Grosz and Sidner in (Moore and Pollack, 1992), has lead to attempts at unification

(Moser and Moore, 1996), but this latter work also conspicuously omits the attentional structure. As discussed in (Reed and Long, 1997b) and explored in detail below in chapter four, the current work integrates topic manipulation into the broader intentional structure, which is itself hierarchically and functionally related to informational structure.

3.1.6 *Belief and Saliency*

The definition of persuasive monologue offered in the §3.1.1, the operationalisation of RST in §3.1.4, and the discussion of functionalism in the previous section all turn on the ability of the system to satisfactorily model the beliefs of the hearer. In particular, those aims of argumentation which focus on altering hearer beliefs generally presuppose some notion of scalarity, in order to express relative changes such as effecting an *increase* in belief in a proposition (Reed *et al.*, 1996). Yet, as discussed in (Galliers, 1992), *inter alia*, there are fundamental problems associated with representing such scalarity. The simplest approach is to attach to beliefs some numerical representation of the degree to which an agent believes them; any deduction process can then employ, say, Bayesian probability to compute new certainty values. This approach – however formulated – suffers from numerous problematic and counter-intuitive results. Firstly, a low value could be construed either as “lack of evidence, or alternatively, plenty of dubious evidence” (Galliers, 1992, p227) – this is the problem of distinguishing ignorance and uncertainty. Secondly, determining accurate probabilities for beliefs is a difficult and often inappropriate task (this problem motivates the more generic framework for medical reasoning which admits both probabilistic and qualitative reasoning presented in (Das *et al.*, 1996)). Finally, numerical values, without further enhancements, permit combination and comparison of beliefs which are intuitively nonsensical (concluding, for example, that ‘I know that today is Tuesday with 10% more certainty than that environmental pollution is morally aberrant’).

Galliers surveys a number of alternative approaches which eschew a restrictive, counter-intuitive numeric representation in favour of deeper, more flexible notions. In particular, the idea that any belief may be believed with complete certainty, but may be more or less corrigible, i.e. subject to retraction. It is much more difficult to induce retraction of a belief that is deeply entrenched (that is, depended upon by a highly populated set of consequents), than of a belief which has less “explanatory power and informational value” (Gärdenfors, 1992) (- a good example might be the difference between changing one’s belief in the existence of evolution, compared with changing one’s belief in the temporal relationship between *Australopithecus afarensis* and *Australopithecus robustus*: the former would require vast changes to one’s belief structure, the latter would call for important but only very local change). The epistemic entrenchment of these beliefs can then be modelled as a separate feature in the model of the hearer, through employment of an ordering relation (based upon Galliers’ *mc*) holding over beliefs.

Following, therefore, (Galliers, 1992) and (Gärdenfors, 1992), the current work employs a dichotomous view of belief certainty modelled upon the BEL (belief) and BMB (mutual belief) operators of (Cohen and Levesque, 1990). The latter of these offers an explicit means of dealing with the notion of mutual belief, used to define the ‘common ground’ upon which an argument is based (i.e. beliefs which both parties hold, and which both parties also know the other to hold). Mutual belief is defined in terms of an infinite regress of nested beliefs: the problem is to pragmatically choose a level

of nesting beyond which 'mutual' belief is to be assumed. In making this choice, it is understood that no matter how many levels a system can cope with, it is always possible to construct a (highly convoluted) example which exceeds the capabilities of that system. From a psychological (and intuitive) point of view, choosing some arbitrary level of nesting by which to define mutuality seems rather implausible. In humans, it would appear that belief nesting is a resource bounded operation with no apparent limit, and it is possible to construct examples involving increasingly deep nesting which make increasingly challenging demands on a hearer (such as the example in Figure 3.3, below). This evidence can be utilised in implementation, such that the fundamental operator BMB is assumed at a naively shallow level of nesting, and then allowed to be quite corrigible in the light of new evidence.

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

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

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

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

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

Figure 3.3 Deep nesting of beliefs

The ability to nest beliefs of a Cohen and Levesque form has been implemented in Wilks's VIEWGEN system, (Lee and Wilks, 1997), (Ballim and Wilks, 1991), upon which the current work draws quite heavily. The characterisation of belief employed in the current work, however, diverges a little from the basic view espoused by Cohen and Levesque, and implemented in VIEWGEN, in permitting the expression of more than just the presence of existence of belief. Using negation as failure, the basic model supports a representation of an agent's belief in a proposition (or its negation) or absence of belief in a proposition (or its negation) using $BEL(Ag, P)$ and $\sim BEL(Ag, P)$, respectively. However, it also becomes crucial to represent a situation in which an agent is *undecided* about the validity of a claim - that is, the agent is aware of a fact but has not yet concluded whether the belief should be in the claim or its negation. This is in contrast to the state $\sim BEL(Ag, P)$, where an agent has no knowledge with regard to a claim. The state of 'undecided' is represented as $BEL(Ag, ?P)$.

A straightforward use of Cohen and Levesque's modal operators simplifies the integration of the belief models with the planning process used to guide the construction of argument. Plan operators used in argument generation (and NLG in general) based on advanced planning architectures rely heavily upon pre- and post-conditions composed, in part, of epistemic facts: Maybury's communicative acts, for example, employ BELIEVE and KNOW-ABOUT (Maybury, 1992); Reithinger's idiosyncratic MB, BGP, etc. (Reithinger, 1991); and Hovy's influential work (Hovy, 1991).

operationalising Rhetorical Structure Theory (RST) relations (Mann and Thompson, 1988) using Cohen and Levesque's BEL and BMB (and this approach is characteristic of all systems which draw upon RST – (Moore and Paris, 1989), (Paris, 1991), etc.). Although the current work follows this last precedent (though with RST playing a rather different role, as discussed in detail above in §3.1.4), the integration with the planning mechanism is somewhat different. Following (Moore and Paris, 1994), the distinction is drawn between communicative goals and linguistic goals – the former are characterised by their use of epistemic pre- and post-condition specifications, and represent intentions of the speaker. The latter, which arise from fulfilment of communicative goals, represent rhetorical relations or speech acts which can be directly realised into text.

The flat structure – mapping a communicative goal directly to a linguistic goal – is expanded in *Rhetorica*, such that communicative goals form the substrate of the planning process at an abstract level. Thus one communicative goal can be fulfilled by one or more further communicative goals, and so on. At the stage where conventional approaches such as (Moore and Paris, 1994) resolve linguistic goals, *Rhetorica* has an intermediate layer of communicative goals which express the intention of the speaker to make particular facts *salient* to the hearer. Further planning is then invoked to decide upon appropriate linguistic goals to achieve the saliency; the subsequent realisation of those linguistic goals is beyond the scope of the current work for reasons put forward in §3.1.3, above.

Introducing goals of saliency as an interface between other communicative goals and linguistic goals offers a number of advantages over more traditional approaches. In the first place, it neatly accommodates the need for avoiding redundancy in discourse. In addition to generic dicta for conciseness such as Grice's maxim of quantity (Grice, 1975), there are also specific considerations for argument generation. The *Modus Brevis* form of argument (Sadock, 1977) whereby the inference step in a Modus Ponens is omitted (discussed in relation to argument structure in (Cohen, 1987)), is an example of generic enthymematic contraction discussed by Aristotle (1926). In a syllogism, consisting by definition of two premises and one conclusion, it is frequent in natural language to find one component omitted (and, moreover, that including the omitted component would render the text cumbersome and repetitive). By introducing an intermediate level between communicative and linguistic goals, it is possible to effect heuristic control over which parts of an argument to make salient and which to leave implicit. Examples of the phenomenon and its production are given below in chapters four and five. Secondly, the approach also admits the generation of the reverse phenomenon. Marcu (1996a) notes that "Contrary to NLG wisdom, a system capable of generating persuasive text will also have to generate information that is known to the hearer". Thus a communicative goal of persuading the hearer of *p* could be fulfilled by making several facts – including some already known to be in the model of the hearer's beliefs – salient to the hearer. Equally, it is quite possible for several goals of saliency to refer to the same proposition at various stages in the text, causing repetition of information in the text. Introducing redundancy into text has been recognised for discourse in general (Walker, 1992) – informationally redundant utterances (IRU's) functioning either to offer evidence or to centre propositions. It is thus an important advantage that the model encompass generation of IRU's. Finally, as explored in detail in the next two chapters, the division of communicative goals into those of belief and those of saliency facilitates precise control over the ordering of components, and in

particular, enables the conclusion of an argument to be ordered with respect to its premises (an important factor in determining persuasive effect, as discussed in (Hovland, 1957) *inter alia*).

3.2 Architecture

The key component of the *Rhetorica* system is the hierarchical planner based on AbNLP (Fox and Long, 1995). The planning process is conceptually divided into four stages of abstraction, as shown in Figure 3.4.




Figure 3.4 Architecture of the *Rhetorica* system

At the lowest levels, the syntactic and morphological structure is determined. This functionality comes within the remit of the LOLITA system (Smith, 1994), which takes as input a SemNet semantic network (Shiu *et al.*, 1996). The structure of this network can be directly mapped from an Rhetorical Structure (Mann and Thompson, 1988) tree, with appropriate enhancements (including tags expressing various stylistic parameters such as those discussed in (Hovy, 1990), and below in chapter four). Consider, for example, the RS tree in Figure 3.5(a) – which can be mapped on to the SemNet representation in Figure 3.5(b) (the example represents an amalgamation of those presented in (Reed *et al.*, 1997a) and (Mann and Thompson, 1988)). The details of the mechanism by which this can be accomplished, and evaluation of its flexibility and expressive power, are beyond the scope of this work, since for the reasons previously put forward in §3.1.3, the current work is focusing entirely on the higher levels of structure determination (and as explained in §3.1.2, the assumption that this lower level functionality can be decoupled in this way is fully justified).




Figure 3.5 (a) Sample RST tree and (b) equivalent SemNet representation

It is the higher levels of functionality, the Argument Structure (AS) and Eloquence Generation (EG) levels, which form the focus of the current work, but before characterising the remit of these levels of the framework, the infrastructure provided by the AbNLP planner needs discussing.

Both the AS and EG levels conduct their processing under the auspices of a planning mechanism closely related to AbNLP (Fox and Long, 1995), (Fox and Long, 1996), an abstraction-based non-linear 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). Figure 3.6 shows a typical operator, with conventional precondition and postcondition lists, and an encapsulated body comprising six partially ordered goals (further justification and discussion of operator design can be found in chapter four).

```

MP (H, P, X)
Shell: Precond:     $\exists X: X, (X \rightarrow P)$ 
                   $\sim \text{BEL}(\text{Ag}, X)$ 
                  Add:     $\text{BEL}(\text{Ag}, X)$ 
Body: Goals:      t0: PUSH_TOPIC (arg(X,P))
                  t1: BEL (Ag, X)
                  t2: IS_SALIENT (Ag, X, arg(X,P))
                  t3: BEL (Ag,  $X \rightarrow P$ )
                  t4: IS_SALIENT (Ag,  $X \rightarrow P$ , arg(X,P))
                  t5: POP_TOPIC (arg(X,P))

```

Figure 3.6 AbNLP operator characterising Modus Ponens

On building a complete 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).

In addition to the putative computational advantages of AbNLP over more traditional techniques (predominantly those based upon NOAH, (Sacerdoti, 1974)), the approach also exploits the intuitive structure of an argument – that, as discussed above in §3.1.1, an argument is composed of subarguments, in which statements stand as premises supporting conclusions – and these latter then stand as premises in further subarguments, and so on. It is the organisation of this structure which forms the major responsibility of the AS level.

Eventually, arguments inevitably rely upon premises which are unsupported: it is also the AS level which determines at what point it is reasonable to assume that the hearer believes, or will accept a particular claim (or, for the third situation mentioned in §3.1.1, whether or not it is appropriate to use a subargument which cannot be further supported). In addition, if there are a number of potential supports for a statement, the AS level both identifies and selects from the alternatives. The problems of thus determining both the breadth and depth of content in an argument are explored in detail in chapter four.

Using planning solely to determine the content of an argument can be criticised on the grounds that the apparent simplicity of the problem does not warrant such an advanced technique (Reed *et al.*, 1997). Information is conveyed to the hearer in a monotonic way; that is, it is never desirable for the speaker of a monologue to make a statement and then at a later stage to claim its negation: this absence of true retraction in a completed monologue has been discussed above in §3.1.1. Indeed, the criticism can be generalised to discourse in general: it is extremely rare to find examples where a speaker straightforwardly performs such a move, as it is likely to lead the hearer to question the speaker's rationality or truthfulness (of course, there are stylistic, figurative and idiomatic examples where assertion of claim and counter-claim are perfectly acceptable – irony and humour, for example, occasionally license such moves). The result of this monotonicity in information conveyance is that it seems unnecessary to consider the order in which information is presented, since all that it required is a

monotonic accumulation (albeit an accumulation of information selected on the basis of the hearer's beliefs). From this point of view, the planning process employs operators with empty 'delete' lists, and, without the potential for conflict, the planning is trivial. Such a criticism, however, ignores the problem of coherency: a discourse must flow naturally from one topic to the next and must 'hang together' (§3.1.4). As discussed in chapter four, the role of planning at the AS level is in coordinating topic manipulation and focus.

In addition to a notion of textual coherency, there is the related issue of ensuring that the argument itself is coherent: it is this aspect of coherency which Cohen examines (Cohen, 1987), and which forms another key role of the AS level. Although Cohen restricts her analyses to arguments composed of only a handful of clauses, the general principles she establishes seem to be generalisable. As a bare minimum, arguments must be arranged such that their integrity is preserved: by drawing on the work of Cohen, in addition to the diagrammatic analyses of Freeman (1991), etc., this notion of *subargument integrity* can be summarised (as in (Reed *et al.*, 1996)) thus:

- (1) An argument consists of one or more premises and exactly one conclusion.
- (2) A premise can be a subargument (which itself consists of one or more premises and exactly one conclusion: the conclusion then stands as the premise in the superargument).
- (3) A subargument is an integral unit whose components cannot be referred to from elsewhere, nor can the conclusion of a subargument rest upon premises extraneous to that subargument.
- (4) The only exception to (3) is where a conclusion in a distant subargument is restated locally as a premise.

From this base level of integrity, however, there remains wide scope for alternative orderings – both amongst premises, and between premises and their conclusion – which are crucial to the coherency of an argument: this forms a key area of investigation in chapter four, and of the work discussed in (Reed and Long, 1997c). Furthermore, the ordering of components within an argument has a major impact not only on the resultant coherency, but equally on the resultant persuasive impact of that argument: this has been assumed and exploited by rhetoric (e.g. (Blair, 1838, pp420–440)), and demonstrated in social psychology (Hovland, 1957a); (McGuire, 1969). It is therefore necessary in determining an appropriate ordering to consider both the effect on coherency and on persuasive strength. The first consideration comes within the remit of the AS level; the second within the remit of the EG level.

With content determined on the basis of the hearer model, and arranged to maximise coherency and persuasive effect, there remains a range of problems associated with integrating a variety of additional situational, stylistic and hearer-dependent facets to render the text effective and persuasive: intuitively, these aspects might be considered to be the chief components of eloquence. The conglomeration of these features thus constitutes the remaining functionality of the EG level. The means by which the various features affect the surface textual form of the argument are examined in detail below in chapter five; here, the features themselves are summarised and their influence described, expanding work presented in (Reed *et al.*, 1996) and (Reed, 1997).

One vitally important parameter affecting the generation of an argument is the relationship

which the speaker wishes to create or maintain with the hearer. This relationship is established through stylistic rather than structural means, and is not necessarily divorced from other aims: if a hearer accepts the speaker's authoritative stance, for example, the speaker may be able to use the relationship to reinforce his statements. Consider, by way of example, two differing relationship stances taken by automatic advice-givers. The Smoking Letters project (Reiter *et al.*, 1997) offers advice on how to give up smoking, tailored to the individual; the letters are seen by patients to originate from the GP's office. As a result, the advice offered carries with it the weight of authority, and this can be called upon in structuring both form and content of the argument. Though not strictly argumentative, many on-line help systems are using more intelligent approaches in recognising what users can and can't do and tailor their advice appropriately. However, such systems are increasingly adopting a relationship of 'friend by your side' rather than the traditional 'computer expert', which can be intimidating to novices. Adopting this approach necessitates a careful use of language to avoid a situation in which users simply don't respect the advice offered.

The technical and general competence of the hearer are also important parameters to be considered at the outset. General competence determines the hearer's ability to understand complex argumentation (and to some extent, complex grammar); technical competence enables the argument to be pitched at the right level, and affects the choice of appropriate vocabulary. Relatedly, structural limits to various aspects of argumentation, such as the maximum number of subarguments contributing to a conclusion, and the length of each, are in part determined by the capabilities of the hearer. Blair emphasises this point:

“...against extending arguments too far, and multiplying them too much. This serves to render a cause suspected, than to give it weight. An unnecessary multiplicity of arguments both burdens the memory and detracts from the weight of that conviction which a few well-chosen arguments carry.” (Blair, 1838, p432).

The investment that the interlocutors have in the outcome of an argument also heavily affects its reception and hence, construction. Thus if the speaker is perceived by the hearer to have significant potential gain in winning an argument – or potential loss in losing – (regardless, of course, of whether or not the perception is accurate), the hearer may be more sceptical. In addition, if the speaker in fact has a significant investment in winning an argument, this would clearly also necessitate more diligent argumentation.

The primary goals of an argument also have a continuing effect on the generation process. The various aims of argument – epistemic, behavioural and emotive – highlighted by (Vorobej, 1997) and discussed above in §3.1.1 each place differing demands on the structure and expression of an argument. Arguments aiming primarily to change belief often explicitly indicate a clear logical structure (consider, for example, the Socratic method); those which try to increase or decrease the hearer's certainty in an existing belief often involve a greater accumulation of supports; those which aim to induce new belief (i.e. commencing from a point at which the hearer is uninformed) are frequently shorter. Arguments focusing on belief change also need to be sensitive to how the hearer's beliefs are *grounded* – i.e. how they have been arrived at, and how they are maintained (Reed *et al.*, 1997a). Beliefs that are deeply entrenched (Gärdenfors, 1992) in the hearer's knowledge – i.e. beliefs which if

removed would have a massive influence on the remaining belief set – will be much more difficult for a speaker to alter: the hearer holds a bias towards these beliefs. Such bias presents two problems. In the first place, a speaker who intends to go up against hearer bias must ensure that she employs a stronger, more cogent set of counterarguments and supports than would normally be required. In addition, there are also secondary effects: if a speaker were to embark upon an argument against a deeply entrenched belief, it may engender a sceptical reaction, which then prejudices the hearer against further arguments. The notion of hearer scepticism is an important factor in argument construction – if the speaker is aware that during all or part of her argument, the hearer is maintaining high levels of scepticism, she must be much more diligent in the construction of that argument. As discussed in more detail in (Reed and Long, 1997b) and in chapter four, this loose, intuitive notion of diligence can be shown to be amenable to a computational reading. By way of example of the effect of scepticism on argument structure, it is interesting to note that an assumption of scepticism often leads to a thin-end-of-the-wedge argument – Blair, for example, notes that

“... the orator conceals his intention concerning the point he is to prove, till he has gradually brought his hearers to the designed conclusion. They are led on, step by step, from one known truth to another, till the conclusion be stolen upon them, as the natural consequence of a chain of propositions... It is a very artful method of reasoning; may be carried on with much beauty, and is proper to be used when hearers are much prejudiced against any truth, and by imperceptible steps must be led to conviction.” (Blair, 1838, p429).

Arguments to induce behaviour often have closely bound aims which could be characterised as emotive – common examples are engendering enthusiasm and solidarity. Such arguments often employ a number of stylistic devices to meet their emotive aims: Martin Luther King’s “I Have a Dream” speech is an accomplished example of argument where the primary aims are emotive and are associated with further aims of behaviour inducement. The most memorable device in his speech, is of course the repetition of ‘I have a dream’, though he also employs a number of other refrains: ‘One hundred years later...’, ‘Now is the time...’, ‘We can never be satisfied...’, ‘Go back to...’, ‘Let freedom ring’, etc. Other emotive aims also make characteristic use of various techniques: arguments to induce guilt seem often to involve a pattern of numerous short arguments; arguments to impress use a wide vocabulary range, high structural complexity, and rich use of metaphor.

Finally, the modality in which the argument is presented will have a very great impact on the construction process. Thus verbal oration will be organised with clearer indications of structure, a greater use of repetition and summary, and lower limits on argument complexity. There may also be absolute physical limits on the amount of time available for presentation. If argument is to be presented textually rather than orally, additional problems of formatting, layout and graphical arrangement become important. Again, column-inches or page-limits may impose strict restrictions on length. Interestingly, in considering the various characteristic aspects of verbal and textual modes, it might be conjectured that the intended mode could be identified from some later presentation of the same argument in a different mode. It is transparently clear, for example, when reading the Martin Luther

King speech, that it was intended for verbal presentation – and this inference comes not from a couple of minor references in the text to the situational context, but rather from the structure of the text itself. Indeed, after conducting an informal survey, a strong tendency has been identified amongst readers to subvocalise the text at a very low speed close to that of spoken text.

In effecting improvements on an argument, the EG level has control over a number of facets: through the planning machinery, the EG can modify the structure of an argument – introducing new components, pruning existing components and determining order over these components. In addition, the EG level can mark segments for particular styles of realisation – where PAULINE (Hovy, 1990) has global parameters expressing formality, emotional charge, etc., *Rhetorica* can, through the EG level, dynamically tag small parts of a text with combinations of such parameters. Finally, the EG level can introduce clue words which indicate the structure of an argument (*because*, *so*, etc.). These clue words have been investigated from an analytic viewpoint, both within computational linguistics (Cohen, 1987) and argumentation theory (Snoeck Henkemans, 1997), and have also been shown to play an important role in argument comprehension (Cohen, 1987). The mechanics of tagging and clue word introduction, and the relationship between clue words and the broader category of cues (which have been widely studied in NLG and computational linguistics (Knott, 1996), (Knott and Dale, 1994), (Grote *et al.*, 1997)) are examined in detail in chapter five.

The AS and EG levels can thus be seen to have two reasonably clearly defined roles: the AS to decide upon the content and ensure it is partially ordered to avoid incoherence; the EG to further resolve that partial order and introduce stylistic and lexical components, all to improve the persuasive effect of the argument. The issue of ordering is complicated by the fact that ordering components in a configuration which maximises coherency may in some cases be detrimental to the persuasive effect, and vice versa. This interaction demands a careful analysis of the two phenomena, and, consequently, a well defined interface between the two levels.

Although in Figure 3.4, the EG level is shown as subordinate to the AS level, it is clearly inappropriate for the AS to plan an argument in its entirety and then to pass control to the EG to perform minor modifications. There are a number of problems with such a scenario. In the first place, the EG would be unable to effect reorderings between any components other than the terminal nodes (i.e. propositions in a subargument), and would lose the important ability to reorder between abstract units (e.g. between two consecutive subarguments). Secondly, the input to the EG level is a flat, non-hierarchical structure (the AS simply produces a partially ordered plan); without the hierarchical information, it is impossible to introduce appropriate clue words between components of an argument (since those components are not necessarily identifiable in the plan). Similarly, without the hierarchical structure available, the EG is unable to remove appropriate segments of the plan, nor introduce new segments in appropriate places. The processing of the AS and EG levels is thus not pipelined but interleaved: by exploiting the construction in AbNLP of complete abstract plans (i.e. plans comprising operators whose bodies are unrefined), control passes from AS to EG immediately prior to each refinement. The AS level has conceptual priority over the EG level – in other words, considerations of coherency take priority over those of persuasive effect, since an argument must, at the very least be

coherent. Thus at the AS level, an abstract plan of the discourse is created with a partial order over its components; this is passed to the EG level which then amends the plan, adding or removing components and altering the ordering.

There is the potential in this framework for a highly undesirable result, namely, that processing at the EG level might violate the coherency carefully constructed at the AS level. By placing under the remit of the EG the potential to utterly revoke AS level planning, there is no advantage in separating the two phases at all (since the EG completely subsumes the functionality of the AS). In order to maintain the intuitions (i) that content structuring for coherency is separate from structuring for persuasive effect, and (ii) that coherency takes precedence over persuasion, the framework needs to be refined. This refinement also needs to take account of the fact that neither the issues of persuasive effect nor of coherency are dichotomous – as mentioned in (Reed and Long, 1997b), the *Rhetorica* system can be seen as mediating between the constraints and maximising both coherency and persuasive effect (rather than ensuring either).

To deal with these issues, the partial order planned at the AS level is augmented by modalities, with each ordering indicated as either *hard* or *soft* – those of the former class are inviolable at the EG level; those of the latter will stand only if the EG does not deem it necessary to select an alternative. The hard/soft distinction is also employed to qualify decisions of content – operators included in the plan as a ‘hard’ decision may not subsequently be pruned by the EG. In this way, the AS level can make both demands and suggestions on the basis of coherency constraints, and the EG can then work within the hard demands, and decide to follow or modify the soft suggestions. The framework thus captures the intuitions of separability of the two tasks, whilst facilitating the interaction and feedback between them, and maintaining the theoretical integrity of each level.

3.3 Worked Example

In order to clarify the issues discussed above, set them in the AbNLP environment, and demonstrate the AS-EG flow of control, a small example is examined and the planning process engaged in its production is analysed. The example is taken from (Reed and Long, 1997b), which, in (Reed and Long, 1997d) is discussed in relation to its original textual form (the original argument was taken from the corpus). Only the propositional structure is examined here.

The initial situation is one in which the hearer is assumed to have no relevant beliefs, and the speaker holds *a* to be true, supported by *b*, *c*, and *d*; *c* and *d* are further supported by *e* and *f*, respectively. This structure is shown in Figure 3.7 below.




Figure 3.7 Sample argument structure

The process is initiated with two goals, $BEL(ag_0, a)$, expressing the aim that the hearer, agent ag_0 , should believe the proposition a , and $IS_SALIENT(ag_0, a, _)$, expressing the aim that a should be salient to the hearer. The third parameter on the second goal represents the context in which a proposition is to be made salient; its role is explored in more detail in chapter four. The underscore represents the initial situation (corresponding closely to the *mutual initial discourse context*). It is goals of belief which give rise to structural planning; those of saliency introduce flexibility into the way information is expressed (as discussed in §3.1.6, above).

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), and play a key role in argumentation (Reed and Long, 1997a). 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 discourse planning has generally been approached through the use of some 'for-all' operation. Maybury (1993), for example, makes explicit use of $\forall$ 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 fulfil a BEL goal. (Although UCPOP (Penberthy and Weld, 1992) offers a principled solution to the iteration problem, other features – principally its nonhierarchical nature – make it unsuitable for argument planning).

The first BEL goal, then, is fulfilled by three Modus Ponens operators, viz. $MP(ag_0, a, b)$, $MP(ag_0, a, c)$, and $MP(ag_0, a, d)$, and the $IS_SALIENT$ goal is fulfilled by a corresponding $MAKE_SALIENT(ag_0, a, _)$. The AS level employs a general heuristic to introduce a soft constraint that the expression of the conclusion – i.e. the $MAKE_SALIENT$ goal – should precede its supports (this heuristic is based upon Blair's (1838, p.429) distinction between *analytic* and *synthetic* argument; the former, where the conclusion is clearly stated at the outset, is to be preferred, *ceteris paribus*). The result of this first round of planning is the abstract plan in Figure 3.8.

t0: $MP(ag_0, a, b)$	[hard]
t1: $MP(ag_0, a, c)$	[soft]
t2: $MP(ag_0, a, d)$	[soft]
t3: $MAKE_SALIENT(ag_0, a, _)$	[hard]
(t3 < t0)	[soft]
(t3 < t1)	[soft]
(t3 < t2)	[soft]

Figure 3.8 First AS plan with partial order

On the generation of this plan, control passes to the EG level, where a heuristic employing knowledge of the persuasive strength of the various propositions fixes the position of the conclusion. Because the conclusion a is noted in the knowledge base to be rather weak, it is inappropriate to have it

positioned at the head of the argument: better to have it succeed a strong subargument and then be further supported. In this case, the heuristic fixes its position between t_0 and t_1 . The first soft ordering suggestion from the AS is thus revoked.

With no further structural manipulation (and, in this example, no further stylistic issues) to be considered, the EG returns control. The abstract plan is now complete and undergoes refinement, opening up each of the bodies of the MPs (the MAKE_SALIENT operator, as a primitive, remains in the plan unchanged with respect to content or position). This results in the list of goals in Figure 3.9(a) (in which the partial order and hard/soft indication have been omitted for clarity). Control reverts to the AS level, where further heuristic manipulation of the structure is effected to improve coherency. In particular, a potential problem is detected in the final MP argument (the argument from d). With the position of the conclusion fixed between the first two subarguments, the coherency of the plan in Figure 3.9(a) 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 ($d \rightarrow a$) is expressed before d itself, thus indicating to the hearer the relevance of d to the conclusion. This reordering is effected in Figure 3.9(b).

(a)	(b)
PUSH_TOPIC (arg(b,a))	PUSH_TOPIC (arg(b,a))
BEL (h, b)	BEL (h, b)
IS_SALIENT (h, b, arg(b,a))	IS_SALIENT (h, b, arg(b,a))
BEL (h, $b \rightarrow a$)	BEL (h, $b \rightarrow a$)
IS_SALIENT (h, $b \rightarrow a$, arg(b,a))	IS_SALIENT (h, $b \rightarrow a$, arg(b,a))
POP_TOPIC (arg(b,a))	POP_TOPIC (arg(b,a))
MAKE_SALIENT (h, a, _)	MAKE_SALIENT (h, a, _)
PUSH_TOPIC (arg(c,a))	PUSH_TOPIC (arg(c,a))
BEL (h, c)	BEL (h, c)
IS_SALIENT (h, c, arg(c,a))	IS_SALIENT (h, c, arg(c,a))
BEL (h, $c \rightarrow a$)	BEL (h, $c \rightarrow a$)
IS_SALIENT (h, $c \rightarrow a$, arg(c,a))	IS_SALIENT (h, $c \rightarrow a$, arg(c,a))
POP_TOPIC (arg(c,a))	POP_TOPIC (arg(c,a))
PUSH_TOPIC (arg(d,a))	PUSH_TOPIC (arg(d,a))
BEL (h, d)	BEL (h, $d \rightarrow a$)
IS_SALIENT (h, d, arg(d,a))	IS_SALIENT (h, $d \rightarrow a$, arg(d,a))
BEL (h, $d \rightarrow a$)	BEL (h, d)
IS_SALIENT (h, $d \rightarrow a$, arg(d,a))	IS_SALIENT (h, d, arg(d,a))
POP_TOPIC (arg(d,a))	POP_TOPIC (arg(d,a))

Figure 3.9 Result of refinement (a) before, and (b) after ordering

Reordering thus occurs in two stages: once at the goal fulfilment phase, and once at the refinement phase. This is necessitated by the requirement that two forms of reordering should be possible: (i) between the supports for a conclusion, and (ii) between supports and their conclusion. This distinction is explained in (Reed and Long, 1997d) by consideration of orderings within an operator body, such as that of MP (in Figure 3.6, above). Ordering between the supports for a conclusion

corresponds to arranging the order of the two pairs t_1-t_2 and t_3-t_4 – i.e. whether X precedes $X \rightarrow P$ or vice versa. These goals are posted at refinement, necessitating ordering at that time. In contrast, positioning a conclusion corresponds to arranging the order within each of those pairs – i.e. for t_1-t_2 whether making X salient should precede or succeed the argument bringing the hearer to believe X , and similarly for t_3-t_4 , whether making $X \rightarrow P$ salient should precede its support or not. In order to account for placing the conclusion amongst *multiple* supports, this latter form of ordering needs to occur not between the `BEL` and `IS_SALIENT` goal, but between the *supports* for the `BEL` and the `IS_SALIENT` goal. As these supports are generated at goal fulfilment, this second form of ordering has to occur separately.

The small reordering illustrated in Figure 3.9 completes the work of the AS at this level of abstraction: control passes to the EG once more for further reorderings for persuasive effect. In this example, there are none, and the processing for this level of abstraction is complete. The AS then takes as input the set of partially ordered goals in Figure 3.9(b), and performs another round of planning to fulfil those goals. The processing cycle thus comprises five distinct phases, as shown in Figure 3.10, below: first, AS level planning, goal fulfilment, and ordering to maximise coherency; next, EG level processing to effect further reorderings (within the hard constraints laid down at the previous stage); then refinement of the operators selected; next AS reordering with the new body goals; and finally, EG reordering, again within the AS constraints.




Figure 3.10 High level processing cycle.

This double interleaving of AS and EG processing (once with goal fulfilment and once with refinement) ensures that the interplay between coherency-based and persuasion-based considerations are taken into account during both the required ordering phases, endowing the framework with a rich model of the high level generation process.

The broad summary of a round of planning for a small example given in this section has outlined the functionality of the AS and EG levels, and indicated the means by which the flow of

control passes between them. The next chapter examines the functionality of the AS level in much more detail, exploring the role of the topic manipulators, and enumerating the logical and quasi-logical operators available to the planning process, and the various heuristics for influencing the structure of an argument to improve its coherency.

IV

Focus and Order

There are two key tasks to be carried out in organising the structure of an argument. The first is to construct a plan of the components and their interrelationships to ensure that the resultant structure is coherent. The result of this task, however, is underspecified with respect to a complete, optimal order, as pointed out by (Suthers, 1994). The second task thus addresses the issue of resolving a complete ordering over plan components. This chapter examines the two tasks in turn. For the first, the underlying notion of *focus* is introduced and explored in the context of related work within both argument generation and NLG as a whole. Following this, the operators (deductive, inductive and fallacious) available to the planning mechanism are introduced, and then the planning mechanism itself is explained in some detail. The role of intentions and the attentional state within the planning process is then addressed, before the discussion explicitly returns to the notion of focus, and how the topic stack, integrated with the planning machinery, plays a key role in the development of focus in a text. The second section explores the second task, and presents an ordering mechanism which relies on a wide-ranging body of heuristics. The discussion is subdivided into two parts, the first addressing techniques for ordering components with the aim of improving coherency, the second, those aimed at improving the persuasive effect of an argument. Further subdivision is arranged on the basis of the source of the heuristics: the fields of rhetoric and psychology contribute the majority of the heuristics implemented, but these are supplemented by further heuristics suggested by intuition. Finally, the interaction between the heuristics is explored through an examination of AS-EG communication and the use of strict and defeasible constraints. The functionality presented in the earlier sections is then drawn together in a single extended example, the text of which is taken from the corpus. In working the example, the generative capability of the system is carefully analysed and compared to the structure of the original text.

4.1 Maintaining Focus

As a coherent discourse progresses, the attention of the hearer is directed from one piece of information

to the next. Aspects of the process of thus directing hearer attention – *focusing* – have been extensively studied from a variety of viewpoints, so incorporating some aspect of its control into a formal model necessitates a clear identification of the phenomenon to be tackled. Derr and McKeown (Derr and McKeown, 1984) summarise a number of influential linguistic theories which deal with focusing and its manifestation, including those drawing distinctions between topic and comment, given and new information, and theme and rheme. Though exhibiting important differences, these various theories all detail a common notion that a given utterance is ‘about’ something – that some information (be it explicit or implicit) is more important and more central than other information conveyed in the same utterance. The way in which the flow of such focused information is developed through a discourse is then one of the chief determinants of the resultant coherency of that discourse.

Unsurprisingly, therefore, a similar notion of focus also forms a key component in most theories in computational linguistics, of which four are of particular pertinence to the current discussion. First, Rhetorical Structure Theory (RST) (Mann and Thompson, 1988) is founded upon the concept of nuclearity, which, as discussed above in §3.1.4, explicitly requires bipartite asymmetry in almost all rhetorical relations. Under this asymmetry, one piece of information functions as a nucleus to which additional, satellite, information can be attached. Mann and Thompson offer in support of this assumption an analysis of a text from which two artificial ‘summaries’ are created – one containing only the satellites from the original, the other only the nuclei. The former is blatantly incoherent and nonsensical; the latter is impoverished, but clearly semantically coherent. The way in which nuclear information is presented is thus clearly related to the linguistic notion of focus (though Mann and Thompson’s account itself does not directly handle focus in a discourse).

The second major attempt to characterise focusing is offered by McCoy and Cheng (McCoy and Cheng, 1991) in their construction of *focus trees*. Their primary aim is to characterise the priming effects by which earlier utterances set up a particular context for interpretation of later utterances. At a given stage in a discourse, the participants have a partially constructed discourse focus tree, and maintain an indication of the current node in that tree. This state then determines which focus shifts are licensed and which incoherent (thus, for example, if the current focus is an object, the candidates for focus shift include the attributes and prominent actions of that object). One useful result of the work is in characterising the ‘layers of focusing’ identified by McKeown (1985, p67): the constraints on focus shift active at the level of discourse referents (such as a balloon, in her example) are also active at the lower level of the attributes of that level. Thus in her example, it is more coherent to group together comments pertaining to the colour of the balloon and then comments pertaining to its size (in the same way that all the comments regarding the balloon are grouped together). It is surprising that neither McKeown nor McCoy and Cheng extend this notion of layered focusing to embrace the distinction drawn by Grosz (1977) between *immediate* and *global* focus – the former being active from one utterance to the next, the latter on a larger scale. Though McCoy and Cheng implicitly seem to assume such a glissando from coarse to fine scale, they do not make any explicit claim to the fact, though a number of references to the “low level focusing” of (Sidner, 1979) and (Grosz *et al.*, 1995), discussed below, seem to suggest that they view their contribution closer to the global end of the focus scale. Nevertheless, like much of the work on focusing, one of the major applications of their work is in

generating anaphoric reference – by identifying the current focus of attention, it is possible to determine appropriate anaphora, and in particular, to pronominalize effectively. Perhaps the single most influential theory of focus shift designed expressly for resolving anaphora forms the third on the list: centering theory.

Grosz, Joshi and Weinstein (Grosz *et al.*, 1995) explain the importance of maintaining the focus of attention with reference to the “inference load placed upon the hearer” – by choosing linguistic resources which are appropriate to the current focus (e.g. through the use of pronominalization) the cognitive effort required by the hearer to extract the speaker’s intended meaning is minimised. They define the term *center* as “those entities serving to link [an] utterance to other utterances in the discourse segment that contains it”. Finally, by defining the means of linking centers between utterances, constraints can be expressed for minimising hearer inference load. Centering theory is thus squarely positioned at the ‘immediate’ end of the focus scale; related work by Grosz and Sidner (1986) complements centering theory by dealing with broader components of discourse.

Although (Grosz and Sidner, 1986) does not solely concern itself with focus of attention, it does examine the broad attentional structure in discourse, and its relation to the intentional and informational aspects. The basis for Grosz and Sidner’s account is the notion of decomposing a discourse into *discourse segments*, each fulfilling a single intention. The overall intention of a discourse – the *discourse purpose* – is contributed to by the various discourse segments (and in turn, each *discourse segment purpose* can be further supported by other discourse segments). The various discourse segments are associated with *focus spaces*, which indicate salient, contextually relevant information (a similar notion of focus spaces is also employed in McKeown’s TEXT system, (1985)). The discussion returns to the relationship between attention and intention in more detail in §4.1.3, below.

With important differences between these theories which aim to characterise essentially the same phenomenon, there has been considerable interest in trying to unify aspects of the theories, and in particular, in forming a bridge between those concerned with lower level, immediate focusing and those concerned with features at the global end of the scale. With the convenience and consequent popularity of RST, it is unsurprising that the two foremost attempts have been in unifying RST and focus trees on the one hand, and RST and Grosz and Sidner (1986) on the other. Hovy and McCoy (Hovy and McCoy, 1989) produced promising results in the first area, though the outcome was still underspecified with respect to coherency, and the work seems not to have been developed further. More recently, Moser and Moore (1996) performed a preliminary analysis in the second area, equating RST nuclearity with Grosz and Sidner dominance, though there are a number of problems associated with this approach, chief of which are (i) that the result can be inflexible because of the lack of a one to one mapping between intentions and rhetorical relations and (ii) intentional information is often left implicit in a text, whereas RST nuclearity enforces explicit realisation.

These approaches to the analysis and generation of focused text, and the interrelations between the approaches can be summarised as in Figure 4.1, which also gives some indication of the scale of the text with which the theories are primarily concerned (in terms of Grosz’s immediate/global distinction).




Figure 4.1 Major theories of focus

In addition to the four theories summarised above, three further merit brief mention because of their direct focus upon mechanisms of topic shift and coherency within argumentation in particular. Firstly, Huang's (1994) PROVERB generation system which intermixes conventional top down planning with localised bottom up processing, following the recognition in psycholinguistics of the partially "unplanned, spontaneous" nature of language. Responsibility for focus change – and thus resulting coherency – is divided between the two processing forms. Secondly, Cohen's (1987) work concentrating on analysis of argumentation identifies numerous heuristics and maxims which ensure appropriate development of focus through an argument – and the discussion returns to her work in more detail in §4.2, below. Finally, Reichman's (1985) treatment of topic using *context spaces*, which roughly correspond to 'subarguments'. Each context space constitutes a single move, such as a claim, a support, a challenge, a concession, etc. Movement between context spaces is restricted through constraints expressed in an ATN, which permits appropriate suspension and subsequent resumption of context spaces (such that, for example, on resuming a particular context space, further support could be adduced for the claim made in that space). Reichman's work was not implemented.

The handling of focus flow in *Rhetorica* is closest in spirit to that of Reichman's work, but with several important departures. Firstly, the attentional control is explicitly distinguished from both intentional and informational processing, following (Moore and Pollack, 1992), (Moser and Moore, 1996), (Young and Moore, 1994) etc. This approach, as mentioned in (Grosz and Sidner, 1986, p176), is in contrast contrast with Reichman's work, where the context spaces themselves are seen as playing informational roles of support or attack. Secondly, the concurrent development of argument structure and topic space is founded upon the well researched mechanics of argument rather than on intuitive insight. Finally, the topic manipulation is carried out explicitly by operators which are introduced in the planning process. The means by which such manipulation is effected within the planning framework is examined in the next section.

4.1.1 AS Level Planning Operators

The operators employed at the Argument Structure (AS) level follow the proforma offered by AbNLP (Fox and Long, 1995), comprising in the shell, an unordered precondition list, an unordered postcondition list, and an unordered delete list, and in the body, a partially ordered list of goals. The preconditions are typically purely epistemic, expressing the state of knowledge of the speaker and presumed state of knowledge of the hearer. The postconditions are similarly constructed, generally expressing beliefs to which the hearer will be brought through application of the operator. The current work thus almost exclusively considers only goals of hearer belief change. That is not to say that the framework as a whole is unable to deal with other intentions (such as KNOW-ABOUT, and others listed in Table 1 of (Moore and Paris, 1994, p670)²⁰), but rather that the structure of argument is predominantly characterised by intentions of the form $BEL(H, P)$, and that other intentions are typically found in roles subordinate to the main structure of the argument. Thus the realisation of these various intentions is more appropriately carried out at a lower level (and planned directly to RST, following Moore and Paris). Finally, operator delete lists are typically empty. This is as a characteristic feature of the domain: as discussed in (Reed *et al.*, 1997), and in §3.2, above, communication can be seen as a ‘monotonic’ process, whereby there is a continual accumulation of knowledge in the hearer’s mind as a result of the speaker’s utterances. For the application of one operator to ‘clobber’ the results of another would be for one piece of text to conclude the reverse of another (or at the very least to induce a belief change in the opposite direction). Since both pieces of text are ultimately aiming at the same thing (i.e. there’s no blatant retraction in the product of a monologue, §3.1.1), the result would be incoherent. (Although counter-counterargumentation plays an important role, the counterargument to be attacked is not claimed to be true by the speaker, rather she claims that some party other than herself holds the counterargument to be true). The Modus Ponens operator given in Figure 3.6 in the previous chapter is a typical example, exhibiting each of these three features – epistemic pre- and postconditions and null delete list.

Almost all operators at the AS level include in their body two topic manipulation goals, a `PUSH_TOPIC` and a `POP_TOPIC`. Following most of the work in the area (with one notable exception, discussed below in §4.1.4, below), *Rhetorica* employs the notion of a *topic stack*, onto which new topics are pushed, and completed topics popped: the conceptual mechanics of this notion are trivial (but the reader is referred to (McKeown, 1985, p57) for details if necessary), and the implementation details are similarly simple, (again, as described in §4.1.4). Thus most operator bodies set up goals for an appropriate topic to be added to the stack and then removed. As might be expected, these actions have a fixed position with respect to the other goals in the body, namely, that the `PUSH_TOPIC` occurs before, and the `POP_TOPIC` after, all of the other goals in the body. With n other goals in the body, this generally requires only $2n$ orderings to be recorded in the body (since there are usually no other fixed orderings within operator bodies). Thus the Modus Ponens operator of Figure 3.6 requires eight orderings to ensure that the topic manipulation goals are fixed at the outset and conclusion of the planning under that operator.

²⁰ Though in the table offered by Moore and Paris, both $(BEL \text{ ?h ?proposition})$ and $(PERSUADED \text{ ?h ?proposition})$ would be subsumed by the goals active at the upper levels of the *Rhetorica* framework.

The choice of operators implemented at the AS level has been influenced by a number of factors. The rules of inference are clear candidates for operationalisation: moves such as Modus Ponens are clearly vital components of any argument – though, as noted in (Grosz and Sidner, 1986, p201), it is inappropriate to view the implication step as one of conventional material implication. The relationship is rather one of *support* – the hearer must be brought to believe that (given the current context and domain of discourse) the first proposition warrants, in part, concluding the second. Even given this weaker, predicate-based reading of a Modus Ponens argument, it is still unclear that any of the other rules of inference (which are, after all, formally redundant) should be necessary. The answer lies in the second consideration, which is entirely empirical – the reason that the AS level needs to be able to employ other rules of inference is that such argument forms occur naturally. Modus Tollens, for example, is perfectly common, as demonstrated in this simple example:

“Radioactive elements disintegrate and eventually turn into lead. If matter has always existed there should be no radioactive elements left. The presence of uranium, etc. is scientific proof that matter has not always existed.” (the argument comes from a pamphlet of the Worldwide Church of God, cited in (Fisher, 1988)).

Further, there is a variety of evidence which suggests that Modus Tollens in fact occupies a crucial position in human reasoning (Ohlsson and Robin (1994) cite examples not only from psychology, artificial intelligence and empirical observation, but also by reference to classic examples of Euclid, Galileo, etc.)

Disjunctive Syllogisms are also found reasonably often (though many apparent examples of Disjunctive Syllogism are often cases of *false dichotomy*, explained in Appendix A), but the remaining rules of inference are found very rarely. For this reason, only the three logical argument forms, MP, MT and DS, are available to the AS level. Before examining the operator definitions for each of these, two digressions are warranted: one to explain the absence of the Conjunction and HS operators from this list, and one to justify the need for three versions of each operator.

Conjunction might very well be expected to figure on the list of AS operators. Its absence results from the implicit assumption of multiplicity of support, and from the notion of multiple subarguments. As discussed in §2.2.3 and (Reed and Long, 1997a), a claim can be supported by either a single proposition, or by several propositions either independently (in a convergent structure) or in tandem (in a linked structure). The logical rule of conjunction is thus subsumed by the structure of convergent and linked argument constructed through the normal planning process (though as explained in §3.3, deriving multiple supports is in itself a nontrivial task given the iteration problem).

Hypothetical Syllogism is also a natural candidate for employment, since intuitively it seems appropriate to be able to claim “if *a* then *b*, and if *b* then *c*, so *a* implies *c*”. The reason that this form has not been identified in a sample of natural language is because it is supplanted by a close relative – “*a*, and if *a* then *b*, so *b*; if *b* then *c*, therefore *c*”. In other words, natural language generally concludes implicitly the intermediate premise (using a standard MP), so what may at first glance appear to be HS is more generally a chain of two MP arguments in a sorites.

The list of available logical rules thus comprises only MP, MT, and DS, which, it might be assumed, justify the implementation of three corresponding operator definitions. However, there is a

complication in that the operators are being applied not in a classical situation of bimodal truth values, but in an environment of belief. Even after eschewing scalarity, such an environment demands that four possible values for an agent's belief in regard to a particular proposition be distinguished: belief, disbelief, indecision and unawareness – this last can be maintained only as a nested belief – e.g. the speaker believing that the hearer is unaware of a proposition. This analysis of agent belief is discussed more fully in §3.1.6 and (Reed *et al.*, 1997a). Thus for each of the three rules of inference, it is necessary to characterise three different epistemic start states: one in which the hearer disbelieves the conclusion, one in which the hearer is undecided, and another in which he is simply unaware of the conclusion. Though the bodies of the operators for each of the three epistemic states are the same (each situation demands that in MP, for example, the premise and its support are identified and made salient), it is nevertheless important to distinguish the different forms to permit a variety of techniques to ensure successful persuasion (an argument to persuade a hearer who is heavily biased against a conclusion will clearly be quite different in nature from one to persuade a hearer who is ill informed – this is discussed in (Reed, 1997), and the current discussion returns to the point in chapter five). The nine operators are shown in Figure 4.2 (the orderings constraining the body goals – enforcing initial and terminal positions for PUSH_TOPIC and POP_TOPIC respectively – are omitted for clarity).

MP (Ag, X, P)

```

Shell: Precond:   $\exists X: X, (X \rightarrow P)$ 
                  BEL (Ag,  $\sim P$ )          ***
                Add:  BEL (Ag, P)
Body:  Goals: t0: PUSH_TOPIC (arg(X,P))
          t1: BEL (Ag, X)
          t2: IS_SALIENT (Ag, X, arg(X,P))
          t3: BEL (Ag,  $X \rightarrow P$ )
          t4: IS_SALIENT (Ag,  $X \rightarrow P$ , arg(X,P))
          t5: POP_TOPIC (arg(X,P))

```

MT (Ag, X, P)

```

Shell: Precond:   $\exists X: X, (arg(X,P))$ 
                  BEL (Ag,  $\sim P$ )          ***
                Add:  BEL (Ag, P)
Body:  Goals: t0: PUSH_TOPIC (arg(X,P))
          t1: BEL (Ag,  $\sim X$ )
          t2: IS_SALIENT (Ag,  $\sim X$ , arg(X,P))
          t3: BEL (Ag,  $\sim P \rightarrow X$ )
          t4: IS_SALIENT (Ag,  $\sim P \rightarrow X$ , arg(X,P))
          t5: POP_TOPIC (arg(X,P))

```

DS (Ag, X, P)

```

Shell: Precond:   $\exists X: X, (X \vee P)$ 
                  BEL (Ag,  $\sim P$ )          ***
                Add:  BEL (Ag, P)
Body:  Goals: t0: PUSH_TOPIC (arg(X,P))
          t1: BEL (Ag,  $\sim X$ )
          t2: IS_SALIENT (Ag,  $\sim X$ , arg(X,P))
          t3: BEL (Ag,  $X \vee P$ )
          t4: IS_SALIENT (Ag,  $X \vee P$ , arg(X,P))
          t5: POP_TOPIC (arg(X,P))

```

For each operator, the line indicated by "****" changes in the three epistemic states: each operator shown is the 'hearer disbelieves conclusion' situation. The 'hearer is undecided' situation is characterised by this line reading

BEL (Ag, ?P)
and 'hearer unaware' by
 $\sim$ BEL (Ag, P) $\wedge$
 $\sim$ BEL (Ag, $\sim P$)

Figure 4.2 The nine deductive operators at the AS level

In addition to these deductive operators, the AS level also employs pseudo-deductive operators, by means of which counter-counterargumentation structures can be developed. The importance of including such refutation in an argument has been conclusively demonstrated in social psychology (Hovland *et al.*, 1965), to the extent that it is incorporated into modern theories of rhetoric such as (Simons, 1976). The operators required to effect the generation of such structure are closely related to the notions of conflict explored in (Haggith, 1996), and draw upon the distinction between rebutting and undercutting counterarguments, identified in (Toulmin, 1958) and popular in computational approaches such as (Das *et al.*, 1996). Given the situation portrayed in Figure 4.3, in which the speaker believes p because of a , and also disbelieves b because of d and e , and the hearer believes $\sim p$ supported by b and c , a number of options are available to the speaker. The conventional MP operator discussed above can be employed to support p by a – this is rebuttal. In addition, the hearer's belief in $\sim p$ can be undercut by arguing against one of its supports, namely, given the arguments available to the speaker, b .




Figure 4.3 Sample scenario for rebutting and undercutting

There are thus no new operators for rebutting, since those in Figure 4.2 already fulfil that role. Undercutting, however, requires two new operators, one which characterises a refutation of a premise (UCP), and one which characterises a refutation of the validity of an inference (UCI). The operator definitions are shown below in Figure 4.4. There are several points to note about these definitions. First, that they are fairly loosely constrained, since the speaker is not obliged to believe the falsity of the hearer's premise, but merely be able to persuade the hearer of that falsity (though the speaker is somewhat constrained by rules of terminal goal fulfilment discussed in §4.1.2 below – in particular, the $BEL(H, P)$ goal is prohibited from fulfilment by any means other than substantial support). Secondly, in the UCP operator, it is necessary to make sure that the hearer is aware that the premise supports his conclusion – but clearly, the speaker doesn't want to offer any further support for the inference, hence the absence of a belief goal corresponding to $IS_SALIENT(Ag, X \rightarrow P, \sim X)$. Lastly, a similar issue faces the UCI operator – the τ_1 goal expresses the need to make the premise salient before attacking it. This is the key to counter-counterargumentation: it is necessary to state the counterarguments. The goal is particularly interesting both from a realisation point of view (which can exploit the information that x is being made salient in the context of $\sim x$), and an ordering point of view (whether statement should precede refutation or not – and then whether UCP/I argumentation should precede other, pro, support or not is a major issue of debate in psychology, as explored further in §4.3).

The definitions of UCP and UCI very accurately characterise the discussion in (Whately, 1855, pp38-9) (which is rather too long to here quote verbatim) in which it is claimed that rebutting “is less strictly and properly called Refutation, ... since it might have been employed equally well had the opposite argument never existed”, and that attacks can be made “either in the denial of the Premises, or an Objection against the conclusiveness of the Reasoning”.

UCP (Ag, X, P)

```

Shell: Precond:  BEL (Ag, ~P)
                   $\exists X$ : BEL (Ag, X), BEL (Ag,  $X \rightarrow \sim P$ )
                  P
      Add:        BEL (Ag, P)
                  BEL (Ag, ~X)
Body: Goals: t0: PUSH_TOPIC (arg(~X, P))
              t1: IS_SALIENT (Ag, X, arg(~X, P))
              t2: BEL (Ag, ~X)
              t3: IS_SALIENT (Ag, ~X, arg(~X, P))
              t4: IS_SALIENT (Ag,  $X \rightarrow \sim P$ , arg(~X, P))
              t5: POP_TOPIC (arg(~X, P))

```

UCI (Ag, X, P)

```

Shell: Precond:  BEL (Ag, ~P)
                   $\exists X$ : BEL (Ag, X), BEL (Ag,  $X \rightarrow \sim P$ )
                   $\sim(X \rightarrow \sim P)$ 
      Add:        BEL (Ag, P)
                  BEL (Ag,  $\sim(X \rightarrow \sim P)$ )
Body: Goals: t0: PUSH_TOPIC (arg( $\sim(X \rightarrow \sim P)$ , P))
              t1: IS_SALIENT (Ag, X, arg( $\sim(X \rightarrow \sim P)$ , P))
              t2: BEL (Ag,  $\sim(X \rightarrow \sim P)$ )
              t3: IS_SALIENT (Ag,  $\sim(X \rightarrow \sim P)$ , arg( $\sim(X \rightarrow \sim P)$ , P))
              t4: POP_TOPIC (arg( $\sim(X \rightarrow \sim P)$ , P))

```

Figure 4.4 The operators UCP and UCI

The deductive operators, however, do not offer the full range of argument forms found in natural text. One major omission is the class of inductive operators, including analogy, inductive generalisation, and causal relation. The framework is designed to admit all these operators, but the first and last are beyond the scope of the current work, which concentrates solely upon inductive generalisation.

Inductive generalisation is a particularly interesting topic for study for a number of reasons. The first is the frequency with which various naively statistical and probabilistic arguments are employed in natural language. More importantly, though, are the problems faced in argumentation theoretic analyses of inductive generalisation. Freeman (1991) examines the problems in depth, and, building on Toulmin's (1958) work, and its criticisms (particularly (Cooley, 1959) on the inability of Toulmin schemas to appropriately handle inductive generalisation) comes to a well justified conclusion that inductive generalisation should be treated as a convergent arrangement. His argument rests in a large part on the distinction between the 'ground adequacy' and 'relevance' questions mentioned earlier in §3.1.1: in analysing any argument as dialogical, the analyst can look at any two premises and infer that some imaginary opponent had asked a question after the first premise to elicit the second. If that question was 'Can you give me another reason?' (ground adequacy), the resulting structure is convergent, whereas if that question was 'Why does the premise support the conclusion?' (relevance), the resulting structure is linked. Clearly, an inductive generalisation is based on a number of premises between which an imaginary opponent continually asks the ground adequacy question. The reason, Freeman claims, that inductive generalisation may be intuitively mistaken for a linked structure is that each premise in itself lends only very weak support to the conclusion, and that this generally results in assumption of linkage.

In following Freeman's attractive account of inductive generalisation, it may appear that the required convergent structure can be fully accounted for in the existing framework, by allowing the standard iterative fulfilment of goals of belief discussed above in §3.3. However, Freeman's account, because it is analytic, omits the rather obvious fact that premises in an inductive generalisation have something in common with each other and with the conclusion. That a premise in an inductive generalisation is related to the conclusion *in some respect* cannot be handled simply by iterating through all available supports for an argument, since there is no way to select all and only those premises which support the conclusion in the given respect. Furthermore, it is important that the inductive generalisation itself is seen as a unit, since it is quite inappropriate for subsequent ordering heuristics to be at liberty to intersperse various deductive premises for a conclusion in the midst of the inductive premises (or further, that if there exist two or more inductive generalisations supporting the same conclusion – each employing a different common attribute – it is inappropriate to mix premises from the various arguments). Seeing the whole inductive generalisation as a unit enables appropriate scoping for reordering: the premises within the unit can be reordered wholly within the unit, and the unit itself can be moved around wholesale with respect to the other premises. An inductive generalisation is thus viewed in the current work as a premise. This is illustrated in the diagrammatic argument notation as a *phantom node*, as shown in Figure 4.5.




Figure 4.5 The hierarchical structure of an inductive generalisation

In order to generate this structure within the existing hierarchical planning framework, it becomes necessary to generate the inductive generalisation premise along with all the other premises for a conclusion, and then, after refinement, to generate the individual premises within the inductive argument (which consequently happens concurrently with supports being identified for the other premises which are at the level of the inductive generalisation). In the scenario illustrated in Figure 4.5, for example, the first round of planning identifies that there are three supports for the conclusion p , namely, a Modus Ponens argument from each of b and c , and an inductive generalisation. After an appropriate order is determined for these three, refinement opens up the bodies of the operators, the supports for b and c are identified, and the inductive generalisation is fleshed out to include a_1 through a_n . The process of building an inductive generalisation thus involves two different operators: the IG operator, which identifies that an inductive generalisation is appropriate, and the ISUP operator, which is used to select each inductive premise. To prevent an inductive generalisation from being considered

at every turn, the precondition list on IG states that there must exist at least one premise which can be used inductively – this is a bare minimum since an inductive generalisation employing a single premise is clearly very weak. Strengthening the notion of inductive generalisation is a trivial task of increasing the minimum number of premises which must exist for the application of IG to be licensed.

The complete definitions for IG and ISUP are given below in Figure 4.6. A single new function is required to express the common feature of the premises and conclusion which license the inductive generalisation – this is implemented as a simple function call to `HAS_PROPERTY` which determines whether or not a given property holds for a given proposition. This simplification follows the approach of Cohen (1989) (which has subsequently been adopted in a number of areas) of assuming the existence of an ‘oracle’ which can be polled with a request for information. To permit actual implementation, the knowledge that a proposition has some property is hand coded explicitly in the knowledge base. In both IG and ISUP, the notion of ‘support’ is thus eschewed altogether and simply remains implicit in the fact that propositions are the same in respect *R*. It is not necessary to introduce a new notion of support.

The body of the IG operator includes a single goal of belief (which the various premises are then selected to fulfil), which expresses the aim of bringing the hearer to believe that an inductive argument exists concluding *p* from a body of support which is similar to *p* in respect *r*. As elsewhere, this goal of belief is complemented by a goal expressing that the existence of an inductive generalisation (to *p* from premises which have property *r*) is made salient to the hearer. The motivation for including this saliency goal, and the ways in which it may become realised are discussed in chapter five.

IG (Ag, P)

```
Shell: Precond:  ∃R: HAS_PROPERTY (P, R)
                  ∃X: HAS_PROPERTY (X, R)
Add:            BEL (Ag, P)
Body: Goals: t0: PUSH_TOPIC (IG (R, P))
              t1: BEL (Ag, IG (R, P))
              t2: IS_SALIENT (Ag, IG (R, P), IG (R, P))
              t3: POP_TOPIC (IG (R, P))
```

ISUP (Ag, P, R)

```
Shell: Precond:  ∃X: HAS_PROPERTY (X, R)
Add:            BEL (Ag, IG(R, P))
Body: Goals: t0: PUSH_TOPIC (HAS_PROPERTY(X, R))
              t1: BEL (Ag, X)
              t2: IS_SALIENT (Ag, X, HAS_PROPERTY(X, R))
              t3: BEL (Ag, HAS_PROPERTY(X, R))
              t4: IS_SALIENT (Ag, HAS_PROPERTY(X, R), HAS_PROPERTY(X, R))
              t5: POP_TOPIC (HAS_PROPERTY(X, R))
```

Figure 4.6 The operators IG and ISUP

The operators available to the AS level – deductive and inductive – still do not offer the necessary flexibility. The final class of operators which are included in the suite at the AS level correspond to fallacious reasoning forms identified in texts such as (Johnson, 1992), (Wilson, 1980),

etc. (in which a ‘gang of eighteen’, to use Woods’ phrase (Woods and Walton, 1989) usually occur – these are listed in Appendix A). To intentionally employ fallacies (which are usually deemed ‘bad arguments’) may seem strange: there are two key reasons why the approach is warranted. In the first place, the aim of the *Rhetorica* system is to create persuasive argument – rather than sound argument. More importantly, however, it is also becoming increasingly clear from recent research in informal logic that so-called fallacies are often perfectly acceptable in certain circumstances – Walton has been compiling a series of monographs examining the nature and reasonable use of each of the fallacies in turn (Walton, 1995), (Walton, 1992), (Walton, 1996), etc.. To excise such reasoning from the model would thus leave a serious omission in its battery of techniques.

Again the current work focuses on a subset of the full list. In considering which fallacies to focus on, it is important to recognise that they fall into a number of groups which each have particular characteristics to be addressed for implementation. There are numerous taxonomies of fallacies ((Whately, 1855), (Locke, 1975), (Copi, 1978), (Johnson, 1992), etc.) and the current work does not aim at contributing to the field; rather, a few broad distinctions important for implementation are introduced. First, there are a variety of ‘ad’s – *ad baculum*, *ad verecundiam*, *ad populum*, etc. which share a number of features, most important of which is that they all involve the use of a premise which is in some sense irrelevant to the conclusion (Willard, 1989). Employment of such fallacies thus rests upon identifying a suitably irrelevant premise (the key word is *suitably* – clearly, premise and conclusion must be similar in a number of respects for the fallacy to be plausible to the hearer). Next, there are a number of fallacies which have a lexical, grammatical or phonological basis – (accent, amphiboly, equivocation, etc.). As these rely on low level features, they are outside the remit of the AS level and therefore beyond the scope of this work (in fact, such fallacies seem to constitute an interesting example in which the pipeline hypothesis of §3.1.2 may be violated). In many accounts, there are formal, structural fallacies (affirming the consequent, denying the antecedent, undistributed middle, etc.) which, unlike the others, are never valid – and are therefore also not considered in this work. The remaining fallacies (slippery slope, *petitio principii*, complex question, red herring, straw man, etc.) have a variety of idiosyncratic semantic constraints which are difficult to formalise without trivialisation. The current work formalises two fallacies, *argumentum ad populum* and *ignoratio elenchi*, from the first and last groups mentioned, respectively.

The *argumentum ad populum* has a variety of readings (discussed by (Copi, 1978, pp93-4); one of the most common – of claiming a proposition is true because many people believe it – forms the use of the term in this work. The formalisation of the fallacy into the AP operator is shown below in Figure 4.7. The key step is in relating the two propositions q , and p : the former must state something of the form ‘many people believe p ’. This relation is verified through the `EXPRESSES_POPULAR_BELIEF_IN` function, which follows the use of `HAS_PROPERTY`, above, in that it is viewed as an oracle, is implemented by explicit representation in the KB, and suffers from a similar set of limitations. The other point of note in the operator description is that τ_3 is rarely fulfilled by further support and τ_4 is rarely textually realised – this is due, of course, to the fact that the support of p by q is the weakest point of the fallacy (so there is unlikely to be support for the link, and it is unlikely to be appropriate to highlight the support). The application of the fallacy would thus be

blocked if the hearer was known not to accept the support between from a proposition of the form q and an associated conclusion of the form p .

AP (Ag, P)

```

Shell: Precond:   $\exists Q: Q, \text{EXPRESSES\_POPULAR\_BELIEF\_IN } (Q, P)$ 
               Add:    BEL (Ag, P)
Body: Goals: t0: PUSH\_TOPIC (arg(Q,P))
               t1: BEL (Ag, Q)
               t2: IS\_SALIENT (Ag, Q, arg(Q,P))
               t3: BEL (Ag,  $Q \rightarrow P$ )
               t4: IS\_SALIENT (Ag,  $Q \rightarrow P$ , arg(Q,P))
               t5: POP\_TOPIC (arg(Q,P))

```

Figure 4.7 The AP operator

The *ignoratio elenchi*, or fallacy of irrelevant conclusion, is characterised by an argument which on the face of it appears to come to one conclusion, but in fact is arguing a slightly different point. A good example from (Copi, 1978, p100), is that in offering a case that a defendant is guilty of murder, the prosecution might build an argument for the heinous nature of murder – even if that argument is quite compelling, the prosecution has nevertheless failed to demonstrate the defendant’s guilt. The key to the *ignoratio elenchi*, then, is the closeness between the apparent and actual conclusions of the argument. This closeness is clearly problematic to characterise, and the IE operator shown in Figure 4.8 employs the HAS_PROPERTY function, both for the sake of representational economy and because it represents a bare minimum for appropriate application of IE. If such a filter on IE admits an inappropriately wide range of conclusion-pseudoconclusion pairs, it could be restricted either by enforcing that the pair must be similar in several respects, or by introducing a new function altogether to determine closeness. As with AP, the claim that the pseudoconclusion supports the actual conclusion – expressed in goals t3 and t4 – is rarely supported further or made explicit in the argument, since it is highly questionable.

IE (Ag, P)

```

Shell: Precond:   $\exists R: \text{HAS\_PROPERTY } (P, R)$ 
                $\exists X: X, \text{HAS\_PROPERTY } (X, R)$ 
               Add:    BEL (Ag, P)
Body: Goals: t0: PUSH\_TOPIC (arg(X,P))
               t1: BEL (Ag, X)
               t2: IS\_SALIENT (Ag, X, arg(X,P))
               t3: BEL (Ag,  $X \rightarrow P$ )
               t4: IS\_SALIENT (Ag,  $X \rightarrow P$ , arg(X,P))
               t5: POP\_TOPIC (arg(X,P))

```

Figure 4.8 The IE operator

Although the current version of *Rhetorica* implements only a couple of fallacious operators, and a single inductive operator in addition to the full set of deductive operators found in natural language, the set is sufficient to illustrate the various control structures which are required in the

planning machinery, including how various types of goals are fulfilled, how preference choices are made over operator selection, how various argument structures are manufactured, and how plans of primitive operators are constructed. It is to these issues that the discussion turns next.

4.1.2 AS Level Planning

As explained in (Reed and Long, 1997b), and in §3.3 above, the initial state for the construction of any argument is a list of two goals, `BEL (ag0, p)` and `IS_SALIENT (ag0, p, _)` for some agent *ag₀* and proposition *p*. Thus, following (Grosz *et al.*, 1995), (Grosz and Sidner, 1986), etc., the initial situation is characterised as consisting of a single primary intention (since the goals of belief and saliency are generally conflated in other research), or *discourse purpose* (Grosz *et al.*, 1995). This initial situation offers a good starting point from which to examine the various aspects of the planning machinery.

Goals of saliency are almost always fulfilled by an appropriate instantiation of the `MAKE_SALIENT` operator (since this is the only operator which lists the goal amongst its postconditions)²¹. The operator is a primitive – i.e. as shown in Figure 4.9, it has an empty body – and therefore features in the final plan produced by the AS level (in conjunction with the EG). The operator also has a run-time role (as well as the plan-time pre- and postconditions) which ensures coherent topic flow and context-sensitive realisation, as explained in §4.1.4 and chapter five, respectively. The *Rhetorica* system does not enforce fulfilment of saliency goals by `MAKE_SALIENT`, and it is possible to envisage a scenario in which the same `IS_SALIENT` goal occurs more than once at a single level of abstraction; this would then license the omission of one of the expected applications of `MAKE_SALIENT`. The current operator definitions given in Figures 4.1, 4.3, 4.4 and 4.5 do not admit such a scenario, but this does not present a problem since any spurious `MAKE_SALIENT` primitives can be pruned by subsequent EG processing.

`MAKE_SALIENT (Ag, P, T)`

```
Shell: Precond:
      Add:      IS_SALIENT (Ag, P, T)
Body: Goals:   none
```

Figure 4.9 The `MAKE_SALIENT` operator

Where goals of saliency lead to actual utterance, goals of belief lead to the structure and content of an argument: the broad means by which this is accomplished are discussed in §3.2 and §3.3, but it remains here to examine how the planning process bottoms out (determination of depth), how subargument options are pruned (determination of breadth), and how structure is formed.

²¹ The qualification here is due to the possibility for saliency goals to be discharged as a result of contextual components of the discourse encounter, particularly where deictic reference is involved (in Sperber and Wilson's terms, it may be unnecessary to communicate a proposition if it is already highly manifest in the mutual cognitive environment (Sperber and Wilson, 1986)). Such forms of goal fulfilment are not currently accounted for in *Rhetorica*.

At the higher levels in the structure of an argument, BEL goals are fulfilled by the various operators discussed in the previous section, but as the plan is refined, these goals increasingly express propositions for which no further support is available. The idea has been introduced (§3.1.1) that such terminal nodes can be arrived at in one of three ways. Firstly, beliefs may simply match those in the model of the hearer's beliefs. This is the most preferable means of terminating an argument, for if its premises are already believed by the hearer, counterargumentation is difficult to formulate, since it is only the support links which are open to dispute (though of course, even if the speaker believes that the hearer believes all an argument's premises, it is not certain that the hearer actually believes those same premises, so even an argument which is 'perfect' as far as the speaker is concerned may still fail to be successful). For an argument to be based wholly upon beliefs held by the hearer (or rather, present in the hearer model) is extremely unusual, for although human reasoners do not maintain the deductive closure of their beliefs, they are nevertheless unlikely to hold an entire set of beliefs licensing a conclusion they do not believe (particularly given that such a set would need to include beliefs about the validity of support between propositions and the unpalatable conclusion). The second means, then, by which goals may be terminally fulfilled is checking that such goals are not in contradiction with the (modelled) beliefs of the hearer – i.e. the goal $BEL(H, P)$ may be satisfied in the absence of any further argumentation simply by the fact that the hearer doesn't believe $\neg p$ (i.e. can't be shown to believe $\neg p$, using negation as failure). A slightly stronger version is that the hearer is known to be unaware of p , and a rather weaker alternative that he is known to be aware of, but undecided with respect to p (if the hearer is aware of a proposition, a proponent's task will be that much more difficult, since a greater range of factors influence the hearer's response – a number of these factors are discussed in regard to the order in which pro and con arguments should be presented, in §4.3.3). Each version can be strengthened by consideration of beliefs which imply p or $\neg p$ – if beliefs exist in the hearer model implying p (e.g. a ; $a \rightarrow p$), or not implying $\neg p$ (e.g. $\neg b$; $b \rightarrow \neg p$), the argument is more likely to succeed in respect of the terminal node p .

These, then, determine the depth of an argument. It is equally important to ensure appropriate breadth – i.e. to employ an optimal number of supports for each belief goal. Clearly when such a goal has only a single possible support, the decision is only between whether to support or to terminate. This decision is decided by a rating heuristic which embodies a preference rating: that termination via hearer belief is preferred to any other support; deductive and strong (i.e. many-premised) inductive support is preferred to other forms of termination; and termination by absence of hearer counterbelief is preferred to fallacious or weak inductive support. Often, however, a belief goal will have a number of possible supports or termination options, and in such cases a more complex heuristic is invoked to determine both between support and termination, and between which and how many of the supports to employ. This heuristic extends the previous preference order to state that multiple deductive supports preclude the employment of any inductive or fallacious arguments (since the latter are more likely to detract from the strength of the former); and that no more than one fallacy should be employed in the support of a proposition (since fallacious argument will generally be weaker, and compounding that weakness is unlikely to result in a strong argument). Further heuristics also enforce other, related constraints, such as the fact that a premise in an inductive generalisation may not itself be supported by another

inductive generalisation based on the same property (thus preventing circularity)²². Coherency heuristics (enforcing, for example, limits on the size of a subargument) clearly also impinge upon these decisions, but discussion of their functionality is deferred to §4.2.

Planning for belief goals, in conjunction with auxiliary heuristics, thus generates the content of an argument. Before examining in detail what the result of this planning actually represents, a brief digression examines the relationship between this structure and the standard diagrammatic notation of, for example, Freeman (1991), summarised in Figure 2.1. The distinction between linked and convergent structure is of particular importance to argumentation in general (Reed and Long, 1997a), and although its role has been emphasised in the current work, no explicit representation of the distinction has yet been identified. Such a representation, however, lies implicitly in the structure of operator bodies, and in particular, in the distinction between the belief and saliency goals relating to the premise and those relating to the support between premise and conclusion. Recall from the previous section that convergent structure results from the implicit question ‘Can you give me another reason?’ and linked from ‘Why does the premise support the conclusion?’; in Modus Ponens (and similarly for the other operators), it is the BEL (H, P) goal which ‘poses’ the first question to the planner, and the complementary BEL $(H, X \rightarrow P)$ goal which poses the second. Thus, quite simply, the first goal gives rise (through multiple instantiation) to convergent structure, and the latter (through one-or-more instantiation) to linked structure. Indicating this simple relationship between the structure of the AS plan and conventional argumentation theoretic analyses leaves open the question of how that structure is to be interpreted in terms of linguistic theory. Clearly, the various saliency goals represent in some sense the surface content of the argument, but as identified in (Grosz and Sidner, 1986) and widely adopted in the NLG community, linguistic structure also involves intentional and attentional facets. The next section therefore addresses the relationship of how the AS plan is related to such aspects of linguistic structure.

4.1.3 *Intention and Attention*

The need for explicit planning with attentional information has not generally been recognised in NLG research. There seem to be several contributing factors to this phenomenon. In the first place, the few examples of topic manipulation functioning generatively in an NLG system ((Huang, 1994) and (Hovy and McCoy, 1989)) use it to support the planning architecture as a means of improving coherency (particularly on a very local basis, compared to the larger scale work of the planning machinery). No attempt is made to integrate attentional operators into the planning framework (even Hovy and McCoy (1989) attempt only to glue together two monolithic theories). Secondly, there appears to be an assumption that topic control itself is unsuitable for a planning operationalisation – partly because of the assumed localisation of its effect (anaphora, etc.), and partly because the extra machinery required to handle focus explicitly is essentially very simple ((Blandford, 1993) is a typical example of this approach).

This shifting of topic control from within the remit of the planning subsystem to some distinct

²² Without this constraint, if a and b were used to support c in an inductive generalisation based on their common property p , it would be possible for a then to be supported by b and c in another inductive generalisation. Of course, an inductive argument based upon a different property could still be used to further support any of a , b or c .

module is surprising given the work of Grosz and Sidner (1986) and its popularity. They maintain that “any discourse is a composite of three distinct but interacting components: the structure of the actual sequence of utterances in the discourse; a structure of intentions; an attentional state”, p176. To account for only two of the three in a unified manner is surprising (as a case in point, Moser and Moore’s (1996) recent attempt to unify the work of Grosz and Sidner with the more prevalent RST paradigm makes no mention of attentional structure whatsoever).

The relation between attentional and intentional structure plays a particularly prominent role in the Grosz and Sidner theory, and with the recognition of the crucial role played by intentionality in NLG (Young and Moore, 1994), (Moser and Moore, 1996), etc., it is this relation which forms a key area of the current work. In (Grosz and Sidner, 1986), discourse segments each have a single purpose; the discourse segment purposes (DSPs) are then arranged in a hierarchy (such that one DSP contributes to or is dominated by another). In addition, the theory permits expression of linearisation constraints – that one DSP must precede (i.e. be satisfied earlier than) another. The notions of dominance and satisfaction-precedence constitute the intentional component of (Grosz and Sidner, 1986). Their characterisation of attentional state is based upon the notion of focus spaces, which comprise the salient information for a given discourse segment (including the DSP). Focus spaces are stacked such that entities in the topmost focus space are more salient than those of lower focus spaces.

The *Rhetorica* system implements a model of discourse generation which integrates focus control, intentional structure and informational content into a unified planning framework. The relationship in the model between the intentional and attentional structures created by the planning mechanism is similar to that proposed by Grosz and Sidner, in a number of respects. Firstly, although *Rhetorica* has no explicit notion of discourse segment, it remains implicit in a more generic form, as the body of an operator. The DSP – i.e. the purpose of an operator – is the belief goal on that operator’s postcondition list (in fact, this is a slight finesse since as Young and Moore (1994) explain, it would prohibit any distinction between intended and side-effects in cases where an operator has more than one postcondition; the intention is thus the goal posted rather than that fulfilled, and this information could be recorded additionally – but since all implemented operators are designed in such a way that all postconditions must have been intended, the current definition and implementation suffices). As explained in §3.2, this definition admits two important features of intentionality – first that one intention can give rise to another (compared, in (Moore and Pollack, 1992) to communicative goals being mapped immediately onto linguistic goals), and second that goals of saliency are also intentional (it is not until a later stage that they are mapped on to linguistic goals).

Furthermore, the definition also admits one further feature: goals of topic shift are intentional. Thus it is claimed that speakers usually intend to make the hearer aware of the structure of the argument (though, as for the *IS_SALIENT* goal, mentioned above, *PUSH_TOPIC* and *POP_TOPIC* are not enforced by the *Rhetorica* framework to be fulfilled by corresponding *PUSH_TOPIC* and *POP_TOPIC* operators, so it is possible to create a situation where attentional structure does not form part of the speaker’s intentional goals – but such a situation is beyond the scope of the current operator definitions). This is consistent both with Grosz and Sidner’s account, and with intuition – if content intentions are to be fulfilled, the hearer must understand how the content is interrelated. The

PUSH_TOPIC and POP_TOPIC goals are fulfilled by operators of the same name, whose definitions – like that of MAKE_SALIENT – are trivial, as can be seen from Figure 4.10. Notice that, whereas MAKE_SALIENT fulfils the IS_SALIENT goal, both PUSH_TOPIC and POP_TOPIC are fulfilled by operators of the same name.

```

PUSH_TOPIC (T)
  Shell: Precond:
        Add:      PUSH_TOPIC (T)
  Body:  Goals:   none

POP_TOPIC (T)
  Shell: Precond:
        Add:      POP_TOPIC (T)
  Body:  Goals:   none

```

Figure 4.10 The PUSH_TOPIC and POP_TOPIC operators

The intentional goals PUSH_TOPIC and POP_TOPIC thus have a twofold role: one in expressing the aim of conveying attentional structure to the hearer, and one in actually manipulating the topic stack. The distinction between these two roles and its relation to the concept of the topic stack is explored in the next section.

4.1.4 The Topic Stack

Following almost all other work in the area ((Grosz and Sidner, 1986), (Hovy, 1990), (McKeown, 1985), etc.) *Rhetorica* employs a notion of a stack on which to record the current focus of attention. The only exception is the work of (McCoy and Cheng, 1991) (extended in (Hovy and McCoy, 1989)) which employs a richer notion of *focus trees*. These structures are built and traversed during a discourse, with various rules constraining valid moves at any given turn. The strength of the model lies in the increased flexibility afforded by the tree structure: although depth first traversal (i.e. that which could be mimicked by a stack) is expected, alternative moves are available to the generator. There are two key reasons why the flexibility offered by the focus tree structures is not appropriate for the *Rhetorica* system. In the first place, as discussed above in chapter one, textual argument is generally more structured than other, freer forms of naturally occurring text, and this structure can be exploited to simplify the generation process. More importantly, however, it seems that the flexibility offered by focus tree structures can in fact be matched by the combination of the AbNLP planning regime, the operationalisation of focus control, and the use of partial ordering discussed below.

Rhetorica differs from much other work in discourse planning, by virtue of the recognition of the important notion of *disjunctive constraints*. Following the well established principle of least commitment in planning, it is desirable to avoid specifying an absolute order amongst operators when such an order is not demanded by the planning constraints (Sacerdoti, 1975). In *Rhetorica*, therefore, the order between premises and conclusion is not determined by the planning process itself – rather, orderings may (or may not) be imposed by the various heuristics discussed in §4.2 and §4.3. Thus

during the planning process, it is necessary to record the few constraints on ordering using a conventional partial order notation. For example, on every occasion that an operator's body is opened up and the goals therein posted, the partial order is updated to reflect the need for the `PUSH_TOPIC` goal to precede, and `POP_TOPIC` to succeed all the other goals in that body. There is one crucial result which follows from such a situation, namely that, by itself, the machinery does not enforce coherency. Consider, for example, the sample argument in Figure 3.7. There is no constraint which prevents the interleaving of the subarguments listed in Figure 3.9, to result in multiple consecutive `PUSH_TOPIC` goals, `PUSH_TOPIC` followed immediately by `POP_TOPIC`, or saliency goals from one argument being interspersed in another – all of which result in incoherent arguments.

To enforce the appropriate constraints to ensure the basic integrity of argument units would involve the use of disjunctive constraints, which are best explained by a slightly larger example taken from (Reed and Long, 1997d), in which a conclusion is supported by two premises, each of which is supported by a further two premises, Figure 4.11. Given the rules of ordering scope presented in §3.2, a number of orders need to be resolved: (i) the order between *d*, *e*, and *b*; (ii) the order between *f*, *g*, and *c*; (iii) the order between the *b* subargument unit, *c* subargument unit and their conclusion *a*. Notice that this means amongst other things that the unordered triple (*d*, *e*, *b*) is itself a member of a larger unordered triple. Thus the orderings between, say, *d* and *f* can only be expressed as a (complex) disjunction of possible scenarios. Such disjunctive constraints cannot be expressed in conventional partial order notation.




Figure 4.11 The limits of partial order

One natural approach to solving this problem is to propose an extension to the conventional representation of partial order. Aside from implementation issues and the need to prove anew soundness and completeness, there is a more fundamental problem in adopting the approach. After summary inspection, a general worst case analysis of the number of orderings required to fully express an unconstrained argument comprising *m* levels of support, at which each proposition is supported by *n* premises (so Figure 4.11 shows an argument with *n* = 2 and *m* = 2) is

$$(n+1)! \sum_{i=0}^{m-1} n^i = (n+1)! \frac{1-n^m}{1-n}$$

This has horrendous complexity $O(n!)^{O(n^m)}$. By way of an example, an argument of *n* = 5, *m* = 2 (the

sort of size which might constitute a small newspaper article), with just twenty-five terminal propositions has a number of orderings in the magnitude 10^{17} . Any means of representing and manipulating this sort of intractable information is therefore an unattractive route to follow.

An alternative approach is to follow further the spirit of least commitment, and defer specification of the partial order itself – or more precisely, the partial order holding between propositions in separate subarguments and levels of abstraction. Thus the constraints enforcing that PUSH_TOPIC and POP_TOPIC goals are endpoints in a subargument are maintained in the plan specification, and these orderings along with any others that may be heuristically introduced are propagated down to lower layers of abstraction during refinement operations. All other constraints are left unspecified, thus avoiding intractable data. One particularly important class of orderings which are thus omitted are those which enforce the integrity of subargument units – i.e. those which group the components of one subargument together. In the example of Figure 4.11, the specified orderings between the various push and pop actions admit an ordering of terminal nodes such as *e, f, d, g* which results in an incoherent argument (Reed and Long, 1997b).

How then does *Rhetorica* avoid the problems of disjunctive constraints whilst still ensuring basic coherency specified by subargument integrity? The solution arises from consideration of a seemingly unrelated generic planning problem which is particularly noticeable in discourse planning: the distinction between a plan and its execution. Arguing analogously to McDermott (1982), the discourse planning process culminates in a piece of text, which, as far as the planner is concerned and barring any nondeterministic events in the world, *has already* conveyed its meaning to the hearer (or persuaded the hearer in the case of *Rhetorica*). In designing the primitive operators in *Rhetorica* (i.e. just MAKE_SALIENT, PUSH_TOPIC and POP_TOPIC) a distinction was drawn between those pre- and postconditions which would function in the planning process, and those which would be dependent upon the state of the world at runtime. Thus in addition to the MAKE_SALIENT operator presented in Figure 4.9, there is a runtime equivalent, shown below in Figure 4.12 (as a primitive, it has no body)²³. The important feature is the precondition, specifying that the context parameter on the saliency goal which generated the application of the MAKE_SALIENT is the same as the current topic, sitting at the top of the focus stack.

```
MAKE_SALIENT (Ag, P, T)
```

```
Shell: Precond: TOPIC_STACK ([T | R])
Add:
```

Figure 4.12 The runtime specification for MAKE_SALIENT

In conjunction with the actions of the PUSH_TOPIC and POP_TOPIC operators, whose runtime definitions (Figure 4.13) are unremarkable, the precondition of the MAKE_SALIENT operator is sufficient to ensure subargument integrity. Furthermore, by preserving subargument integrity at every

²³ The [A | B] notation is a result of implementing in Prolog, and is read as a list comprising A as its head and B as its tail (which may be an empty list).

level of abstraction, but introducing no new constraints, there is an effective partial order over the entire plan which is precisely equivalent to that required.

```

PUSH_TOPIC (T)
  Shell: Precond:  TOPIC_STACK (S)
            Add:    TOPIC_STACK ([T | S])
            Del:    TOPIC_STACK (S)

```

```

POP_TOPIC (T)
  Shell: Precond:  TOPIC_STACK ([T | S])
            Add:    TOPIC_STACK (S)
            Del:    TOPIC_STACK ([T | S])

```

Figure 4.13 The runtime specification for PUSH_TOPIC and POP_TOPIC

Clearly, the approach suffers a slight problem inasmuch as the final plan of primitives generated by the combined efforts of the AS and EG levels is still an abstract one, requiring further refinement through the RST and syntactic levels of the framework. As such, it is hardly applicable to introduce a notion of ‘runtime’ for such a plan. The reason for describing and implementing the post-planning effects in the way illustrated in Figures 4.9 and 4.10 is simply that in the current work, *Rhetorica* is not concerned with sub-EG planning, as justified in §3.2. Rather, its remit is to produce a fully specified, fully ordered plan of primitives. Thus when the partially specified plan of primitives is completed, it is ‘executed’ to enforce the constraints imposed by MAKE_SALIENT and the topic stack, and to effect an arbitrary ordering on any remaining freedoms. In a fully implemented system, the same approach could easily be employed, fixing abstract structure absolutely, before passing control to the lower levels. Instead of this, however, a more sophisticated approach is envisaged, whereby operators which map from the various intentional goals (i.e. primitives at the AS/EG) to RST relations employ the topic stack dynamically. Thus rather than resolving order arbitrarily, it is the RST operators (or, conceivably, even more refined operators) which determine ordering choices explicitly, progressively constraining the plan further and further through reference to a topic stack which is created on the fly. Further exploration of these issues is beyond the scope of the current work.

It is important to note that whichever approach is taken to resolving the partial order, such resolution occurs after the plan is produced. This in turn means that the notion of a topic stack has two facets: one at plan-time control and one at runtime. This concurs neatly with the observation by (McCoy and Cheng, 1991) that representations of focus in the speaker and hearer can be used to check for coherency: a maximally coherent argument is one in which the two focus structures are identical. Such a definition relies on the runtime notion of a topic stack, and contrasts (as McCoy and Cheng point out, p120) with the Grosz and Sidner account which views the focus structure as lying in the discourse itself – a view more consistent with the plan-time notion of the topic stack (and a view more commonly adopted by discourse planning theorists).

4.2 Ordering for Coherency

The key source from which to determine appropriate heuristics for ordering components within an argument is intuition – though in defence of a rather weak source, it is for the most part the intuition (and experience) of the greatest authorities: Aristotle, Quintilian, Whately, etc. Thus it is rhetoric (and its modern-day guises such as communication theory) which furnish the necessary suggestions. It is surprising that it is so difficult to find appropriate heuristics elsewhere – particularly in linguistics, where it might be expected that the appropriate rules on discourse structure would have been formulated. However, it seems that linguistics (including pragmatics and discourse analysis) is preoccupied with a smaller-scale notion of structure. Even work which purports to mix linguistics and rhetoric (Gray, 1977), or in the most abstract linguistic theories (of which RST is a good example), the structures analysed are simply not abstract enough. Consider, for example, the canonical orderings identified in RST (e.g. that the nucleus of an elaboration generally precedes the satellite): as discussed in §3.1.4, argument structure is built around more abstract structures (e.g. that an MP can employ evidence, antithesis or justification), so any orderings which generalise no further than individual relations is of little use. Other linguistic and argumentation theoretic characterisations of the notion of coherency in argumentation (Cohen, 1987), (Cohen, 1984), (Walton and Krabbe, 1995), (Carlson, 1983), etc. also offer little in the way of direction for the current considerations since the constraints on coherency there proposed have already been accounted for in the planning process itself (i.e. through the concept of subargument integrity), primarily through the constraints imposed by the topic stack explained above.

The next two sections, then, cover the heuristics which have been extracted from various rhetoric sources, and those which have been identified after examination of a small corpus and seem to form intuitively correct orderings. These heuristics are available to the AS level as post-processing operations in a blackboard phase active on the completion of the planning and reordering identified in Figure 3.10. That is, at the culmination of AS planning, and again after reordering, any heuristics which are triggered effect the appropriate change to the plan, introducing new ordering constraints. The fact that the heuristics are available at two distinct phases permits their design to include either goals (available after the reordering phase) or operators (available after the planning phase), though usually a given operator will be active at only one of the phases. Clearly, some of these heuristics are of greater value than others (some are highly specific and are to be preferred to others which are more default in nature). The preference ordering over the heuristics could be implemented explicitly, but the *Rhetorica* system exploits the inherent database ordering in Prolog such that earlier heuristics are applied first (and thus if the conditions for application of later heuristics no longer hold, those later heuristics are not applied).

4.2.1 Heuristics from Rhetoric

Although rhetoric texts are predominantly concerned with issues regarding the success of an argument in convincing its audience – i.e. with persuasive heuristics – most also discuss a number of features which might be more appropriately classed as constraining componential ordering to improve coherency. The first such heuristic is Blair's suggestion that arguments of a similar nature should be grouped together (Blair, 1838, p430). He claims that "all arguments whatever are directed to prove one

or other of these three things; that something is true; that it is morally right or fit; or that it is profitable and good.”, and that arguments in each class should be presented together. To determine which class a given proposition falls into would require significant additional information in the knowledge base, and such information would inevitably be recorded in an ad hoc manner. Furthermore, it is not entirely clear that the three classes represent the most appropriate taxonomy (Reed *et al.*, 1996), (Reed *et al.*, 1997a). Instead, *Rhetorica* follows the spirit of Blair’s suggestion, making use of the HAS_PROPERTY function introduced above. Given two propositions *a* and *b*, the HCR1 heuristic ensures that no other statement comes between them²⁴. Of course, on reflection this again poses a problem for the conventional partial order notation (it would be extremely computationally expensive to introduce and then manage constraints between every other component of the abstract plan and the propositions in question). The same solution is adopted here as was proposed in the previous section. By introducing a new subargument expressing that *a* and *b* are related and may not be separated, the same MAKE_SALIENT topic machinery will ensure their cohesion. In this way, the action of HCR1 is similar to the introduction of an IG phantom node (except, of course, that the propositions *a* and *b* are still supporting their conclusion deductively). The definition of HCR1 is given below in Figure 4.14. It is based entirely on primitives, since it is only applicable at the post-planning phase (refinement will not offer any new premises to be grouped, as the premises within an operator body will necessarily already be grouped together).

HCR1

```

Trigger:      tm: MAKE_SALIENT (Ag, P, C)
              tn: MAKE_SALIENT (Ag, Q, C)
              HAS_PROPERTY (P, R)
              HAS_PROPERTY (Q, R)
Update:      tx: PUSH_TOPIC (arg(R,C))
              ty: POP_TOPIC (arg(R,C))
              (tx < tm < ty)
              (tx < tn < ty)

```

Figure 4.14 The definition of HCR1: grouping topics

As mentioned in §3.2 above, *Rhetorica* maintains a representation of a variety of facets of the hearer. Following the characterisation offered in PAULINE (Hovy, 1990) and HealthDoc (DiMarco *et al.*, 1995), these facets are represented as parameters with a small range of discrete values. One of these parameters is an indication of the hearer’s general competence, which Whately indicates is crucial for determining an appropriate level of complexity:

“The less enlightened the hearers, the harder, of course, it is to make them comprehend a long and complex train of Reasoning; so that sometimes, the Arguments, in themselves the most cogent, cannot be employed at all with effect.”
(Whately, 1855, §7 p58).

²⁴ A generic naming convention is used, so that HCR1 is the first Heuristic improving Coherency based on a Rhetoric source.

Blair too identifies the problem in a more general sense:

“... extending arguments too far, and multiplying them too much ... serves rather to render a cause suspected, than to give it weight. An unnecessary multiplicity of arguments both burdens the memory and detracts from the weight of that conviction which a few well-chosen arguments carry.” (Blair, 1838, p432).

Following this advice, HCR2 is designed to curb excess in the face of low hearer competence, by pruning weaker arguments (the notion of argument strength is discussed in §4.3.1, below). The specification of HCR2 is given in Figure 4.15, in which the functions `count` and `remove_arguments` are assumed to be defined, and the threshold value is fixed and stored explicitly. Though the current implementation of *Rhetorica* uses a value derived from a small survey of arguments in the corpus, the potential for developing a richer, more dynamic means of determining this value (and others like it active in other heuristics) is discussed in chapter seven. The single parameter to both `count` and `remove_arguments` is an abstraction of the various argument types (MP, IG, etc.), and thus refers to any of the planned supports for a given proposition. HCR2 is only applicable at the post-planning phase, since the goals produced by refinement are not subject to such pruning (they are assumed to be necessary to the application of a given argument form).

HCR2

Trigger:	<code>parameter (general_competence, Ag, low)</code> <code>count (argument(Ag, X, P)) > fixed_value</code>
Update:	<code>remove_argument (argument(Ag, X, P))</code>

Figure 4.15 The definition of HCR2: reduce breadth

4.2.2 Heuristics from Intuition

Even after consideration of the intuitions from the authorities in rhetoric, there remain a number of heuristics which are intuitively appropriate, but not elsewhere documented, perhaps because of their lack of import, or because they are so straightforward. First, is the heuristic alluded to in §3.3, by which an MP subargument distanced from the conclusion it supports (for example, by a large intervening subargument) may have the order of its components switched; this is HCI1. The evaluation of the intervening distance is performed by the function `distance_between` the result of which is again thresholded at an arbitrary, fixed value.

HCI1

```

Trigger:      ta: MAKE_SALIENT (Ag, P, C)
              tb: PUSH_TOPIC (X → P)
              tc: BEL (H, X)
              td: IS_SALIENT (H, X, X → P)
              te: BEL (H, X → P)
              tf: IS_SALIENT (H, X → P, X → P)
              tb < tc
              tb < td
              tb < te
              tb < tf
              distance_between (ta, tb) > fixed_value
Update:       tc < te
              tc < tf
              td < te
              td < tf

```

Figure 4.16 The definition of HCI1: reminding of implication

4.2.3 Hard and Soft Constraints

The variety of heuristics presented in the three preceding sections provide a good basis for discussing in more depth than was presented in §3.3 the introduction and ramifications of the constraint classes ‘hard’ and ‘soft’. As with the heuristics, the abstract operators used in planning at the AS level are implicitly ordered in the Prolog database. The list of these operators runs in order of preference from the deductive operators whose preconditions describe that the hearer is unaware of the conclusion, all the way through to fallacy operators. This means that the first applicable operator selected is always the ‘optimal’ operator. Consequently, the AS planning mechanism can attach a hard constraint, i.e. expressing necessity, to the first operator selected for a given goal, and allow any other applicable operators to be introduced with soft constraints, i.e. made optional (since additional convergent supports are not essential). Thus on completion of a round of planning at the AS and EG levels, there will be a number of operators (abstract and primitive), some of which are marked as mandatory (hard), others as optional (soft). At this stage, the various heuristics are applicable, and may introduce new ordering constraints. Some, such as HCR2, will actually remove components of the plan – clearly, such deletion will only be permitted for optional components of the plan. Heuristics can introduce both hard and soft ordering constraints; the latter can be overturned by later processing, the former cannot. At refinement, all ordering constraints are propagated to the members of the operator bodies, and then once again the heuristics can be applied, again overturning only soft constraints, and introducing new constraints of either kind (the algorithm employed to propagate orderings to body goals thus represents a simplification of the rich mechanism provided in the original AbNLP planner (Fox and Long, 1995)).

4.3 Ordering for Persuasive Effect

Like the heuristics designed to improve coherency, those concentrating on persuasive effect also draw heavily upon rhetoric, but in addition also draw upon psychology, and in particular the social psychology of attitude change. Before examining the various heuristics in detail however, it is necessary to elaborate on a distinction mentioned above between two means of appraising the strength of an argument – for many of the heuristics make significant use of the notion.

4.3.1 Persuasive and Inferential Strength

The strength of an argument plays an important role in determining where it should be placed in the discourse as a whole, yet it is far from clear exactly what the notion connotes. As mentioned in (Reed and Long, 1997b), the *Rhetorica* system follows (Freeman, 1991) in distinguishing two facets of argument strength – inferential and persuasive:

“Inferential strength is a completely normative issue. How well does this inference satisfy the canons of deductive or inductive logic? Persuasive force is studied empirically. We can measure it by developing a tool to determine just how more strongly (or less strongly) an audience accepts a claim after hearing a particular argument for that claim. Inferential strength and persuasive force are two separate and independent features of argument.” p243.

Perelman and Olbrechts-Tyteca (1969) makes a similar claim, “The strength of arguments varies therefore with the audience and the object of the argumentation”, p461, though he also claims (p499) that referring to the ‘strength of an argument’ presupposes that such a feature is independent of, for example, the position in which the argument occurs (and indeed Perelman claims that such independence does not exist in this respect). In the current work, it is assumed that arguments do have some intrinsic strength, but that this can be improved upon by appropriate positioning.

It is generally the persuasive force of an argument which is referred to by the various authors mentioned in the next three sections when proposing presentation techniques based upon argument strength. Unfortunately, it is also persuasive strength which is the more difficult to model. There are two alternatives available for characterising persuasive strength. The first is to extend the representation of beliefs to enable support relationships to be adumbrated with some notion of how plausible the support link might be. There are numerous problems with this approach, central amongst which is its independence from hearer beliefs – the notion of persuasion (for both Freeman, and to an even greater extent, Perelman) is heavily dependent upon the intended audience for assessment. The alternative, adopted in *Rhetorica*, is to calculate persuasive strength on the basis of hearer beliefs and the likelihood of the particular argument to effect the desired attitude change. Clearly, this definition still admits a wide range of approaches: in *Rhetorica*, persuasive strength is assessed on the basis of (i) the epistemic state of the hearer in regard to the conclusion of the given argument, and (ii) the class of argument used. Under the first point, using an argument in a scenario in which the hearer is unaware of the conclusion is persuasively stronger than employing an argument where the hearer is undecided, which is in turn is stronger than an argument used against an established hearer counter-belief. Under the second point, a ‘pro’ deductive argument (MP, MT, DS) is persuasively stronger (i.e., in this context, is more likely to have the intended effect on the hearer) than refutation (UCP, UCI), which is stronger than an inductive argument, which is stronger than a fallacious argument. This definition is still clearly a simplification of a very complex problem, but the reference to hearer beliefs and the ease of integration with the existing framework make it both attractive from a pragmatic point of view, and also a reasonable approximation for employment by the variety of heuristics discussed below.

4.3.2 *Heuristics from Rhetoric*

One of the most ubiquitous maxims, occurring in (Quintilian, 1960, VII, i. 10 & 16-17), (Blair, 1838, p430); (Whately, 1855, p43); (Perelman and Olbrechts-Tyteca, 1969, p499) and in numerous modern text books on rhetoric and persuasive communication (e.g. (Simons, 1976)), is due to Cicero: “ut augeatur semper, et increseat oratio” – that arguments should increase in strength in climax order. There are, however, a number of caveats to this general rule. First Blair:

“But this rule is not always to be followed. For if he distrusts his cause, and has but one material argument on which to lay the stress, putting less confidence in the rest, in this case, it is often proper for him to place this material argument in the front; to preoccupy the hearers early” p431.

Next, Whately:

“It will be advisable, however, (and by this means you means you may secure the last advantage) when the strongest Arguments naturally occupy the foremost place, to *recapitulate in a reverse order*; which will destroy the appearance of anti-climax”, p43.

Although there are a number of further contextual and psychological factors (discussed below and in (Reed *et al.*, 1996), (Reed, 1997)) which may mitigate against the general rule, it is this conventional rule proposed by rhetoricians which is captured by HPR1 (Figure 4.17), into which exceptions have not yet been incorporated. The function `all_arguments` returns all the various argument(*x*) and `order_by_persuasion` simply evaluates each of the *x* to determine its persuasive force (as defined above), and then places those *x* in ascending order.

HPR1

Trigger: `all_arguments (Ag, X, P)`
 Update: `order_by_persuasion (all_arguments (Ag, X, P))`

Figure 4.17 The definition of HPR1: climax ordering

Related to the ordering suggested by strength of individual subarguments is that concerning the conclusion – again, the position of the conclusion has attracted the attention of most authors on the subject from Cicero (cited in (Perelman and Olbrechts-Tyteca, p498)) onwards. Blair terms the two primary alternatives *synthetic*,

“when the point to be proved is fairly laid down, and one argument after another is made to bear upon it, till the hearers be fully convinced”,

as opposed to *analytic*,

“when the orator conceals his intention concerning the point he is to prove, till he has gradually brought his hearers to the designed conclusion. They are led on, step by

step, from one known truth to another, till the conclusion be stolen upon them”
(Blair, 1838, p429).

The former, ‘conclusion-first’ option is by far the more common, with the analytic construction being reserved for situations in which the hearer is particularly sceptical. A hybrid form, where the conclusion is placed amongst the premises is rarer still, being the result, in dialogue, of an analytic argument which was not accepted by the hearer, and which is then further bolstered. (It should be noted that there is not a unanimous consensus – (Brooks and Warren, 1972) claim that the conclusion-last construction is more common, but their claim is unsupported, so the current work defers to the earlier authors). The work of experimental psychologist Darnell also offers limited support for the conclusion-first arrangement (Darnell, 1963) (though his evidence suggests only that bare comprehension is improved through use of this ordering: he did not look at subsequent persuasive effect). HPR2 characterises the usual situation of conclusion-first, admitting two exceptions: (1) that the hearer is highly sceptical, (2) that the hearer disbelieves the conclusion. HPR2 is only active at the post-refinement stage, because conclusions and supports are not available at any other time.

HPR2

Trigger: ~(parameter(scepticism(Ag, high))
 ~(BEL (Ag, ~P))
 tm: MAKE_SALIENT (Ag, P, C)
 tn: argument(Ag, X, P)
Update: tm < tn

Figure 4.18 The definition of HPR2: conclusion first

If the conclusion – or any other proposition – is obvious (i.e. the hearer already believes it), it should come at the beginning of the argument: “A Proposition that is *well known* should in general be stated at once”, (Whately, 1855, p36). This is simply characterised by HPR3, which, through multiple instantiation of *y* can introduce the necessary ordering between the known belief and all others:

HPR3

Trigger: BEL (Ag, X)
 tm: MAKE_SALIENT (Ag, X, P)
 tn: MAKE_SALIENT (Ag, Y, P)
Update: tm < tn

Figure 4.19 The definition of HPR3: well known first

Lastly, with the conclusion fixed, it remains to determine an appropriate place for counter-counterarguments. Whately’s suggestion, that

“Refutations of Objections should generally be placed in the midst of the other
Arguments, but nearer the beginning than the end.” (Whately, 1855, p38),

is not terribly helpful; instead, the current work follows the bolder stance of Quintilian and Perelman

that such counter-counterargumentation should simply come first (in most situations). This is easily incorporated through HPR4:

HPR4

Trigger: tm: UCP (Ag, X, C)
 tn: argument (Ag, Y, C)
 Update: tm < tn

Figure 4.20 The definition of HPR4: refutations first

HPR4 requires two variants – the one shown in Figure 4.20, and a corresponding version involving UCI rather than UCP (since the initial position of refutation is demanded regardless of whether it is premise or inference which is being refuted). Again, multiple application of the operator ensures that all refutation arguments are placed before all others.

4.3.3 Heuristics from Psychology

There is a large body of social psychology research dealing with the effect that message structure (including ordering alternatives) has on persuasion, due mostly to a preoccupation with the primacy effect identified by Lund (1925). One work is of particular importance: Hovland *et al.*'s collection of papers entitled *The Order of Presentation in Persuasion*, (Hovland, 1957b). The work of Janis and Feierabend (Janis and Feierabend, 1957) in that volume builds on earlier work (Hovland, 1965) (originally published in 1949) that demonstrated the utility of including counter-counterargumentation, and supports the primacy effect, inasmuch as subjects who were presented with 'con' arguments first evinced less attitude change than those who were presented with 'pro' arguments first (this is the reverse of the suggestion from rhetoric mentioned above and characterised in HPR4). Several other researchers, however, offer alternative results; chief amongst these are Hass and Linder (1972), who make several important observations. First is that presenting and countering counterarguments when the available refutation is weak has – unsurprisingly – a detrimental effect on the persuasiveness of the argument as a whole. This in itself is worthy of being addressed, and is formalised into HPP1, below. As in the aforementioned HCR2, the threshold value is assumed to be fixed.

HPP1

Trigger: UCP (Ag, X, P)
 persuasive_strength (UCP (Ag, X, P)) < *fixed_value*
 Update: remove_argument (Ag, X, P)

Figure 4.21 The definition of HPP1: no weak refutations

Note that the threshold may in fact be quite low, since subsequent EG processing may decline to realise the (weak) support for a counterargument: this mode of acknowledging but not explicitly refuting counterarguments can add surprising support to the weight of an argument (Hovland *et al.*, 1957). Hass and Linder also report on a series of careful experiments which, roughly, demonstrate that under most

scenarios, the refutation-first order is the one most likely to effect attitude change, thus supporting the application of HPR4.

The placement of counter-counter argumentation is not the only guide discussed in the previous section which has also been the topic for study in social psychology. The position of the conclusion has also been addressed – also with mixed results. Apart from cross-cultural studies indicating that persuasive effect of conclusion placement may be affected by cultural background, (Triandis, 1971, p188), there are few disadvantages of the conclusion-first structure (namely, that it may alienate those who disagree with the conclusion, or prejudice a hearer that recognises he is to be persuaded). In contrast, there are several advantages, including increasing comprehensibility, focusing hearer attention, etc. (McGuire, 1969, p212). The social psychology evidence thus also supports HPR2, including its restrictions (specifically, that it is inapplicable for situations of hearer scepticism).

Finally, McGuire (McGuire, 1969) also discusses the efficacy of climax ordering, and although there is clearly no consensus, there exists extensive support for the approach implemented by HPR1.

4.4 Pulling it Together

Before examining an extended example which integrates the planning and heuristic considerations put forward in the preceding sections of this chapter, a brief digression is required to introduce the preference ordering between the heuristics discussed above. In most cases, there is unlikely to be conflict since the operators (a) act at different times in the processing cycle (i.e. either after planning or refinement) and (b) have, for the most part, quite specific triggers. However, such conflict is possible particularly involving the more general heuristics such as HPR2 (conclusion first), so the preference order remains reasonably important. Since the implementation requires complete specification (because of the implicit ordering in the Prolog database), the order given below is complete rather than partial, even where that completeness is unnecessary (e.g. between heuristics active at different phases). In order of preference, the full list of heuristics is as shown in below:

HCR2: reduce breadth
 HCI1: reminding of implication
 HCR1: grouping topics
 HPR3: well known first
 HPP1: no weak refutations
 HPR4: refutations first
 HPR2: conclusion first
 HPR1: climax ordering

Figure 4.22 The preference order amongst heuristics from most to least preferred

As mentioned in §3.3 above, this ordering ensures that coherency heuristics take priority over those improving persuasive effect. To eradicate redundant processing, it also ensures that the heuristics which remove content (HCR2 and HPP1) precede those heuristics which reorder components which are liable for deletion.

To demonstrate the role of the various planning and heuristic operations discussed above, an extended worked example is offered. To find a single naturally occurring argument which exhibits the results of a majority of the operations in a manageable unit is a tall order – instead, one long argument from the corpus²⁵ has been used as the basis for creating one of a more appropriate length which illustrates a reasonable number of the planning and heuristic operations available in *Rhetorica*. Although abridged, the extract remains true to the original in several important respects: (i) the order of components has not been changed; (ii) edited parts are clearly excisable from the argument structure; (iii) punctuation and formatting have been retained; (iv) selected extracts are included verbatim. The argument to be used for analysis, referred to hereafter as the *vegetarianism argument*, is given below:

Pigs are tethered in barren stalls, so deprived of stimulation that they often go mad. They are forced to have as many litters of piglets as their bodies will stand but they are not allowed to mother them.

But it's fine because they spend their lives in buildings sheltered from the elements with all the food they need and no animal wants more than that. Well, that's what we're told.

If a person did the same thing to dogs he or she would be prosecuted for cruelty – and yet a pig is equally as intelligent as a dog. So why the difference? Why do we have two sets of rules, one for a dog and one for a pig? There is no logical explanation.

The only excuse for visiting such suffering on other creatures can be ignorance. But that exit is closed to us because we do know the truth. We might pretend we don't, we might say that to give up pork²⁶ is pointless as one person's abstinence doesn't make any difference. But really we know it's wrong and that knowledge makes us push the reality into the darker reaches of our minds.

²⁵ The argument is an excellent example of modern day eloquence by Michael Mansfield QC, available at <http://www.viva.org.uk/Viva!Guides/justice.html>. The argument is also reproduced in full in Appendix B.

²⁶ Notice that the word *pork* replaces the phrase *battery eggs* found in the original argument. This is because an intervening paragraph in which the topic shifts from the suffering of pigs to that of battery hens occurs in the original argument. With this paragraph deleted in the pursuit of brevity in the analysis, the lexical swap is required to maintain the coherency of the whole. The modification has no significant impact on the analysis of the structure, and subsequent discussion of the generation of that structure.

Although analyses of natural arguments are notoriously equivocal, the following analysis of propositional structure offers a reasonably close and reliable characterisation:

- (a) Eating meat is wrong
- (b) Animals raised for meat suffer
- (c) Pigs are tethered in barren stalls so deprived of stimulation that they often go mad
- (d) No animal wants more than to live in a building sheltered from the elements with all the food they need
- (e) If a person did the same thing (c) to dogs, he or she would be prosecuted for cruelty
- (f) A pig is equally intelligent as a dog
- (g) There is no logical explanation for having two sets of rules, one for a dog and one for a pig
- (h) The only excuse can be ignorance
- (i) We do know the truth (even though we may pretend...)

The relationships between these propositions is shown below in Figure 4.23. This diagram (in common with most argumentation theoretic diagramming techniques) omits the rules of inference which are employed at the various nodes. In Figure 4.23, all links are Modus Ponens, except for the argument from *d* to *b* which is an undercutting argument (i.e. UCP), and that from $\sim h$ to *a*, which is a DS.




Figure 4.23 The argument structure for the vegetarianism argument

The system is initiated with a knowledge base containing all the relevant facts and inferences, in addition to a model of relevant beliefs held by the hearer. This knowledge base is enumerated in Figure 4.24.

Speaker Beliefs

a	b	c	d
e	f	g	h
i	$b \rightarrow a$	$c \rightarrow b$	$d \rightarrow \sim b$
$e \rightarrow \sim d$	$g \rightarrow (e \rightarrow \sim d)$	$f \rightarrow g$	$a \vee h$
$\sim h$	$i \rightarrow \sim h$		

Speaker's Model of Hearer Beliefs

$\text{bel}(ag_0, ?a) \quad \text{bel}(ag_0, d) \quad \text{bel}(ag_0, d \rightarrow \sim b)$

Figure 4.24 The initial knowledge base for the vegetarianism argument

The aim for the *Rhetorica* system, then, is to map from the knowledge base in Figure 4.24 (which, of course, would represent only a fragment of the complete KB of an applied system) to a representation of the structure in Figure 4.23. The planning process commences with the goals $\text{BEL}(ag_0, a)$ and $\text{IS_SALIENT}(ag_0, a, \text{arg}(a, _))$. The former is fulfilled by instantiation of two operators, $\text{MP}(ag_0, b, a)$ and $\text{DS}(ag_0, h, a)$; the latter by a corresponding MAKE_SALIENT . At this stage, the AS completes processing, and control passes to the EG, where the various heuristics are applicable. A number of these heuristics draw upon the contextual parameters describing the discourse encounter: an appropriate parameterisation of the context for the vegetarianism argument is given in Figure 4.25.

```
scepticism(ag0, high)
general_competence(ag0, high)
technical_competence(ag0, high)
investment(s, low)
medium(text)
```

Figure 4.25 Context parameters for the vegetarianism argument

At this stage, only one heuristic, HPR1 (climax ordering), is applicable (since, working down the preference order in Figure 4.22, there are too few arguments for HCR2, no implication steps for HCI1, no common property between subarguments for HCR1, no beliefs held by the hearer for HPR3, no refutations for HPP1, and high scepticism prohibiting HPR2). Given that the oracle evaluating argument strength returns that $\text{DS}(ag_0, h, a)$ is stronger than $\text{MP}(ag_0, b, a)$, the latter is constrained to precede the former by HPR1. This completes EG functionality, and control passes back to the AS, at which point, refinement is carried out.

The body goals of the MP and DS are posted, with the constraints active at the previous level of refinement (i.e. in this case, just that the MP precedes the DS) propagated to the post-refinement goals. Control passes back to the EG, where no new heuristic modification is effected, so control passes directly back to the AS where planning recommences (this represents one cycle of the flow illustrated in Figure 3.10).

Of the two goals in the MP body, only one gets fulfilled by subargument support – the other,

a: eating meat wrong; b: animals suffer; c: tethered pigs go mad; d: no animal wants more than food and shelter; e: doing the same to dogs is prosecutable; f: pigs as intelligent as dogs; g: no explanation for two sets of rule; h: excuse is ignorance; i: we do know the truth

$BEL(ag_0, b \rightarrow a)$, is met because, using negation as failure, $\sim BEL(ag_0, \sim(b \rightarrow a))$ is true. As discussed in §4.1.2, this weak form of terminal goal fulfilment is only suitable for situations in which no further support exists. A similar termination also satisfies the major premise of the DS subargument (i.e. $BEL(ag_0, a \vee h)$). The remaining two belief goals both give rise to further support. In the MP, $BEL(ag_0, b)$ gives rise to two supports: $MP(ag_0, c, b)$ and $UCP(ag_0, d, b)$. In the DS, the sole support for the minor premise is $MP(ag_0, i, \sim h)$. With fulfilment of all saliency goals by appropriate instantiations of `MAKE_SALIENT`, this completes the round of planning, and control passes to the EG.

Two of the heuristics which might now potentially be relevant are barred: `HCI1` reminds of implication if there is a large amount of intervening material, and such material is not present in the current plan; `HPP1` removes a UCP if it is weak, whereas the current UCP is strong enough (as determined by the argument strength evaluation oracle) to resist `HPP1`. The `HPR4` (refutations first) heuristic, is triggered however, and introduces a constraint that the $UCP(ag_0, d, b)$ must precede $MP(ag_0, c, b)$. With no further heuristics applicable, control reverts to the AS level for refinement.

The state resulting from refinement is shown in Figure 4.26. The orderings between goals are omitted for clarity, but the order in which goals are listed in the figure complies with the constraints on the partial order. The main subargument units are indicated by braces to the left of the partial plan.




Figure 4.26 Midway status of the planning process for the vegetarianism argument

a: eating meat wrong; b: animals suffer; c: tethered pigs go mad; d: no animal wants more than food and shelter; e: doing the same to dogs is prosecutable; f: pigs as intelligent as dogs; g: no explanation for two sets of rule; h: excuse is ignorance; i: we do know the truth

With saliency goals fulfilled appropriately, and conditions for belief goal satisfaction met, most of the argument is fixed – the only area to be resolved is the support for $BEL(ag_0, \sim d)$. The AS processing determines this support in a straightforward manner, with the goal being fulfilled by $MP(ag_0, e, \sim d)$, and then, after refinement, the $BEL(ag_0, e \rightarrow \sim d)$ goal being fulfilled by another $MP, MP(ag_0, g, e \rightarrow \sim d)$. This latter component, supporting the major premise of the earlier MP , gives rise to the linked structure shown in Figure 4.23 (all other support to this point has been convergent). Finally, when the last MP is refined, the last component of the argument, $MP(ag_0, f, g)$, is introduced to support $BEL(ag_0, g)$. The final plan of primitives is shown in Figure 4.27.

```

PUSH_TOPIC (arg(b, a))
PUSH_TOPIC (arg( $\sim d$ , b))
MAKE_SALIENT ( $ag_0$ , d, arg( $\sim d$ , b))
PUSH_TOPIC (arg(e, d))
MAKE_SALIENT ( $ag_0$ , e, arg(e, d))
MAKE_SALIENT ( $ag_0$ ,  $e \rightarrow d$ , arg(e, d))
PUSH_TOPIC (arg(g,  $e \rightarrow d$ ))
MAKE_SALIENT ( $ag_0$ , g, arg(g,  $e \rightarrow d$ ))
PUSH_TOPIC (arg(f, g))
MAKE_SALIENT ( $ag_0$ , f, arg(f, g))
MAKE_SALIENT ( $ag_0$ ,  $f \rightarrow g$ , arg(f, g))
POP_TOPIC (arg(f, g))
MAKE_SALIENT ( $ag_0$ ,  $g \rightarrow (e \rightarrow d)$ , arg(g,  $e \rightarrow d$ ))
POP_TOPIC (arg(g,  $e \rightarrow d$ ))
POP_TOPIC (arg(e, d))
MAKE_SALIENT ( $ag_0$ ,  $\sim d$ , arg( $\sim d$ , arg( $\sim d$ , b)))
MAKE_SALIENT ( $ag_0$ ,  $d \rightarrow \sim b$ , arg( $\sim d$ , b))
POP_TOPIC (arg( $\sim d$ , b))
PUSH_TOPIC (arg(c, b))
MAKE_SALIENT ( $ag_0$ , c, arg(c, b))
MAKE_SALIENT ( $ag_0$ ,  $c \rightarrow b$ , arg(c, b))
POP_TOPIC (arg(c, b))
MAKE_SALIENT ( $ag_0$ , b, arg(b, a))
MAKE_SALIENT ( $ag_0$ ,  $b \rightarrow a$ , arg(b, a))
POP_TOPIC (arg(b, a))
PUSH_TOPIC (arg( $\sim h$ , a))
PUSH_TOPIC (arg(i,  $\sim h$ ))
MAKE_SALIENT ( $ag_0$ , i, arg(i,  $\sim h$ ))
MAKE_SALIENT ( $ag_0$ ,  $i \rightarrow \sim h$ , arg(i,  $\sim h$ ))
POP_TOPIC (arg(i,  $\sim h$ ))
MAKE_SALIENT ( $ag_0$ ,  $\sim h$ , arg( $\sim h$ , a))
MAKE_SALIENT ( $ag_0$ ,  $a \vee h$ , arg( $\sim h$ , a))
POP_TOPIC (arg( $\sim h$ , a))
MAKE_SALIENT ( $ag_0$ , a, arg(a, _))

```

Figure 4.27 Final plan of primitives for vegetarianism argument

This partially specified plan not only describes the structure of the vegetarianism argument, it also leaves unconstrained choices open, such that the sub-EG processing has the flexibility to realise the clauses in the most appropriate way (given interclause and lexical constraints). Furthermore, the plan is also perfectly suited to driving the introduction of a number of surface features which are closely related to the abstract structure of the argument. The following chapter explores the range of such devices, and the means by which features of the plan can be exploited in their generation.

a: eating meat wrong; b: animals suffer; c: tethered pigs go mad; d: no animal wants more than food and shelter; e: doing the same to dogs is prosecutable; f: pigs as intelligent as dogs; g: no explanation for two sets of rule; h: excuse is ignorance; i: we do know the truth
--

V

Surface Features

The skeletal form of the argument created at the AS level lends itself to an assault on several important and interesting problems. The first concerns the use of clues, words or phrases which serve to linguistically mark relations between spans of text. A survey and analysis of current research into the use and possible generation of such clues (at both clausal and argument levels) is provided, before the presentation of one component of the EG level wherein inter-argument clues are introduced into the text plan in a clear, principled manner. The model is then extended to additionally handle a small set of punctuation and formatting devices, which can be viewed in a similar way. The second issue addressed at the EG level is the generation of enthymemes – arguments with components left implicit. The use of goals of saliency is shown to be crucial in building a model which can heuristically create enthymematic argument. The third role of the EG level is in introducing ‘rhetorical’ text: insubstantial, nonfunctional, yet highly effective spans of text which serve to perform premise establishment and qualified quantification. Related to this issue of text which does not directly play a role in the functional structure of an argument, is the fourth and final component of EG functionality: *affect*. Several heuristics are presented which mark the need for loaded lexical realisation, so as to subtly influence how a claim will be received by the hearer. These four components of the EG level are then demonstrated and critiqued with respect to the vegetarianism argument introduced in §4.4, illustrating the AS and EG levels working in tandem to account for a range of features in large scale, real world persuasive text

5.1 Clues

Once the structure and content of an argument has been planned at an abstract level, it remains to introduce appropriate cue phrases to signal the structure to the hearer (following (Grosz and Sidner, 1986), the current work distinguishes *clue words* from *cue phrases*: the former refers to indications of large-scale argument structure, the latter is a superset of the former, encompassing all indications of discourse structure). Such cue phrases have been recognised at a variety of levels of analysis, both with

a specific bent on argumentation, and across natural language in general, e.g. (Grosz and Sidner, 1986), (Birnbaum, 1982), (Hobbs, 1982), (Reichman-Adar, 1984), (Wilson, 1980), (Johnson, 1992), (Brooks and Warren, 1972) and (Snoeck Henkemans, 1997). Of this work, however – which is almost exclusively analytic rather than generative – two studies are of particular importance²⁷. At an interclause level, Knott's (1996) empirical study, motivated by the need for a means of determining an appropriate set of coherence relations, characterises the specificity of some two hundred cue phrases by analysis of their semantic role (through the linguistically well established approach of a substitution test). Though the majority of these are beyond the scope of the current work (either because they do not signal argument structure, or because they relate clauses at a much lower level), a number are considered for handling within the *Rhetorica* system in §5.1.1, below. At a more abstract level, Cohen's work (1984), (1987) offers a taxonomy of clues employed for signalling various argument structures, and it is this taxonomy (again constructed for an analytic theory) which forms the chief basis for clue generation, explored in §5.1.2.

5.1.1 Interclause cue phrase generation

Knott (1996, chapter six) offers a first cut at a feature based taxonomy of cue phrases employing eight independent binary dimensions: SOURCE OF COHERENCE, ANCHOR, PATTERN OF INSTANTIATION, FOCUS OF POLARITY, POLARITY, PRESUPPOSITIONALITY, MODAL STATUS and RULE TYPE. Of these, two are of particular import to the current work. The first is POLARITY, which Knott characterises (p101) in terms of the causal or consequential rule implicitly connecting two premises; in cases where this rule is defeated (Knott *et al.*, 1997), the polarity of the relation between the two clauses is negative, otherwise it is positive. The stereotypical (or, in Knott's work, the most abstract) cue in situations of negative polarity is *but*. From both Knott's definition of negative polarity, and from its associated cue *but*, it is clear that the link between the components in a Modus Tollens (MT) could be expressed in these terms. Indeed, in realising an MT it is difficult to avoid using *but* or something similar – in particular arrangements of the components of the MT²⁸. An MT comprises three parts, the implication step, the denial of the consequent, and concluding the reverse of the antecedent. There are thus six possible arrangements of these components (ignoring for the moment the possibility of further support for each of the three claims, or for elements to be left implicit). The canonical order shown in Example 5.1 thus has the various permutations illustrated in Examples 5.2 through 5.6:

²⁷ Reichman, though both recognising the importance of clue words, and building parts of a model which was to be generative, did not address the generation of clue words explicitly.

²⁸ Interestingly, Knott seems not to have considered the impact of component ordering on the resultant applicability of cue phrases: he seems implicitly to assume that if two clauses are related by a particular relation then a single cue phrase (or set of phrases) is applicable. In the light of the evidence of Examples 5.1 through 5.6, it would seem that the assumption would only hold if separate relations were defined for different ordering scenarios.

If it's been raining, it'd be wet, but it isn't wet so it hasn't been raining. (5.1)

If it's been raining, it'd be wet, but it hasn't been raining because it isn't wet. (5.2)

It hasn't been raining because if it had been, it'd be wet but it isn't. (5.3)

It hasn't been raining because it's not wet, and if it had been raining it would be (5.4)

It's not wet, so it hasn't been raining, because if it had been, it would be wet. (5.5)

It's not wet, and if it had been raining it would be, so it hasn't been raining. (5.6)

The corresponding propositional readings of these statements are shown in 5.1a through 5.6a:

$P \rightarrow Q, \sim Q, \sim P$ (5.1a)

$P \rightarrow Q, \sim P, \sim Q$ (5.2a)

$\sim P, P \rightarrow Q, \sim Q$ (5.3a)

$\sim P, \sim Q, P \rightarrow Q$ (5.4a)

$\sim Q, \sim P, P \rightarrow Q$ (5.5a)

$\sim Q, P \rightarrow Q, \sim P$ (5.6a)

A cursory glance shows that the negative polarity indicator, *but*, only occurs in Examples 5.1, 5.2 and 5.3 – the only examples in which $P \rightarrow Q$ precedes $\sim Q$ (though the former does not necessarily have to immediately precede the latter, as demonstrated by Example 5.2). Of course, it is not obligatory for *but* (or something similar such as *however*) to be realised, and indeed a more general cue such as *and* is also applicable in some cases (such as Example 5.2) – the problem of determining whether or not to realise pieces of information (both clues and argument content) is addressed in §5.2. Regardless of whether or not a clue is realised however, it is only where $P \rightarrow Q$ precedes $\sim Q$ that such a clue is licensed at all (the use of *but* in Examples 5.4 through 5.6 would be incoherent). Thus, the negative polarity is an intrinsic feature of the occurrence of $P \rightarrow Q$ prior to $\sim Q$ in an MT.

Knott's account relies upon the notion of a rule being defeated – this defeat then results in the negative polarity relationship. It is important to note that the *but* in an MT is not indicating the defeat of the major premise $P \rightarrow Q$. Indeed, the upholding of this premise is crucial for the successful application of MT. Rather, it seems that there is a meta-level rule which is being defeated. Uttering $P \rightarrow Q$ carries with it a weak defeasible rule that the speaker actually holds P and Q to be true. It is *this* rule which is then defeated, and the resultant negative polarity indicated by *but*. (It is for this reason that the fallacious argument $P \rightarrow Q, \sim P$ therefore $\sim Q$ – the fallacy of denying the antecedent – also characteristically employs *but* between the first two premises).

The question of from where the various features arise is not addressed in Knott (as it to be expected, given his analytic stance). For a generation system to account for cues, however, it is inappropriate simply to consider adjacent clauses, as Knott does. Rather, the feature needs to be introduced as a result of the planning process – if the feature is a result of a particular operator applied at some level of abstraction, then the feature itself need to be introduced at that point, rather than inferred at some later stage. Negative polarity holds between the two premises of an MT, so whenever an MT is applied, the planner must introduce negative polarity between the two premises (which may, of course, after further planning turn out to be complex subarguments).

Responsibility for introducing an indication of an appropriate cue phrase lies with the EG level. The functionality is incorporated in heuristic structures which are active during EG processing

(i.e. the second and last boxes in the control flow diagram Figure 3.10). The notation used to specify cue phrases forms an interface between the EG and processing at sub-EG levels – and ultimately at the lexical level. As these do not fall within the remit of the current work, it is difficult to anticipate latter requirements and confidently define a specific interface. Furthermore, the pipeline hypothesis has inhibited consideration of introducing lexical information at more abstract, ‘earlier’ stages of the generation process; partially constraining lexical choice on the basis of abstract structure seems to have attracted very limited attention (Rubinoff (1992), for example, introduces explicit feedback from lexical levels to planning, rather than constraining lexicalisation during planning). As a result, the current work cannot simply adopt the approach of a precedent.

Instead, an assumption is made that subsequent planning will be able to interpret the feature POLARITY itself. Thus the EG level does not determine the lexeme to be used, rather, it specifies an abstraction of a group of applicable lexemes, thus maintaining a proper degree of abstraction. This is not only consistent with Knott’s work (such that if a number of features held between argument components, then an appropriate compound marker would be introduced by the EG), but also with recent work on lexical generation, particularly Meteer (1991) and Panaget (1994), (1997). In his abstraction-based approach which unifies lexical and grammatical resources, Panaget (1994) includes one example of a clause modifier which could function as a cue phrase – *so that* (p133). His characterisation permits multiple ideational concepts to be associated with a particular abstract linguistic resource (i.e. lexical or grammatical unit); and if Knott’s features can be construed as ideational categories, then there appears to be no problem in unifying the two approaches, and exploiting such a unification in the specification of cue phrases introduced by the EG.

The negative polarity relationship holding between the two premises of an MT is introduced by CLUE-MT, shown in Figure 4.7. The trigger matches the set of body goals of an MT, when they are constrained such that the implication premise precedes the denial of the consequent (as required for negative polarity). The update introduces a new goal to be placed between the two MT premises. That goal is to make salient to the hearer the feature holding between the two premises. Thus CLUE takes two parameters, specifying feature and value, and the MAKE_SALIENT of which it is part represents a goal for subsequent planning to introduce a cue phrase of an appropriate feature type. It is thus the binary predicate CLUE which forms the EG/sub-EG interface, with the values that CLUE can take following Knott’s characterisation.

CLUE-MT

Trigger: ta: PUSH_TOPIC ($\sim P \rightarrow X$)
 tb: IS_SALIENT (Ag, $\sim X$, $\sim P \rightarrow X$)
 tc: IS_SALIENT (Ag, $\sim P \rightarrow X$, $\sim P \rightarrow X$)
 td: POP_TOPIC ($\sim P \rightarrow X$)
 (ta < tb < td)
 (ta < tc < td)
 (tc < tb)

Update: tn: MAKE_SALIENT (Ag, CLUE(polarity, negative), $\sim P \rightarrow X$)
 (tc < tn < tb)

Figure 5.1 The CLUE-MT heuristic

The intersection of the POLARITY feature with another on Knott's list, SOURCE OF COHERENCE is also important in realising argument. In any situation where a rule is defeated (including the defeasible rule of speaker adherence to antecedent and consequent of an implication used in MT), the speaker intends

“to inform the reader that some general rule in the reader's model of the world is defeated in the situation being described” (Knott, 1996, p104)

This Knott terms negative polarity with a semantic source of coherence (following (Sanders *et al.*, 1992)), and is contrasted with a negative polarity and pragmatic source of coherence, where

“the writer's aim is that some general rule in the world itself *actually* fails in the present instance” *ibid.*

In the latter case only, *but* can be replaced by the construction *admittedly...but*. The use of both *but* and *admittedly...but* features in the realisation of counter-counterargumentation, and it might be expected therefore that the pragmatic/semantic distinction might play an important role.

The distinction between semantic and pragmatic negative polarity is illustrated by Knott in his examples 6.10 (pragmatic) and 6.11 (semantic), shown in Figure 5.2 (notice that Knott employs ‘#’ to indicate a semantically unlicensed cue – rather than the more usual linguistic reading of ‘probably ill-formed’).

United have some key players injured; $\left\{ \begin{array}{l} \textit{admittedly...but} \\ \surd \textit{but} \end{array} \right\}$ they're still bound to win. (6.10)

Mary was behaving oddly. $\left\{ \begin{array}{l} \textit{but} \\ \# \textit{admittedly...but} \end{array} \right\}$ she didn't eat any of it. (6.11)

Figure 5.2 Examples 6.10 and 6.11 from (Knott, 1996, p103)

In the first example, *admittedly...but* and *but* are both acceptable; in the latter, *admittedly...but* is not. This, Knott claims, is a result of the defeat of a (real-world) causal rule for the former, and of a (hearer-resident) deductive rule for the latter.

Unfortunately, this seems to be something of a weak spot in Knott's characterisation. The first problem is with (6.11), which intuitively seems rather weak, particularly as an extra sentence of context is introduced. It is unclear exactly how to distinguish the rules in the hearer's model of the world and the world itself. In Knott's examples, the problem is confounded by the inclusion of a much more significant difference which may, in fact, account for the *but / admittedly...but* distinction rather more definitively. Knott mentions that *admittedly...but* "has a specifically argumentative flavour" (p103). A more incisive reformulation of this claim is that the locus of attention is the conclusion – it is the conclusion which forms the functional aim of the text fragment (this is often indicated by stress patterns in an utterance, i.e. *Admittedly P, but not C*). This is in clear contrast to (6.11) where the locus is the defeated rule itself – the speaker is uttering something like *P but not C; Interesting, huh?* This is manifest in Knott's example by the presence of the first sentence – a conclusion based on the defeat of the inference rule. Reformulations of Knott's two examples may serve to clarify this point.

$$\begin{array}{l} \text{(Mary is anorexic)} \\ \text{She ordered a pizza} \end{array} \left\{ \begin{array}{l} \text{admittedly...but} \\ \checkmark \text{ but} \end{array} \right\} \text{ she didn't eat any of it} \quad (5.7)$$

$$\begin{array}{l} \text{United are renowned for their} \\ \text{underhand player transfers. They} \\ \text{have some key players injured} \end{array} \left\{ \begin{array}{l} \text{but} \\ \# \text{ admittedly...but} \end{array} \right\} \text{ they're still bound to win} \quad (5.8)$$

In (5.7) stress patterns and the parenthetic sentence contextualise the utterance to ensure the locus of attention is the conclusion (if that context were available from non-linguistic sources, neither stress nor the parenthetic phrase would be necessary). With the locus so fixed, the *admittedly...but* construction becomes valid. Similarly, in (5.8), the first sentence emphasises that the locus is the defeated rule. Importantly, the rule itself does not change from Knott's original to these reformulations – i.e. the same rule is defeated in (6.10) as is defeated in (5.8) and the same rule is defeated in (6.11) as is defeated in (5.7). The rule has not shifted from the hearer's model of the world to the world itself, nor in the opposite direction. By Knott's analysis therefore, the licensed cues should not change – and yet, examples (5.7) and (5.8) demonstrate that they do. This works well for an analysis of the negative polarity in MT. Under Knott's characterisation, it would be necessary to determine whether the inference rule 'uttering an implication suggests belief in antecedent and consequent' is part of the hearer's model of the world, or whether it is a rule of the world itself. There would seem to be reasonable arguments for taking either line (or a combination of both); eschewing either completely would be rather unpalatable. Instead, using a locus-based approach, it is necessary only to determine where the locus of attention lies in an MT. It is quite clear that the locus of attention is the conclusion (i.e. the consequent of the MT) and not the rule ('uttering an implication ... etc.'). The only exception would be a highly contrived example of the form (5.9) – where, as expected, *admittedly...but* is prohibited. In standard cases of MT, however, uttering *Admittedly P → Q, but not Q...* is perfectly acceptable.

$$\begin{array}{l} \text{The fact that my stating an implication} \\ \text{does not necessarily imply I believe} \\ \text{the consequent is interesting. A implies B} \end{array} \left\{ \begin{array}{l} \text{but} \\ \# \text{ admittedly...but} \end{array} \right\} \text{ not B. (so A)} \quad (5.9)$$

A similar situation is found in counter-counterargumentation, between attacking premises of counterarguments, and attacking the support that premises have for conclusions in counterarguments – as discussed in §4.1.1, these two techniques are characterised by the AS level operators UCP and UCI respectively. In a situation where the hearer (as modelled by the speaker) maintains an argument that it is wet on the basis that it has been raining, and that the latter implies the former, the speaker might invoke either UCP or UCI to generate the arguments in Examples 5.10 and 5.11²⁹:

$$\text{Raining makes the } \left\{ \begin{array}{l} \text{admittedly...but} \\ \sqrt{\text{but}} \end{array} \right\} \begin{array}{l} \text{it hasn't been raining,} \\ \text{so it's not wet.} \end{array} \quad (5.10)$$

$$\text{It's been raining } \left\{ \begin{array}{l} \text{but} \\ \sqrt{\text{admittedly...but}} \end{array} \right\} \begin{array}{l} \text{that doesn't mean the} \\ \text{ground is wet} \end{array} \quad (5.11)$$

In both cases, the locus of attention is the conclusion, and so in both cases, either *but* or *admittedly...but* is appropriate. Identifying the negative polarity in UCI and UCP enables the construction of two further heuristics, CLUE-UCP and CLUE-UCI, which introduce appropriate goals of saliency with these features specified. Figure 5.3 gives the definitions for the two heuristics. Note that as with CLUE-MT, CLUE-UCP requires the implication step to precede the refutation of the precedent to license the introduction of a negative polarity feature; CLUE-UCI similarly requires the step which is dissonant with hearer beliefs – this time the implication step itself – to succeed that which is consonant.

²⁹ Without further support, of course, Example 5.10 is an example of the fallacy of denying the antecedent. Clearly, however, UCP is generally used in conjunction with other supports for the conclusion. Example 5.11 is even weaker – there is a strong expectation for information about the abnormality (that the ground is covered, sheltered, etc.)

this respect, the work of Cohen closely resembles that of Knott). The relationship definitions are based upon notions of hierarchical dependence and subargument units which are similar to those used in *Rhetorica* (and although not identical, the differences are not important to the current discussion). These categories are shown in Figure 5.4 below.

CATEGORY	RELATION: P to S	EXAMPLE
parallel	brother	First, Second
inference	son	As a result
detail	father	In particular
summary	multiple sons	In sum
reformulation	son (& father)	In other words
contrast	brother or father	On the other hand

Figure 5.4 Clue categories, after (Cohen, 1987, p15)

Parallel clues express the relation holding between consecutive subarguments which contribute to the same conclusion (and can thus occur in any ordering arrangement). Generation of such clues relies on identifying an appropriate pattern of goals in a partially specified abstract plan – the appropriate pattern in this case is simply the occurrence of multiple `PUSH_TOPIC` goals with parameter `arg(X, Y)` between the `PUSH_TOPIC` and `POP_TOPIC` of a superargument (i.e. whose parameters are `arg(Y, Z)`). This pattern is detected by the heuristic `CLUE-PARALLEL` which introduces an appropriate saliency goal for a parallel clue, as shown in Figure 5.5. Note that when a particular pattern of primitives license the introduction of a clue, it is *not* the case that those primitives are deleted from the plan. Although clues can be seen as fulfilling in part goals such as `PUSH_TOPIC` and `POP_TOPIC` (by expressing the speaker's intention to make the hearer aware of the structure of an argument), it is important to bear in mind that subsequent, sub-EG planning may further fulfil the goals with sentence delimitation, parenthesis, etc., and the topic manipulation goals need therefore to persist through the EG level processing.

CLUE-PARALLEL

```

Trigger:      ta: PUSH_TOPIC (arg (X, P))
              tb: PUSH_TOPIC (arg (A, X))
              tc: POP_TOPIC (arg (A, X))
              td: PUSH_TOPIC (arg (B, X))
              te: POP_TOPIC (arg (B, X))
              tf: POP_TOPIC (arg (X, P))
              (ta < tb < tf)
              (ta < tc < tf)
              (ta < td < tf)
              (ta < te < tf)

Update:      tn: MAKE_SALIENT (Ag, CLUE(parallel), arg(X, P))
              (ta < tn < tf)

```

Figure 5.5 The CLUE-PARALLEL heuristic

The position of the goal expressing clue introduction is not fixed between the two subarguments tb - tc and td - te , for not all parallel clues are found between the subarguments they connect. Cohen's example of *first...second* is a case in point, though an entire phrase such as *There are three reasons for this...* is also an example of a non-interposed parallel clue. It is lexico-grammatical constraints which force the positioning of a particular parallel clue in a particular position with respect to the subarguments it relates, and these constraints are enforced at lower levels in the planning framework. Such constraints cannot be considered until the exact choice of clue word – e.g. whether to use *first...second* or *furthermore* – has been made, and this decision is also deferred: the EG level simply identifies the category of clue required. Of course, the position of the clue is constrained inasmuch as the context parameter of the `MAKE_SALIENT` goal enforces realisation of the clue outside the subarguments tb - tc and td - te , but within the scope of the goals `PUSH_TOPIC (arg (X, P))` and `POP_TOPIC (arg (X, P))`. This provides precisely the required flexibility, whilst not incurring any significant extra computation for the planning machinery.

Although not made explicit in Cohen's work, it would appear that lexical marking is not the only form of clue to indicate parallel structuring, particularly in written text. At coarser-grained levels (i.e., in a hierarchical planning framework, at a higher level of abstraction), the division of an argument into subarguments is often accompanied by formatting conventions; in the preceding discussion the use of the phrase *There are three reasons for this...* might be followed by a numbered or bulleted list. The generation of such features is discussed in §5.1.3

Inference clues can be generated in a similar way to parallel clues, but their introduction relies on a rather more localised pattern. The pattern specification licensing inference clue introduction is simple: an expressed premise immediately preceding its associated conclusion. Given that any conclusion is a belief goal / saliency goal pair, the pattern is equivalent to the occurrence of the belief goal immediately preceding the saliency goal. It is difficult (within the conventional partial order notation adopted in *Rhetorica*) to specify that one goal *immediately* precedes another: instead, the constraint is enforced by the focusing control imposing immediate proximity of conclusion to supporting argument(s) discussed above in chapter four. The only under-specification is in whether the conclusion precedes or follows a given subargument; an inference clue is then only licensed by pre-order (i.e. premise-then-conclusion). This specification is captured by `CLUE-INFERERENCE`, shown in Figure 5.6.

CLUE-INFERERENCE

```

Trigger:      ta: BEL (Ag, X)
              tb: IS_SALIENT (Ag, X, arg (X, P))
              (ta < tb)

Update:       tn: MAKE_SALIENT (Ag, CLUE(inference), arg(X, P))
              (ta < tn < tb)

```

Figure 5.6 The `CLUE-INFERERENCE` heuristic

Generation of detail clues occurs with the reverse pattern used by CLUE-INFERENCE – i.e. with hybrid or post-order, between the conclusion and the first premise (or subargument) which it precedes. As with inference clues, the required immediacy constraints are imposed not by the clue introduction heuristic, but by the focus control mechanisms active at run-time. The introduction of detail clues can thus be performed by CLUE-DETAIL, shown in Figure 5.7.

CLUE-DETAIL

Trigger: ta: IS_SALIENT (Ag, X, arg (X, P))
 ta: BEL (Ag, X)
 (ta < tb)

Update: tn: MAKE_SALIENT (Ag, CLUE(detail), arg(X, P))
 (ta < tn < tb)

Figure 5.7 The CLUE-DETAIL heuristic

Clues of summary are currently beyond the scope of the *Rhetorica* system for a number of reasons. The first is that they require the generation of content which the hearer not only knows (this much is accounted for through the distinction between saliency and belief – beliefs known to be held by the hearer may still be made salient in the construction of an argument) but furthermore, that the hearer has already been told during the course of the argument. This repetition of information (albeit in a reduced form in a summary) is not accounted for in the planning process, since it violates the basic definition of the hierarchical structure of argument structure, that any single claim has a single functional role to play, and is supported and made salient to the hearer. Repeating information also presents a problem for the implicit assumption of the monotonicity of the communication process discussed above and in (Reed *et al.*, 1997): once a claim has been conveyed, it is unnecessary to convey that claim again. This in turn involves an assumption that the saliency of propositions persists – for at least as long as is convenient for the speaker (though this unrealistic assumption is addressed in part by ordering and enthymeme heuristics discussed below). Finally, the use of summary clues is also related to the gross scale organisation of orations proposed in Cicero etc. (see e.g. (Blair, 1838)), whereby the argumentative part characterised by *Rhetorica* is preceded by stages of introduction, division, narration, and followed by stages of pathetic appeal and conclusion. There is significant repetition between these stages (particularly division, narration, argumentative and conclusion), which might not only be characterised intuitively as summary, but also make characteristic use of the clues of summary discussed by Cohen. Only the argumentative part of persuasive text is considered in the current work: in practice, this limits its generative capability to shorter texts (though still accounting for multi-paragraph text usually not considered in NLG).

Clues of reformulation suffer from a similar set of problems to those of summary, given that they too introduce a repetition of material. Tackling the generation of reformulation, however, is simplified somewhat by their very localised introduction – in Cohen’s analysis, reformulation of some individual claim occurs immediately after the initial statement itself. Thus in the *Rhetorica* system, both the initial claim and its reformulation are seen as the result of a single MAKE_SALIENT goal. This

indicates that it is not generally the activity of the higher levels which introduces reformulation; rather, it is a lower level process which realises a particular saliency goal (in conjunction, perhaps with additional information such as the hearer's competence) as not one but two clauses, connected by a reformulation clue. Such a process might be triggered by additional information attached to the MAKE_SALIENT goal indicating the importance of the premise requiring expression (such that a crucial claim might be reformulated).

Less frequently, clues of reformulation are also found in patterns which could be accounted for by higher level processing. The typical structure is one in which the conclusion of an argument is presented before its supporting premises and then again as a reformulation at the close. This arrangement shares many of the features of a summary – and suffers from the same problems.

Finally, Cohen associates clues of contrast with a loose set of relations which seem to have been characterised much more precisely by Knott in his definition of POLARITY. The means by which *Rhetorica* employs polarity (and its intersection with the source of coherence) have been explored in the previous section.

Unfortunately, Cohen draws little upon work in argumentation theory, and as a result her taxonomy does not recognise the importance of the distinction between linked and convergent argumentation (Reed and Long, 1997a). The conflation of the two forms represents a particular problem for clues indicating parallel structure, where usual examples such as *first...second...* are almost always inappropriate for convergent arrangements³⁰ (consider, for example, *George said it's true, and he's reliable, so it must be so* – introducing a parallel clue, *Firstly, George said it's true; secondly, he's reliable, so all in all it must be true*, seems to border on incoherency). Identification of clues bound to various argument forms is a major programme of research, and is consequently beyond the scope of the current investigation; early results of one major project in this area are offered in (Snoeck Henkemans, 1997).

5.1.3 Punctuation and formatting as clues

Punctuation is a crucial component of written text. It functions to clarify text by delimiting semantic units, a process which does not just clarify meaning, but plays a vital role in comprehension (compare, for example, the growing use of punctuation analysis as an aid in parsing technology: (Briscoe, 1996); (Jones, 1994a); (Peh and Ting, 1996); (Simard, 1996)). Nunberg (1990) describes the role of punctuation as one of indicating *text-categories*: in the same way that the lexical content of a text is based upon the lexical categories, so the graphological appearance is based upon text categories. Thus Nunberg proposes the category of *text-sentence*, which is typically delimited by a capital letter at the head, and a full stop at the tail. A text-sentence does not necessarily coincide with a lexical sentence (of standard linguistic analysis): a text sentence may comprise several lexical sentences (conjoined, for example, by a semicolon), or may alternatively not even contain a single lexical sentence³¹.

Text generation has not formally distinguished between these two forms of grammar (i.e. the textual and the lexical) – generation seems to have relied upon the linguistics of speech (despite the fact

³⁰ A brief survey seems to suggest that *furthermore* is one of a rather small set of parallel clues which can serve in linked structures.

³¹ Nunberg's example (p22) offers three text sentences of which only one is classed as a lexical sentence: *The L9000 delivers everything you wanted in a luxury sedan. With more power. At a price you can afford.*

that realisation is most commonly textual). As Nunberg persuasively argues, it is important to see written text as a linguistic system in itself – not just a transcript of speech (the role of punctuation, for example, is not one of indicating intonation patterns). Despite failing to make this distinction, NLG has nevertheless been able to approach the problem of producing effective punctuation because many systems maintain an abstract representation of communicative intent. This representation is free of structural and textual constraints – these constraints can then be introduced as the planning process is conducted, and the communicative intentions are realised into surface textual form. As the planning proceeds, both lexical and textual structure can be imposed – and the latter can include punctuation.

At the lowest levels of realisation where lexical and morphological structure is determined, punctuation such as hyphens and virgules (e.g. *and/or*) are introduced. Above this sits further grammatical functionality, which Hovy and Wanner (1996) have proposed as the *sentence generation level*: it is here that delimiting and separating commas are introduced, along with dash interpolation and parentheses. More abstract again, the inter-clause relations of RST link semantic units of text, and can give rise to punctuation appropriate to this level in the hierarchy. Full stops, colons and semicolons are all generated at this level, due to their dependence upon the notion of ‘clause’. However, it is also clear that they are rather close to the border with the previous level, at which it has been proposed that text-sentence content delimitation is performed (Hovy and Wanner, 1996). Given the recursive manner in which RST is applied, it is unsurprising that it can also be used to account for a number of other, ‘higher-level’ forms of punctuation. For example, Hovy and Arens (1991) discuss the generation of enumerated lists from consecutive SEQUENCE relations, and of itemised lists from clauses connected by a JOINT relation. They also suggest that section headings are the result of the IDENTIFICATION relation and appendices and footnotes of the BACKGROUND relation. Lastly, they note that some structures are the result of concepts not explicitly handled by RST – quotation, for example, relying upon PROJECTION.

Other punctuation forms, however, cannot adequately be accounted for by structure determined at RST or lexical levels of planning. For the same reasons that RST or other approaches based on a set of coherence relations are unable to account for the internal structure of an argument, they are also unable to generate explicit indication of that structure in the form of punctuation. Responsibility for the introduction of such features must therefore lie with the more abstract levels where the gross scale argument structure is determined. Since the punctuation introduced at these levels is functioning to indicate structure to the hearer, it is appropriate to view punctuation as a form of clue – and this is precisely how the implementation is handled. Three types of punctuation are particularly common for realising abstract structural intentions: breaks, footnotes and quotations.

The use of chapter and section breaks is termed by Hovy and Arens (1991) *separation*, referring to the vertical displacement often associated with these formatting conventions. However, the organisation of chapters, sections and subsections relies upon more than the vertical distance that such headings are afforded (since headings may or may not be numbered, and may or may not be in-line, and may or may not be in bold or italic face, and so on). In the text of extended argument (such as a monograph), the chapter and section breaks correspond to focus shifts at the highest levels of abstraction: these shifts are marked in the plan by the PUSH_TOPIC and POP_TOPIC goals, so it is these

which form the basis of the trigger for the various punctuation heuristics. The use of the topic manipulation goals (which as mentioned above, are communicative goals of intention) to generate formatting is precisely the notion implicit in the analysis of chapter and section breaks offered by (Pascual and Virbel, 1996), in which the use of such formatting expresses the underlying architecture of the text, which they term a *metadiscourse*. A metadiscourse consists of statements made by the speaker, though the use of formatting, regarding the structure of the discourse itself (e.g. *The author creates a text identified as X; The author organizes X into two parts identified by A and B; The author assigns the level of chapter to A and B; etc.*).

The RST based approach for introducing breaks proposed by Hovy and Arens suffers from two key problems. In the first place, RST is unable to account for the intentionality identified by Pascual and Virbel as lying behind high level discourse breaks, as discussed at length in §3.1.4 above. Secondly, although such breaks may occur at a single level in the hierarchical text plan, they do not occur with a single type of relation; rather, they can be associated with almost any relation which happens to be included at the appropriate level in the hierarchy. Thus the generation of the metadiscourse proposed in (Pascual and Virbel, 1996) and (Pascual, 1996) relies upon a representation more abstract than RST.

The plan of primitives produced by *Rhetorica* has the potential for characterising the metadiscourse, and thereby generating appropriate high level breaks. The crucial problem is to decide precisely what plan structure licenses the introduction of a break (and, then, what kind of break is most appropriate). In the current implementation, the problem is simplified in two respects: (i) only one break is considered – a standard paragraph break, and (ii) the break is associated with a fixed level of argument nesting. After inspection of several natural arguments collected in the corpus, it has been determined that paragraph breaks are typically associated with focus shifts occurring between arguments which are distanced from the terminal propositions by about three levels of subargumentation. Clearly, this figure is easily altered, but on the basis of the small sample, the rule is implemented by PUNC-BREAK, shown in Figure 5.8.

PUNC-BREAK

```

Trigger:      ta: PUSH (arg(P, Q))
              tb: PUSH (arg(X, P))
              tc: PUSH (arg(Y, X))
              td: PUSH (arg(Z, Y))
              te: POP  (arg(Z, Y))
              tf: POP  (arg(Y, X))
              tg: POP  (arg(X, P))
              (ta < tb < tc < td < te < tf < tg)

Update:      tn: MAKE_SALIENT (Ag, break, arg(X, P))
              (ta < tn < tg)

```

Figure 5.8 The PUNC-BREAK heuristic

The saliency goal uses the unique atom break to indicate the need for a paragraph break; the context parameter ensures that it occurs at the correct level of abstraction (i.e. avoiding the break occurring

within or between the subarguments). The fact that the `MAKE_SALIENT` goal is an intentional goal also permits the result to be read in a similar way to that proposed by Pascual – the abstract goal `tn` might be read something like “the author identifies a new component supporting *P*, and this component is at the same level as *X*”.

Footnotes also play an important role in indicating the structure of an argument – particularly a large argument with a rich structure of subargument support. Clearly, footnotes play a number of roles, including indicating terminological usage and referencing: the current work concentrates specifically upon their use for including non-essential information in a discourse. Footnotes are primarily used to avoid disrupting the coherence of an argument through inclusion of digressive, supplementary or background information, and it is this use which is captured by the heuristics presented here.

Hovy and Arens (1991) suggest that footnotes and other *offpage* formatting (appendices, side bars, marginal notes, etc.) are the result of the RST relation `BACKGROUND` – as are parentheses – with a greater or lesser amount of material included in the Satellite. However, actual footnote usage does not seem to be restricted in this way: offpage material may also be related through `ELABORATION`, `ANALOGY`, etc.), but in such a way that the resultant text (both in the body of the text and in the footnote) is coherent. There are thus two distinct issues to be tackled in footnote generation: (i) the importance or centrality of information to the current argument and (ii) the potential impact on coherency of eschewing footnote data. Within the current *Rhetorica* framework, the former is rather less amenable to analysis than the latter. In order to evaluate the importance of information, the system must represent such information in its knowledge base and then attach some indication to each saliency goal. This issue is discussed further in chapter seven; here, the problem is decoupled and formulated as an oracle which can be polled quite specifically for an evaluation. Thus the `PUNC-FN1` heuristic shown in Figure 5.9 determines footnote applicability on the basis of the result returned from the importance oracle.

PUNC-FN1

Trigger:	<code>ta: PUSH (arg(X, P))</code> <code>tb: MAKE_SALIENT (Ag, A, arg(X, P))</code> <code>tc: POP (arg(X, P))</code> <code>~∃ PUSH (arg(Y, X))</code> <code>importance(A) < fixed_value</code> <code>(ta < tb < tc)</code>
Update:	<code>tm: MAKE_SALIENT (Ag, footnote-start, arg(X, P))</code> <code>tn: MAKE_SALIENT (Ag, footnote-end, arg(X, P))</code> <code>(ta < tm < tn < tc)</code>

Figure 5.9 The `PUNC-FN1` heuristic

The formulation of `PUNC-FN1` correctly enforces the conclusion of a footnote to lie in the main text and the subargumentation in the offpage unit.

The second issue in footnote generation, coherency, is more amenable to characterisation, but like paragraph break introduction, relies on an ad hoc approximation of applicability. It seems that

several rules impinge on the decision over whether or not to place information in a footnote: (i) that the length of a footnote should not exceed one or two claims, (ii) that the information in a footnote should not play any further role in the main-text argument, (iii) that footnotes should not be introduced when the number of subarguments at the same level of abstraction is small. These general rules are fixed precisely by PUNC-FN2, which sets the maximum number of footnote claims at two, and the number of sibling components at four. These figures are open to dispute, but seem to represent a fair approximation for heuristic use. (Note that PUNC-FN2 uses square brackets in their conventional BNF sense of indicating optional components).

PUNC-FN2

```

Trigger:   ta: PUSH (arg(X, P))
           tb: PUSH (arg(A, X))
           tc: PUSH (arg(B, X))
           td: PUSH (arg(C, X))
           te: PUSH (arg(D, X))
           tf: MAKE_SALIENT(Ag, E, arg(D, X))
           [tg: MAKE_SALIENT(Ag, F, arg(D, X))]
           th: POP (arg(D, X))
           ti: POP (arg(X, P))
           (ta < tb < tc < td < te < tf < th < ti)

Update:    tm: MAKE_SALIENT (Ag, footnote-start, arg(D, X))
           tn: MAKE_SALIENT (Ag, footnote-end, arg(D, X))
           (te < tm < tn < th )

```

Figure 5.10 The PUNC-FN2 heuristic

5.2 *Explicit and Implicit Saliency*

An abstract plan comprises an ordered list of propositions to be made salient to the hearer under various constraints of attention shift. A claim does not, however, necessarily need to be expressed in order for it to be salient – this is the assumption in most work which employs a notion of saliency. Walker (1996) offers a good example of this assumption: in building her model of language generation based upon the hearer's cognitive abilities, she acknowledges the need to express information which may already be salient (the discussion returns to this phenomenon in the next section). Determining whether or not a proposition is already salient is clearly a complex task – Walker implicitly defines a notion of saliency based upon entailment and presupposition, i.e. a proposition *p* is salient iff (i) *p* has been uttered, (ii) *p* has been entailed by an utterance *u*, or (iii) *p* is presupposed by an utterance *u*. This is problematic in itself (since there is no principled route to determining whether some proposition is entailed or presupposed by another – other than by laboriously representing such facts explicitly), but in the context of argumentation, it is insufficient. This can be demonstrated by the endemic use of enthymemes, recognised – as indicated by the word's etymology – in classical texts on rhetoric (Aristotle, 1926). In classical analyses, an enthymeme is a syllogism (an argument of exactly two premises and one conclusion) with one component left implicit. Clearly, this implicit step is salient to the hearer (for otherwise the hearer would be unable to follow the argument). One particular form of enthymematic argument is found so commonly in natural argument, that it has led both to extensive

investigation in social psychology (Triandis, 1971, pp187-188); (McGuire, 1969, pp208-210), and also to separate linguistic characterisation and terminology: omitting one of the steps of a Modus Ponens or Modus Tollens is characterised by Sadock (1977) as *Modus Brevis*. Cohen (1987) employs Sadock's distinction in her computational analysis of argument structure, distinguishing each of the possible forms (where any one or two of the three components are omitted) in each of her 'rule of inference frames'. Importantly, she notes that particular forms of Modus Brevis are far more common than others (in a Modus Ponens, for example, the major premise $P \rightarrow Q$ is usually omitted). It is the responsibility of the EG to capture such enthymematic contraction in the most appropriate way: given the context of each subargument, the EG must determine which, if any, of the saliency goals might be completely removed from the plan.

The rules of enthymematic contraction are dependent upon a large number of factors. Firstly, they cannot violate hard constraints imposed by earlier planning at either AS or EG levels – if a saliency goal has been introduced with a hard constraint, the EG is not at liberty to delete it. Secondly, there are structural constraints imposed by considerations of argument coherency: the implication step of an MP, for example, may be the most common example of enthymematic contraction, but it is non-optional in situations where the MP is distanced from its conclusion by intervening argumentation, as in Figure 3.9, above. Thirdly, the beliefs of the hearer modelled by the speaker affect the potential for contraction, with each of various epistemic modalities associated with a different likelihood of enthymematic contraction: if a particular claim is already believed by the hearer, then it is likely to be left implicit; if the hearer is unaware of a claim, or is undecided (i.e. aware, but uncommitted to belief or disbelief) on the status of a claim, then it is less likely to be omitted unless it represents a canonical contraction (e.g. the omission of the major premise of a Modus Ponens); and finally, if a claim is disbelieved by the hearer then it is highly unlikely that leaving the claim implicit will result in persuasive argument. Fourthly, a related consideration of the hearer's epistemic state will affect the decision over whether or not to omit a particular claim. As discussed above in §3.1.6, beliefs that are deeply entrenched in the hearer's knowledge – i.e. beliefs for which alteration would involve changes in many other beliefs as a consequence – need to be made explicit (as well as also needing to be supported by a great deal of cogent argumentation). Fifthly, there are also the non-epistemic features of the hearer: if the hearer finds complex argument difficult to follow (i.e. the general competence parameter has a low value) it is less likely that omission of steps in the argument will be successful. Sixthly, there are other more general parameters describing the hearer which affect not only enthymeme contraction, but the precision of argument in general – if the hearer is particularly sceptical, or has reason to believe that the speaker has a lot to gain from winning the argument, the generation process has less freedom in omitting information, and needs rather to adopt a more diligent approach. Finally, other attributes of the wider context in which the argument is situated, such as media-specific constraints restricting column inches or presentation minutes, will affect decisions of content omission – if an argument needs to be short, one option is to omit intermediate steps.

It is this range of influences which need to be characterised within the EG level processing to determine whether any given saliency goal should be removed from the plan, on the understanding that it would be salient to the hearer even if not made explicit. Clearly, this is a complex task, and the

heuristics presented below offer a first cut, utilising a number of simplifications to facilitate an appropriate integration of the various features.

The first enthymeme contraction heuristic, EC-EPISTEMIC2, is founded upon the idea that if a hearer believes a proposition then, other things being equal, that proposition can be left implicit. The 'other things' which need to be considered are the general competence and scepticism of the hearer. One further constraint introduced is that at least two components of a syllogism must be expressed. This is somewhat stronger than the conventional restriction, which ensures the inclusion of at least one component, but, as pointed out by Cohen (1987), the omission of two components from an argument is much rarer than the omission of one. It is not clear what additional contextual features license a dramatic contraction where two components are omitted, so the current work restricts implementation to this more restrictive notion of enthymeme. When EC-EPISTEMIC2 is applied to a particular subargument in which inclusion of each component is registered as a soft constraint, the removal of one of those components results in the reregistering of the remaining components as hard constraints. This ensures that enthymeme contraction is only performed once on any given argument. (Notice that no extra explicit machinery is required to differentiate contraction of premises and conclusions, since the latter are simply handled by the same heuristics acting at a higher level of abstraction).

EC-EPISTEMIC2

```

Trigger:  ta: PUSH (arg(X, P))
          tb: MAKE_SALIENT (Ag, A, arg(X, P))           [soft]
          tc: MAKE_SALIENT (Ag, B, arg(X, P))
          td: POP (arg(X, P))
          te: MAKE_SALIENT (Ag, P, arg(P, Q))
          (ta < tb < td)
          (ta < tc < td)
          BEL (Ag, A)
          parameter(hearer_competence) > fixed_value1
          parameter(hearer_scepticism) < fixed_value2

Update:   DELETE (tb)
          tc: MAKE_SALIENT (Ag, B, arg(X, P))           [hard]
          te: MAKE_SALIENT (Ag, P, arg(P, Q))           [hard]

```

Figure 5.11 The EC-EPISTEMIC2 heuristic

As reflected by the numbering, EC-EPISTEMIC2 has a dual in EC-EPISTEMIC1 in which the claim removed is constrained to be of the form $a \rightarrow b$, for as mentioned above, it is the major premise of both an MP and an MT which is more susceptible to enthymematic contraction. The numbering also indicates that the omission of the major premise is to be preferred to the omission of the minor premise (that is, the enthymeme contraction heuristics are listed numerically in the Prolog database).

Related to the first two epistemically founded heuristics are two further which are rather weaker: EC-EPISTEMIC3 and EC-EPISTEMIC4 represent a similar licensing of contraction in situations where the hearer does not believe a claim (i.e. is unaware or undecided) – clearly, though, for EC-EPISTEMIC3 and EC-EPISTEMIC4, the thresholds of hearer competence (and scepticism) which license heuristic application are significantly higher (lower) than for EC-EPISTEMIC1 and EC-EPISTEMIC2.

The next heuristic, EC-LENGTH, widens the remit of enthymeme contraction to admit propositions upon which the hearer is undecided or unaware, but is only applied in the face of constraints on the absolute length of the argument. Again, EC-LENGTH needs two versions, one for deleting a premise (EC-LENGTH1, shown in Figure 5.12), and another for deleting a conclusion (EC-LENGTH2, with a similar modification to the *bel* goals and constraint updates).

EC-LENGTH1

Trigger:	ta: PUSH (arg(X, P))	
	tb: MAKE_SALIENT (Ag, A, arg(X, P))	[soft]
	tc: MAKE_SALIENT (Ag, B, arg(X, P))	
	td: POP (arg(X, P))	
	te: MAKE_SALIENT (Ag, P, arg(P, Q))	
	(ta < tb < td)	
	(ta < tc < td)	
	BEL (Ag, A) $\vee$ BEL (Ag, ?A) $\vee$ $\sim$ BEL (Ag, A)	
	parameter(length_constraint) > fixed_value	
Update:	DELETE (tb)	
	tc: MAKE_SALIENT (Ag, B, arg(X, P))	[hard]
	te: MAKE_SALIENT (Ag, P, arg(P, Q))	[hard]

Figure 5.12 The EC-LENGTH1 heuristic

Thus the current characterisation of enthymeme contraction considers most of the features listed above – the only exception is the dictum that deeply entrenched hearer beliefs should not be left implicit. As this rule presupposes the possibility that claims disbelieved by the hearer might be left implicit, the entrenchment rule is currently ignored (since leaving *p* implicit when the hearer is known to believe $\sim p$ is licensed by neither EC-EPISTEMIC nor EC-LENGTH).

5.3 Rhetoric

The content of an argument determined by the AS level is comprised of units which play a functional role in fulfilling the intentions of the speaker. Unsurprisingly, the view that such units can completely specify a given text has been predominant in NLG, and it is only recently that suggestions have been made that nonfunctional units – components which are in some way redundant – also have a part to play. There are two broad classes of such components: those which are generic, and those which are specific to persuasive text. Both of these classes seem to fit well under the traditional notion of ‘rhetoric’.

The former class has been investigated under the heading of *informationally redundant units* (IRUs) by Walker (1996). She distinguishes (p184) three separate roles that an IRU may play in a discourse: (i) “to provide evidence supporting beliefs about mutual understanding and acceptance”; (ii) “to manipulate the focus of attention of the discourse participants by making a proposition salient”; (iii) “to augment the evidence supporting beliefs that certain inferences are licensed”. The first class, attitude IRUs, are employed to indicate assent in a dialogue – and seem to be unique to dialogue, with no equivalent in monologue. Attitude IRUs are therefore not addressed in the current *Rhetorica* system. The second class, consequence IRUs, seem in fact to be the complement to the saliency issues

discussed in the previous section. For Walker, the uttering of the two premises of, for example, a Modus Ponens, logically entail their conclusion. As a result, the text will often then omit to make that conclusion explicit (as discussed by Sadock (1977)). Walker goes on to propose that the conclusion might still be made explicit – as an IRU. This contrasts with the approach adopted in the *Rhetorica* system, whereby a deep representation of the text plan includes each of the three components of a Modus Ponens, but that during heuristic processing, items may be deselected under certain conditions. These conditions which license the deletion of saliency goals from a plan in *Rhetorica* form a complementary analogue to those which permit the introduction of consequence IRUs in Walker's theory – how "hard" the inference is, how important it is, etc. The logical omniscience assumption made (though acknowledged to be unrealistic) in (Walker, 1996), is not made in *Rhetorica*, and so the inclusion of all parts of a syllogism is not seen in the current work as an IRU. The third and final class in Walker's taxonomy are the attention IRUs which, she claims, can be used to shift the current topic of the dialogue. The claim seems to involve two distinct points, one of which is subsumed in the current *Rhetorica* framework, and one of which violates the underlying assumptions of hierarchical argument structure discussed above. The first is that propositions – even those known to already be believed by the hearer – must be made salient to the hearer (this has been discussed in (Marcu, 1996a) and (Reed and Long, 1997b)). This is incorporated in AS processing, and although such saliency goals are at a greater risk of being pruned from the plan, they are not considered to be redundant since the focus is upon creating a connected series of salient claims – which does not exclude claims already believed by the hearer. The second is that claims may be *repeated* to bring propositions back into focus. One of the central tenets of the planning in *Rhetorica* is that such repetition is unnecessary if the argument is well structured: this second use of IRUs is thus prohibited.

In summary, then, Walker's attitude IRUs are beyond the scope of the current work because they are unique to dialogue, consequence IRUs are not considered redundant, but are licensed in a similar way, and attention IRUs are either effected by saliency and topic manipulation planning at the AS level, or else violate the principles on which it is founded and are therefore not considered. There are, however, further textual units which seem redundant on a wider definition of functionality – that a given claim must form a part of the textual structure – or in the current work, of the argument structure. Classical works on rhetoric and public speaking offer an enormous range of advice of which only a relatively small part pertains to the construction of the core argument; a great deal of time is spent on more generic suggestions, which amount to far more than a catalogue of anodyne comments of the "Unaccustomed as I am to public speaking..." ilk. The current work examines only a small number of these suggestions – importantly, *Rhetorica* concentrates solely upon the argumentative part of persuasive discourse, leaving the phases of introduction, division, narration, pathetic and peroration for future work. From Blair's (1838) discussion of these sections, it is clear that it is the pathetic part which is the primary location of nonfunctional textual units. Some examples, however, do occur within the argumentative phase which forms the remit of the *Rhetorica* system. These 'rhetorical' components of an argument do not offer support through any of the means (deductive, inductive or fallacious) established at the AS level. Rather, they represent a body of 'tricks' for improving an argument's persuasive effect. One prime example of their use is in establishing terminal premises. Rescher

(Rescher, 1997) suggests, indeed, that rhetoric (as distinguished from rational argumentation) is always employed in this way, since there is no alternative for establishing premise validity (appealing to existing hearer beliefs is in his view a rhetorical appeal). A typical example of a rhetorical appeal for premise establishment is something of the form *I wouldn't insult your intelligence by...* or *I realise you understand that....* One famous example is offered in Malthus' *Essay on the Principle of Population*, (in (Fisher, 1988, p30)) where, his initial claims are prefaced by "I think I may fairly make two postulata...".

Such rhetorical components are very common, and thus need to be accounted for in any large-scale model of argumentation. The introduction of such components could proceed in one of two ways – either by marking a requirement for the generation of such a unit at a high level (where notions of support and subargument are available), or else introducing the lexical phrase at a high level, and then constraining subsequent processing to fit in with that lexical decision. The former approach is adopted for two reasons – the first is consistency with the remainder of the framework (particularly clue introduction discussed in §5.1, and lexical tagging discussed in §5.4), and the second is that a high level plan with a number of lexical constraints cannot be guaranteed to have a grammatically correct realisation that can be found by lower level processing (this problem is analogous to Panaget's argument, discussed in §3.1.2). The conditions under which a rhetorical premise establishment phrase is introduced are simple – that the hearer is undecided with regard to the value of a proposition and is not sceptical, that a similar component has not been used recently, and that no non-rhetorical support for the claim is available. These constraints for the introduction of premise establishment by a rhetorical unit are formulated in RHETORIC-PE:

RHETORIC-PE

```

Trigger:    ta: MAKE_SALIENT (Ag, P, arg(P, Q))
            ~∃ MAKE_SALIENT (Ag, X, arg(X, P))
            BEL (Ag, P) ∨ BEL (Ag, ?P) ∨ ~BEL (Ag, ~P)
            distance_between(ta, tb) > fixed_value
            parameter(scepticism) < fixed_value

Update:     tb: MAKE_SALIENT (Ag, rhetoric-pe(P), arg(P, Q))

```

Figure 5.13 The RHETORIC-PE heuristic

The other use of rhetorical text components in *Rhetorica* is in the generation of qualifiers: phrases which indicate a higher or lower level of confidence in a claim or inference link. As explained in §3.1.6, the strength of a belief is not explicitly represented in *Rhetorica* – the qualifiers introduced by the following heuristics are based upon the type of argument involved and, relatedly, the anticipated reception of that argument. Thus conclusions which are known to be in contradiction with beliefs held by the hearer may be introduced with a less confrontational qualifier such as *it seems to me that* or *and so probably*, etc. Similarly, claims which are based on deductive arguments, and which lead to conclusions which the hearer does not already disbelieve often involve strong qualifiers such as *it's clearly the case...* or *without a doubt....* These qualifiers do not function as clues since they are not indicating argument structure and are consequently not the result of the topic manipulation operators in

through one of the seven stylistic dimensions he proposes (namely, slant). The other dimensions (formality, euphemism, archaic, floridity, abstractness and force) follow in spirit the work on computational stylistics of DiMarco, Green, Hirst, *et al.*, (DiMarco *et al.*, 1992), (Green and DiMarco, 1996), but as a primarily lexically-oriented problem, falls beyond the scope of the current work.

The use of loaded phrases has also been explored in Hovy's PAULINE system (Hovy, 1990) through the broad concept of *affect*. The PAULINE system distinguishes three values of affect: GOOD, BAD and NEUTRAL. Each concept represented in the PAULINE knowledge base is either inherently marked with one of these values (Hovy's example is that in a neutral context, the concept ARREST has the affect marking BAD), or else may occur in the system's lists of sympathies or antipathies. In characterising the notion of affect, the *Rhetorica* system differs from PAULINE in a number of important respects. In the first place, PAULINE employs three *topic collection plans* (namely, CONVINCE, DESCRIBE and RELATE), whereas only one of these (CONVINCE) has a parallel in *Rhetorica*. The process of topic collection (which for Hovy's CONVINCE is given at (Hovy, 1990, p185)), is not seen in the current work as a rigid, predefined specification, but rather, as a crucially important planning task – this task represents the complete functionality of the AS level. Secondly, where sympathies and antipathies need to be specified manually by the user in PAULINE, *Rhetorica* makes the assumption that all beliefs included in the argument have the affect value NEUTRAL, and all contradictory beliefs held by the hearer have the affect value BAD. Finally, where PAULINE includes an array of slanting techniques for injecting affect values into text, the *Rhetorica* system teases the techniques apart (with topic collection being planned at the AS level, clue-based structure introduction being handled by the EG level, etc.) leaving only Hovy's last point – “simply *enhance or mitigate the topic* with sensitive aspect – just say it and allow subsequent realization decisions to give it the appropriate slant” – to be dealt with here. It is proposed that this point is rather less simple than Hovy claims, both with respect to realisation (which is not explored in the current work) and to introduction (which is).

The introduction of appropriate loading onto saliency goals is dependent upon consideration of the persuasive intentions of the speaker in the context of the abstract representation, since the appropriate affect of a given claim is dependent upon its role in the structure of the argument. The introduction of loading must be integrated into the generation of that argument structure, and needs, therefore, to be co-ordinated by the EG level. Introduction of affect is performed through a markup of the appropriate saliency goal represented as *affect* (tx, value) where value is one of the possible degrees of affect – in the current work, Hovy's broad categorisation of affect into GOOD, BAD and NEUTRAL is adopted because, as he points out, the “three values are sufficient to give the program interesting behaviour” (p182). Where affect is not explicitly specified as either GOOD or BAD, NEUTRAL is assumed by default. As in Hovy's work, a number of propositions may initially be registered as conveying good or bad affect – these are essentially concepts which the speaker believes will meet with general agreement in respect of their affect (examples might be liberty, justice, war, cruelty, etc.). Any saliency goals introduced in the planning phases will inherit the affect of the propositions they make salient, if such affect is specified.

There are a range of situations which typically demand clausal realisation with a particular

slant, and a number of these are formulated in the following heuristics. The first, AFFECT-UC, characterises the use of negative affect in mentioning the claims in a counterargument. Thus the undercut premise in a UCP argument, or undercut inference in a UCI argument, will be realised with an affect marking of BAD. The appropriate saliency goal can be identified by exploiting the fact that such goals have a characteristic form – namely, that a proposition is being made salient in the context of an argument which employs the negation of that proposition as its premise.

AFFECT-UC

Trigger: ta: MAKE_SALIENT (Ag, X, arg(~X, P))
 Update: affect(ta, bad)

Figure 5.15 The AFFECT-UC heuristic

The AFFECT-GOOD series of heuristics control a broad introduction of positive affect markers under contextual conditions similar to those proposed by Hovy (p183). AFFECT-GOOD1 specifies that indicating the saliency of a proposition believed by the hearer can be marked with GOOD. AFFECT-GOOD2 specifies that any proposition, whether believed or disbelieved by the hearer (or for which the hearer is undecided or unaware) can be marked with good if the relationship to be maintained or created between speaker and hearer is not close (a high value for this figure indicates closeness).

AFFECT-GOOD1

Trigger: ta: MAKE_SALIENT (Ag, X, arg(X, P))
 BEL (Ag, X)
 Update: affect(ta, good)

AFFECT-GOOD2

Trigger: ta: MAKE_SALIENT (Ag, X, arg(X, P))
 relationship < *fixed value*
 Update: affect(ta, good)

Figure 5.16 The AFFECT-GOOD heuristics

The affect of particular saliency goals also interacts with other goals, particularly with regard to appropriate cues and clues. A good example of such interaction is in the disjunctive syllogism construction, wherein situations of bad affect associated with one component of the disjunction, and good or neutral with the other, the two premises are linked with a negative polarity cue if the major premise precedes the minor premise. This can be demonstrated through consideration of the simple Examples (5.12) to (5.15). In the first, both disjuncts are presumed to carry bad affect, in the third, both have good affect, and in the second and fourth, one carries good, the other bad (and in the second, good precedes bad, in the fourth, vice versa).

Films are always either $\left\{ \begin{array}{l} \surd \text{and} \\ \# \text{but} \end{array} \right\}$ this one isn't long,
too long or too violent so it'll be violent (5.12)

Films are always either $\left\{ \begin{array}{l} \surd \text{and} \\ \surd \text{but} \end{array} \right\}$ this one isn't informative,
informative or violent so it'll be violent (5.13)

Films are always either $\left\{ \begin{array}{l} \surd \text{and} \\ \# \text{but} \end{array} \right\}$ this one isn't informative,
informative or thrilling so it'll be thrilling (5.14)

Films are always either $\left\{ \begin{array}{l} \surd \text{and} \\ \surd \text{but} \end{array} \right\}$ this one isn't violent,
violent or informative so it should be informative (5.15)

The involvement of affect in the clue between DS components is captured in CLUE-DS, shown in Figure 5.17.

CLUE-DS

Trigger: ta: PUSH_TOPIC (arg(X, P))
 tb: IS_SALIENT (Ag, X $\vee$ P, arg(X, P))
 tc: IS_SALIENT (Ag, $\sim$ X, arg(X, P))
 td: POP_TOPIC (arg(X, P))
 (ta < tb < td)
 (ta < tc < td)
 (tb < tc)

Update: tn: MAKE_SALIENT (Ag, CLUE(polarity, negative), arg(X, P))
 (tb < tn < tc)

Figure 5.17 The CLUE-DS heuristic

As *Rhetorica* is not directly concerned with realisation issues in general, it also defers decisions over how to realise affect itself until subsequent sentence and lexical planning. The processing required to map from an indication of affect to its manifestation in text (through devices such as juxtaposition and word choice) are discussed in (Hovy, 1990, pp187-190).

5.5 Pulling it Together

The example examined in §4.4 is here continued and expanded using the range of heuristics deleting goals and introducing cues, clues, punctuation, nonfunctional text and affect markings, discussed in §5.1 to §5.4. The plan in Figure 4.27 (or rather, the intermediate abstract plans generated during its construction, such as that in Figure 4.26) licenses the application of a number of the heuristics discussed above.

Starting with cue introduction, the clearest example of cue use is *but* in the phrase *The only excuse for visiting such suffering on other creatures can be ignorance. But that exit is closed to us...* Analysing the structural context shows that the construction is based on DS – the key is in the use of the word *only*³³ – either there exists this one excuse, or else we are visiting suffering on animals.

³³ The generation of *only* is clearly the result of lower level realisation issues as it is used to relate the two-propositions of the compound disjunction *ignorance $\vee$ cruelty*.

With affect on suffering registered as bad, and none registered on ignorance, the CLUE-DS heuristic is licensed, instantiated as shown in Figure 5.18.

CLUE-DS (arg(h, a))

Trigger: ta: PUSH_TOPIC (arg(h, a))
 tb: IS_SALIENT (Ag, ~h, arg(~h, a))
 tc: IS_SALIENT (Ag, a ∨ h, arg(~h, a))
 td: POP_TOPIC (arg(~h, a))
 (ta < tb < td)
 (ta < tc < td)
 (tc < tb)

Update: tn: MAKE_SALIENT (Ag, CLUE(polarity, negative), arg(~h,a))
 (tc < tn < tb)

Figure 5.18 The CLUE-DS instantiation in the vegetarianism argument

The other cue introduction dependent upon a particular operator occurs earlier in the argument, with the counter-counterargument between $\sim d$ and b . With the implication step, $d \rightarrow \sim b$, preceding the premise, $\sim d$, in the UCP (this is enforced by HCR1), the CLUE-UCP heuristic is triggered, introducing a negative polarity, pragmatic cue (of which *Admittedly...but* was the example given above). Subsequent contraction, through the application of EC_EPISTEMIC3, removes the implication $d \rightarrow b$, and subsequent realisation is then forced to use the more general cue *but* at the start of paragraph two: *But it's fine because they spend their lives ...*

Notice however, that by this account alone, the sentence should read something like *But animals want more than just food and shelter* – expressing $\sim d$. The actual realisation expresses d – though the intention of the speaker is that the hearer read the claim as $\sim d$: this is an example of irony. Although a characterisation of this complex, subtle phenomenon (Ito and Takizawa, 1995) is beyond the scope of this work, its introduction at least can be accounted for at this point in the text by consideration of AFFECT-UC, which indicates that the denied premise should, if realised, be expressed with bad affect. It seems reasonable that such bad affect could be realised through the use of irony.

There are also several situations where the application of clue heuristics lead to the introduction of clues in the text. Interestingly, however, explicit clues of inference and detail are surprisingly rare in the argument. The primary cause for this is that many of the conclusions which form components of the argument structure (a , b , and d) are not explicitly realised. This means that the saliency goals associated with these claims have been removed by one of the EC heuristics (this is discussed further below), and as these heuristics are preferred over those of clue introduction, CLUE-DETAIL and CLUE-INFERENCE will not fire (as they both require the presence of the appropriate saliency goal). As a result, there are only two points where such clues would be expected. The first is in concluding g following the expression of f . The link between these two claims is unfortunately clouded by the involvement of two rhetorical questions, the use of which is a particularly troublesome problem in the analysis of argument (Slot, 1993). Ignoring, however, the repetition (the two questions are making the same propositional point), and the use of questions to make the claim, the leading *so* (in *So*

a: eating meat wrong; b: animals suffer; c: tethered pigs go mad; d: no animal wants more than food and shelter; e: doing the same to dogs is prosecutable; f: pigs as intelligent as dogs; g: no explanation for two sets of rule; h: excuse is ignorance; i: we do know the truth

why the difference?) still remains, as can be demonstrated in a rewording which paraphrases the rhetorical questions: *So there is no logical explanation for there being two sets of rules*. This is propositionally equivalent, coherent, and precisely to be expected given the CLUE-DETAIL heuristic application.

The second example of clue introduction occurs between the subargument from *i* to its conclusion $\sim h$. This argument exhibits pre-order (unlike the previous clue introduction) and is therefore liable to trigger CLUE-INFERENCE: the result becomes manifest as the lexeme *because* found between the statement of $\sim h$ (*But that exit is closed to us ...* – where *that exit* is an anaphoric reference to *h*, the claim to ignorance) and its justification *i* (*... we do know the truth*).

Next are clues associated with punctuation. There are three paragraph breaks in the extract, between the four paragraphs expressing (i) MP(*c*, *b*), (ii) UCP(*d*, *b*) (iii) MP(*e*, *d*) and MP(*g*, (*e*→*d*)), (iv) DS(*h*, *a*). One of these – the last – is generated by the application of PUNC-BREAK, instantiated as shown in Figure 5.19.

PUNC-BREAK

```

Trigger:      ta: PUSH (arg(b, a))
               tb: PUSH (arg(d, b))
               tc: PUSH (arg(g, d))
               td: PUSH (arg(f, g))
               te: POP  (arg(f, g))
               tf: POP  (arg(g, d))
               tg: POP  (arg(d, b))
               (ta < tb < tc < td < te < tf < tg)

Update:       tn: MAKE_SALIENT (Ag, break, arg(b, a))
               (ta < tn < tg)

```

Figure 5.19 The PUNC-BREAK heuristic instantiated in the vegetarianism argument

The remaining two breaks however, cannot be explained through application of PUNC-BREAK. There are two possible routes to explaining their generation. The first is to relax the conditions of PUNC-BREAK to permit introduction of breaks at lower levels of nesting – reducing the levels of nesting to two would neatly account for the required breaks. Although the article as a whole seems to use rather short paragraphs, extending the applicability of PUNC-BREAK would often lead to inappropriately short paragraphs (particularly when information is left implicit). An alternative route is to view paragraph breaks as one of the possible realisations for other types of clues. Adopting this approach offers a less drastic solution, and is also consistent with the realisation model of Meteer (1991) and Panaget (1994), discussed above (since there is no reason why formatting devices such as paragraph breaks should not characterised – at least in part – as abstract linguistic resources). For the first paragraph break, the clue is parallel: in the earlier paragraph, the argument is from *c* to *b*, in the later, from $\sim d$ to *b*. The second paragraph break separates the conclusion *d* from its subsequent supports, the *e-f-g* complex; the clue is thus one of inference.

Of course, the disadvantage of this approach is that further functionality is being deferred to

a: eating meat wrong; b: animals suffer; c: tethered pigs go mad; d: no animal wants more than food and shelter; e: doing the same to dogs is prosecutable; f: pigs as intelligent as dogs; g: no explanation for two sets of rule; h: excuse is ignorance; i: we do know the truth
--

sub EG processing; the decision seems warranted however by consideration of the alternatives the realisation process has available. For the first paragraph break, there exists the standard range of parallel clues – one possible result would be *animals raised for meat suffer, firstly because pigs are tethered in barren stalls, etc., and secondly because all animals want more than just food and shelter*. However, as mentioned above, the *but* introduced as a result of the DS construction, and the bad affect resulting in irony – both of which are sub EG realisation issues – prevent the use such parallel clues (it is unclear, for example, how the above construction could be accommodated after the but+irony decision). To identify the parallel structuring, the only option left open (to sub EG planning) is the introduction of a paragraph break.

The complexity – and particularly the ironical slant – of the expression of *~d* similarly bars the employment of a simple inference clue such as *because* or *for* between *~d* and the *e-f-g* complex (*But it's fine because no animal wants more than food and shelter. For if a person did the same thing to dogs etc.* is substantially incoherent).

The example also exhibits the application of a range of enthymematic contraction heuristics: the markup of the argument shown below demonstrates, when compared with the list of MAKE_SALIENT plan of primitives in Figure 4.26, that barely half of all the steps in the argument are actually realised into text:

[Pigs are tethered in barren stalls, so deprived of stimulation that they often go mad. They are forced to have as many litters of piglets as their bodies will stand but they are not allowed to mother them.]_c

But [it's fine because they spend their lives in buildings sheltered from the elements with all the food they need and no animal wants more than that.]_{-d} Well, that's what we're told.

[If a person did the same thing to dogs he or she would be prosecuted for cruelty]_e – and yet [a pig is equally as intelligent as a dog.]_f So why the difference? [Why do we have two sets of rules, one for a dog and one for a pig? There is no logical explanation.]_g

[The only excuse for visiting such suffering on other creatures can be ignorance.]_a ∨ _h But [that exit is closed to us]_{-h} because [we do know the truth.]_i We might pretend we don't, we might say that to give up pork³⁴ is pointless as one person's abstinence doesn't make any difference. But really we know it's wrong and that knowledge makes us push the reality into the darker reaches of our minds.

³⁴ Notice that the word *pork* replaces the phrase *battery eggs* found in the original argument. This is because an intervening paragraph in which the topic shifts from the suffering of pigs to that of battery hens occurs in the original argument. With this paragraph deleted in the pursuit of brevity in the analysis, the lexical swap is required to maintain the coherency of the whole. The modification has no significant impact on the analysis of the structure, and subsequent discussion of the generation of that structure.

a: eating meat wrong; b: animals suffer; c: tethered pigs go mad; d: no animal wants more than food and shelter; e: doing the same to dogs is prosecutable; f: pigs as intelligent as dogs; g: no explanation for two sets of rule; h: excuse is ignorance; i: we do know the truth

The majority of the implicit claims are the result of EC-EPISTEMIC3, i.e. the contraction of implication steps in cases where the hearer has no reason to disbelieve the inference. This accounts for $b \rightarrow a$, $c \rightarrow b$, $e \rightarrow d$, $f \rightarrow g$, $g \rightarrow (e \rightarrow d)$, and $i \rightarrow \sim h$. The inference step $d \rightarrow \sim b$ (in the UCP argument) is left implicit by the application of EC-EPISTEMIC1 (because the hearer believes this step). This leaves three further implicit claims unaccounted for: b , d and a .

Although syntactically, the first sentence of the second paragraph appears to be conveying d , as explained above, the function of the sentence is in fact to convey the opposite, $\sim d$, through the use of irony. The phrase *Well, that's what we're told* emphasises this irony. Leaving d implicit could be explained simply through the application of EC-EPISTEMIC2 (omitting information known to the hearer), but it seems more likely that in fact it is a result of the use of irony: it would be clumsy – and probably incoherent – to utter d twice, once to indicate its role in the argument held by the hearer, and once to act as the defeat of that very argument.

This irony also seems to be responsible for the enthymematic contraction of b : if $\sim d$ was expressed directly, it seems reasonable to propose that that paragraph could close with some statement such as *There seems little doubt that these animals suffer horrendously*. But with $\sim d$ expressed through an ironic statement, such a claim seems out of place. To adequately account for this contraction, therefore, a fuller account of the generation and role of irony in persuasive text is required, and such an account is beyond the scope of the current work.

There is no single reason motivating the decision to leave a implicit: rather, there seem to be a number of related factors. In the first place, the argument is an extract from a much longer piece, and a serves as the conclusion for the entire text, and amongst complex closing rhetoric, the author comes very close to stating a directly (e.g. *[by becoming vegetarian] you are no longer responsible for most of the daily cruelties...*). In addition however, it seems that the argument from b to a (something that is cruel is morally wrong) is assumed by the author to be almost a truism, and that claiming b is tantamount to claiming a . This is probably an example of the stronger form of enthymematic contraction not currently handled by *Rhetorica*, whereby two (from three) components of an argument may be left implicit. The use of such contraction is presumably restricted to cases such as this, where the speaker can reliably judge that the hearer can infer an entire argument from a single (minor) premise.

Finally, the argument also offers an example of the successful use of rhetorical premise establishment, through the application of RHETORIC-PE. The claim i , that *we do know the truth*, is not supported by any propositional components, and yet a significant amount of text follows the initial expression of the claim. Of particular interest are the phrases *We might pretend we don't*, and *But really we know it's wrong* – these are offering superficial support for i , but are clearly insubstantial. These phrases are typical examples of rhetorical appeals in support of basic premises, and are the result of the introduction of the saliency goal `MAKE_SALIENT(ag0, rhetoric-pe(i), arg(i, ~h))` by the RHETORIC-PE heuristic.

The extended example presented in this chapter serves to motivate each of the components of the *Rhetorica* system, but it does not provide concrete examples of the input, processing and subsequent output of the system. Three such examples are provided in the next chapter.

In drawing together the various control aspects – both those acting on the deep structure of an argument, and those reflecting the demands of coherency and persuasive effect on the surface structure – it has become clear that there are many areas which need further work to enrich the model and lead its generative capacity towards text which is more natural and more persuasive. An overview and discussion of these areas, and the ramifications of their current under-development is set out in chapter seven.

VI

System Output

To demonstrate the functionality of *Rhetorica*, three examples of system processing are presented, using arguments from the corpus as a basis: by analysing the structure of naturally occurring arguments, the system input is realistic and system output can be compared to a real-world standard. The three arguments were selected for their length (long enough to manifest interesting structure, but short enough to avoid uninformative repetition and tedious analysis), but not for any a priori suitability for system processing.

The first argument selected was the 'tourist facility signs' argument from *The Guardian*, Letters to the Editor, Saturday 19th October 1996. It is reproduced below, annotated with a mark-up to indicate premise units (note that here, as elsewhere, all punctuation is retained):

Since December 1995, when the Department of Transport relaxed the type of premises that could be sign-posted on motorways and trunk roads, increasing numbers of "tourist facility" signs have appeared. They signpost facilities such as pubs, restaurants, shops and nurseries on an apparently permanent basis.

[Such signs add to roadside clutter]_c and, through [their fixture on the supports holding road signs]_d, [distract motorists]_e by [competing with essential information]_f. The problem is aggravated when [the signs have the colour and form of the traditional "temporary" AA signs]_g – [which commanded attention for the information they carried]_h.

[There appears to be nothing to stop any concern from joining the rush to promote itself in this way]_i. [Such signs' siting and design needs much more careful thought]_j. The highway is no place for advertisements masquerading as traffic signs.

The analysis of this argument, using the mark-up above, is shown in Figure 6.1. The two sentences of the first paragraph, and the last sentence of the last paragraph are not components of the argumentative phase of the text, but rather, stand as exordium (or possibly narration) and peroration, respectively. During the reconstruction of the argumentative part of the text, one enthymeme was identified in which the premise was omitted. This premise forms the proposition *l*, "Tourist facility signs command attention".




Figure 6.1 Structure of the 'tourist facility signs' argument

The beliefs of the speaker which capture this structure are thus as shown in Figure 6.2. The hearer is assumed to be of high technical competence, have low scepticism, and to hold no beliefs with regard to the subject matter (i.e. the *Rhetorica* system assumes that the hearer will accept all inferences and premises as there is no information in the hearer model to the contrary). The speaker also represents the fact that there is a topical similarity between propositions *c* and *i* (property *r* might be characterised as 'clutteredness').

<code>bel(speaker, c).</code>	<code>bel(speaker, supports(g,l)).</code>
<code>bel(speaker, f).</code>	<code>bel(speaker, supports(h,supports(g,l))).</code>
<code>bel(speaker, d).</code>	<code>bel(speaker, supports(f,e)).</code>
<code>bel(speaker, g).</code>	<code>bel(speaker, supports(l,f)).</code>
<code>bel(speaker, h).</code>	<code>bel(speaker, supports(i,j)).</code>
<code>bel(speaker, l).</code>	<code>bel(speaker, supports(e,j)).</code>
<code>bel(speaker, e).</code>	<code>bel(speaker, has_property(c, r)).</code>
<code>bel(speaker, i).</code>	<code>bel(speaker, has_property(i, r)).</code>
<code>bel(speaker, j).</code>	
<code>bel(speaker, supports(c,j)).</code>	<code>parameter(scepticism, low).</code>
<code>bel(speaker, supports(d,f)).</code>	<code>parameter(general_competence, high).</code>

Figure 6.2 System beliefs and parameters for the 'tourist facility signs' argument

The full system processing for the tourist facility signs argument is given in Appendix C: it is there that intermediate output of the developing plan and its partial order from one level of abstraction to the next, and the firing of various heuristics is recorded. Here, in Figure 6.3, only the final output comprising a fully ordered, fully instantiated list of primitive operators is shown.

```

Complete plan of primitives:
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
9.push_topic(v_topic=supports(c,j),v_parent=supports(r,j))
11.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
14.pop_topic(v_topic=supports(c,j),v_parent=supports(r,j))
15.push_topic(v_topic=supports(i,j),v_parent=supports(r,j))
17.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
20.pop_topic(v_topic=supports(i,j),v_parent=supports(r,j))
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
21.push_topic(v_topic=supports(e,j),v_parent=null_topic)
59.make_salient(v_agent=ag0,v_proposition=break,v_topic=supports(e,j))
23.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
36.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(e,j))
28.push_topic(v_topic=supports(f,e),v_parent=supports(e,j))
30.make_salient(v_agent=ag0,v_proposition=f,v_topic=supports(f,e))
51.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(f,e))
37.push_topic(v_topic=supports(d,f),v_parent=supports(f,e))
39.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,f))
42.pop_topic(v_topic=supports(d,f),v_parent=supports(f,e))
50.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(f,e))
43.push_topic(v_topic=supports(l,f),v_parent=supports(f,e))
45.make_salient(v_agent=ag0,v_proposition=l,v_topic=supports(l,f))
60.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(l,f))
52.push_topic(v_topic=supports(g,l),v_parent=supports(l,f))
54.make_salient(v_agent=ag0,v_proposition=g,v_topic=supports(g,l))
61.push_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
63.make_salient(v_agent=ag0,v_proposition=h,v_topic=supports(h,supports(g,l)))
66.pop_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
57.pop_topic(v_topic=supports(g,l),v_parent=supports(l,f))
48.pop_topic(v_topic=supports(l,f),v_parent=supports(f,e))
33.pop_topic(v_topic=supports(f,e),v_parent=supports(e,j))
26.pop_topic(v_topic=supports(e,j),v_parent=null_topic)

laden with affect on the following goals:
[]

```

Figure 6.3 Final *Rhetorica* output for the ‘tourist facility signs’ argument

In order to examine this output in a more intuitive format (and to perform system evaluation, as discussed in chapter seven), it is possible to generate text associated with this structure by hand. Although manual, the process is essentially one of assuming the original textual formulation to represent canned text for each proposition: the only alterations are those necessary to maintain valid syntactic form. All other changes – punctuation, ordering, clues, and affect – are the result of elements of the final plan of primitives generated by *Rhetorica*. Other features of a text (for example, components which form exordium and peroration) are left entirely unchanged. Thus the plan in Figure 6.3 is realised into the following text:

Since December 1995, when the Department of Transport relaxed the type of premises that could be sign-posted on motorways and trunk roads, increasing numbers of "tourist facility" signs have appeared. They signpost facilities such as pubs, restaurants, shops and nurseries on an apparently permanent basis.

The siting and design of tourist facility signs needs much more careful thought. In the first place, such signs add to roadside clutter – there seems to be nothing to stop any concern from joining the rush to promote itself in this way.

More importantly, these signs distract motorists, because they compete with essential information through their fixture on the supports holding road signs. The problem is aggravated by tourist facility signs commanding attention from drivers because the signs have the colour and form of the traditional AA “temporary” signs which carry truly important information.

The highway is no place for advertisements masquerading as traffic signs.

There are several points to note about the realisation shown above. First, the detail clue at goal 8 is realised as *In the first place*. Second, the break clue at goal 59 is realised as a straightforward paragraph break between the second and third paragraphs. Third, the detail clue at goal 36 is realised simply as *because*. Fourth, following the text of the original, the detail clue at goal 51 is realised as *through*. Fifth, again as in the original, the parallel clue at 50 is realised as *The problem is aggravated by*. Lastly, the detail clue at 60 is realised as *because*.

The second argument selected was the ‘Clare Short’ argument from *The Guardian*, Letters to the Editor, Saturday 19th October 1996, reproduced below:

For Clare Short, the wait is over. Sadly, however, if Toby Graham had not searched for her, she would still be suffering her painful loss in silence, [as are many other birth parents]_b. [The Contact Register is only of limited value]_c as [so few know of its existence]_d. [Is it not time that the law concerning contact between adults after adoption is revised to enable the birth parent to have identifying information once the adoptee is 18 or possibly 25?]_e

The structural analysis of this argument is shown in Figure 6.4, which involves an enthymeme, premise *a*: “Contact can only be made if the adoptee searches out their birth parents”. The first sentence and the first half of the second play the role of exordium and are therefore excluded from the analysis, and included wholesale in both artificially generated arguments.




Figure 6.4 Structure of the 'Clare Short' argument

(Note that Figure 6.4 makes use of two nonstandard diagrammatic devices: (i) the '=' serving simply to link a proposition and its negation, and (ii) the dotted arrow – from $\sim c$ to a indicating a support arc assumed to be maintained by the audience).

From this, the beliefs in Figure 6.5 are easily inferable (as with the preceding argument, the hearer is assumed not to hold beliefs on the issue, and has the same high competence and low scepticism).

```

bel(speaker, e) .
bel(speaker, a) .
bel(speaker, supports(a, e)) .
bel(speaker, b) .
bel(speaker, supports(b, a)) .
bel(speaker, c) .
bel(speaker, supports(d, c)) .
bel(speaker, d) .
bel(ag0, supports(not(c), not(a))) .

```

Figure 6.5 System beliefs for the 'Clare Short' argument

The final output of the *Rhetorica* system running with the input in Figure 6.5 is shown in Figure 6.6 (and complete, level by level, output can be found in Appendix C). Again, the output concludes with a completely ordered plan, this time also including a list of affect laden goals, due to an application of the AFFECT-UC heuristic.

```

Complete plan of primitives:
2.make_salient(v_agent=ag0,v_proposition=e,v_topic=null_topic)
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(a,e),v_parent=null_topic)
6.make_salient(v_agent=ag0,v_proposition=a,v_topic=supports(a,e))
27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(a,e))
19.push_topic(v_topic=supports(c,a),v_parent=supports(a,e))
20.make_salient(v_agent=ag0,v_proposition=not c,v_topic=supports(c,a))
28.push_topic(v_topic=supports(d,c),v_parent=supports(c,a))
30.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,c))
33.pop_topic(v_topic=supports(d,c),v_parent=supports(c,a))
24.pop_topic(v_topic=supports(c,a),v_parent=supports(a,e))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(a,e))
13.push_topic(v_topic=supports(b,a),v_parent=supports(a,e))
15.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
18.pop_topic(v_topic=supports(b,a),v_parent=supports(a,e))
9.pop_topic(v_topic=supports(a,e),v_parent=null_topic)

laden with affect on the following goals:
[[20,bad]]

```

Figure 6.6 Final *Rhetorica* output for the 'Clare Short' argument

The text associated with this output is:

For Clare Short, the wait is over. Sadly, however, if Toby Graham had not searched for her, she would still be suffering her painful loss.

It is high time that the law concerning contact between adults after adoption is revised to enable birth parents to have identifying information once the adoptee is 18 or perhaps, 25. This change is required because at the moment, contact can only be made if the adoptee searches out their birth parents. The Contact Register is hopelessly limited as so few know of its existence, leaving many birth parents to suffer in silence.

Again there are several points of note in this realisation. First, the detail clue at goal 12 is realised as *because* (and requires some grammatical reformulation to accommodate it – namely the phrase *This change is required*). The detail clue at goal 27 is not easily realised given syntactic constraints (i.e. those constraints which impinge upon the functioning of any realisation component in regard to the generation of text on the basis of the hard, and in this case, soft plan components generated by *Rhetorica*). As a result, the clue was left simply as a sentence break. The parallel clue of 26 is similarly constrained. The bad affect on goal 20, introduced by the UCP argument in 11, is realised using *hopelessly*.

The final argument selected is the 'Irish census' question from the Letters page of *The Guardian*, from Wednesday, 23rd October, 1996:

Tomorrow, the Office of National Statistics will take a decision affecting Irish people in Britain for the next 15 years. It concerns the inclusion of an Irish category in the Ethnic Group question in the census.

[It has become increasingly apparent since 1991 that the lack of such census information places Britain's Irish community at a serious disadvantage]. [The

planning of social services, housing, healthcare and many other support services depends on such accurate census statistics]_c. [Irish people in Britain experience poorer health, higher rates of mortality and economic disadvantages which are passed from one generation to the next]_d. [These disadvantages are striking when compared to the indigenous white population and when compared to other economic minorities]_e.

[We would like to state our support for a separate category for Irish people in the Ethnic Group question in the 2001 census]_a. We would urge the ONS to put an end to the anomalous situation where the largest ethnic-minority group in Britain is invisible.

The analysis of this text yields the structure in Figure 6.7:




Figure 6.7 Structure of the 'Irish census' argument

This leads to the epistemic characterisation in Figure 6.8:

```

bel (speaker, e) .
bel (speaker, c) .
bel (speaker, supports (e, d) ) .
bel (speaker, d) .
bel (speaker, supports (d, b) ) .
bel (speaker, b) .
bel (speaker, supports (c, b) ) .
bel (speaker, a) .
bel (speaker, supports (b, a) ) .
  
```

Figure 6.8 System beliefs in the 'Irish census' argument

When these are used as input to *Rhetorica*, the final output produced is as in Figure 6.9 (again complete output can be found in Appendix C):

Complete plan of primitives:

```

2.make_salient(v_agent=ag0,v_proposition=a,v_topic=null_topic)
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(b,a),v_parent=null_topic)
6.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(b,a))
13.push_topic(v_topic=supports(d,b),v_parent=supports(b,a))
15.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,b))
34.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(d,b))
28.push_topic(v_topic=supports(e,d),v_parent=supports(d,b))
30.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,d))
33.pop_topic(v_topic=supports(e,d),v_parent=supports(d,b))
18.pop_topic(v_topic=supports(d,b),v_parent=supports(b,a))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(b,a))
19.push_topic(v_topic=supports(c,b),v_parent=supports(b,a))
21.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,b))
24.pop_topic(v_topic=supports(c,b),v_parent=supports(b,a))
9.pop_topic(v_topic=supports(b,a),v_parent=null_topic)

laden with affect on the following goals:
[]

```

Figure 6.9 Final *Rhetorica* output for the 'Irish census' argument

The resultant text is as follows:

Tomorrow, the Office of National Statistics will take a decision affecting Irish people in Britain for the next 15 years. It concerns the inclusion of an Irish category in the Ethnic Group question in the census.

We would like to offer our wholehearted support for a separate category for Irish people in the Ethnic Group question in the 2001 census. Our primary reason is that it has become increasingly apparent since 1991 that the lack of such census information places Britain's Irish community at a serious disadvantage, as there is evidence to suggest that Irish people in Britain experience poorer health, higher rates of mortality and economic disadvantages which are passed from one generation to the next – and these disadvantages are striking when compared to the indigenous white population and when compared to other economic minorities. In addition, the planning of social services, housing, healthcare and many other support services also depends on the collection of such accurate statistics.

We would urge the ONS to put an end to the anomalous situation where the largest ethnic-minority group in Britain is invisible.

There are four clues realised in this text. First, detail at 12, realised as *Our primary reason*. Second, detail at 27, realised as *as there is evidence to suggest*. Third, detail at 34 realised simply as the em dash followed by *and*. Finally, the parallel clue at 26, realised using *In addition*.

These three examples demonstrate not only a range of functionality of the *Rhetorica* system, but also that the output is structurally well formed and can be associated with coherent text. The next chapter in part addresses whether *Rhetorica* output is also persuasive, by tackling the key problem of system evaluation.

VII

Peroration

7.1 Appraisal

The evaluation of natural language processing systems presents a range of problems and has become a field of research in its own right (Galliers and Sparck Jones, 1993). Evaluation of the *Rhetorica* system is confounded by several further problems. In the first instance, it is not currently applied in any specific domain: Galliers and Sparck Jones make repeated use of case studies – often in information retrieval – to examine evaluation methods. In such situations, there are a number of metrics which can be employed in judging the efficacy of a system. Without an application, evaluative study becomes that much more difficult. Secondly, as Galliers and Sparck Jones point out (pp98-99), evaluation for NLG systems is particularly difficult – the only possible route they mention is “to evaluate NLG systems in the context of their application”, leaving more generic theories of text organisation difficult to evaluate. *Rhetorica* suffers further from the fact that it does not generate the surface text, but an annotated deep structure which then requires realisation by subsequent processing not addressed in this work. Evaluating *Rhetorica* on the basis of output text thus becomes more difficult and less reliable.

There are a number of questions which can be asked in trying to determine the level of *Rhetorica*'s performance. The two key issues which have run throughout the work, motivating design decisions at a variety of levels, are the notions of coherency and persuasive effect, and these can be employed in examining system output. Does *Rhetorica* produce text which is coherent? Does *Rhetorica* produce *only* text which is coherent? The first question is the easier to establish: in the two examples discussed (the first in §3.3, the second – the vegetarianism argument – in §4.4 and §5.5), the result of the system is extremely close to the structure of the original natural language; since these are coherent, *Rhetorica* has demonstrated that it is at least capable of generating coherent structure. But does *Rhetorica* ensure that the structure it generates is coherent? There are two problems which beset attempts to answer this question: (i) as discussed in §3.1.4, there is no definition for coherency, so, it is impossible to reason a priori about the coherency of structures *Rhetorica* may produce – instead, a

highly labour intensive (and non-automatable) task of generate-and-test is called for; and (ii) coherency is not dichotomous, but scalar (also discussed in §3.1.4), so results can at best be graded, and at worst – and in practice – only be compared with other, natural, coherent versions with the same propositional content. Because such validation is so expensive, only limited work has been conducted in this direction. However, these limited results are promising: both with the examples discussed above, and in several other small arguments taken from the corpus, *Rhetorica* creates structures which are not only coherent, but also highly similar to those of the natural arguments. Finally, the standard which *Rhetorica* must achieve in order to create only coherent structure is actually rather low: as Cohen (1987) points out, introduction of appropriate clue words can repair significant structural impairment (indeed, Cohen tentatively suggests that any incoherency in structure can be rectified by clue word repair). Structural coherency, then does not in itself represent a suitable vehicle for system evaluation.

Since *Rhetorica* focuses on building arguments to persuade, the persuasive effect of a text might be employed as an evaluative metric. But clearly, this approach too suffers from major problems. Quite apart from problems equivalent to those which beset coherency (particularly that it is the structure not the form which must be evaluated), there are even greater difficulties in assessing persuasive effect. As explained in §2.2.1, persuasive effect as construed by both Perelman (Perelman and Olbrechts-Tyteca, 1969) and Freeman (1991) is contingent upon a specific audience; any evaluation must therefore be with respect to some audience. Even if an appropriately detailed characterisation of an audience were available, it still remains doubtful whether an objective assessment of persuasive effect (on that audience) would be possible. One option might be to consider whether or not the argument is successful in altering the beliefs of its audience (assuming that is its goal); this is too crude an estimate however, since (i) it is possible for a highly persuasive argument to fail to effect an absolute belief change, and (ii) it does not admit of intermediate levels of persuasion, and it seems inappropriate to consider persuasive effect as simply present (successful belief change) or absent (no belief change). As with coherency, it is possible to peek at the persuasive capabilities of *Rhetorica* by comparing system output with arguments known to have been persuasive, but, as with coherency, the limited scope and arbitrariness of such validation limits the utility of the approach (though, as with coherency, in the few examples examined, *Rhetorica* performs extremely well).

One final avenue open for investigation is a more empirical, more objective assessment of the overall effectiveness of an argument: direct experimentation. By studying the response of a large sample to a variety of textual arguments, some generated by *Rhetorica*, some extracted from a natural environment, it should be possible to assess the competence of the *Rhetorica* system. One of the key problems with this attractive approach is in the complexity of conducting experimental studies dealing with such abstract phenomena as ‘persuasiveness’ (witness, for example, the conflicting surveyed in (McGuire, 1969)). Constructing an assay to evaluate the functionality of *Rhetorica* whilst eliminating as many confounding features as possible – and then conducting a large scale experiment with concomitant statistical analysis – is not only ambitious but also premature. In the first place, it is unclear what would constitute a good (exhaustive, weighty) list of external parameters to be controlled for. Secondly, there is a broader issue that NLG systems have not traditionally been evaluated by human centred experimentation, and so the utility and applicability of the approach is far from certain.

For these reasons, the evaluation of the *Rhetorica* system carried out in the current work centres upon a preliminary pilot study, which, in addition to examining the efficacy of the system, also aims to uncover related phenomena and to explore the potential for experimental psychology research within the domain of NLG evaluation, with a view to motivating, justifying and delimiting a full scale investigation to be executed within the programme of future work.

The Pilot *Rhetorica* Evaluation Study (PRES) follows a simple design rubric: an argument taken from the corpus is analysed to elicit its internal structure and to infer the original beliefs of the speaker (with regard both to the arena of discussion, and to the beliefs and parametric characterisation of the audience). These beliefs and parameters then form the input to *Rhetorica*, which constructs the deep structure of an argument appropriate to that input. The final stage is to construct, by hand, the surface text of that deep structure. It is during this process that bias could unintentionally be introduced by the experimenter, so to minimise this possibility, realisation was restricted almost entirely to the original wording, except where this was prohibited by syntactic constraints. Sentence boundaries and other punctuation devices were also maintained, except where clue realisation could not be performed in their presence. The aim of this stage is thus to have available two arguments on the same topic: the original and one generated by *Rhetorica*, which may have (i) potentially different orderings of premises and conclusions; (ii) potentially different enthymematic contraction; (iii) potentially different use of clue phrases; (iv) potentially different punctuation breaks; (v) potentially differing content (in that the *Rhetorica*-generated argument may eliminate subarguments due to breadth or depth constraints).

In addition to these two versions of the argument, a third is also generated, again using the *Rhetorica* system, but employing only the planning mechanism in conjunction with argument integrity constraints – i.e. with all EG level heuristics inoperative. The textual realisation process is then the same as above. Finally, the entire process is repeated on two further arguments from the corpus. The three arguments exhibit a number of important differences: they have different levels of emotiveness (argument one, for example, concerns road signs; argument two, the rights of birth parents in cases of adoption); they are of different length (both textually and in terms of the number of functional units); they are of different complexity (both in terms of depth, breadth and range of subarguments); and they involve premises of different sizes. All arguments, however, are assumed to be aimed at a similar, broad audience (specifically, the arguments were all taken from the ‘Letters’ page of *The Guardian* appearing during one week in October 1996) with a high level of technical competence and low level of scepticism. The arguments, their analyses and their *Rhetorica* generation, are those presented above in chapter six; the argument variants generated with EG functionality disabled are given in Appendix C.

Thus the subject is presented with three variants of each of three arguments: the original from *The Guardian* (this is termed ‘Orig’ in the analysis), the argument generated by the full *Rhetorica* system (abbreviated ‘FullRhet’), and finally, the argument generated by *Rhetorica* without EG heuristic functionality (termed ‘NoFrills’). The subject is then provided with the following instructions: “... you are asked to rate the texts on the basis of how persuasive you find them. Following each text is a box: please enter a number between 0 and 9 to indicate whether you found the argument highly persuasive (9) or totally hopeless (0).” No further information is provided on how to perform the assessment. The

subjects are, unusually, quite diverse: in total, the 34 participants included computer science academics, CS students, non-CS academics, healthcare professionals, and university clerical staff. This wide range is important and appropriate – to take, for example, just a single cohort of students is not only unrealistic, but also quite different from the audience for whom the argument was originally intended (namely, the diverse readership of *The Guardian*).

In addition to factors such as emotiveness and complexity, which are varied across the arguments, it is also likely that primacy and recency effects would play a role in subject's assessment of the arguments (though, on the basis of (Hovland and Mandell, 1957), it is not certain what that effect would be). For this reason, the order of presentation of the arguments was varied: for the first set of arguments, that generated by *Rhetorica* (FullRhet) preceded the original (Orig), which in turn preceded the minimal *Rhetorica* argument (NoFrills); for the second, the order was Orig, NoFrills, FullRhet, and for the third, NoFrills, Orig, FullRhet. In this way, the primacy effect, recency effect, and argument ordering effect were all randomised to the fullest extent given only three replicates.

The raw results are given in Appendix C, and are summarised below in Figure 7.1. The results have been presented in two forms: firstly, the raw figures provided by the subjects, and secondly, the rankings based on those figures. Although the translation from the former to the latter involves a loss of information, it eliminates the skew introduced by individual subjects who employ an unusually high or unusually low range of marks.

	Full Rhet	Orig	NoFrills	Overall
Total	649	562	605	1816
Mean	6.36	5.51	5.93	5.93
StDev	1.94	1.84	1.75	1.87
#times 1st	52% (53)	21% (21)	19% (19)	
#times 2nd	19% (19)	27% (28)	37% (38)	
#times 3rd	20% (20)	38% (39)	28% (29)	

Figure 7.1 Summary of results from PRES as raw data and rankings

Though the sample size was small (a total of 34 subjects) – as is appropriate to a pilot study – the results are encouraging not only for *Rhetorica*, but also for the approach in evaluative technique in general, and in terms of identifying factors which need to be taken into account in a full study.

The raw data has several pertinent features. The first is the high standard, and the relatively small difference in the mean scores of the argument types. These figures are summarised in the graph in Figure 7.2, from which it can be seen that even if the variance were smaller, the difference between averages of the argument types is unlikely to be significant. Nevertheless, it is clearly demonstrated that the text produced by *Rhetorica* – with or without EG level processing – is at least as good as the original arguments. The results suggest that the maintenance of subargument integrity is thus sufficient for acceptable textual argument.




Figure 7.2 Summary of raw results from PRES

The high variance and consequent non-significant differences between argument types is to be expected, however, because subjects were given minimal guidance on how to evaluate arguments, and were at liberty to interpret the notion of 'persuasiveness' in any way they saw fit. A more insightful means of comparison between the argument types is by consideration of the order in which subjects ranked arguments in a given set. Figure 7.3 illustrates the aggregate rank profile as a percentage, i.e. the frequency with which a particular argument type was rated as the best, middle, or worst in its set.




Figure 7.3 Summary of results from PRES analysed as rankings

These ranked results show several important features. The first, and most striking, is the high frequency with which the arguments generated by the full *Rhetorica* processing are rated the best in their set: these arguments were selected by subjects as the best in 52% of cases, as compared with subjects choosing the original as best in 21% of cases, and the EG-deactivated arguments as best in 19%³⁵. A similar, but less striking relationship is found amongst those arguments rated as ‘middle’ – the arguments generated by the version of *Rhetorica* with EG level processing deactivated were evaluated ‘middle’ in 37% of cases, as compared with 27% for the original arguments and 19% for those generated by the full *Rhetorica*. Finally, original arguments were more commonly rated as the worst in a set: 38% (compared with 20% for the full *Rhetorica* and 28% for the restricted system).

It seems from these figures that those arguments generated by the full *Rhetorica* system performed significantly better than their natural, original counterparts at appearing persuasive. Furthermore, it is clearly demonstrated that the employment of the various forms of heuristic processing at the EG level does significantly contribute to the persuasiveness of a text – since the full *Rhetorica* system far outperforms the version without such processing. These are particularly encouraging results, as it suggests that the rich rhetorical and psychological techniques exploited by *Rhetorica* enhance a text beyond what may be achieved by a naïve human author.

This is not to say that PRES has definitively proved the capabilities of the *Rhetorica* system: as a pilot study, an analysis of the shortcomings and future extensions of PRES are almost as important as the results gained in the first instance. It is to this analysis that the next section is in part addressed.

7.2 Future work

Before examining the limitations and future extensions to the *Rhetorica* system itself, the discussion focuses upon the shortcomings of PRES, and the measures which could be taken to address those shortcomings in a larger, more rigorous study.

The first set of problems is to identify factors which may affect the outcome so that those factors may be controlled for. The first such factor to consider is the effect of primacy and recency. The PRES attempts to minimise these effects by randomising the order across argument sets, however with only three replicates, the effects could potentially impact the results. In analysing the data, however, it seems that primacy and recency play little or no role, as evinced by the summary in Figure 7.4: if such effects were impacting the results in PRES, a correlation would be expected between position and rating/ranking (positive if the dominant effect was recency, negative if primacy). Such a correlation is clearly not present to any significant degree (though cursory inspection might suggest a very slight recency effect). Given the identification of the influence of primacy and recency on experimental results in previous studies, however, it would still be desirable to eliminate their effects entirely. This could be achieved quite simply in a larger study by offering only one argument from each group to any given subject (necessitating a larger sample size). It is assumed to be unlikely that primacy or recency would affect the reception of arguments in differing domains (e.g. the presentation of a ‘pro’ argument concerning road signs is unlikely to effect the reception of a ‘con’ argument concerning abortion).

³⁵ The small discrepancy – these percentages should sum to 100% – is due to the occurrence of equally rated arguments: such ratings were ignored in the counts.




Figure 7.4 Summary of results on primacy/recency effects in PRES

Another, practical, problem with the PRES comes as a result of the small corpus: although *Rhetorica* includes a rich, wide range of interacting arguments forms and heuristics, only a limited subset of these were actually identified in the portion of the corpus employed for the PRES. A more rigorous study would benefit from being preceded by a wider corpus study, from which it would then be possible to extract examples of a more diverse range of argument forms and heuristics (e.g. Modus Tollens, Inductive Generalisation, and their associated heuristics). The section of the corpus employed in the PRES may also be susceptible to slight criticisms of artificiality, since pieces on the 'Letters' page are subject to editing before appearing in print. This editing process may improve or potentially damage (through over-zealous enthymematic contraction) an author's original argument. A wider corpus study would eliminate this problem (though in defence of the original approach, the arguments which appear in print are still persuasive monologues aiming to alter beliefs or behaviour of a particular audience, and are therefore encompassed within the remit of the *Rhetorica* system).

Relatedly, a full scale study would also need to address audiences of different types – i.e. audiences about whom an argument's author maintains different assumptions. This would not only test the hearer-sensitive aspects of *Rhetorica*, but also widen the generality of the investigation and, consequently, its results.

The single greatest problem with the PRES, however, concerns its means of eliciting subjects' judgements. Simply asking explicitly for an estimate of the persuasive effect of a text carries with it a number of problems. In the first place, subjects may have a preconceived notion of what constitutes a persuasive text (e.g. that a technical piece is more persuasive than a nontechnical text), and this may have only a limited correspondence with what is actually persuasive to those subjects (i.e. what actually effects a change in their belief, behaviour, etc.). One approach to dealing with this problem is to excise any measure of persuasion, and instead assess directly the degree to which change in belief (or behaviour, etc.) has been effected. This in itself barely simplifies the problem, but if the assessment is

carried out by testing the subject's commitment to a particular course of action, then it may be possible to assess belief change whilst avoiding the need for subjective introspection. One example of the means by which this might be executed is through the use of arguments competing for limited resources over which the subject has control (e.g. the subject playing the role of treasurer, choosing between various proposals).

Finally, a greater understanding of the factors affecting the persuasiveness of a text may perhaps be gained through requesting from the subjects a more or less extensive commentary on their evaluation (be it of persuasive effect or commitment to some decision). Comments referring to clue word usage, premise length/complexity and contextual information, and the effect that these features had on evaluation, were all volunteered by various subjects in PRES. A more careful collection and analysis of such data could provide valuable insight into the way in which persuasive effect can be enhanced.

Turning now to the *Rhetorica* system itself, there is one important simplification exploited by the current implementation of *Rhetorica* which forms an important component of future work. A variety of fixed values are employed to characterise a range of thresholds and heuristic limits, such as the number of subarguments permitted to support a given claim, the distance between a subargument and its conclusion required to generate a reminder of the implication link, etc. Using a static value for these attributes intuitively seems too rigid: it would be more appropriate to determine values dynamically given the context in which the heuristics are being applied. In particular, settling upon some value for one variable may affect the value adopted for another. For example, if the breadth of an argument is severely constrained (i.e. the number of subarguments supporting a given claim is restricted to only a few) then it seems more likely that the depth of any one of the argument chains could be permitted to be longer.

This interaction between various variables can be explained through consideration of the effort required by the hearer to follow an argument, and the *cognitive load* imposed upon the hearer by various tactics in the argument. The hearer's ability to understand an argument will call upon a range of faculties, but all told, there is some fixed, finite limit to the hearer's capabilities. It seems reasonable to suggest that given this fixed bound to the hearer's cognitive resources, various presentation techniques will demand a greater or lesser portion of those resources: if one approach (say, severely restricting argument breadth) uses less, then more may be available for another (say, argument depth). This claim is consistent with results in cognitive psychology, demonstrating that tasks which place a load on centralised cognitive processing detract from the ability to perform further concurrent processing (Kahneman, 1973), (Baddeley, 1986). It also follows in the spirit of Sperber and Wilson's theory of relevance which rests upon a key assumption that one fact is more relevant to the hearer than another if the hearer is called upon to perform less processing in order to derive that fact (Sperber and Wilson, 1986) – in the current work, similar assumptions are made, through the definitions of the heuristics, of the relationships between the demands placed upon the hearer's cognitive processing and the resultant coherency and persuasive strength (though the relationships are not, as in Sperber and Wilson, a direct mapping from more/less processing to less/more salient).

Viewing the variables in this way as an interconnected mesh of constraints on the ultimate cognitive load placed on the hearer has three important results: the wide range of influence that the notion of cognitive load has on *Rhetorica*, the potential for defining a notion of style, and the relation to a computational definition of resource allocation.

The first point is simply the observation that anticipating cognitive load is a task which pervades an enormous variety of processes in *Rhetorica*. Recall, for example, from §5.2, that during enthymeme contraction, it is difficult to establish how many components of a subargument need to be left explicit (or conversely, how many may be left implicit). *Rhetorica* currently restricts the scope of enthymeme contraction to a maximum of one implicit component. Cohen (1987) suggests that two components may be left implicit. It is not difficult, however, to construct an example of a sorites containing a syllogism in which *all three* components are omitted – particularly where additional information is available from the context. The argument in (6.1) provides an example:

- (6.1) *There's a big black thunderstorm over our normal route; if the road there is flooded, we'd only get through by taking the back route. So we'd better take the back route.*

In example (6.1), the sorites runs through four propositions, viz. thunderstorm, rain, flooded, back-route. In the second syllogism *rain, rain → flooded, flooded*, all three components are left implicit. In order to imbue *Rhetorica* with the flexibility to leave an arbitrary number of syllogism components implicit, richer, more context-sensitive heuristic control is required. This control could be provided by the approach currently adopted in *Rhetorica*, enhanced through the use of a dynamic evaluation of the potential for enthymeme contraction based on estimation of cognitive load (such that, for example, if the surrounding elements of a sorites provided sufficiently strong causal indication of both minor premise and conclusion of a syllogism, then the cognitive load placed on the hearer through omitting all three components of that syllogism may be acceptable).

In addition to enthymeme contraction (i.e. in the EC-EPISTEMIC and EC-LENGTH heuristics), and argument breadth and depth decisions (i.e. in HCR2), estimation of cognitive load also plays a role in all the other heuristics – either in determining values used within the heuristics as they currently stand (such as calculating inter subargument distance in HCI1 – reminding of implication), or in providing a principled justification for the use of heuristics such as HPP1 (no weak refutations) and HPR4 (refutations first) – and perhaps in the process also tightening up the scope of their application.

The second role for the notion of cognitive load is in defining 'style'. The supra-linguistic features (i.e. those which are above the level of the lexical and grammatical concerns discussed in (Sandell, 1977)) by which an author's style can be characterised are much the same as the variables which underpin the various heuristics (such as the minimum inter-subargument distance which demands a clue linking premise to conclusion in HCI1). If estimation of cognitive load is seen as a means of calculating the values of these variables dynamically, then a characterisation of style can be built using the functions which perform this calculation. Characterising style in terms of these functions also abstracts from the resultant values themselves in precisely the way suggested by intuition: an author's style will be tempered to the particular extra-linguistic context of the argument – i.e. it will be

modified by the various parameters (hearer scepticism, hearer competence, the medium in which the argument is presented, etc.) defining the discourse encounter. Nevertheless, the idea that an author has a particular style relies on the assumption that her style will manifest itself in any context. So by defining style in terms of the *functions* rather than the *values* used in heuristics, it becomes possible to talk about an author's style independently of any given context (in which particular values are calculated).

The last advantage of adopting an approach based on cognitive load is the potential for tying the estimation of cognitive load to the process of resource allocation. As mentioned in (Reed and Long, 1997b), resource allocation can be employed within the AbNLP framework for restricting the planning resources available to a particular part of the abstract plan. In *Rhetorica*, this translates into laying down limits on the depth and breadth of subargumentation permitted in support of a particular claim. At present, the technique is under-exploited, mainly because there are no principled grounds available for making appropriate estimates of the resources appropriate for supporting a particular claim. Introducing estimates of cognitive load, and imposing upper limits on the value that load can take offers a means of effecting principled resource allocation between parts of an argument.

Together, these three opportunities arising from a cognitive load based analysis make the approach highly attractive as an avenue for future development. In order to employ the notion of cognitive load computationally, a number of issues will need to be addressed. First and foremost, quantitative or qualitative base values will need to be attached to various argumentative strategies: which places a greater load on the hearer's abilities, a serial configuration or a convergent configuration, for example? Then, rules will be required for combining these various values (perhaps simple addition would be inappropriate in some situations). Then there remains the major problem of how to integrate these base values with parameters describing the context of the argument: low hearer competence will clearly increase the weightings of the various estimates (based on an assumption that lower hearer competence is equivalent to claiming that the capacity of the hearer's centralised cognitive process is small). The medium in which the argument is to be expressed will also have a significant impact on the calculations: the cognitive load placed on the hearer by n subarguments might be claimed to be, say, proportional to n in the case where the argument is conveyed orally. In the case of textual presentation, it might be claimed that this falls to $k \log n$ as a result of textual devices such as bullets, numbering, and other formatting which allow the reader to skip back to a previous point if necessary, and to see the structure of the subargument at a glance. It is clearly important to be able to represent this sort of difference and account for it in calculations of cognitive load.

The utilisation of the concept of cognitive load thus represents an extension to the current work which fits well into the existing framework, but introduces a significant amount of new machinery for estimation and integration with context parameters. There are also a number of other extensions to the work which are more derivative in nature and require the introduction of less new functionality.

The most obvious of these derivative extensions is the completion of the set of operators. Thus work should focus on characterising the causality and analogy inductive operators, and a fuller range of fallacy operators. As discussed below, there are far-reaching problems with the design of the inductive operators, and in the absence of a perfect, principled solution, a simplified pragmatic approach – such

as formed the foundation of the inductive generalisation operator – would need to be designed. There are fewer problems posed by the design of the fallacy operators, other than the preliminary decision over how many to recognise as distinct fallacies – lists range from Locke's four (Locke, 1975) to more recent lists numbering around four dozen (Lowder, 1997). Perhaps the most appropriate pragmatic solution to this problem is to go for the 'gang of eighteen' proposed by Woods and Walton (Woods and Walton, 1989) *inter alia*, but in making this choice, it would be necessary to identify which forms of fallacy could not be characterised or uniquely identified, and what the ramifications of such limitations would be.

Similarly straightforward is the development of a wider set of heuristics. The current set of thirty heuristics plays two roles in the *Rhetorica* system: firstly, it allows the generation of valid, coherent, and reasonably persuasive argument structures, as demonstrated by the results achieved from the experiment discussed in §6.1. Secondly, the set illustrates the range of heuristics permissible – the use of threshold values, the role played by context parameters, the preconditions based on goals or operators, etc. The set is by no means intended to be exhaustive, since new heuristics could be developed on the basis of research in each of the contributory fields. Four areas, in particular, have the potential for offering many new heuristics. In the first place, the rules set down in tomes of classical rhetoric have barely been touched upon – Quintilian (1960) is perhaps the best example: thirteen volumes of precisely the sort of rules characterised in the heuristics in chapter four. Secondly, there is the psychology of interpersonal relations, and the role, currently not represented in the body of heuristics, of the face goals discussed in §2.2.2. Thirdly, the work of Knott (1996) and others in characterising cue words offers a rich source of information which in the current work has been only very lightly tapped. A fuller exploitation of this research could lead to the development of a more complete set of clue word heuristics. Finally, there is the empirical source: in analysing how the structure of natural arguments came about, rules can be formulated which can be generalised and implemented. This last route, though laborious, has the potential to lead to a vast catalogue of heuristics which would always be subject to incremental revision.

A similarly direct extension of the current work is the consideration of other stages of argument. The original remit of the *Rhetorica* system was to investigate the argumentative part of persuasive text. By the classical analysis, this leaves the exordium, division, narration, pathetic and peroration. Although conceptually, the generalisation is simple, the design of a system to cope with all six stages is faced with a number of important issues. On an architectural level, should the six parts of the argument be planned sequentially or in tandem? How are constraints managed between the various phases? At an implementation level, what is the impact on the planning process of the need to repeat information: to what extent does this need violate assumptions of communicative monotonicity, an important simplifying assumption? What is the role of Walker's (1996) informationally redundant units in the other phases of argument? The ability to generate introductions and conclusions is clearly related to summarisation work (such as (Robin, 1994), (Marcu, 1997b), and the SUMMARIST project at ISI) – to what extent can the parallels be identified and exploited?

In addition to these fairly clearly defined extensions to the current work, there are also a number of areas which would benefit from further exploration (and for which precise, implementation .

aims could not be set out until such exploration had been undertaken). A good example is the relationship between the inductive operators, and between them and their deductive counterparts. Inductive generalisation and analogy (in common with some of the fallacious operators, such as IE, introduced in §4.1.1) both rely upon some concept of similarity, yet this concept is poorly understood. *Rhetorica* avoids the issue by demanding that such similarity be explicitly recorded in the database using the HAS_PROPERTY predicate. The work of (Long and Garigliano, 1994) or its foundations might be adduced in supporting *Rhetorica* with a richer model of similarity which might then offer a more competent handling of argument structures employing analogy and inductive generalisation. A similar issue faces the link between causality (and reasoning based upon such causality) and the notion of support used in the deductive operators. There is clearly some commonality between the two concepts, and yet their interrelation is unclear – the tacit assumption in the design of the *Rhetorica* operators has been that support links are a superset of causality links, and that a causality-argument operator would be invoked only in situations where the causality link was being established through co-correlative means.

7.3 Contributions

The primary contributions of the current work are a result of the novel design of the *Rhetorica* system, and the benefits that that design offers. In the first place, the distinction between the logical structure of an argument and the various additional filigree which enhances persuasive effect is crucial. The distinction is similar to that recognised in argumentation theory between logic and rhetoric (Rescher, 1997), but, apart from being the first computational handling of the distinction, differs from the view in argumentation theory in two respects – (i) that argument structure includes some component ordering (normally classified as rhetoric); (ii) that argument structure includes inductive and fallacious techniques (the latter are almost always viewed as rhetorical devices). The separation of the phase of creating the structure of an argument from that of then augmenting it with further ‘rhetoric’ is an important modularisation of functionality, enabling distinct processing remits to be defined. Although conceptually the process is a two phase pipeline, in refining the model, the shortcomings of the generic NLG pipeline model have become clear. To overcome these problems, a slightly different approach is proposed: a ‘cyclic pipeline’, whereby the phases of argument structuring and eloquence generation strictly proceed later generation tasks, but between which incremental, cyclic processing is permitted whereby at a given level of abstraction, the AS determines hard and soft constraints on the plan, and the EG can then introduce new constraints or alter those which the AS recorded as soft. Processing then returns to the AS at the next level of abstraction. This cyclic pipeline is a direct result of the hierarchical nature of the planning process, which itself reflects the hierarchical structure of argument.

A related contribution is the signalling that an argument which is logically sound is not the same as one which is persuasive. Even if the logical content of an argument (either in purely logical terms, or in the wider sense of ‘argument structures’, used here to include inductive and fallacious strategies) can be determined, this alone will rarely suffice to achieve the desired affect on a human hearer. Several factors contribute to this point, including (i) the uncertainty and incompleteness of the dialogue situation (a speaker cannot know all the hearer’s beliefs on the matter at hand, nor can she be

sure that those beliefs of the hearer's of which she is aware actually represent an accurate model of the hearer); (ii) the impact of the 'social context' on the communication (including parameters such as the general and technical competence of the hearer, his scepticism, his view of the speaker's investment in winning the argument, and so on); and (iii) the need for connecting information in the hearer's mind (the hearer, as a human, is not an idealised, rational judge, capable of immediately surmising the deductive closure of some set of premises, but must be led carefully along a coherent chain of reasoning from one point to the next). These problems – and particularly those concerning uncertainty and incompleteness – have been shown to lead to counter-intuitive results, even in the simplified computer-computer interaction occurring in 'distributed defeasible reasoning', §2.2.3. In addition, the fact that logical content may, by itself, be insufficient to effect belief change in a human audience, leads to an important phenomenon which cannot be accounted for in classical logic, namely, the use of multiple subarguments supporting a single conclusion. A computational analysis of multiple subargumentation forms a key contribution, particularly as it constitutes a very basic assumption in the *Rhetorica* planning framework (that goals may be fulfilled one or more times). The particular species of multiple subargumentation represented in an inductive generalisation is also characterised in a novel way which clarifies its role with respect to sibling and parent arguments.

One of the key tenets of the thesis, then, is that both the logical and rhetorical components of argument need to be accounted for in a generation system if that system is to generate text which can be effective faced with uncertain and incomplete information about a human, imperfectly rational, hearer. One of the means by which a system can deal with such scenarios is by distinguishing coherency and persuasive effect: a wide class of component arrangements may be coherent, whereas only some small subset will also be highly persuasive: a key contribution of *Rhetorica* is to demonstrate how *multiple* orderings of units in coherent text structure need to be maintained in order that a subset can be selected on persuasion grounds. The advantages of this approach also carry through to lower levels of abstraction, both as the AS/EG plan is refined (so that the class of coherent, persuasive orderings at one level of abstraction may be subsequently restricted given what is optimally coherent at the next level of abstraction) and then as control passes to lower levels in the framework (so that, for example, lexicalisation can itself force some orderings between components).

At a more practical level, the work also delivers formal characterisations of a range of argument forms – in addition to those deductive operators found to be present in natural language (primarily modus ponens, modus tollens and disjunctive syllogism), there are also novel implementations of inductive generalisation (relying upon several technical points, including the definition of a phantom node), and the fallacies of *argumentum ad populum* and *ignoratio elenchi* (relying upon the availability of rich information in the knowledge base). Importantly, all these potentially disparate forms of argumentation are formalised in a uniform manner, and are drawn together into a single coherent planning framework. In addition, the characterisation of these argument forms offers an appealing analogy between the ground adequacy and relevance questions used in argumentation theory, and the two belief goals (representing major and minor premises) found in the bodies of the deductive operators.

There are also a number of technical aspects of the work which represent contributions to the

field. First is the handling of disjunctive constraints. Although a number of planning systems have adopted a variety of approaches to the problem (e.g. UCPOP (Penberthy and Weld, 1992)), *Rhetorica* represents the first discourse planner (a) to clearly identify the role played by disjunctive constraints (namely, to express the set of coherent ordering relations over components of a text), and (b) to develop a technique of 'deferred commitment' by which a distinction is drawn between the plan-time and run-time constraints on operator applicability. Through this approach, the management of a plan structure implicitly involving disjunctive constraints becomes tractable. Secondly, control of the topic stack is managed explicitly – the push and pop operations form planning operators which in turn fulfil topic manipulation goals. Such explicit handling offers a number of advantages, including, importantly, a basis by which to motivate the introduction of particular clue phrases (which can then be seen as fulfilling the intentions represented by the AS/EG topic manipulation goals). This generation of clue phrases marks a third key contribution: rather than viewing clues as lexical manifestations of strictly inter-clause relations, the current work views them as links between arbitrarily large units within the text, such that a 'therefore' may function to link a conclusion to (the sum of) numerous earlier premises – rather than simply to the immediately preceding premise. The introduction of these clues (or rather, abstractions that classify the set of appropriate clue phrase realisations) occurs *at the same level of abstraction as the units of text which they relate*, and is thus in contrast to other generation work which, due to the assumption of their inter-clausal function, defers generation of clues until the lexical level. A final technical contribution is in clearly distinguishing the role played by goals of belief and those of saliency, and relating both to recent work on the role of intentions in NLG. Although the notion of saliency has been identified in text generation, (Walker, 1996), the current work develops its usage such that goals of saliency form the basis of all subsequent realisation – goals of belief (traditionally realised into text) are responsible for the structure of the text.

These contributions – of the logical/rhetorical distinction, the cyclic pipeline control structure, the identification of problems in purely logical accounts of argument, the distinction between persuasion and coherency, the characterisation of multiple subarguments, the uniform formalisation of deductive, inductive and fallacious planning operators, the deferred commitment approach to handling disjunctive constraints, the explicit handling of topic manipulators, the introduction of clue phrases at high levels of abstraction, the belief/saliency distinction, and the experimental testing of efficacy – are emphasised by considering similar NLG work. The four most relevant systems are those of Hovy (1990) on the role of style in text generation, Elhadad (1992a) and Maybury (1993) on the generation of textual argument, and Zukerman *et al.* (Zukerman *et al.*, 1996) on the generation of argument structure. None of these systems distinguish explicitly between the logical and rhetorical components of an argument. It is assumed in both (Elhadad, 1992a) and (Maybury, 1993) that the structure of an argument can be directly realised into text without need for any additional work; in (Zukerman *et al.*, 1996) it is unclear exactly how the structure of an argument is realised into text, but it is assumed that there is a direct correspondence between structure and surface form (the emphasis in her work is squarely upon the pre-linguistic organisation of argument). Hovy's account offers the richest support for the influence of rhetorical and extra-linguistic factors on textual form, but it does not clearly delineate such factors from other structural and stylistic components of the generator.

In regard to the control of the planning process, it is again only Hovy's PAULINE which addresses the issue in a similar way: by identifying the need for feedback between structural and lexical levels, he proposes an interleaved approach; *Rhetorica* adopts a similar approach, not between lexical and nonlexical issues, but rather, between much more abstract levels of processing. In addition, *Rhetorica* extends the paradigm by using a two tiered notion of constraint, permitting some constraints to be revoked and leaving others irrevocable.

In identifying the need for more than just a logical account of argument, the target audience is not the wider NLG community (for whom issues of uncertain, incomplete information, demands on linear coherence, etc., are standard), but the narrower community focusing on natural language argument (in addition to other argumentation communities). Thus work of Elhadad and Maybury involves strong assumptions about the reliability of the hearer model, but is otherwise competent at handling extra-logical argument (though this point is not directly addressed in either work). The problem is that large scale structure in either model is poorly co-ordinated. Zukerman's NAG, in contrast, uses a logical approach, and benefits as a result in having a rich model of argument structuring. However, the work is restricted to an entirely logical account, and thus suffers from serious shortcomings, such as its inability to handle multiple subargumentation.

The distinction between persuasion and coherency does not appear to have been addressed even in work focusing upon argument generation. This stems, in the case of Zukerman, from an implicit assumption that logically correct is equivalent to persuasive, and in the case of Elhadad, from a lack of representation of the effect of an argument on the hearer. For Maybury, as representative of much of NLG, the assumption is that text which is correctly planned (i.e. coherent) using various operators (which may have persuasive overtones, such as 'CONVINCE_BY_CAUSE_AND_EVIDENCE') will have the desired effect. *Rhetorica* represents the first system to tackle coherency as a separate issue from persuasive effect.

The formalisation of deductive, inductive and fallacious operators is not in itself new – Maybury, Sycara (1989), and others approach the problem, but none offer a characterisation which is flexible (in that it can accommodate any additional operators) and uniform between the operator types. Furthermore, the conventional deductive operators have not previously been explicitly represented as operators, but rather, have been assumed in the definitions of a wider set of operators embodying particular instances (such as the aforementioned convincing by offering causal evidence, from Maybury).

The need for maintaining multiple coherent orderings (and the subsequent requirement in an abstract planning framework for the maintenance of disjunctive constraints) has not been recognised in NLG, seemingly for two related reasons: first, due to an assumption that any coherent ordering is good enough (since there is no other metric for success), and second, that there is a single optimally coherent ordering (or, more usually a single coherent – as opposed to incoherent – ordering). By enriching a model of discourse to include an additional metric for success – persuasive effect – and assuming coherency to be a scalar rather than dichotomous phenomenon, the handling of multiple options becomes necessary.

Handling of topic manipulation through the use of explicit plan operators also marks a

departure from conventional NLG practice in which it is only – to use terminology from (Grosz and Sidner, 1986) – the rhetorical and intentional components which are used in the planning process. This is true of Hovy, Maybury and Elhadad – amongst many others. The advantage in also planning with attentional information, apart from the intuitive appeal of representing and manipulating all three components of the Grosz and Sidner model in a uniform way, is that clue word generation can be easily incorporated in a principled manner. Although clue word introduction is not addressed directly in any of Hovy, Elhadad, Maybury or Zukerman, it is emerging as a new area of investigation in NLG (consider, for example, that the COLING-ACL meeting of 1998 hosts a workshop devoted to the subject), though seems tied to a view that clues connect clauses. That clues function to link larger structures within the text is a key claim of the current work, and is utilised in motivating clue introduction at abstract levels in the planning process, in contrast to being left as a lexical realisation problem.

Almost all NLG research – and certainly that represented by Hovy, Maybury and Elhadad – relies upon planning processes based on a derivative of Cohen and Levesque's formalisation of agent beliefs (Cohen and Levesque, 1990). Thus the recognition and operational employment of a notion of saliency is quite rare, and has not previously been used, in tandem with belief, to codify the speaker's intentions – i.e. to distinguish between the speaker's goal of making the hearer believe some fact, and her goal of making him aware of that fact.

Finally, the experimental technique, albeit as a pilot version, represents an attempt to rigourise evaluative methods in NLG, which have been characterised by subjectivity and lack of generalisable results. Systems such as those of Hovy, Maybury and Elhadad, and their derivatives, undergo testing to establish that they are sufficiently capable to perform in a given application. Such testing makes it difficult to establish whether or not the techniques proposed in each system are generally applicable. As *Rhetorica* has no application domain, testing efficacy can only be testing its ability in a generic sense. The current work therefore contributes to the problem by indicating one possible route to improving standards of evaluation in NLG: the use of carefully designed experiments in cognitive psychology.

In considering the contributions in comparison with closely related work, another, critical, function of the work becomes relevant. Although the vast array of NLG work based upon RST has suffered limited criticism (Hovy, 1993), the current work offers an array of problems from which the approach suffers (if it is to be viewed as a panacea), and offers a well justified argument for adopting a view whereby RST plays a role at an intermediate level of processing, but which requires larger-scale direction from more abstract levels.

In conclusion, perhaps the single clearest contribution – clear even from the foregoing list – is the unified integration of views of argument from a diverse range of sources: classical and modern rhetoric, argumentation theory and informal logic, computational linguistics, agent theory, social psychology, and corpus study. The *Rhetorica* framework offers an environment in which dicta and observations from these various sources can be combined and brought to bear on the process of generating argument, and in which the various facets of argument – logical and fallacious, deductive and inductive, heuristic and definitive – can be represented in a uniform way. In adopting such an

integrative approach, it is perhaps unsurprising that the current work also offers small contributions to fields other than NLG, such as the exploration of argument as a means of agent communication in distributed AI, and the development of a definition for persuasive monologue in argumentation theory. It is this inter-disciplinary nature which has provided *Rhetorica* with techniques for handling many of the complexities of the discourse situation – extra-linguistic, contextual, hearer-dependent techniques which are crucial in generating textual arguments which are not only coherent, but also persuasive.

Appendices

Appendix A. The ‘Gang of Eighteen’: A Common List of Fallacies

This list is a summary of (Johnson, 1992, pp237-267).

Appeal to Authority, *Argumentum ad Verecundiam*. Testimony of an authority is used as evidence where that authority is not appropriate for the current argument.

Appeal to the People, *Argumentum ad Populum*. Arguing for some conclusion on the basis that many people believe it to be so.

Appeal to Force, *Argumentum ad Baculum*. Argument based upon a threat.

Appeal to Pity, *Argumentum ad Misericordiam*. Argument relying upon the pity of the audience.

Appeal to Ignorance, *Argumentum ad Ignorantium*. Arguing that a conclusion is true on the basis of an absence of counter information.

***Ad Hominem*.** An attack against a proponent rather than her argument. (Three subcategories are often distinguished: the abusive *ad hominem*, attacking the proponent’s character, the circumstantial *ad hominem*, attacking the proponent on the basis of assumed vested interest, and the *tu quoque*, accusing the proponent of the same shortcomings as are ascribed to the opponent).

False Cause, *Post hoc ergo propter hoc*. Arguing for a causal relationship purely on the basis of temporal precedence.

Slippery Slope. An argument that one thing will lead to another, that to a third, etc.

False Dichotomy. Employing a mistaken assumption that there exist only two alternatives (one of which is then negated). If there really are only two alternatives, this ‘fallacy’ becomes equivalent to disjunctive syllogism.

Equivocation. Utilises an ambiguity in a word or phrase.

Hasty Generalisation. Drawing a generalisation on the basis of a small or unrepresentative sample.

Fallacy of Composition. Employs reasoning that some property of a collection of parts holds for the whole.

Fallacy of Division. Employs reasoning that some property of a whole holds for its parts.

False Analogy. An analogy failing to meet minimum standards of the criteria for inductive strength: (i) that the relevance of the properties the items have in common to the proposed property; (ii) that there are no relevant dissimilarities.

Begging the Question. An argument which requires at some stage its conclusion to be assumed as a premise.

Straw Man. An opponent misrepresents (and then attacks) the proponent's position.

Red Herring. Drawing attention away from the issue by introducing another, unrelated (but often emotive) issue.

Inconsistency. Reasoning from inconsistent premises.

Appendix B. The *vegetarianism* argument in full

Reproduced with permission from <http://www.viva.org.uk/Viva!Guides/justice.html>

Justice for All

by Michael Mansfield QC

Michael Mansfield is almost certainly the best – and the best known – criminal defence barrister in Britain.

He has represented and won many of the big headline cases – the Birmingham Six, Tottenham Three, Cardiff Three, Judith Ward.

He is outraged by injustice and is motivated by a desire for a more just and equitable world.

He is a vegetarian and a trustee of Viva!

An agonised debate wracked society about a couple of hundred years ago. Church, government and the middle classes searched their consciences, argued publicly and offered evidence to back their views. However, it was a one-sided argument and they soon reached agreement, with a collective sigh of relief. It was official – black people didn't have souls. More than that, they would benefit from the discipline of hard labour provided for them by Christian masters.

The outcome was – business as usual. Ship's architects played with little models of prostrate humans, jiggling and juggling them until every available space on every possible deck was filled. They produced new vessels that could carry even greater numbers of slaves. It translated into millions of human beings being chained, motionless, side by side in the festering dark of a ship's hold for weeks on end. Over two million died and were simply tossed into the Atlantic's waves.

In West Africa, slave traders continued to plunder villages, yoking the inhabitants by the neck in long lines of misery. In the colonies, plantation owners continued to divide families, beat, abuse and exhaust the people over whom they had complete and total control. In Bristol and other slave transporting sea towns, the monied counted their dividends from this trade in human degradation.

So long as we persuaded ourselves that black people were beasts then we could do to them as we wished. Few people posed the question of whether we should be handing out such barbarity even to beasts. And not a lot has changed when you think about it. The overriding concern in that early debate was to defend the right to make money. In order to do that almost anything could be excused.

The argument these days is not about souls but about consciousness, awareness, the ability of animals to feel or fear. But it is motivated by entirely the same morality. Pigs are tethered in barren stalls, so deprived of stimulation that they often go mad. They are forced to have as many litters of piglets as their bodies will stand but they are not allowed to mother them.

But it's fine because they spend their lives in buildings sheltered from the elements with all the food they need and no animal wants more than that. Well, that's what we're told.

If a person did the same thing to dogs he or she would be prosecuted for cruelty – and yet a pig is equally as intelligent as a dog. So why the difference? Why do we have two sets of rules, one for a dog and one for a pig? There is no logical explanation except that we eat pigs after we've been cruel to them and industry makes a great deal of money from their suffering.

Doesn't it seem extraordinary that we can cram sheep into lorries in three tiers, knowing that

the top tier will urinate and defecate on the sheep below and that tier will do the same to the sheep below them. We then drive them around Europe without food, water or rest and talk about a humane society. Could anyone with no financial interest conceive of such a way to treat any feeling creature? But, of course, they feel differently to us! Or so we're told.

Trawlers criss-cross the world's oceans, their trawl boards crushing everything on the sea bed that has the misfortune to be in their path. Much of what is caught is either thrown back dead or turned into fish meal to feed other animals. Those to be eaten have a knife thrust into them and are disembowelled while still alive. But that's all right because fish don't feel pain! Who could ever have conceived of such an excuse and who would have believed that so many people would be taken in by it.

Watch a chicken as it forages around a field. It struts and scratches and runs and bathes in sand. Its feet thrust the earth to one side in search of bugs and beetles and seeds and other tasty things to eat. Its feathers shine and glint in the sunshine. What intelligence could ever have conceived of taking these restless creatures and cramming them five to a cage little bigger than a microwave oven? The beak is clipped, the feathers fall out and the pathetic creatures are bred to produce twenty-five times more eggs than they would naturally. Instead of bugs and beetles and seeds we grind up their own kind and even their droppings and feed that to them. Every attempt to end the battery practice is resisted. Well, chickens aren't really animals, are they?

The only excuse for visiting such suffering on other creatures can be ignorance. But that exit is closed to us because we do know the truth. We might pretend we don't, we might say that to give up battery eggs is pointless as one person's abstinence doesn't make any difference and we might even say what the battery owners say – that it's of no consequence. But really we know it's wrong and that knowledge makes us push the reality into the darker reaches of our minds. If we don't think about it, it doesn't exist.

The world has existed for nearly five billion years during which time various life forms have developed and evolved. From the earliest sea life to the most complex mammals such as apes, they have all had one thing in common – they have lived within their environment, part of it and dependent upon it.

On the African plains, a cheetah chases after a Tomson's gazelle. It has evolved to live largely on gazelles and it chases because it has no choice – it's that or die. Both are closely matched in speed and turning ability, ensuring that it is largely the weak, the old or the injured who are caught. Similarly a cheetah which is slow will not survive.

This selection process ensures that only the fit and healthy survive to pass their genes on to their offspring. It is a finely-tuned balancing act which has gone on for millions of years.

Into this fine web of existence came human beings. In evolutionary terms we have been here for little more than a twinkling of light. But already we have begun to tear and break the individual strands which go to make up the web of life. Our entire teaching, both political and religious, is to place us above and beyond the rules by which all other animals live, as though they simply don't apply to us, as though we are not animals. But the smartest Porsche car, the best stereo system or the biggest swimming pool can't alter the fact that we are animals and we are governed by the same rules which govern all animals. To ignore these rules is to court disaster. And that is what we are doing – ignoring

the rules.

As a species, we have looked at the world and said that nothing matters but us. All the glories and wonders are there to be exploited and if they can't be exploited then they count for very little. We destroy without knowing the long-term effects of such actions. And even when we do know, we continue to destroy because today is much more important than tomorrow. It is by today's achievements, today's profit margins, today's boasts that we are judged.

The very philosophy which has brought the planet teetering to the edge of destruction is, we are told, the same philosophy which will save us. Like the practice of bleeding in the 18th century, the cure for our haemorrhage, they say, is to prize open even wider the severed arteries of life. Allow those with the power even greater licence to make money by destroying ever more and the cure will be found. Profit is now the global penicillin. Perhaps the most frightening aspect of this voracious appetite is that it can never be satisfied. The more it gets the more it wants.

Profit, however, knows no morality other than to be successful in financial terms. In pursuit of that success, humans and other animals are exploited – ever more demanded of them. Knowledge, truth and understanding have ceased to be signposts to the future and have become minor obstacles to be circumvented. For instance, we know that smoking is the biggest avoidable killer yet every high street is littered with adverts for cigarettes. We know that a vegetarian diet is much healthier than a meat-based diet but it is the livestock farmer who receives all the subsidies. We know that poverty destroys people but the gap between rich and poor grows ever wider.

There is an age-old practice of leaders refusing to accept responsibility for their actions but lack of integrity has become an epidemic disease. They pass the blame for their failures down the line – to single parents, to students, to travellers, to squatters, to immigrants, to other nationalities. Part of the process is to stereotype each of these groups so they cease to be a collection of ordinary humans beings with the normal range of feelings and hopes and fears.

So what has this to do with animals? Everything. It is all part of the same process of denial. By pretending that animals have no real desire for freedom, to procreate naturally, to mother their young or even to experience pleasure by lying in the sun, then it is so much easier to use them purely as commodities – like so much iron or coal or steel – and ignore their pain and fear. It is a philosophy from which human animals are not excluded.

In the poorer parts of the world we ignore suffering on a massive scale and in the process conveniently forget our own history of colonisation. It was they who provided the wealth which we now squander and in return we destroyed their agriculture and social fabric. We now hand out a few paltry pounds and pretend that we are helping to cure the problem which we were instrumental in creating. We still control their economies and demand fodder for livestock while their children die from hunger.

Despite the disproportionate volume of the world's wealth which we control, we are seemingly incapable or unwilling to solve the problems of poverty and inequality – even on our own doorstep. When care and concern for our own species is so stingingly withheld, what hope is there for animals? If you refuse to accept this morality of self interest and violence, which binds this all together, and choose to exercise your right to protest – you will almost certainly be portrayed as the violent one. New laws .

can be used to prevent even the most peaceful demonstration and so collective action against the violence of factory farming or fox hunting can now be prevented, allowing the perpetrators of the real violence to continue unhindered.

We, as humans, obviously believe we have the right to determine everything, who and what shall live and die. We slaughter owls, hawks, crows and magpies so that grouse or pheasants can be reared in large numbers. We then slaughter them by sending lead shot ripping through their flesh – and call it sport. We destroy rabbits as vermin and then demonise the foxes who live on them. We then hunt the foxes. We gas badgers because they might have TB; we trap and kill rooks because we don't like their habits; chase hares with dogs for entertainment; do anything we like to rats and mice; shoot pigeons in their tens of thousands. We determine which animals we will eat and deny them everything, we determine which will be labelled vermin and try to annihilate them; we allow others the comfort of our hearth.

Across the globe we chase whales and destroy them for cultural reasons. We destroy dolphins and seals because they dare to eat the same food as us; and there is hardly a species which will not be exterminated if their interests and ours collide.

It seems we are incapable of understanding that every living creature has its part to play in maintaining the glorious fabric of our wonderful world. We pretend that only we can maintain the balance by determining what shall live and what shall die. It seems we never stop and look around us to witness the appalling mess we have made – always deluding ourselves that we know what we are doing. None of the animals which we slaughter, even those we demonise as vermin, pose any threat to the survival of the planet. It is not they which threaten its existence but us. In this maniacal juggling act we have begun to drop the balls.

The only hope we have is to fundamentally reassess our role in things and our attitude to the planet and living creatures who share it with us. When a calf is prodded and dragged into the killing pen, wide-eyed and terrified with the stench of blood and death in its nostrils, there is no compassion as the captive bolt shatters its forehead. When the slaughterer's hand grabs the muzzle of a lamb to stifle its bleating and applies the knife to its throat, there is no compassion. And without compassion there is little hope for any of us.

Make no mistake, becoming a vegetarian is an important act. Instantly, you are no longer a part of this insanity. You are no longer responsible for most of the daily cruelties handed out to farm animals. You are taking the first step in allowing the planet to heal itself but it is much more than that. It is a political act and a clear expression of a belief in a different way of doing things, a different kind of world – as better world.

Appendix C. The Pilot Rhetorica Evaluation Study (PRES)

C.1 The Arguments

Textual and structural analyses of the three arguments are given in chapter six. Here, Figures C.1, C.3 and C.5 give the full *Rhetorica* system processing showing each level of abstraction, each partial plan and its ordering constraints, and each heuristic firing. In addition, the generation of ‘NoFrills’ arguments (i.e. arguments generated with EG functionality disabled) is also explicated (in Figures C.2, C.4 and C.6).

The full processing for the ‘tourist facility signs’ argument is shown in Figure C.1:

```
| ?- example1.
Applying HCR1: Grouping Topics
Applying HPR2: Conclusion First
Applying CLUE-DETAIL
Applying HPR1: Climax ordering

Complete plan of primitives:
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
3.argModusPonensEl(v_agent=ag0,v_conclusion=j,v_premise=c,v_parent=supports(r,j))
4.argModusPonensEl(v_agent=ag0,v_conclusion=j,v_premise=i,v_parent=supports(r,j))
5.argModusPonensEl(v_agent=ag0,v_conclusion=j,v_parent=null_topic,v_premise=e)
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)

[[start,1],[1,end],[start,2],[2,end]]
Applying EC-EPISTEMIC3
Applying EC-EPISTEMIC3
Applying EC-EPISTEMIC3
Applying HPR2: Conclusion First

Complete plan of primitives:
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
9.push_topic(v_topic=supports(c,j),v_parent=supports(r,j))
11.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
14.pop_topic(v_topic=supports(c,j),v_parent=supports(r,j))
15.push_topic(v_topic=supports(i,j),v_parent=supports(r,j))
17.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
20.pop_topic(v_topic=supports(i,j),v_parent=supports(r,j))
21.push_topic(v_topic=supports(e,j),v_parent=null_topic)
27.argModusPonensEl(v_agent=ag0,v_conclusion=e,v_parent=supports(e,j),v_premise=f)
23.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
26.pop_topic(v_topic=supports(e,j),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)

[[21,25],[25,26],[21,24],[24,26],[21,23],[23,26],[21,22],[22,26],[15,19],[19,20],[15,18],[18,20],[
15,17],[17,20],[15,16],[16,20],[9,13],[13,14],[9,12],[12,14],[9,11],[11,14],[9,10],[10,14],[2,8],[
8,6],[2,9],[2,15],[2,21],[6,15],[20,7],[6,9],[14,7],[start,1],[start,21],[start,end],[start,6],[1,
end],[26,end],[start,end],[7,end],[start,2],[2,end]]
Applying EC-EPISTEMIC3
Applying HPR2: Conclusion First
Applying CLUE-DETAIL
Applying HPR1: Climax ordering

Complete plan of primitives:
36.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(e,j))
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
9.push_topic(v_topic=supports(c,j),v_parent=supports(r,j))
11.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
14.pop_topic(v_topic=supports(c,j),v_parent=supports(r,j))
15.push_topic(v_topic=supports(i,j),v_parent=supports(r,j))
17.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
20.pop_topic(v_topic=supports(i,j),v_parent=supports(r,j))
21.push_topic(v_topic=supports(e,j),v_parent=null_topic)
28.push_topic(v_topic=supports(f,e),v_parent=supports(e,j))
34.argModusPonensEl(v_agent=ag0,v_conclusion=f,v_parent=supports(f,e),v_premise=d)
35.argModusPonensEl(v_agent=ag0,v_conclusion=f,v_parent=supports(f,e),v_premise=l)
30.make_salient(v_agent=ag0,v_proposition=f,v_topic=supports(f,e))
33.pop_topic(v_topic=supports(f,e),v_parent=supports(e,j))
```

```

23.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
26.pop_topic(v_topic=supports(e,j),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)

[[28,32],[32,33],[28,31],[31,33],[28,30],[30,33],[28,29],[29,33],[23,28],[21,26],[21,26],[21,24],[
24,26],[21,23],[23,26],[21,22],[21,28],[22,26],[33,26],[15,20],[15,20],[15,18],[18,20],[15,17],[17
,20],[15,16],[16,20],[9,14],[9,14],[9,12],[12,14],[9,11],[11,14],[9,10],[10,14],[2,8],[8,6],[2,9],
[2,15],[2,21],[6,15],[20,7],[6,9],[14,7],[start,1],[start,21],[start,end],[start,6],[1,end],[26,en
d],[start,end],[7,end],[start,2],[2,end]]
Applying EC-EPISTEMIC3
Applying EC-EPISTEMIC3
Applying CLUE-PARALLEL
Applying HPR2: Conclusion First
Applying CLUE-DETAIL

```

Complete plan of primitives:

```

51.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(f,e))
50.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(f,e))
36.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(e,j))
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
9.push_topic(v_topic=supports(c,j),v_parent=supports(r,j))
11.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
14.pop_topic(v_topic=supports(c,j),v_parent=supports(r,j))
15.push_topic(v_topic=supports(i,j),v_parent=supports(r,j))
17.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
20.pop_topic(v_topic=supports(i,j),v_parent=supports(r,j))
21.push_topic(v_topic=supports(e,j),v_parent=null_topic)
28.push_topic(v_topic=supports(f,e),v_parent=supports(e,j))
37.push_topic(v_topic=supports(d,f),v_parent=supports(f,e))
39.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,f))
42.pop_topic(v_topic=supports(d,f),v_parent=supports(f,e))
43.push_topic(v_topic=supports(l,f),v_parent=supports(f,e))
49.argModusPonensEl(v_agent=ag0,v_conclusion=l,v_parent=supports(l,f),v_premise=g)
45.make_salient(v_agent=ag0,v_proposition=l,v_topic=supports(l,f))
48.pop_topic(v_topic=supports(l,f),v_parent=supports(f,e))
30.make_salient(v_agent=ag0,v_proposition=f,v_topic=supports(f,e))
33.pop_topic(v_topic=supports(f,e),v_parent=supports(e,j))
23.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
26.pop_topic(v_topic=supports(e,j),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)

```

```

[[43,47],[47,48],[43,46],[46,48],[43,45],[45,48],[43,44],[44,48],[37,41],[41,42],[37,40],[40,42],[
37,39],[39,42],[37,38],[38,42],[23,36],[36,28],[30,37],[30,43],[28,33],[28,33],[28,31],[31,33],[28
,30],[30,33],[28,29],[28,43],[28,37],[29,33],[48,33],[42,33],[23,28],[21,26],[21,26],[21,24],[24,2
6],[21,23],[23,26],[21,22],[21,28],[22,26],[33,26],[15,20],[15,20],[15,18],[18,20],[15,17],[17,20]
,[15,16],[16,20],[9,14],[9,14],[9,12],[12,14],[9,11],[11,14],[9,10],[10,14],[2,8],[8,6],[2,9],[2,1
5],[2,21],[6,15],[20,7],[6,9],[14,7],[start,1],[start,21],[start,end],[start,6],[1,end],[26,end],[
start,end],[7,end],[start,2],[2,end]]
Applying EC-EPISTEMIC3
Applying PUNC_BREAK
Applying CLUE-DETAIL

```

Complete plan of primitives:

```

60.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(l,f))
59.make_salient(v_agent=ag0,v_proposition=break,v_topic=supports(e,j))
51.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(f,e))
50.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(f,e))
36.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(e,j))
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
9.push_topic(v_topic=supports(c,j),v_parent=supports(r,j))
11.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
14.pop_topic(v_topic=supports(c,j),v_parent=supports(r,j))
15.push_topic(v_topic=supports(i,j),v_parent=supports(r,j))
17.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
20.pop_topic(v_topic=supports(i,j),v_parent=supports(r,j))
21.push_topic(v_topic=supports(e,j),v_parent=null_topic)
28.push_topic(v_topic=supports(f,e),v_parent=supports(e,j))
37.push_topic(v_topic=supports(d,f),v_parent=supports(f,e))
39.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,f))
42.pop_topic(v_topic=supports(d,f),v_parent=supports(f,e))
43.push_topic(v_topic=supports(l,f),v_parent=supports(f,e))
52.push_topic(v_topic=supports(g,l),v_parent=supports(l,f))
54.make_salient(v_agent=ag0,v_proposition=g,v_topic=supports(g,l))
58.argModusPonensEl(v_agent=ag0,v_conclusion=supports(g,l),v_parent=supports(g,l),v_premise=h)
57.pop_topic(v_topic=supports(g,l),v_parent=supports(l,f))
45.make_salient(v_agent=ag0,v_proposition=l,v_topic=supports(l,f))
48.pop_topic(v_topic=supports(l,f),v_parent=supports(f,e))
30.make_salient(v_agent=ag0,v_proposition=f,v_topic=supports(f,e))
33.pop_topic(v_topic=supports(f,e),v_parent=supports(e,j))
23.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
26.pop_topic(v_topic=supports(e,j),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)

```

```
[52,56],[56,57],[52,55],[55,57],[52,54],[54,57],[52,53],[53,57],[30,51],[51,37],[45,52],[42,50],[
50,43],[43,48],[43,48],[43,46],[46,48],[43,45],[45,48],[43,44],[43,52],[44,48],[57,48],[37,42],[37
,42],[37,40],[40,42],[37,39],[39,42],[37,38],[38,42],[23,36],[36,28],[30,37],[30,43],[28,33],[28,3
3],[28,31],[31,33],[28,30],[30,33],[28,29],[28,43],[28,37],[29,33],[48,33],[42,33],[23,28],[21,26]
,[21,26],[21,24],[24,26],[21,23],[23,26],[21,22],[21,28],[22,26],[33,26],[15,20],[15,20],[15,18],[
18,20],[15,17],[17,20],[15,16],[16,20],[9,14],[9,14],[9,12],[12,14],[9,11],[11,14],[9,10],[10,14],
[2,8],[8,6],[2,9],[2,15],[2,21],[6,15],[20,7],[6,9],[14,7],[start,1],[start,21],[start,end],[start
,6],[1,end],[26,end],[start,end],[7,end],[start,2],[2,end]]Applying EC-EPISTEMIC3
```

Complete plan of primitives:

```
60.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(l,f))
59.make_salient(v_agent=ag0,v_proposition=break,v_topic=supports(e,j))
51.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(f,e))
50.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(f,e))
36.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(e,j))
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
9.push_topic(v_topic=supports(c,j),v_parent=supports(r,j))
11.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
14.pop_topic(v_topic=supports(c,j),v_parent=supports(r,j))
15.push_topic(v_topic=supports(i,j),v_parent=supports(r,j))
17.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
20.pop_topic(v_topic=supports(i,j),v_parent=supports(r,j))
21.push_topic(v_topic=supports(e,j),v_parent=null_topic)
28.push_topic(v_topic=supports(f,e),v_parent=supports(e,j))
37.push_topic(v_topic=supports(d,f),v_parent=supports(f,e))
39.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,f))
42.pop_topic(v_topic=supports(d,f),v_parent=supports(f,e))
43.push_topic(v_topic=supports(l,f),v_parent=supports(f,e))
52.push_topic(v_topic=supports(g,l),v_parent=supports(l,f))
54.make_salient(v_agent=ag0,v_proposition=g,v_topic=supports(g,l))
61.push_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
63.make_salient(v_agent=ag0,v_proposition=h,v_topic=supports(h,supports(g,l)))
66.pop_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
57.pop_topic(v_topic=supports(g,l),v_parent=supports(l,f))
45.make_salient(v_agent=ag0,v_proposition=l,v_topic=supports(l,f))
48.pop_topic(v_topic=supports(l,f),v_parent=supports(f,e))
30.make_salient(v_agent=ag0,v_proposition=f,v_topic=supports(f,e))
33.pop_topic(v_topic=supports(f,e),v_parent=supports(e,j))
23.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
26.pop_topic(v_topic=supports(e,j),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)
```

```
[61,65],[65,66],[61,64],[64,66],[61,63],[63,66],[61,62],[62,66],[45,60],[60,52],[21,59],[59,26],[
52,57],[52,57],[52,55],[52,61],[55,57],[66,57],[52,54],[54,57],[52,53],[53,57],[30,51],[51,37],[45
,52],[42,50],[50,43],[43,48],[43,48],[43,46],[46,48],[43,45],[45,48],[43,44],[43,52],[44,48],[57,4
8],[37,42],[37,42],[37,40],[40,42],[37,39],[39,42],[37,38],[38,42],[23,36],[36,28],[30,37],[30,43]
,[28,33],[28,33],[28,31],[31,33],[28,30],[30,33],[28,29],[28,43],[28,37],[29,33],[48,33],[42,33],[
23,28],[21,26],[21,26],[21,24],[24,26],[21,23],[23,26],[21,22],[21,28],[22,26],[33,26],[15,20],[15
,20],[15,18],[18,20],[15,17],[17,20],[15,16],[16,20],[9,14],[9,14],[9,12],[12,14],[9,11],[11,14],[
9,10],[10,14],[2,8],[8,6],[2,9],[2,15],[2,21],[6,15],[20,7],[6,9],[14,7],[start,1],[start,21],[sta
rt,end],[start,6],[1,end],[26,end],[start,end],[7,end],[start,2],[2,end]]
```

Complete plan of primitives:

```
60.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(l,f))
59.make_salient(v_agent=ag0,v_proposition=break,v_topic=supports(e,j))
51.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(f,e))
50.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(f,e))
36.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(e,j))
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
9.push_topic(v_topic=supports(c,j),v_parent=supports(r,j))
11.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
14.pop_topic(v_topic=supports(c,j),v_parent=supports(r,j))
15.push_topic(v_topic=supports(i,j),v_parent=supports(r,j))
17.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
20.pop_topic(v_topic=supports(i,j),v_parent=supports(r,j))
21.push_topic(v_topic=supports(e,j),v_parent=null_topic)
28.push_topic(v_topic=supports(f,e),v_parent=supports(e,j))
37.push_topic(v_topic=supports(d,f),v_parent=supports(f,e))
39.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,f))
42.pop_topic(v_topic=supports(d,f),v_parent=supports(f,e))
43.push_topic(v_topic=supports(l,f),v_parent=supports(f,e))
52.push_topic(v_topic=supports(g,l),v_parent=supports(l,f))
54.make_salient(v_agent=ag0,v_proposition=g,v_topic=supports(g,l))
61.push_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
63.make_salient(v_agent=ag0,v_proposition=h,v_topic=supports(h,supports(g,l)))
66.pop_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
57.pop_topic(v_topic=supports(g,l),v_parent=supports(l,f))
45.make_salient(v_agent=ag0,v_proposition=l,v_topic=supports(l,f))
48.pop_topic(v_topic=supports(l,f),v_parent=supports(f,e))
30.make_salient(v_agent=ag0,v_proposition=f,v_topic=supports(f,e))
33.pop_topic(v_topic=supports(f,e),v_parent=supports(e,j))
23.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
26.pop_topic(v_topic=supports(e,j),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)
```

```

with partial order
[[61,64],[64,66],[61,63],[63,66],[61,62],[62,66],[45,60],[60,52],[21,59],[59,26],[52,55],[52,61],[
55,57],[66,57],[52,54],[54,57],[52,53],[53,57],[30,51],[51,37],[42,50],[50,43],[43,46],[46,48],[43
,45],[43,44],[44,48],[57,48],[37,40],[40,42],[37,39],[39,42],[37,38],[38,42],[23,36],[36,28],[28,3
1],[31,33],[28,30],[28,29],[29,33],[48,33],[21,24],[24,26],[21,23],[21,22],[22,26],[33,26],[15,18]
,[18,20],[15,17],[17,20],[15,16],[16,20],[9,12],[12,14],[9,11],[11,14],[9,10],[10,14],[2,8],[8,6]
,[2,21],[6,15],[20,7],[6,9],[14,7],[start,1],[1,end],[26,end],[7,end],[start,2]]

Complete plan of primitives:
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)
8.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
6.push_topic(v_topic=supports(r,j),v_parent=null_topic)
9.push_topic(v_topic=supports(c,j),v_parent=supports(r,j))
11.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
14.pop_topic(v_topic=supports(c,j),v_parent=supports(r,j))
15.push_topic(v_topic=supports(i,j),v_parent=supports(r,j))
17.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
20.pop_topic(v_topic=supports(i,j),v_parent=supports(r,j))
7.pop_topic(v_topic=supports(r,j),v_parent=null_topic)
21.push_topic(v_topic=supports(e,j),v_parent=null_topic)
59.make_salient(v_agent=ag0,v_proposition=break,v_topic=supports(e,j))
23.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
36.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(e,j))
28.push_topic(v_topic=supports(f,e),v_parent=supports(e,j))
30.make_salient(v_agent=ag0,v_proposition=f,v_topic=supports(f,e))
51.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(f,e))
37.push_topic(v_topic=supports(d,f),v_parent=supports(f,e))
39.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,f))
42.pop_topic(v_topic=supports(d,f),v_parent=supports(f,e))
50.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(f,e))
43.push_topic(v_topic=supports(l,f),v_parent=supports(f,e))
45.make_salient(v_agent=ag0,v_proposition=l,v_topic=supports(l,f))
60.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(l,f))
52.push_topic(v_topic=supports(g,l),v_parent=supports(l,f))
54.make_salient(v_agent=ag0,v_proposition=g,v_topic=supports(g,l))
61.push_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
63.make_salient(v_agent=ag0,v_proposition=h,v_topic=supports(h,supports(g,l)))
66.pop_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
57.pop_topic(v_topic=supports(g,l),v_parent=supports(l,f))
48.pop_topic(v_topic=supports(l,f),v_parent=supports(f,e))
33.pop_topic(v_topic=supports(f,e),v_parent=supports(e,j))
26.pop_topic(v_topic=supports(e,j),v_parent=null_topic)

laden with affect on the following goals:
[]

```

Figure C.1 *Rhetorica* processing for the 'tourist facility signs' argument

When *Rhetorica* was run a second time with all EG heuristics deactivated, the final, fully ordered plan of primitives is shown in Figure C.2 (the intermediate, abstract-plan steps are trivially inferable from the detailed account in Figure C.1).

```

6.push_topic(v_topic=supports(c,j),v_parent=null_topic)
8.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,j))
10.make_salient(v_agent=ag0,v_proposition=supports(c,j),v_topic=supports(c,j))
11.pop_topic(v_topic=supports(c,j),v_parent=null_topic)
12.push_topic(v_topic=supports(i,j),v_parent=null_topic)
14.make_salient(v_agent=ag0,v_proposition=i,v_topic=supports(i,j))
16.make_salient(v_agent=ag0,v_proposition=supports(i,j),v_topic=supports(i,j))
17.pop_topic(v_topic=supports(i,j),v_parent=null_topic)
18.push_topic(v_topic=supports(e,j),v_parent=null_topic)
25.push_topic(v_topic=supports(f,e),v_parent=supports(e,j))
33.push_topic(v_topic=supports(d,f),v_parent=supports(f,e))
35.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,f))
37.make_salient(v_agent=ag0,v_proposition=supports(d,f),v_topic=supports(d,f))
38.pop_topic(v_topic=supports(d,f),v_parent=supports(f,e))
39.push_topic(v_topic=supports(l,f),v_parent=supports(f,e))
46.push_topic(v_topic=supports(g,l),v_parent=supports(l,f))
48.make_salient(v_agent=ag0,v_proposition=g,v_topic=supports(g,l))
53.push_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
55.make_salient(v_agent=ag0,v_proposition=h,v_topic=supports(h,supports(g,l)))
57.make_salient(v_agent=ag0,v_proposition=supports(h,supports(g,l)),v_topic=supports(h,supports(g,l)))
58.pop_topic(v_topic=supports(h,supports(g,l)),v_parent=supports(g,l))
50.make_salient(v_agent=ag0,v_proposition=supports(g,l),v_topic=supports(g,l))
51.pop_topic(v_topic=supports(g,l),v_parent=supports(l,f))
41.make_salient(v_agent=ag0,v_proposition=l,v_topic=supports(l,f))
43.make_salient(v_agent=ag0,v_proposition=supports(l,f),v_topic=supports(l,f))
44.pop_topic(v_topic=supports(l,f),v_parent=supports(f,e))

```

```

27.make_salient(v_agent=ag0,v_proposition=f,v_topic=supports(f,e))
29.make_salient(v_agent=ag0,v_proposition=supports(f,e),v_topic=supports(f,e))
30.pop_topic(v_topic=supports(f,e),v_parent=supports(e,j))
20.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,j))
22.make_salient(v_agent=ag0,v_proposition=supports(e,j),v_topic=supports(e,j))
23.pop_topic(v_topic=supports(e,j),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=j,v_topic=null_topic)

```

Figure C.2 *Rhetorica*-minus-EG processing for the ‘tourist facility signs’ argument

The plan in Figure C.2 is realised into the following text:

Since December 1995, when the Department of Transport relaxed the type of premises that could be sign-posted on motorways and trunk roads, increasing numbers of "tourist facility" signs have appeared. They signpost facilities such as pubs, restaurants, shops and nurseries on an apparently permanent basis.

Tourist facility signs add to roadside clutter. There appears to be nothing to stop any concern from joining the rush to promote itself in this way. The signs are fixed to the supports holding road signs. The signs also have the colour and form of traditional “temporary” AA signs, which carry truly important information. Tourist facility signs similarly command the attention of drivers. These signs compete with essential information, and distract motorists.

The highway is no place for advertisements masquerading as traffic signs.

Realising text in the absence of any goals expressing clue phrases is difficult, as flowing text makes frequent recourse to such connectives: the prohibition unavoidably leads to jumpy, disjointed text such as that shown above. The plan created by *Rhetorica* naturally includes each of the three components of each Modus Ponens – and in particular, includes the major premise (since no enthymematic contraction is performed). Realising text which makes explicit the major premise in every Modus Ponens would be hopelessly cumbersome, and would undoubtedly prejudice subjects, and make the PRES less informative. Therefore, the during realisation, every major premise was omitted, in an attempt to make the comparison fairer.

The full processing for the second, ‘Clare Short’ argument is shown in Figure C.3:

```

| ?- example2.
Applying HPR2: Conclusion First

Complete plan of primitives:
3.argModusPonensEl(v_agent=ag0,v_conclusion=e,v_parent=null_topic,v_premise=a)
2.make_salient(v_agent=ag0,v_proposition=e,v_topic=null_topic)

[[start,1],[1,end],[start,2],[2,end]]
Applying EC-EPISTEMIC3
Applying HPR4: Refutations first
Applying HPR2: Conclusion First
Applying CLUE-DETAIL

Complete plan of primitives:
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(a,e),v_parent=null_topic)
10.argModusPonensEl(v_agent=ag0,v_conclusion=a,v_parent=supports(a,e),v_premise=b)

```

```

11.argUCP(v_agent=ag0,v_conclusion=a,v_parent=supports(a,e),v_premise=c)
6.make_salient(v_agent=ag0,v_proposition=a,v_topic=supports(a,e))
9.pop_topic(v_topic=supports(a,e),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=e,v_topic=null_topic)

[[4,8],[8,9],[4,7],[7,9],[4,6],[6,9],[4,5],[5,9],[2,4],[start,1],[start,4],[1,end],[9,end],[start,2],[2,end]]
Applying AFFECT-UC
Applying AFFECT-GOOD1
Applying EC-EPISTEMIC1
Applying EC-EPISTEMIC3
Applying EC-EPISTEMIC4
Applying CLUE-PARALLEL
Applying CLUE-DETAIL

```

Complete plan of primitives:

```

27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(a,e))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(a,e))
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(a,e),v_parent=null_topic)
13.push_topic(v_topic=supports(b,a),v_parent=supports(a,e))
15.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
18.pop_topic(v_topic=supports(b,a),v_parent=supports(a,e))
19.push_topic(v_topic=supports(c,a),v_parent=supports(a,e))
20.make_salient(v_agent=ag0,v_proposition=not c,v_topic=supports(c,a))
25.argModusPonensEl(v_agent=ag0,v_conclusion=c,v_parent=supports(c,a),v_premise=d)
24.pop_topic(v_topic=supports(c,a),v_parent=supports(a,e))
6.make_salient(v_agent=ag0,v_proposition=a,v_topic=supports(a,e))
9.pop_topic(v_topic=supports(a,e),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=e,v_topic=null_topic)

[[19,23],[23,24],[19,22],[22,24],[19,21],[21,24],[19,20],[20,24],[13,17],[17,18],[13,16],[16,18],[13,15],[15,18],[13,14],[14,18],[2,12],[12,4],[6,19],[24,13],[4,9],[4,9],[4,7],[7,9],[4,6],[6,9],[4,5],[4,19],[4,13],[5,9],[24,9],[18,9],[2,4],[start,1],[start,4],[1,end],[9,end],[start,2],[2,end]]
Applying EC-EPISTEMIC3

```

Complete plan of primitives:

```

27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(a,e))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(a,e))
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(a,e),v_parent=null_topic)
13.push_topic(v_topic=supports(b,a),v_parent=supports(a,e))
15.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
18.pop_topic(v_topic=supports(b,a),v_parent=supports(a,e))
19.push_topic(v_topic=supports(c,a),v_parent=supports(a,e))
20.make_salient(v_agent=ag0,v_proposition=not c,v_topic=supports(c,a))
28.push_topic(v_topic=supports(d,c),v_parent=supports(c,a))
30.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,c))
33.pop_topic(v_topic=supports(d,c),v_parent=supports(c,a))
24.pop_topic(v_topic=supports(c,a),v_parent=supports(a,e))
6.make_salient(v_agent=ag0,v_proposition=a,v_topic=supports(a,e))
9.pop_topic(v_topic=supports(a,e),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=e,v_topic=null_topic)

```

```

[[28,32],[32,33],[28,31],[31,33],[28,30],[30,33],[28,29],[29,33],[6,27],[27,13],[24,26],[26,13],[19,24],[19,24],[19,24],[19,21],[19,28],[21,24],[33,24],[19,20],[20,24],[13,18],[13,18],[13,16],[16,18],[13,15],[15,18],[13,14],[14,18],[2,12],[12,4],[6,19],[24,13],[4,9],[4,9],[4,7],[7,9],[4,6],[6,9],[4,5],[4,19],[4,13],[5,9],[24,9],[18,9],[2,4],[start,1],[start,4],[1,end],[9,end],[start,2],[2,end]]

```

Complete plan of primitives:

```

27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(a,e))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(a,e))
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(a,e),v_parent=null_topic)
13.push_topic(v_topic=supports(b,a),v_parent=supports(a,e))
15.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
18.pop_topic(v_topic=supports(b,a),v_parent=supports(a,e))
19.push_topic(v_topic=supports(c,a),v_parent=supports(a,e))
20.make_salient(v_agent=ag0,v_proposition=not c,v_topic=supports(c,a))
28.push_topic(v_topic=supports(d,c),v_parent=supports(c,a))
30.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,c))
33.pop_topic(v_topic=supports(d,c),v_parent=supports(c,a))
24.pop_topic(v_topic=supports(c,a),v_parent=supports(a,e))
6.make_salient(v_agent=ag0,v_proposition=a,v_topic=supports(a,e))
9.pop_topic(v_topic=supports(a,e),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=e,v_topic=null_topic)

```

with partial order

```

[[28,31],[31,33],[28,30],[30,33],[28,29],[29,33],[6,27],[27,13],[24,26],[26,13],[19,21],[19,28],[21,24],[33,24],[19,20],[20,24],[13,16],[16,18],[13,15],[15,18],[13,14],[14,18],[2,12],[12,4],[6,19],[4,7],[7,9],[4,6],[4,5],[5,9],[18,9],[start,1],[1,end],[9,end],[start,2]]

```

Complete plan of primitives:

```

2.make_salient(v_agent=ag0,v_proposition=e,v_topic=null_topic)
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(a,e),v_parent=null_topic)
6.make_salient(v_agent=ag0,v_proposition=a,v_topic=supports(a,e))
27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(a,e))

```

```

19.push_topic(v_topic=supports(c,a),v_parent=supports(a,e))
20.make_salient(v_agent=ag0,v_proposition=not c,v_topic=supports(c,a))
28.push_topic(v_topic=supports(d,c),v_parent=supports(c,a))
30.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,c))
33.pop_topic(v_topic=supports(d,c),v_parent=supports(c,a))
24.pop_topic(v_topic=supports(c,a),v_parent=supports(a,e))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(a,e))
13.push_topic(v_topic=supports(b,a),v_parent=supports(a,e))
15.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
18.pop_topic(v_topic=supports(b,a),v_parent=supports(a,e))
9.pop_topic(v_topic=supports(a,e),v_parent=null_topic)

laden with affect on the following goals:
[[20,bad]]

```

Figure C.3 *Rhetorica* processing for the 'Clare Short' argument

Again, *Rhetorica* was also run on the same input with EG processing inactive: the output is given in Figure C.4:

```

Complete plan of primitives:
4.push_topic(v_topic=supports(a,e),v_parent=null_topic)
12.push_topic(v_topic=supports(b,a),v_parent=supports(a,e))
14.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
16.make_salient(v_agent=ag0,v_proposition=supports(b,a),v_topic=supports(b,a))
17.pop_topic(v_topic=supports(b,a),v_parent=supports(a,e))
18.push_topic(v_topic=supports(c,a),v_parent=supports(a,e))
19.make_salient(v_agent=ag0,v_proposition=not c,v_topic=supports(c,a))
25.push_topic(v_topic=supports(d,c),v_parent=supports(c,a))
27.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,c))
29.make_salient(v_agent=ag0,v_proposition=supports(d,c),v_topic=supports(d,c))
30.pop_topic(v_topic=supports(d,c),v_parent=supports(c,a))
21.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,a))
22.make_salient(v_agent=ag0,v_proposition=supports(not c,not a),v_topic=supports(c,a))
23.pop_topic(v_topic=supports(c,a),v_parent=supports(a,e))
6.make_salient(v_agent=ag0,v_proposition=a,v_topic=supports(a,e))
8.make_salient(v_agent=ag0,v_proposition=supports(a,e),v_topic=supports(a,e))
9.pop_topic(v_topic=supports(a,e),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=e,v_topic=null_topic)

```

Figure C.4 *Rhetorica*-minus-EG processing for the 'Clare Short' argument

The text from Figure C.4 is as follows:

For Clare Short, the wait is over. Sadly, however, if Toby Graham had not searched for her, she would still be suffering in silence, as are many other birth parents. The Contact Register should help parents find their children. Few people know of its existence: it is of limited value. Contact can only be made if the adoptee searches out their birth parents. Is it not time that the law concerning contact between adults after adoption is revised to enable the birth parent to have identifying information once the adoptee is 18 or possibly 25?

The full processing for the third, 'Irish census' argument is shown in Figure C.5:

| ?- example3.

Applying HPR2: Conclusion First

Complete plan of primitives:

```
3.argModusPonensEl(v_agent=ag0,v_conclusion=a,v_parent=null_topic,v_premise=b)
2.make_salient(v_agent=ag0,v_proposition=a,v_topic=null_topic)
```

```
[[start,1],[2,end],[start,2],[1,end]]
```

Applying EC-EPISTEMIC3

Applying HPR2: Conclusion First

Applying CLUE-DETAIL

Applying HPR1: Climax ordering

Complete plan of primitives:

```
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(b,a),v_parent=null_topic)
10.argModusPonensEl(v_agent=ag0,v_conclusion=b,v_parent=supports(b,a),v_premise=d)
11.argModusPonensEl(v_agent=ag0,v_conclusion=b,v_parent=supports(b,a),v_premise=c)
6.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
9.pop_topic(v_topic=supports(b,a),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=a,v_topic=null_topic)
```

```
[[4,8],[8,9],[4,7],[7,9],[4,6],[6,9],[4,5],[5,9],[2,4],[start,1],[start,4],[2,end],[start,2],[1,end],
[9,end]]
```

Applying EC-EPISTEMIC3

Applying EC-EPISTEMIC3

Applying CLUE-PARALLEL

Applying HPR2: Conclusion First

Applying CLUE-DETAIL

Complete plan of primitives:

```
27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(b,a))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(b,a))
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(b,a),v_parent=null_topic)
13.push_topic(v_topic=supports(d,b),v_parent=supports(b,a))
25.argModusPonensEl(v_agent=ag0,v_conclusion=d,v_parent=supports(d,b),v_premise=e)
15.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,b))
18.pop_topic(v_topic=supports(d,b),v_parent=supports(b,a))
19.push_topic(v_topic=supports(c,b),v_parent=supports(b,a))
21.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,b))
24.pop_topic(v_topic=supports(c,b),v_parent=supports(b,a))
6.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
9.pop_topic(v_topic=supports(b,a),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=a,v_topic=null_topic)
```

```
[[19,23],[23,24],[19,22],[22,24],[19,21],[21,24],[19,20],[20,24],[13,17],[17,18],[13,16],[16,18],[
13,15],[15,18],[13,14],[14,18],[2,12],[12,4],[6,13],[6,19],[4,9],[4,9],[4,7],[7,9],[4,6],[6,9],[4,
5],[4,19],[4,13],[5,9],[24,9],[18,9],[2,4],[start,1],[start,4],[2,end],[start,2],[1,end],[9,end]]
```

Applying EC-EPISTEMIC3

Applying CLUE-DETAIL

Complete plan of primitives:

```
34.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(d,b))
27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(b,a))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(b,a))
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(b,a),v_parent=null_topic)
13.push_topic(v_topic=supports(d,b),v_parent=supports(b,a))
28.push_topic(v_topic=supports(e,d),v_parent=supports(d,b))
30.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,d))
33.pop_topic(v_topic=supports(e,d),v_parent=supports(d,b))
15.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,b))
18.pop_topic(v_topic=supports(d,b),v_parent=supports(b,a))
19.push_topic(v_topic=supports(c,b),v_parent=supports(b,a))
21.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,b))
24.pop_topic(v_topic=supports(c,b),v_parent=supports(b,a))
6.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
9.pop_topic(v_topic=supports(b,a),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=a,v_topic=null_topic)
```

```
[[28,32],[32,33],[28,31],[31,33],[28,30],[30,33],[28,29],[29,33],[6,27],[27,13],[15,28],[18,26],[2
6,19],[19,24],[19,24],[19,22],[22,24],[19,21],[21,24],[19,20],[20,24],[13,18],[13,18],[13,16],[16,
18],[13,15],[15,18],[13,14],[14,18],[33,28],[14,18],[33,18],[2,12],[12,4],[6,13],[6,19],[4,9],[4,9],[4,7],
[7,9],[4,6],[6,9],[4,5],[4,19],[4,13],[5,9],[24,9],[18,9],[2,4],[start,1],[start,4],[2,end],[start
,2],[1,end],[9,end]]
```

Complete plan of primitives:

```
34.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(d,b))
27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(b,a))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(b,a))
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(b,a),v_parent=null_topic)
13.push_topic(v_topic=supports(d,b),v_parent=supports(b,a))
28.push_topic(v_topic=supports(e,d),v_parent=supports(d,b))
30.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,d))
```

```

33.pop_topic(v_topic=supports(e,d),v_parent=supports(d,b))
15.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,b))
18.pop_topic(v_topic=supports(d,b),v_parent=supports(b,a))
19.push_topic(v_topic=supports(c,b),v_parent=supports(b,a))
21.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,b))
24.pop_topic(v_topic=supports(c,b),v_parent=supports(b,a))
6.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
9.pop_topic(v_topic=supports(b,a),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=a,v_topic=null_topic)

with partial order
[[15,34],[34,28],[28,31],[31,33],[28,30],[30,33],[28,29],[29,33],[6,27],[27,13],[18,26],[26,19],[1
9,22],[22,24],[19,21],[21,24],[19,20],[20,24],[13,16],[16,18],[13,15],[13,14],[14,18],[33,18],[2,1
2],[12,4],[4,7],[7,9],[4,6],[4,5],[5,9],[24,9],[start,1],[start,2],[1,end],[9,end]]

Complete plan of primitives:
2.make_salient(v_agent=ag0,v_proposition=a,v_topic=null_topic)
12.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=null_topic)
4.push_topic(v_topic=supports(b,a),v_parent=null_topic)
6.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
27.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(b,a))
13.push_topic(v_topic=supports(d,b),v_parent=supports(b,a))
15.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,b))
34.make_salient(v_agent=ag0,v_proposition=clue(detail),v_topic=supports(d,b))
28.push_topic(v_topic=supports(e,d),v_parent=supports(d,b))
30.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,d))
33.pop_topic(v_topic=supports(e,d),v_parent=supports(d,b))
18.pop_topic(v_topic=supports(d,b),v_parent=supports(b,a))
26.make_salient(v_agent=ag0,v_proposition=clue(parallel),v_topic=supports(b,a))
19.push_topic(v_topic=supports(c,b),v_parent=supports(b,a))
21.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,b))
24.pop_topic(v_topic=supports(c,b),v_parent=supports(b,a))
9.pop_topic(v_topic=supports(b,a),v_parent=null_topic)

laden with affect on the following goals:
[]

```

Figure C.5 *Rhetorica* processing for the ‘Irish census’ argument

The processing carried out with the EG level deactivated is shown in Figure C.6

```

Complete plan of primitives:
4.push_topic(v_topic=supports(b,a),v_parent=null_topic)
12.push_topic(v_topic=supports(d,b),v_parent=supports(b,a))
25.push_topic(v_topic=supports(e,d),v_parent=supports(d,b))
27.make_salient(v_agent=ag0,v_proposition=e,v_topic=supports(e,d))
29.make_salient(v_agent=ag0,v_proposition=supports(e,d),v_topic=supports(e,d))
30.pop_topic(v_topic=supports(e,d),v_parent=supports(d,b))
14.make_salient(v_agent=ag0,v_proposition=d,v_topic=supports(d,b))
16.make_salient(v_agent=ag0,v_proposition=supports(d,b),v_topic=supports(d,b))
17.pop_topic(v_topic=supports(d,b),v_parent=supports(b,a))
18.push_topic(v_topic=supports(c,b),v_parent=supports(b,a))
20.make_salient(v_agent=ag0,v_proposition=c,v_topic=supports(c,b))
22.make_salient(v_agent=ag0,v_proposition=supports(c,b),v_topic=supports(c,b))
23.pop_topic(v_topic=supports(c,b),v_parent=supports(b,a))
6.make_salient(v_agent=ag0,v_proposition=b,v_topic=supports(b,a))
8.make_salient(v_agent=ag0,v_proposition=supports(b,a),v_topic=supports(b,a))
9.pop_topic(v_topic=supports(b,a),v_parent=null_topic)
2.make_salient(v_agent=ag0,v_proposition=a,v_topic=null_topic)

```

Figure C.6 *Rhetorica*-minus-EG processing for the ‘Irish census’ argument

Finally, the text of this argument is given below:

Tomorrow, the Office of National Statistics will take a decision affecting Irish people in Britain for the next 15 years. It concerns the inclusion of an Irish category in the Ethnic Group question in the census.

The disadvantages of Britain's Irish community are striking when compared to the indigenous white population and when compared to other economic minorities. Irish people in Britain experience poorer health, higher rates of mortality and economic disadvantages which are passed from one generation to the next. The planning of social services, housing, healthcare and many other support services depends on the collection of accurate census statistics. It has become increasingly apparent since 1991 that the lack of such census information places Britain's Irish community at a serious disadvantage. We would like to offer our wholehearted support for a separate category for Irish people in the Ethnic Group question in the 2001 census.

We would urge the ONS to put an end to the anomalous situation where the largest ethnic-minority group in Britain is invisible.

The nine arguments were presented to subjects as a web page, with each group of three arguments on a single topic demarcated as a 'set' (set one was thus the three versions of the 'tourist facility signs' argument, etc.). Within each set, the order of arguments was randomised thus:

Tourist facility signs:	Full Rhet, Orig, No Frills
Clare Short::	Orig, No Frills, Full Rhet
Irish census:	No Frills, Orig, Full Rhet

The guidance text was limited to the following:

The questionnaire is comprised of three sets of texts, each text in a set representing one version of a "letter to the editor" of a national newspaper. For each set, you are asked to rate the texts on the basis of how persuasive you find them. Following each text is a box: please enter a number between 0 and 9 to indicate whether you found the argument highly persuasive (9) or totally hopeless (0).

No additional guidance was given to subjects either on the web page or orally.

C.2 The Results

The results were collected over a three day period from a variety of colleagues and personal acquaintances who kindly agreed to participate. In the table of results in Figure C.7, subjects' names have been replaced by a broad professional classification. The last three rows indicate the number of times a given argument form was rated first, second, or third. If two or more of a subject's ratings in a given set were identical, no update was made to the ranking accumulators (this leads to percentage totals less than 100% in the analyses).

Figure C.7 The results of PRES

Index of First Authors

- | | | | | | |
|----------------------------------|-----------------------|-------------------|-------------------------|--------------------------|------------------------|
| Abelson | 40 | Farley | 19, 28 | Jones | 24, 122 |
| Aleven | 34 | Fawcett | 19, 47 | Jonsson | 24 |
| Alexy | 28 | Feteris | 18, 28 | Kahneman | 156 |
| Allen | 20, 41 | Fikes | 21 | Kamp | 27, 59 |
| Anscombe | 35 | Finin | 17, 20 | Kass | 17 |
| Antoniou | 33 | Fink | 21 | Kemper | 37 |
| Apostel | 18 | Finocchiaro | 30 | Kienpointner | 30 |
| Appelt | 17, 24 | Fisher | 80, 131 | Knoblock | 22 |
| Aristotle | 27, 126 | Flowers | 35 | Knott | 25, 69, 112, |
| Asher | 27 | Fogelin | 31 | | 115, 118, 159 |
| Austin | 12 | Fox | 21, 23, 26, 27, | Konolige | 17 |
| Bacchus | 22, 65 | | 33, 34, 63, 64, 79, 100 | Kowalski | 34 |
| Baddeley | 156 | Franklin | 37 | Krause | 33, 34 |
| Ballim | 61 | Freeman | 19, 29, 30, 48, | Kuhn | 32 |
| Bavelas | 32 | | 71, 84, 91, 101, 150 | Laasri | 19 |
| Billig | 27, 28 | Gabbay | 33 | Labrou | 20 |
| Binstead | 42 | Galliers | 60, 149 | Laver | 32 |
| Birnbaum | 35 | Gardenfors | 60 | Lee | 61 |
| Blair | 28, 30, 32, 42, 44, | Gazen | 23 | Leopold-Wildburger | 18 |
| | 47, 57, 71, 102, 130 | Geffner | 33 | Levelt | 52, 53, 55 |
| Blandford | 20, 91 | Genesereth | 19 | Levin | 17 |
| Bonarini | 17 | Gerevini | 23 | Locke | 87, 159 |
| Bratman | 20 | Gilbert | 31, 32 | Long | 160 |
| Briscoe | 122 | Girle | 18 | Lorenz | 18 |
| Brockriede | 33, 44 | Goldman | 27 | Loui | 33 |
| Brooks | 103, 112 | Goodman | 17 | Lowder | 159 |
| Brown | 32 | Gordon | 34 | Lund | 104 |
| Calistri-Yeh | 17 | Grasso | 42 | Mackenzie | 19 |
| Carberry | 17 | Gray | 97 | Mann | 24, 27, 36, 52, 55, |
| Carlson | 18, 97 | Green | 32 | | 56, 59, 62, 76 |
| Castelfranchi | 37 | Grice | 13, 18, 62 | Marcu | 25, 27, 56, 58, |
| Cavalli-Sforza | 34 | Grimes | 24 | | 62, 130, 159 |
| Cawsey | 26 | Grosz | 25, 59, 76, 80, | Matthiessen | 55 |
| Chapman | 21 | | 92, 93, 111, 164 | Matwin | 34 |
| Cohen | 5, 14, 24, 36, 57, | Grote | 69 | Maybury | 26, 35, 36, 42, |
| | 59, 60, 62, 66, 69, | Haggith | 34, 82 | | 61, 71, 162 |
| | 78, 86, 97, 112, 118, | Haller | 27 | Mayfield | 17 |
| ... 119, 127, 128, 150, 157, 164 | | Hamblin | 28, 31 | McAllester | 21 |
| Cole | 24, 27 | Hansen | 31 | McConachy | 35 |
| Conklin | 34 | Harsanyi | 37 | McCoy | 17, 25, 76, 93, 96 |
| Cooley | 30, 84 | Hartley | 42 | McDermott | 22, 95 |
| Cooper | 42 | Hass | 104 | McGuire | 66, 105, 127, 150 |
| Copi | 87 | Heidrich | 18 | McKeown | 24, 76, 77, 79, 93 |
| Craig | 32 | Hewitt | 37 | Mellish | 24, 27 |
| Currie | 23 | Hintikka | 17, 18 | Meteer | 27, 51, 52, 55, 114 |
| D'Inverno | 37 | Hirst | 27, 42 | Mooney | 36 |
| Dahlgreen | 56 | Hobbs | 25, 59, 112 | Moore | 25, 36, 55, 59, |
| Dale | 24, 27 | Horacek | 27 | | 62, 71, 78 |
| Daradoumis | 19 | Hovland | 63, 66, 82, 104, 152 | Morik | 17 |
| Darnell | 103 | Hovy | 16, 24, 25, 26, | Moser | 26, 60, 77, 78, 92 |
| Das | 33, 60, 82 | | 53, 55, 56, 59, | Nash | 37 |
| de | 17, 27 | | 61, 63, 69, 71, | Newell | 21 |
| Derr | 76 | | 77, 91, 93, 123, | Nicholas | 57 |
| DiMarco | 27, 42, 98, 133 | | 133, 135, 162, 164 | Nunberg | 122 |
| Doyle | 17 | Huang | 55, 78, 91 | Nwana | 19, 37 |
| Dung | 38 | Inui | 27 | O'Keefe | 32, 44 |
| Eemeren | 29, 30, 48 | Ito | 136 | Ohlsson | 80 |
| Elhadad | 27, 35, 55, 162 | Jackson | 34 | Panaget | 51, 52, 114 |
| Eller | 17 | Jameson | 16 | Paris | 17, 24, 27, 36, 42, 62 |
| Ellis | 56 | Janis | 104 | Parsons | 33, 35, 38 |
| Elvang-Goransson | 33, 38 | Jennings | 19 | Pascual | 124 |
| Erol | 23 | Johnson | 28, 29, 30, | Pednault | 22 |
| Etzioni | 22 | | 86, 112, 167 | Peh | 122 |

- | | | | | | |
|--------------------|-----------------------|-----------------------|----------------|------------------|-----------------|
| Penberthy..... | 21, 71, 162 | Sandell | 157 | Tracy..... | 32 |
| Penman | 32 | Sanders | 115 | Triandis..... | 105, 127 |
| Perelman..... | 28, 43, 101, 102, 150 | Schegloff..... | 18 | Turing | 46 |
| Pilkington | 19, 34 | Scott..... | 25, 27 | Turner | 20 |
| Pollock..... | 33 | Searle..... | 13 | Verheij..... | 34 |
| Prakken..... | 33, 34 | Shadbolt..... | 21 | Vorobej..... | 42, 50, 67 |
| Quine | 56 | Shiu..... | 63 | Vreeswijk..... | 20, 33, 38 |
| Quintilian..... | 102, 159 | Shoham..... | 20 | Wahlster..... | 17, 27 |
| Rambow..... | 24 | Sider..... | 27 | Waldron | 32 |
| Rao..... | 20 | Sidner..... | 76 | Walker | 62, 126, 129, |
| Reed..... | 19, 26, 32, 33, 35, | Sierra | 19 | | 130, 159, 162 |
| | 40, 42, 47, 57, 60, | Sillince..... | 33 | Walton | 19, 31, 36, |
| | 65, 66, 70, 79, 80, | Simard..... | 122 | | 43, 50, 87, 97 |
| | 81, 91, 94, 98, 102, | Simons | 82, 102 | Wellman | 35 |
| | 121, 122, 130 | Slot | 136 | Whately..... | 28, 31, 83, |
| Reichman..... | 59, 78 | Smeltzer | 32 | | 98, 102, 103 |
| Reichman-Adar..... | 112 | Smith | 19, 23, 55, 63 | Wilkins | 21 |
| Reiter | 22, 24, 42, 67 | Snoeck Henkemans..... | 36, 57, 59, | Willard..... | 32, 87 |
| Reithinger | 27, 53, 61 | | 69, 112, 122 | Williams | 24 |
| Rescher | 28, 131, 160 | Sparck | 17 | Wilson | 31, 36, 86, 112 |
| Rich | 17 | Sperber..... | 15, 156 | Woods..... | 28, 87, 159 |
| Rips..... | 32 | Stede | 132 | Wooldridge | 19, 37 |
| Rittel | 34 | Stefik..... | 21, 23 | Yang | 22 |
| Robin | 27, 159 | Strawson | 15 | Young | 22, 26, 78, 92 |
| Rosner..... | 42 | Sussman | 21 | Zlotkin | 37, 38 |
| Rubinoff..... | 114 | Suthers | 75 | Zukerman..... | 35, 162 |
| Sacerdoti..... | 21, 22, 26, 65, 93 | Sycara | 35, 39, 163 | | |
| Sadock | 29, 62, 127, 130 | Toulmin | 29, 48, 82 | | |

Index

AbNLP, 27, 63, 64, 65, 79, 100

Acceptability, 38

Achievement goal, 71

AFFECT-GOOD1, 134

AFFECT-GOOD2, 134

AFFECT-UC, 134

Agent autonomy, 37

AOP, 20

AP, 88

Argument₁, 32

Argument₂, 32

Argumentum ad populum, 87

Canonical orderings, 97

CATO, 34

Centering theory, 77

Clue words, 36, 54

CLUE-DETAIL, 121

CLUE-DS, 135

CLUE-INFERENCE, 120

CLUE-MT, 115

CLUE-PARALLEL, 119

CLUE-UCI, 118

CLUE-UCP, 118

Coherence, 56

Communicative acts, 36

Communicative goals, 26

Conjunction, 80

Context spaces, 78

Cues, 36

Diagrammation, 30

Dialogue games, 17

Dialogue types, 40

Discourse markers, 36

Discourse segments, 77

Disjunctive constraints, 93, 94

Disjunctive Syllogisms, 80

Distributed defeasible reasoning, 38

Dominance, 92

DPOCL, 26

DRAFTER, 42

DRT, 59

DS, 82

EC-EPISTEMIC, 128

EC-LENGTH, 129

EC-QUALIFIER1, 132

EC-QUALIFIER2, 132

EES, 27, 36

Encapsulation, 64

Enveloping structures, 57

EPICURE, 27

Epistemic entrenchment, 60

Explanation, 36

Face goals, 32

Face-to-face persuasive monologue, 44

Fallacies, 31

Focus spaces, 92

Focus trees, 76, 93

FORA, 34

Formulator, 52

FUF, 27, 55

Fully-engaged dialogue, 44

Functional embedding, 40

General audience, 43

Generation gap, 52

Global focus, 76

Goldfish, 42

Gorgias, 27

Graphplan, 23

Ground adequacy question, 48

HAS_PROPERTY, 86, 87, 98

HCI1, 100

HCR1, 98

HCR2, 99

HealthDoc, 27, 42, 98

HPP1, 104

HPR1, 102

HPR2, 103

HPR3, 103

HPR4, 104

HTN, 23

HYPO, 34

Hypothetical Syllogism, 80

IDAS, 24, 42

IE, 88

IG, 86

Ignoratio elenchi, 88

Implicit discussion, 48

Inductive generalisation, 84, 85

IRU, 62

ISUP, 86

Iteration problem, 71

KQML, 19

LA, 33

Linguistic goals, 26

LOLITA, 24, 55, 63

Longbow, 26

Macroplanning, 52

MAKE_SALIENT, 89, 95

Microplanning, 52

Modus Brevis, 29, 62

Modus Ponens, 80

Modus Tollens, 80

MP, 82

MT, 82

Multiple subarguments, 33

Mutual belief, 60

Mutual initial discourse context, 71

- NAG, 35
- Negoplan, 34
- Negotiation, 37
- Negotiation set, 37
- NIGEL, 55
- NOAH, 65
- Non-engaged dialogue, 44
- Nuclearity, 57

- OFFICE-PLAN, 27

- Parallel structures, 57
- Particular audience, 43
- PAULINE, 24, 71, 98
- PENMAN, 27, 55
- PERSUADER, 39
- Persuasive monologues, 42
- Phantom node, 85
- Phatic communication, 32
- Piglit, 42
- Pipeline, 51
- Politeness, 32
- POP_TOPIC, 93, 96
- POPEL, 27, 53
- PROVERB, 55, 78
- PUNC-BREAK, 124
- PUNC-FN1, 125
- PUNC-FN2, 126
- PUSH_TOPIC, 93, 96

- RAGs, 42
- Rebutting, 38, 82
- Refinement, 65
- Relevance, 15

- Relevance question, 48
- Retraction, 50, 79
- Rhetoric, 27
- RHETORIC-PE, 131
- RST, 36, 54, 55, 56, 57, 58, 76, 77, 97

- SAGE, 27
- Satisfaction-precedence, 92
- SemNet, 63
- Sentence planning, 53
- Sentence realisation, 51
- Smoking Letters, 42, 67
- Speech acts, 13
- Standard treatment, 29
- Subargument integrity, 66, 95
- SURGE, 55

- TECHDOC, 42
- TEXT, 24, 77
- Text planning, 51
- Topic stack, 79
- Toulmin schema, 29
- Trains, 41

- UCI, 84
- UCP, 84
- UCPOP, 71
- Undercutting, 38, 82
- User modelling, 16

- Veiled persuasive monologue, 44
- VERBMOBIL, 27
- VIEWGEN, 61

- Zeno Framework, 34

References

- (Abelson, 1959) Abelson, R.P. (1959) "Modes of resolution of belief dilemmas", *Journal of Conflict Resolution* 3, pp343-352
- (Aleven and Ashley, 1992) Aleven, V. & Ashley, K.D. (1992) "Automated Generation of Examples for a Tutorial in Case-Based Argumentation" in *Proceedings of the 2nd International Conference on Intelligent Tutoring Systems*, pp575-584
- (Aleven and Ashley, 1994) Aleven, V. & Ashley, K.D. (1994) "An Instructional Environment for Practising Argumentation Skills" in *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI'94)*, AAAI, pp485-492
- (Alexy, 1989) Adler, R. & MacCormick, N. (trans) (1989) Alexy, R., "A Theory of Legal Argumentation", Clarendon Press, Oxford
- (Allen *et al.*, 1995) Allen, J.F., Schubert, L.K., Ferguson, G., Heeman, P., et al (1995) "The TRAINS project: a case study in building a conversational planning agent", *Journal of Experimental and Theoretical Artificial Intelligence* 7 (1), pp7-48
- (Anscombre and Ducrot, 1983) Anscombre, J.C. & Ducrot, O. (1983) "Philosophie et langage: L'argumentation dans la langue", Pierre Mardaga, Bruxelles
- (Antoniou, 1996) Antoniou, G. (1996) "A comparative survey of Default Logic Variants" in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085)*, Springer, pp15-28
- (Apostel, 1982) Apostel, L. (1982) "Towards a General Theory of Argumentation" in Barth, E.M., Martens, J.L., (eds), *Argumentation: Approaches to Theory Formation*, John Benjamins, Amsterdam, pp93-122
- (Appelt and Pollack, 1992) Appelt, D.E. & Pollack, M.E. (1992) "Weighted Abduction for Plan Ascription", *User Modelling and User Adapted Interaction* 2, pp1-25
- (Appelt, 1985) Appelt, D.E. (1985) "Planning English Sentences", Cambridge University Press, Cambridge
- (Aristotle, 1926) Freese, J.H. (trans) (1926) Aristotle, "The Art of Rhetoric", Heinmann: Loeb Classical Library, London
- (Asher and Lascarides, 1994) Asher, N. & Lascarides, A. (1994) "Intentions and Information in Discourse" in *Proceedings of the Meeting of the Association for Computational Linguistics (ACL'94)*, pp34-41
- (Austin, 1976) Austin, J.L. (1976) "How to do Things with Words (Second Edition)", OUP, Oxford
- (Bacchus and Yang, 1991) Bacchus, F. & Yang, Q. (1991) "The Downward Refinement Property" in *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI'91)*, Sydney, pp286-292
- (Bacchus and Yang, 1992) Bacchus, F. & Yang, Q. (1992) "The Expected Value of Hierarchical Problem-Solving" in *Proceedings of the National Conference on Artificial Intelligence (AAAI-92)*, AAAI, pp369-374
- (Bacchus and Yang, 1994) Bacchus, F. & Yang, Q. (1994) "Downward Refinement and the Efficiency of Hierarchical Problem Solving", *Artificial Intelligence* 71 (1), pp43-100
- (Baddeley, 1986) Baddeley, A.P. (1986) "Working Memory", Oxford University Press, Oxford
- (Ballim and Wilks, 1991) Ballim, A. & Wilks, Y. (1991) "Artificial Believers", Lawrence Erlbaum Associates, Hillsdale, NJ
- (Bavelas *et al.*, 1990) Bavelas, J.B., Black, A., Chovil, N. & Mullett, J. (1990) "Truths, Lies, and Equivocations: The Effects of Conflicting Goals on Discourse", *Journal of Language and Social Psychology* 9 (1-2), pp135-161
- (Billig, 1996) Billig, M. (1996) "Arguing and Thinking: A Rhetorical Approach to Social Psychology: 2nd Edition", Cambridge University Press, Cambridge
- (Binstead *et al.*, 1995) Binstead, K., Cawsey, A. & Jones, R. (1995) "Generating Personalised Patient Information Using the Medical Record" in *Proceedings of the Conference on Artificial Intelligence in Medicine, Europe (AIME'95)*
- (Birnbaum, 1982) Birnbaum, L. (1982) "Argument Molecules: A Functional Representation of Argument Structure" in *Proceedings of the 2nd National Conference on Artificial Intelligence (AAAI-82)*, AAAI, Pittsburgh, PA, pp63-65
- (Blair, 1838) Blair, H. (1838) "Lectures on Rhetoric and Belles Lettres", Charles Daly, London
- (Blair, 1996) Blair, J.A. (1996) "Notes for the Tutorial on Evaluating Arguments at the First International Conference on Formal and Applied Practical Reasoning (FAPR'96)", (unpublished manuscript), Bonn

- (Blair, 1997) Blair, J.A. (1997) "The Limits of the Dialogue Model of Argument" in Proceedings of the OSSA Conference on Argument and Rhetoric, (to appear), St. Catharines, Canada
- (Blandford, 1993) Blandford, A.E. (1993) "An agent-theoretic approach to computer participation in dialogue", *International Journal of Man-Machine Studies* 39, pp965-998
- (Bonarini, 1987) Bonarini, A. (1987) "User Modeling in Person-Machine Conversation: A Computational Approach" in Proceedings of COGNITIVA'87, Paris, pp377-382
- (Bratman *et al.*, 1988) Bratman, M.E., Israel, D.J. & Pollack, M.E. (1988) "Plans & resource-bounded practical reasoning", *Computational Intelligence* 4, pp349-355
- (Briscoe, 1996) Briscoe, E. (1996) "The Syntax and Semantics of Punctuation and its Use in Interpretation" in Jones, B., (eds), *Proceedings of the ACL/SIGPARSE Workshop on Punctuation in Computational Linguistics*, Santa Cruz, CA, pp1-8
- (Brockriede, 1975) Brockriede, W. (1975) "Where is Argument?", *Journal of the American Forensic Association* 11, pp179-182
- (Brockriede, 1977) Brockriede, W. (1977) "Characteristics of Arguments and Arguing", *Journal of the American Forensic Association* 13, pp129-134
- (Brooks and Warren, 1972) Brooks, C. & Warren, R.P. (1972) "Modern rhetoric (Shorter Third Edition)", Harcourt Brace Jovanovich, New York
- (Brown and Levinson, 1987) Brown, P. & Levinson, S.C. (1987) "Politeness: Some universals in language usage", Cambridge University Press, Cambridge
- (Calistri-Yeh, 1991) Calistri-Yeh, R.J. (1991) "Utilizing User Models to Handle Ambiguity and Misconception in Robust Plan Recognition", *User Modelling and User Adapted Interaction* 1, pp289-322
- (Carberry and Pope, 1993) Carberry, S. & Pope, W.A. (1993) "Plan recognition strategies for language understanding", *International Journal on Man-Machine Studies* 39, pp529-577
- (Carlson, 1983) Carlson, L. (1983) "Dialogue games: An approach to discourse analysis", Reidel, Dordrecht
- (Castelfranchi, 1995) Castelfranchi, C. (1995) "Guarantees for autonomy in cognitive agent architectures" in Wooldridge, M.J., Jennings, N.R., (eds), *Proceedings of the Workshop on Agent Theories, Architectures and Languages (ATAL94)*, LNAI Vol. 890, Springer Verlag, Berlin, pp56-70
- (Cavalli-Sforza *et al.*, 1992) Cavalli-Sforza, V., Lesgold, A.M. and Weiner, A.W. (1992) "Strategies for Contributing to Collaborative Arguments" in Proceedings of the 14th Conference of the Cognitive Science Society, pp755-760
- (Cavalli-Sforza *et al.*, 1993) Cavalli-Sforza, V., Moore, J.D. and Suthers, D.D. (1993) "Helping Students Articulate, Support, and Criticize Scientific Explanations" in Proceedings of the World Conference on Artificial Intelligence in Education, pp113-120
- (Cawsey, 1993) Cawsey, A. (1993) "Explanation and Interaction: The Computer Generation of Explanatory Dialogues", MIT Press, Cambridge
- (Chapman, 1987) Chapman, D. (1987) "Planning for Conjunctive Goals", *Artificial Intelligence* 32, pp333-377
- (Cohen and Levesque, 1990) Cohen, P.R. & Levesque, H.J. (1990) "Rational Interaction as the Basis for Communication" in Cohen, P.R., Morgan, J. & Pollack, M.E., (eds), *Intentions in Communication*, MIT Press, Boston, pp221-255
- (Cohen and Perrault, 1988) Cohen, P.R. & Perrault, C.R. (1988) "Elements of a Plan-Based Theory of Speech Acts", *Readings in Distributed AI*, pp169-186
- (Cohen, 1984) Cohen, R. (1984) "A Computational Theory of the Function of Clue Words in Argument Understanding" in Proceedings of the 10th International Conference on Computational Linguistics (COLING'84), pp251-258
- (Cohen, 1987) Cohen, R. (1987) "Analyzing the Structure of Argumentative Discourse", *Computational Linguistics* 13 (1), pp11-24
- (Cohen, 1989) Cohen, R. (1989) "Studying Arguments to Gain Insight into Discourse Structure" in Taylor, M.M., Neel, F. & Bouwhuis, D.G., (eds), *The Structure of Multimodal Dialogue*, Elsevier, pp75-84
- (Cole *et al.*, 1995) Cole, R.A., Mariani, J., Uszkoreit, H., Zaenen, A., Zue., V. (1995) "Survey of the State of the Art in Human Language Technology"
- (Conklin and Begeman, 1988) Conklin, J. & Begeman, M.L. (1988) "gIBIS: A Hypertext Tool for Exploratory Policy Discussion", *ACM Transactions on Office Information Systems* 6(4), pp303-331

- (Cooley, 1959) Cooley, J.C. (1959) "On Mr. Toulmin's Revolution in Logic", *The Journal of Philosophy* 56, pp297-319
- (Cooper et al., 1996) Cooper, R., Fox, J., Farrington, J. & Shallice, T. (1996) "A Systematic Methodology for Cognitive Modelling", *Artificial Intelligence* 85, pp3-44
- (Copi, 1978) Copi, I.M. (1978) "Introduction to Logic: 5th Edition", Macmillan, New York
- (Craig, 1990) Craig, R.T. (1990) "Multiple Goals in Discourse: An Epilogue", *Journal of Language and Social Psychology* 9 (1-2), pp163-170
- (Currie and Tate, 1991) Currie, K. & Tate, A. (1991) "O-Plan: the Open Planning Architecture", *Artificial Intelligence* 52, pp49-86
- (Dahlgreen, 1996) Dahlgreen, K. (1996) "Discourse Coherence and Segmentation" in Hovy, E., Scott, D., (eds), *Computational and Conversational Discourse: Burning Issues - An Interdisciplinary Account: Proceedings of NATO ARW, Maratea, Italy, 1993*, Springer Verlag, pp111-138
- (Dale and Reiter, 1999) Dale, R. & Reiter, E. (1999) "Building Natural Language Generation Systems", Cambridge University Press, Cambridge
- (Dale, 1990) Dale, R. (1990) "Generating Recipes: An Overview of Epicure" in Dale, R., Mellish, C. & Zock, M., (eds), *Current Research in Natural Language Generation*, Academic Press, London, pp229-255
- (Daradoumis, 1996) Daradoumis, T. (1996) "Towards a Representation of the Rhetorical Structure of Interrupted Exchanges" in Adomi, G., Zock, M., (eds), *Trends in NLG: An AI Perspective - Selected Papers from EWNLG'93 (LNAI Vol. 1036)*, Springer Verlag, pp106-124
- (Darnell, 1963) Darnell, D.K. (1963) "The relation between sentence order and comprehension", *Speech Monographs* 30, pp97-100
- (Das et al., 1996) Das, S.K., Fox, J., Krause, P. (1996) "A unified framework for hypothetical and practical reasoning (1): theoretical foundations" in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085)*, Springer Verlag, Berlin, pp58-72
- (Davies and Edwards, 1994) Davies, W.H.E. & Edwards, P. (1994) "Agent-K: An Integration of AOP and KQML" in Finin, T. & Labrou, Y. (eds), *Proceedings of the CIKM'94 Workshop on Intelligent Agents*, Gaithersburg, Maryland.
- (de Kleer, 1986) de Kleer, J. (1986) "An Assumption-Based TMS", *Artificial Intelligence* 28(2), pp127-162
- (de Rosis et al., 1997) de Rosis, F., Grasso, F., & Berry, D.C. (1997) "Strengthening Argumentation in Medical Explanations by Text Plan Revision" in Keravnou, E., Garbay, C., Baud, R. & Wyatt, J., (eds), *Proceedings of the 6th Conference on Artificial Intelligence in Medicine in Europe (AIME'97)*, Springer Verlag, Grenoble, pp349-360
- (Derr and McKeown, 1984) Derr, M.A. & McKeown, K.R. (1984) "Using Focus to Generate Complex and Simple Sentences" in *Proceeding of the 10th International Conference on Computational Linguistics (COLING'84)*, pp319-326
- (DiMarco et al., 1992) DiMarco, C., Green, S.J., Hirst, G., Mah, K. & Makuta-Gilak, M. (1992) "Four Papers on Computational Stylistics", University of Waterloo, Technical Report CS-92-35
- (DiMarco et al., 1995) DiMarco, C., Hirst, G., Wanner, L. & Wilkinson, J. (1995) "HealthDoc: Customizing patient information and health education by medical condition and personal characteristics" in *Working Notes of the Workshop on AI in Patient Education*, Glasgow
- (d'Inverno and Luck, 1996) d'Inverno, M. & Luck, M. (1996) "Understanding Autonomous Interaction" in *Proceedings of the Twelfth European Conference on Artificial Intelligence (ECAI-96)*, John Wiley, pp529-533
- (Doyle, 1979) Doyle, J. (1979) "A truth maintenance system", *Artificial Intelligence* 12(3)
- (Dung, 1995) Dung, P.M. (1995) "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games", *Artificial Intelligence* 77, pp321-357
- (Eemeren and Grootendorst, 1992) Eemeren, F.H. van & Grootendorst, R. (1992) "Argumentation, Communication and Fallacies: A Pragma-Dialectical Perspective", Lawrence Erlbaum Associates, Hillsdale, NJ
- (Eemeren et al., 1993) Eemeren, F.H. van, Grootendorst, R., Jackson, S. & Jacobs, S. (1993) "Reconstructing Argumentative Discourse", University of Alabama Press, Tuscaloosa
- (Eemeren et al., 1996) Eemeren, F.H. van, Grootendorst, R. & Snoeck-Henkemans, F. (1996) "Fundamentals of Argumentation Theory", Lawrence Erlbaum, Mahwah, NJ

- (Elhadad and Robin, 1999) Elhadad, M. & Robin, J. (1999, to appear) "SURGE: A Comprehensive Plug-in Syntactic Realisation Component for Text Generation", Computational Linguistics
- (Elhadad, 1992) Elhadad, M. (1992) "Using Argumentation to Control Lexical Choice: A Functional Unification Implementation", Columbia University PhD Thesis, New York
- (Elhadad, 1992a) Elhadad, M. (1992) "Generating Coherent Argument Paragraphs" in Proceedings of the Conference on Computational Linguistics (COLING'92), Nantes, pp638-644
- (Elhadad, 1995) Elhadad, M. (1995) "Using Argumentation in Text Generation", Journal of Pragmatics 24, pp189-220
- (Eller and Carberry, 1992) Eller, R. & Carberry, S. (1992) "A Meta-Rule Approach to Flexible Plan Recognition in Dialogue", User Modelling and User Adapted Interaction 2, pp27-53
- (Ellis, 1983) Ellis, D.G. (1983) "Language, Coherence, and Textuality" in Craig, R.T., Tracy, K., (eds), Conversational Coherence, Sage, Beverly Hills, CA, pp222-240
- (Elvang-Goransson *et al.*, 1993) Elvang-Goransson, M., Krause, P. & Fox, J. (1993) "Dialectic reasoning with inconsistent information" in Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI'93), pp114-121
- (Erol *et al.*, 1994) Erol, K., Hendler, J. & Nau, D.S. (1994) "HTN Planning: Complexity and Expressivity" in Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI'94), AAAI
- (Erol *et al.*, 1994a) Erol, K., Hendler, J. & Nau, D.S. (1994) "UMCP: A Sound and Complete Procedure for Hierarchical Task-Network Planning" in Proceedings of the Second International Conference on AI and Planning Systems (AIPS'94), Chicago, IL
- (Etzioni *et al.*, 1997) Etzioni, O., Golden, K. & Weld, D.S. (1997) "Sound and efficient closed-world reasoning for planning", Artificial Intelligence 89, pp113-148
- (Farley and Freeman, 1995) Farley, A.M. & Freeman, K. (1995) "Burden of proof in a computational model of argumentation" in Eemeren, F.H. van, Grootendorst, R., Blair, J.A. & Willard, C.A., (eds), Proceedings of the 3rd ISSA Conference on Argumentation (ISSA'94), Volume III, SicSat, Amsterdam, pp410-422
- (Fawcett and Davies, 1992) Fawcett, R.P. & Davies, B.L. (1992) "Monologue as a turn in a dialogue: towards an integration of exchange structure and rhetorical structure theory" in Dale, R., Hovy, E., Rosner, D., & Stock, O., (eds), Aspect of Automated Natural Language Generation, Springer Verlag, Berlin, pp151-166
- (Feteris, 1996) Feteris, E. (1996) "The Analysis and Evaluation of Legal Argumentation from a Pragma-Dialectical Perspective" in Gabbay, D., Ohlbach, H.J., (eds), Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085), Springer Verlag, Berlin, pp151-166
- (Feteris, 1997) Feteris, E.T. (1997) "A Survey of 25 Years of Research on Legal Argumentation", Argumentation 11(3), pp355-376
- (Fikes and Nilsson, 1971) Fikes, R.E. & Nilsson, N.J. (1971) "STRIPS: A new approach to the application of theorem proving to problem solving", Artificial Intelligence 2(3-4), pp189-208
- (Finin *et al.*, 1997) Finin, T., Labrou, Y. & Mayfield, J. (1997) "KQML as an agent communication language" in Bradshaw, J., (eds), Software Agents, MIT Press
- (Finin, 1989) Finin, T.W. (1989) "GUMS - A General User Modelling Shell" in Kobsa, A., Wahlster, W., (eds), User Models in Dialog Systems, Springer-Verlag, Berlin, pp411-430
- (Fink and Yang, 1997) Fink, E. & Yang, Q. (1997) "Automatically selecting and using primacy effects in planning: theory and experiments", Artificial Intelligence 89, pp285-315
- (Finocchiaro, 1980) Finocchiaro, M.A. (1980) "Galileo and the art of reasoning", Reidel, Dordrecht
- (Fisher, 1988) Fisher, A. (1988) "The Logic of Real Arguments", Cambridge University Press, Cambridge
- (Flowers *et al.*, 1982) Flowers, M., McGuire, R. & Birnbaum, L. (1982) "Adversary Arguments and the Logic of Personal Attacks" in Lehnert, W.G., Ringle, M.H., (eds), Strategies for Natural Language Processing, Lawrence Erlbaum Associates, pp275-294
- (Fogelin and Sinnott-Armstrong, 1991) Fogelin, R.J. & Sinnott-Armstrong, W. (1991) "Understanding Arguments: An Introduction to Informal Logic (4th Edition)", Harcourt Brace Jovanovich College Publishers, Fort Worth
- (Fox and Das, 1996) Fox, J. & Das, S. (1996) "A unified framework for hypothetical and practical reasoning (2): lessons from medical applications" in Gabbay, D., Ohlbach, H.J., (eds) Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085), Springer Verlag, Berlin, pp73-92

- (Fox and Long, 1993) Fox, M. & Long, D. (1993) "A Note on Chapman's Modal Truth Criterion" in Filgueiras, M., Damas, L., (eds), Progress in Artificial Intelligence. Proceedings of the 6th Portuguese Conference on AI (EPIA'93), Lecture Notes in AI Vol. 727, Springer Verlag, p307-310
- (Fox and Long, 1995) Fox, M. & Long, D. (1995) "Hierarchical planning using abstraction", IEE Proc.-Control Theory Appl., 142 (3), pp197-210
- (Fox and Long, 1996) Fox, M. & Long, D.P. (1996) "An Efficient Algorithm for Managing Partial Orders in Planning", SIGART Bulletin 7 (4), pp3-9
- (Fox, 1997) Fox, M. (1997) "Natural Hierarchical Planning using Operator Decomposition" in Steel, S., Alami, R., (eds), Proceedings of the 4th European Conference on Planning (ECP'97), Springer Verlag, pp195-208
- (Franklin and Graesser, 1996) Franklin, S. & Graesser, A. (1996) "Is it an Agent, or just a Program? A Taxonomy for Autonomous Agents" in Proceedings of the Workshop on Agent Theories, Architectures and Languages (ATAL'95), Springer Verlag, pp193-206
- (Freeman and Farley, 1994) Freeman, K. & Farley, A. (1994) "Toward Formalizing Dialectical Argumentation" in Proceedings of the 15th Conference of the Cognitive Science Society, pp440-445
- (Freeman, 1991) Freeman, J.B. (1991) "Dialectics and the Macrostructure of Arguments", Foris, Dordrecht
- (Gabbay, 1992) Gabbay, D. (1992) "LDS - Labelled Deductive Systems, 7th Expanded Draft", Imperial College Technical Report, London
- (Galliers and Sparck Jones, 1993) Galliers, J.R. & Sparck Jones, K. (1993) "Evaluating Natural Language Processing Systems", University of Cambridge, Technical Report 291
- (Galliers, 1992) Galliers, J.R. (1992) "Autonomous belief revision and communication" in Gardenfors, P., (eds), Belief Revision, Cambridge University Press, Cambridge, pp220-246
- (Gardenfors, 1992) Gardenfors, P. (1992) "Belief Revision: An Introduction" in Gardenfors, P., (eds), Belief Revision, Cambridge University Press, Cambridge, pp1-28
- (Gazen and Knoblock, 1997) Gazen, B.C. & Knoblock, C.A. (1997) "Combining the Expressivity of UCPOP with the Efficiency of Graphplan" in Proceedings of the European Conference on Planning (ECP'97)
- (Geffner, 1996) Geffner, H. (1996) "A Formal Framework for Causal Modelling and Argumentation" in Gabbay, D., Ohlbach, H.J., (eds), Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085), Springer Verlag, Berlin, pp208-222
- (Genesereth and Ketchpel, 1994) Genesereth, M.R. & Ketchpel, S.P. (1994) "Software Agents", Communications of the ACM 37(7), pp48-53
- (Gerevini and Schubert, 1996) Gerevini, A. & Schubert, L. (1996) "Accelerating Partial-Order Planners: Some Techniques for Effective Search Control and Pruning", Journal of Artificial Intelligence Research 5, pp95-137
- (Gilbert, 1995) Gilbert, M.A. (1995) "Argument and Arguers", Teaching Philosophy 18 (2), pp125-138
- (Gilbert, 1996) Gilbert, M.A. (1996) "Goals in Argumentation" in Gabbay, D., Ohlbach, H.J., (eds), Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085), Springer Verlag, Berlin, pp223-230
- (Girle, 1996) Girle, R.A. (1996) "Commands in Dialogue Logic" in Gabbay, D., Ohlbach, H.J., (eds), Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085), Springer Verlag, Berlin, pp246-260
- (Goldman, 1970) Goldman, A.I. (1970) "A Theory of Human Action", Prentice Hall, Englewood Cliffs, NJ
- (Goodman and Litman, 1992) Goodman, B.A. & Litman, D.J. (1992) "On the Interaction between Plan Recognition and Intelligent Interfaces", User Modelling and User Adapted Interaction 2, pp83-111
- (Gordon and Karacapilidis, 1996) Gordon, T. & Karacapilidis, N. (1996) "The Zeno Argumentation Framework" in Proceedings of the FAPR'96 Workshop on Computational Dialectics, Bonn
- (Gordon, 1994) Gordon, T.F. (1994) "Computational Dialectics" in Proceedings of the Workshop Kooperative Juristische Informationssysteme, GMD Studien Nr. 241, pp25-36
- (Grasso, 1997) Grasso, F. (1997) "Using Dialectical Argumentation for User Modelling in Decision Support Systems" in Jameson, A., Paris, C. & Tasso, C., (eds), User Modeling: Proceedings of the 6th International Conference (UM'97), Springer Wien, New York, pp83-86
- (Gray, 1977) Gray, B. (1977) "The Grammatical Foundations of Rhetoric", Mouton, The Hague

- (Green and DiMarco, 1996) Green, S.J. & DiMarco, C. (1996) "Stylistic Decision-Making in Natural Language Generation" in Adorni, G., Zock, M., (eds), Trends in NLG: An AI Perspective - Selected Papers from EWNLG93 (LNAI Vol. 1036), Springer Verlag, pp125-143
- (Green, 1996) Green, D.W. (1996) "Arguments and mental models: a position paper" in Gabbay, D., Ohlbach, H.J., (eds), Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR96 (Lecture Notes in AI Vol. 1085), Springer Verlag, Berlin
- (Grice, 1957) Grice, H.P. (1957) "Meaning", *Philosophical Review* 66, pp377-388
- (Grice, 1975) Grice, H.P. (1975) "Logic and conversation" in Cole, P., Morgan, J.L., (eds), *Studies in Syntax*, Vol. III, Seminar Press, New York
- (Grimes, 1975) Grimes, J.E. (1975) "The Thread of Discourse", Mouton, The Hague
- (Grosz and Sidner, 1986) Grosz, B.J. & Sidner, C.L. (1986) "Attention, Intentions and the Structure of Discourse", *Computational Linguistics* 12 (3), pp175-204
- (Grosz *et al.*, 1995) Grosz, B.J., Joshi, A.K. and Weinstein, S. (1995) "Centering: A Framework for Modeling the Local Coherence of Discourse", *Computational Linguistics* 21 (2), pp203-225
- (Grosz, 1977) Grosz, B.J. (1977) "The representation and use of focus in dialogue understanding", SRI International AI Center Technical Report 151
- (Grote *et al.*, 1997) Grote, B., Lenke, N. & Stede, M. (1997) "Ma(r)king concessions in English and German", *Discourse Processes* 24(1), pp87-118
- (Haggith, 1995) Haggith, M. (1995) "A meta-level framework for exploring conflicts in multiple knowledge bases" in Hallam, J., (eds), *Hybrid Problems, Hybrid Solutions*, IOS Press, pp87-98
- (Haggith, 1996) Haggith, M. (1996) "A meta-level argumentation framework for representing and reasoning about disagreement", PhD Thesis, University of Edinburgh
- (Haller and Shapiro, 1996) Haller, S.M. & Shapiro, S.C. (1996) "IDP - An Interactive Discourse Planner" in Adorni, G., Zock, M., (eds), Trends in NLG: An AI Perspective - Selected Papers from EWNLG93 (LNAI Vol. 1036), Springer Verlag, pp144-167
- (Haller, 1994) Haller, S.M. (1994) "Recognizing Digressive Questions During Interactive Generation" in Proceedings of the 7th International Workshop on Natural Language Generation, Kennebunkport, Maine, pp181-188
- (Hamblin, 1970) Hamblin, C.L. (1970) "Fallacies", Methuen, London
- (Hansen, 1996) Hansen, H. (1996) "Aristotle, Whately, and the Taxonomy of Fallacies" in Gabbay, D., Ohlbach, H.J., (eds), Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR96 (Lecture Notes in AI Vol. 1085), Springer Verlag, Berlin, pp318-330
- (Harsanyi, 1977) Harsanyi, J.C. (1977) "Rational Behaviour and Bargaining Equilibrium in Games and Social Situations", Cambridge University Press, Cambridge
- (Hartley and Paris, 1996) Hartley, A. & Paris, C. (1996) "Two Sources of Control over the Generation of Software Instructions" in Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL96), pp192-199
- (Hass and Linder, 1972) Hass, R.G. & Linder, D.E. (1972) "Counterargument availability and the effects of message structure on persuasion", *Journal of Personality and Social Psychology* 23 (2), pp219-233
- (Heidrich, 1982) Heidrich, C.H. (1982) "Montague-Grammars for Argumentative Dialogues" in Barth, E.M., Martens, J.L., (eds), *Argumentation: Approaches to Theory Formation*, John Benjamins, Amsterdam, pp191-227
- (Hewitt, 1977) Hewitt, C. (1977) "Viewing control structures as patterns of passing messages", *Artificial Intelligence* 8(3), pp323-364
- (Hintikka and Hintikka, 1982) Hintikka, J. & Hintikka, M.B. (1982) "Sherlock Holmes Confronts Modern Logic: Toward a Theory of Information-Seeking Through Questioning" in Barth, E.M., Martens, J.L., (eds), *Argumentation: Approaches to Theory Formation*, John Benjamins, Amsterdam, pp55-76
- (Hintikka *et al.*, 1996) Hintikka, J., Halonen, I. & Mutanen, A. (1996) "Interrogative Logic as a General Theory of Reasoning", Unpublished manuscript
- (Hintikka, 1962) Hintikka, J. (1962) "Knowledge and Belief", Cornell University Press, Ithaca, NY
- (Hirst *et al.*, 1997) Hirst, G., DiMarco, C., Hovy, E. & Parsons, K. (1997) "Authoring and Generating Health-Education Documents That Are Tailored to the Needs of the Individual Patient" in Jameson, A., Paris, C. & Tasso, C., (eds), *User Modeling: Proceedings of the 6th International Conference (UM'97)*, Springer Wien, New York, pp107-118

- (Hobbs, 1982) Hobbs, J.R. (1982) "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
- (Horacek, 1992) Horacek, H. (1992) "An Integrated View of Text Planning" in Dale, R., Hovy, E., Rosner, D., & Stock, O., (eds), *Aspect of Automated Natural Language Generation*, Springer-Verlag, Berlin, pp29-44
- (Horacek, 1994) Horacek, H. (1994) "How to avoid explaining obvious things" in Cohn, A., (eds), *Proceedings of the 11th European Conference on Artificial Intelligence*, John Wiley & Sons, pp520-524
- (Hovland and Mandell, 1957) Hovland, C.I. & Mandell, W. (1957) "Is There a "Law of Primacy in Persuasion"?" in Hovland, C.I., (eds), *The Order of Presentation in Persuasion*, Yale University Press, New Haven, CT, pp13-22
- (Hovland *et al.*, 1957) Hovland, C.I., Campbell, E.H. & Brock, T. (1957) "The Effects of "Committment" on Opinion Change Following Communication" in Hovland, C.I., (eds), *The Order of Presentation in Persuasion*, Yale University Press, New Haven, CT, pp23-32
- (Hovland *et al.*, 1965) Hovland, C.I., Lumsdaine, A. & Sheffield, F.D. (1965) "Experiments on Mass Communication", Wiley, New York
- (Hovland, 1957) Hovland, C.I. (1957) "Introduction" in Hovland, C.I., (eds), *The Order of Presentation in Persuasion*, Yale University Press, New Haven, CT, pp1-10
- (Hovland, 1957a) Hovland, C.I. (1957) "Summary and Implications" in Hovland, C.I., (eds), *The Order of Presentation in Persuasion*, Yale University Press, New Haven, CT, pp129-157
- (Hovland, 1957b) Hovland, C.I. (ed) (1957), *The Order of Presentation in Persuasion*, Yale University Press, New Haven, CT
- (Hovy and Arens, 1991) Hovy, E.H. & Arens, Y. (1991) "Automatic Generation of Formatted Text" in *Proceedings of the National Conference on Artificial Intelligence (AAAI'91)*, AAAI Press, Anaheim, pp92-96
- (Hovy and Maier, 1993) Hovy, E.H. & Maier, E. (1993) "Parsimonious and profligate: how many and which discourse structure relations?", *Discourse Processes*
- (Hovy and McCoy, 1989) Hovy, E. & McCoy, K.F. (1989) "Focusing your RST: A Step toward Generating Coherent Multisentential Text" in *Proceedings of the 11th Conference of the Cognitive Science Society*, pp667-674
- (Hovy and Wanner, 1996) Hovy, E.H. & Wanner, L. (1996) "Managing Sentence Planning Requirements" in *Proceedings of the ECAI'96 Workshop on Planning and Natural Language Generation*, Budapest, pp53-58
- (Hovy, 1986) Hovy, E.H. (1986) "Putting Affect into Text" in *Proceedings of the 8th Conference of the Cognitive Science Society*, pp669-675
- (Hovy, 1986a) Hovy, E. (1986) "Some Pragmatic Decision Criteria in Generation" in Kempen, G., (eds), *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*, Kluwer, pp3-17
- (Hovy, 1988) Hovy, E.H. (1988) "Planning Coherent Multisentential Text" in *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics (ACL'88)*, pp163-169
- (Hovy, 1990) Hovy, E.H. (1990) "Pragmatics and Natural Language Generation", *Artificial Intelligence* 43, pp153-197
- (Hovy, 1991) Hovy, E.H. (1991) "Approaches to the Planning of Coherent Text" in Paris, C.L., Swartout, W.R. & Mann, W.C., (eds), *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, Kluwer, pp83-102
- (Hovy, 1993) Hovy, E.H. (1993) "Automated discourse generation using discourse structure relations", *Artificial Intelligence* 63, pp341-385
- (Huang and Fiedler, 1996) Huang, X. & Fiedler, A. (1996) "Paraphrasing and Aggregating Argumentative Text Using Text Structure" in *Working Notes of the 8th International Workshop on Natural Language Generation*, Herstmonceux
- (Huang, 1994) Huang, X. (1994) "Planning Argumentative Texts" in *Proceedings of the International Conference on Computational Linguistics (COLING'94)*, Kyoto
- (Inui *et al.*, 1992) Inui, K., Tokunaga, T. & Tanaka, H. (1992) "Text Revision: A Model and Its Implementation" in Dale, R., Hovy, E., Rosner, D., & Stock, O., (eds), *Aspect of Automated Natural Language Generation*, Springer-Verlag, Berlin, pp215-230
- (Ito and Takizawa, 1995) Ito, A. & Takizawa, O. (1995) "Formulation of Irony Based on a Dialogue Model", *International Journal of Psycholinguistics* 11(1), pp75-86

- (Jackson, 1997) Jackson, S. (1997) "Disputation by design" in Proceedings of the OSSA Conference on Argument and Rhetoric, (to appear), St. Catharines, Canada
- (Jameson and Wahlster, 1982) Jameson, A. & Wahlster, W. (1982) "User Modelling in Anaphora Generation: Ellipsis and Definite Description" in Proceedings of the European Conference on Artificial Intelligence (ECAI'82), Orsay, France, pp222-227
- (Jameson, 1989) Jameson, A. (1989) "But What Will the Listener Think? Belief Ascription and Images Maintenance in Dialog" in Kobsa, A., Wahlster, W., (eds), User Models in Dialog Systems, Springer-Verlag, Berlin, pp255-311
- (Janis and Feierabend, 1957) Janis, I.L. & Feierabend, R.L. (1957) "Effects of Alternative Ways of Ordering Pro and Con Arguments in Persuasive Communications" in Hovland, C.I., (eds), The Order of Presentation in Persuasion, Yale University Press, New Haven, CT, pp115-128
- (Jennings *et al.*, 1996) Jennings, N.R., Faratin, P., Johnson, M.J., O'Brien, P., Wiegand, M.E. (1996) "Using Intelligent Agents to Manage Business Processes" in Proceedings of the Conference on Practical Applications of Intelligent Agents and Multi-Agent Technology (PAAM'96), London
- (Johnson and Blair, 1993) Johnson, R.H. & Blair, J.A. (1993) "Logical self-defense", McGraw-Hill, Toronto
- (Johnson, 1992) Johnson, R.M. (1992) "A Logic Book", Wadsworth, Belmont, CA
- (Jones, 1994) Jones, C.E. (1994) "Dialogue Structure Models: An Engineering Approach to Machine Analysis and Generation of Dialogue", PhD Thesis, University of Durham
- (Jones, 1994a) Jones, B.E.M. (1994) "Exploring the Role of Punctuation in Parsing Natural Language Text" in Proceedings of the International Conference on Computational Linguistics (COLING'94), Kyoto, pp421-425
- (Jonsson, 1996) Jonsson, A. (1996) "Natural Language Generation without Intentions" in Proceedings of the ECAI'96 Workshop on Planning and Natural Language Generation, Budapest, pp102-104
- (Kahneman, 1973) Kahneman, D. (1973) "Attention and Effort", Prentice Hall
- (Kamp, 1981) Kamp, H. (1981) "A theory of truth and semantic representation" in Groenendijk, J.A.G., Janssen, T.M.V. & Stokhof, M.B.J., (eds), Formal Methods in the Study of Language, Mathematical Centre Tracts No. 136, University of Amsterdam, Amsterdam, pp277-322
- (Kass and Finin, 1988) Kass, R. & Finin, T. (1988) "Modeling the User in Natural Language Systems", Computational Linguistics 14 (3), pp5-22
- (Kemper and Kemper, 1994) Kemper, R.E. & Kemper, D.R. (1994) "Negotiation Literature: A bibliographic essay, citations and sources", Scarecrow, Metuchen, NJ
- (Kienpointner, 1992) Kienpointner, M. (1992) "How to Classify Arguments" in Eemeren, F.H. van, Grootendorst, R., Blair, J.A. & Willard, C.A., (eds), Argumentation Illuminated, SICSAT, Amsterdam, pp178-188
- (Knoblock *et al.*, 1991) Knoblock, C.A., Tenenber, J.D. & Yang, Q. (1991) "Characterizing Abstraction Hierarchies for Planning" in Proceedings of the National Conference on Artificial Intelligence (AAAI'91), AAAI Press, Anaheim
- (Knoblock, 1994) Knoblock, C.A. (1994) "Automatically Generating Abstractions for Planning", Artificial Intelligence 68 (2)
- (Knott and Dale, 1994) Knott, A. & Dale, R. (1994) "Using Linguistic Phenomena to Motivate a Set of Coherence Relations", Discourse Processes 18(1), pp35-62
- (Knott and Dale, 1996) Knott, A. & Dale, R. (1996) "Choosing a Set of Coherence Relations for Text Generation: A Data-Driven Approach" in Adami, G., Zock, M., (eds), Trends in NLG: An AI Perspective - Selected Papers from EWNLG'93 (LNAI Vol. 1036), Springer Verlag, pp47-67
- (Knott *et al.*, 1997) Knott, A., O'Donnell, M., Oberlander, J. & Mellish, C. (1997) "Defeasible Rules in Content Selection and Text Structuring" in Proceedings of the 6th European Workshop on Natural Language Generation (EWNLG'97), Duisburg, pp50-60
- (Knott, 1993) Knott, A. (1993) "Using Cue Phrases to Determine a Set of Rhetorical Relations" in Rambow, O., (eds), Working Notes of the ACL SIGGEN Workshop
- (Knott, 1996) Knott, A. (1996) "A Data-Driven Methodology for Motivating a Set of Coherence Relations", PhD Thesis, University of Edinburgh
- (Konolige and Pollack, 1989) Konolige, K. & Pollack, M.E. (1989) "Ascribing plans to agents" in Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI'89), Detroit, MI, pp924-930

- (Kowalski and Toni, 1994) Kowalski, R.A. & Toni, F. (1994) "Argument and Reconciliation" in Proceedings of the Legal Reasoning Workshop at the International Symposium on 5th Generation Computer Systems, Tokyo, pp9-16
- (Krause *et al.*, 1995) Krause, P., Ambler, S., Elvang-Goransson, M., Fox, J. (1995) "A Logic of Argumentation for Reasoning under Uncertainty", *Computational Intelligence* 11 (1), pp113-131
- (Kuhn, 1991) Kuhn, D. (1991) "The Skills of Argument", Cambridge University Press, Cambridge
- (Laasri *et al.*, 1992) Laasri, B., Laasri, H., Lander, S. & Lesser, V. (1992) "A Generic Model for Intelligent Negotiating Agents", *International Journal of Intelligent and Cooperative Information Systems* 1 (2), pp291-317
- (Labrou and Finin, 1997) Labrou, Y. & Finin, T. (1997) "Semantics and Conversations for an Agent Communication Language" in Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI'97), Nagoya, Japan, pp584-591
- (Laver, 1976) Laver, J. (1976) "Communicative Functions of Phatic Communion" in Kendon, A., Harris, R.M. & Key, M.R., (eds), *Organization of behavior in face-to-face interaction*, Mouton, The Hague, pp215-238
- (Lee and Wilks, 1997) Lee, M. & Wilks, Y. (1997) "Eliminating deceptions and mistaken belief to infer conversational implicature" in Working Notes of the IJCAI'97 Workshop on Collaboration, Cooperation and Conflict in Dialogue Systems, Nagoya, Japan, pp59-65
- (Leopold-Wildburger, 1982) Leopold-Wildburger, U. (1982) "A Decision-Theoretical Interpretation of Dialogues" in Barth, E.M., Martens, J.L., (eds), *Argumentation: Approaches to Theory Formation*, John Benjamins, Amsterdam, pp273-280
- (Levelt, 1989) Levelt, W.J.M. (1989) "Speaking: from intention to articulation", MIT Press, Cambridge, MA
- (Levin and Moore, 1977) Levin, J.A. & Moore, J.A. (1977) "Dialogue Games: Metacommunication Structures for Natural Language Interaction", *Cognitive Science* 1, pp395-420
- (Locke, 1975) P.H. Nidditch (ed) (1975) Locke, J., "An Essay Concerning Human Understanding", Clarendon Press, Oxford
- (Long and Garigliano, 1994) Long, D.P. & Garigliano, R. (1994) "Reasoning by Analogy and Causality: A model and application", Ellis Horwood, New York
- (Lorenz, 1982) Lorenz, K. (1982) "On the Criteria for the Choice of the Rules of Dialogue Logic" in Barth, E.M., Martens, J.L., (eds), *Argumentation: Approaches to Theory Formation*, John Benjamins, Amsterdam, pp145-157
- (Loui, 1994) Loui, R. (1994) "Argumentation and Arbitration Games", Presented at the AAAI-94 workshop on Computational Dialectics
- (Lowder, 1997) Lowder, J.J. (1997) "Logic and Fallacies", <http://www.infidels.org/news/atheism/logic.html>
- (Lund, 1925) Lund, F.H. (1925) "The psychology of belief. IV. The law of primacy in persuasion.", *Journal of Abnormal Social Psychology* 20, pp183-191
- (Mackenzie, 1979) Mackenzie, J.D. (1979) "Question-Begging in Non-Cumulative Systems", *Journal of Philosophical Logic* 8, pp117-133
- (Mann and Thompson, 1986) Mann, W.C. & Thompson, S.A. (1986) "Rhetorical structure theory: description and construction of text structures" in Kempen, G., (eds), *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics*, Kluwer, pp279-300
- (Mann and Thompson, 1988) Mann, W.C. & Thompson, S.A. (1988) "Rhetorical Structure Theory: A Theory of Text Organisation", *Text* 8 (3), pp243-281
- (Mann, 1983) Mann, W.C. (1983) "An Overview of the PENMAN Text Generation System" in Proceedings of the 3rd National Conference on Artificial Intelligence (AAAI-83), AAAI, Washington, DC, pp261-265
- (Marcu, 1996) Marcu, D. (1996) "Building Up Rhetorical Structure Trees" in Proceedings of the 13th National Conference on Artificial Intelligence (AAAI'96), AAAI, Portland, Oregon
- (Marcu, 1996a) Marcu, D. (1996) "The Conceptual and Linguistic Facets of Persuasive Arguments" in Proceedings of the ECAI'96 Workshop on Planning and Natural Language Generation, Budapest, pp43-46
- (Marcu, 1997) Marcu, D. (1997) "From local to global coherence: A bottom-up approach to text planning" in Proceedings of the 14th National Conference on Artificial Intelligence (AAAI'97), AAAI, Providence, RI
- (Marcu, 1997a) Marcu, D. (1997) "The Rhetorical Parsing of Natural Language Texts" in Proceedings of the 35th Annual Meeting of the Association for Computational Linguistics (ACL'97), Madrid

- (Marcu, 1997b) Marcu, D. (1997) "From discourse structures to text summaries" in Working Notes of the ACL'97 Workshop on Intelligent Scalable Text Summarization
- (Matthiessen and Bateman, 1991) Matthiessen, C.M.I.M & Bateman, J.A. (1991) "Text Generation and Systemic-Functional Linguistics", Pinter, London
- (Matwin *et al.*, 1989) Matwin, S., Szpakowicz, S., Koperczak, Z., Kersten, G.E & Michalowski, W. (1989) "Negoplan: An Expert System Shell for Negotiation Support", IEEE Expert 4 (4), pp50-62
- (Maybury, 1992) Maybury, M.T. (1992) "Communicative Acts for Explanation Generation", International journal of Man-Machine Studies 37, pp135-172
- (Maybury, 1993) Maybury, M.T. (1993) "Communicative Acts for Generative Natural Language Arguments" in Proceedings of the National Conference on Artificial Intelligence (AAAI'93), AAAI, pp357-364
- (Mayfield, 1992) Mayfield, J. (1992) "Controlling Inference in Plan Recognition", User Modelling and User Adapted Interaction 2, pp55-82
- (McAllester and Rosenblitt, 1991) McAllester, D. & Rosenblitt, D. (1991) "Systematic Nonlinear Planning", MIT AI Memo No. 1339
- (McConachy and Zukerman, 1996) McConachy, R. & Zukerman, I. (1996) "Using Argument Graphs to Generate Arguments" in Proceedings of the Twelfth European Conference on Artificial Intelligence (ECAI-96), John Wiley, pp592-596
- (McCoy and Cheng, 1991) McCoy, K.F. & Cheng, J. (1991) "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
- (McCoy, 1986) McCoy, K.F. (1986) "Contextual Effects on Responses to Misconceptions" in Kempen, G., (eds), Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics, Kluwer, pp43-54
- (McDermott, 1982) McDermott, D. (1982) "A Temporal Logic for Reasoning about Processes and Plans", Cognitive Science 6, pp101-155
- (McDermott, 1991) McDermott, D. (1991) "Regression planning", International Journal of Intelligent Systems 6, pp357-416
- (McGuire, 1969) McGuire, W.J. (1969) "The Nature Of Attitudes and Attitude Change" in Lindzey, G., Aronson, E., (eds), The Handbook of Social Psychology (Second Edition) Vol. III, Addison-Wesley, Reading, MA, pp136-314
- (McKeown, 1985) McKeown, K.R. (1985) "Text Generation: Using discourse strategies and focus constraints to generate natural language text", Cambridge University Press, Cambridge
- (Mellish *et al.*, 1996) Mellish, C., Reiter, E. & Levine, J. (1996) "NLG Applications to Technical Documentation: A View Through IDAS" in Adami, G., Zock, M., (eds), Trends in NLG: An AI Perspective - Selected Papers from EWNLG'93 (LNAI Vol. 1036), Springer Verlag, pp368-382
- (Mellish *et al.*, 1998) Mellish, C., O'Donnell, M., Oberlander, J. & Knott, A. (1998) "An Architecture for Opportunistic Text Generation" in Proceedings of the International Workshop on Natural Language Generation (INLG'98), Niagra
- (Meter, 1991) Meter, M.W. (1991) "Bridging the Generation Gap between Text Planning and Linguistic Realization", Computational Intelligence 7 (4), pp296-304
- (Meter, 1994) Meter, M. (1994) "Generating Event Descriptions with SAGE: a Simulation and Generation Environment" in Proceedings of the 7th International Workshop on Natural Language Generation, Kennebunkport, Maine, pp99-107
- (Mooney *et al.*, 1990) Mooney, D.J., Carberry, S. & McCoy, K.F. (1990) "The Generation of High-Level Structure for Extended Explanations" in Proceedings of the International Conference on Computational Linguistics (COLING'90), Helsinki, Vol 2: pp276-281
- (Moore and Paris, 1989) Moore, J.D. & Paris, C.L. (1989) "Planning Text for Advisory Dialogues" in Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL'89), pp203-211
- (Moore and Paris, 1994) Moore, J.D. & Paris, C.L. (1994) "Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information", Computational Linguistics 19 (4), pp651-694
- (Moore and Pollack, 1992) Moore, J.D. & Pollack, M.E. (1992) "A Problem for RST: The Need for Multi-Level Discourse Analysis", Computational Linguistics 18 (4), pp537-544

- (Moore and Swartout, 1991) Moore, J.D. & Swartout, W.R. (1991) "A Reactive Approach to Explanation: Taking the User's Feedback into Account" in Paris, C.L., Swartout, W.R. & Mann, W.C., (eds), *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, Kluwer, pp3-48
- (Moore, 1994) Moore, J.D. (1994) "What Makes Human Explanations Effective?" in *Proceedings of the 15th Conference of the Cognitive Science Society*, pp131-136
- (Morik, 1989) Morik, K. (1989) "User models and Conversational Settings: Modeling the User's Wants" in Kobsa, A., Wahlster, W., (eds), *User Models in Dialog Systems*, Springer-Verlag, Berlin, pp364-385
- (Moser and Moore, 1996) Moser, M. & Moore, J.D. (1996) "Towards a Synthesis of Two Accounts of Discourse Structure", *Computational Linguistics* 22 (3), pp409-419
- (Nash, 1950) Nash, J.F. (1950) "The bargaining problem", *Econometrica* 18, pp155-162
- (Nash, 1953) Nash, J.F. (1953) "Two-person cooperative games", *Econometrica* 21, pp129-140
- (Newell and Simon, 1963) Newell, A. & Simon, H.A. (1963) "GPS, a program that simulates human thought" in Feigenbaum, E.A., Feldman, J., (eds), *Computers and Thought*, Mc-Graw Hill, New York
- (Nicholas, 1995) Nicholas, N. (1995) "Parameters for Rhetorical Structure Theory Ontology" in *Working Papers in Linguistics* 15, University of Melbourne
- (Nunberg, 1990) Nunberg, G. (1990) "The linguistics of punctuation", *CSLI Lecture Notes No.18*, Stanford
- (Nwana, 1996) Nwana, H.S. (1996) "Software Agents: An Overview", *The Knowledge Engineering Review* 11 (3), pp205-244
- (O'Keefe, 1995) O'Keefe, B.J. (1995) "Identity and Influence in Social Interaction", *Argumentation* 9, pp785-800
- (Ohlsson and Robin, 1994) Ohlsson, S. & Robin, N. (1994) "The Power of Negative Thinking: The Central Role of Modus Tollens in Human Cognition" in *Proceedings of the 16th Conference of the Cognitive Science Society*, Atlanta, GA, pp681-686
- (O'Keefe, 1977) O'Keefe, D.J. (1977) "Two Concepts of Argument", *The Journal of the American Forensic Association* 13, pp121-128
- (Panaget, 1994) Panaget, F. (1994) "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
- (Panaget, 1997) Panaget, F. (1997) "Micro-planning: A Unified Representation of Lexical and Grammatical Resources" in *Proceedings of the 6th European Workshop on Natural Language Generation (EWNLG'97)*, Duisburg, pp97-106
- (Paris and VanderLinden, 1996) Paris, C. & VanderLinden, K. (1996) "Building Knowledge for the Generation of Software Documentation" in *Proceedings of the International Conference on Computational Linguistics (COLING'96)*, pp734-739
- (Paris, 1991) Paris, C.L. (1991) "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
- (Paris, 1993) Paris, C.L. (1993) "User Modelling in Text Generation", Pinter, London
- (Parsons and Jennings, 1996) Parsons, S. & Jennings, N.R. (1996) "Negotiation Through Argumentation - A Preliminary Report" in *Proceedings of the International Conference on Multi Agent Systems (ICMAS'96)*, Kyoto, Japan
- (Parsons, 1996) Parsons, S. (1996) "Defining Normative Systems for Qualitative Argumentation" in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085)*, Springer Verlag, pp449-463
- (Parsons, 1997) Parsons, S. (1997) "Normative Argumentation and Qualitative Probability" in Gabbay, D., Kruse, R., Nonnengart, A. & Ohlbach, H.J., (eds), *Proceedings of the 2nd International Conference on Formal and Applied Practical Reasoning (FAPR'97)*, Springer Verlag, pp466-480
- (Pascual and Virbel, 1996) Pascual, E. & Virbel, J. (1996) "Semantic and Layout Properties of Text Punctuation" in Jones, B., (eds), *Proceedings of the ACL/SIGPARSE Workshop on Punctuation in Computational Linguistics*, Santa Cruz, CA, pp41-48
- (Pascual, 1996) Pascual, E. (1996) "Integrating Text Formatting and Text Generation" in Adomi, G., Zock, M., (eds), *Trends in NLG: An AI Perspective - Selected Papers from EWNLG'93 (LNAI Vol. 1036)*, Springer Verlag, pp205-221
- (Pednault, 1989) Pednault, E.P.D. & ADL: Exploring the middle ground between STRIPS and the situation calculus (1989) "" in *Proceedings of the Knowledge Representation Conference (KR'89)*

- (Peh and Ting, 1996) Peh, L.S. & Ting, C.H.A. (1996) "A Divide-and-Conquer Strategy for Parsing" in Jones, B., (eds), *Proceedings of the ACL/SIGPARSE Workshop on Punctuation in Computational Linguistics*, Santa Cruz, CA, pp57-66
- (Penman, 1990) Penman, R. (1990) "Facework and Politeness: Multiple Goals in Courtroom Discourse", *Journal of Language and Social Psychology* 9 (1-2), pp15-38
- (Penberthy and Weld, 1992) Penberthy, J.S. & Weld, D. (1992) "UCPOP: A Sound, Complete, Partial-order Planner for ADL" in *Proceedings of the International Workshop on Knowledge Representation (KR'92)*, pp103-114
- (Perelman and Ohlbrechts-Tyteca, 1969) Perelman, Ch. & Ohlbrechts-Tyteca (1969) "The New Rhetoric", University of Notre Dame Press
- (Pilkington *et al.*, 1992) Pilkington, R.M., Hartley, J.R., Hintze, D. & Moore, D.J. (1992) "Learning to Argue and Arguing to Learn: An Interface for Computer-based Dialogue Games", *Journal of Artificial Intelligence in Education* 3 (3), pp275-295
- (Pollock, 1987) Pollock, J.L. (1987) "Defeasible Reasoning", *Cognitive Science* 11, pp481-518
- (Prakken and Sartor, 1996) Prakken, H. & Sartor, G. (1996) "A System for Defeasible Argumentation, with Defeasible Priorities" in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085)*, Springer Verlag, pp510-524
- (Prakken, 1996) Prakken, H. (1996) "Dialectical proof theory for defeasible argumentation with defeasible priorities" in *Proceedings of the FAPR'96 Workshop on Computational Dialectics*, Bonn
- (Quine, 1960) Quine, W.V.O. (1960) "Word and Object", MIT Press, Cambridge, MA
- (Quintilian, 1960) Page, T.E. (ed) (1960) Quintilian, "Institutio Oratoria", Heinemann, London
- (Rambow and Korelsky, 1992) Rambow, O. & Korelsky, T. (1992) "Applied Text Generation" in *Proceedings of the 3rd Applied Natural Language Processing Conference (ANLP'92)*, Trento
- (Rao and Georgeff, 1992) Rao, A.S. & Georgeff, M.P. (1992) "An Abstract Architecture for Rational Agents" in *Proceedings of the International Workshop on Knowledge Representation (KR'92)*
- (Reed and Long, 1997) Reed, C.A. & Long, D.P. (1997) "Persuasive Monologue" in *Proceedings of the OSSA Conference on Argument and Rhetoric*, St. Catharines, Canada
- (Reed and Long, 1997a) Reed, C.A. & Long, D.P. (1997) "Multiple Subarguments in Logic, Argumentation, Rhetoric and Text Generation" in *Proceedings of the 2nd International Conference on Formal and Applied Practical Reasoning (FAPR'97)*, Springer Verlag, Berlin, pp496-510
- (Reed and Long, 1997b) Reed, C.A. & Long, D.P. (1997) "Content Ordering in the Generation of Persuasive Discourse" in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI'97)*, Nagoya, Japan, pp1022-1027
- (Reed and Long, 1997c) Reed, C.A. & Long, D.P. (1997) "Collaboration, Cooperation and Dialogue Classification" in *Working Notes of the IJCAI'97 Workshop on Collaboration, Cooperation and Conflict in Dialogue Systems*, Nagoya, Japan, pp73-78
- (Reed and Long, 1997d) Reed, C.A. & Long, D.P. (1997) "Ordering and Focusing in an Architecture for Persuasive Discourse Planning" in *Proceedings of the 6th European Workshop on Natural Language Generation (EWNLG'97)*, Duisburg
- (Reed *et al.*, 1996) Reed, C.A., Long, D.P. & Fox, M. (1996) "An Architecture for Argumentative Discourse Planning" in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96 (Lecture Notes in AI Vol. 1085)*, Springer Verlag, pp555-565
- (Reed *et al.*, 1997) Reed, C.A., Long, D.P. & Fox, M. (1997) "Context and focusing in argumentative discourse planning" in *Proceedings of the International Conference on Formalising and Modelling Context (CONTEXT'97)*, Rio de Janeiro, pp88-99
- (Reed *et al.*, 1997a) Reed, C.A., Long, D.P., Fox, M. & Garagnani, M. (1997) "Persuasion as a Form of Inter-Agent Negotiation" in Lukose, D., Zhang, C., (eds), *Proceedings of the 2nd Australian Workshop on DAI (Lecture Notes in AI Vol. 1286)*, Springer Verlag, pp120-136
- (Reed, 1997) Reed, C.A. (1997) "Representing and Applying Knowledge for Argumentation in a Social Context", *AI and Society* 11(3-4), pp138-154
- (Reed, 1998) Reed, C.A. (1998) "Dialogue Frames in Agent Communication" in Demazeau, Y., (eds), *Proceedings of the Third International Conference on Multi Agent Systems (ICMAS'98)*, IEEE Press, pp246-253

- (Reichman, 1985) Reichman, R. (1985) "Getting Computers to Talk Like You and Me", MIT Press, Cambridge, MA
- (Reichman-Adar, 1984) Reichman-Adar, R. (1984) "Extended Person-Machine Interface", *Artificial Intelligence* 22, pp157-218
- (Reiter *et al.*, 1992) Reiter, E., Mellish, C. & Levine, J. (1992) "Automatic Generation of On-Line Documentation in the IDAS Project" in *Proceedings of the 3rd Applied Natural Language Processing Conference (ANLP92)*, Trento
- (Reiter *et al.*, 1995) Reiter, E., Mellish, C. & Levine, J. (1995) "Automatic Generation of Technical Generation", *Applied Artificial Intelligence* 9 (3), pp259-287
- (Reiter *et al.*, 1997) Reiter, E., Cawsey, A., Osman, L. & Roff, Y. (1997) "Knowledge Acquisition for Content Selection" in *Proceedings of the 6th European Workshop on Natural Language Generation (EWNLG97)*, Duisburg, pp117-126
- (Reiter, 1978) Reiter, R. (1978) "On closed world data bases" in Gallaire, H., Minker, J., (eds), *Logic and Data Bases*, Plenum Press, New York, pp55-76
- (Reithinger, 1991) Reithinger, N. (1991) "POPEL - A Parallel and Incremental Natural Language Generation System" in Paris, C.L., Swartout, W.R. & Mann, W.C., (eds), *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, Kluwer, pp179-199
- (Rescher, 1977) Rescher, N. (1977) "Dialectics: A controversy-oriented approach to the theory of knowledge", Statue University of New York Press, Albany
- (Rescher, 1997) Rescher, N. (1997) "Rhetoric and Rational Argumentation" in *Proceedings of the OSSA Conference on Argument and Rhetoric*, (to appear), St. Catharines, Canada
- (Rich, 1989) Rich, E. (1989) "Stereotypes and User Modeling" in Kobsa, A., Wahlster, W., (eds), *User Models in Dialog Systems*, Springer-Verlag, Berlin, pp35-51
- (Rips, 1994) Rips, L.J. (1994) "The Psychology of Proof", MIT Press, Cambridge, MA
- (Rittel and Webber, 1973) Rittel, H.W. & Webber, M.M. (1973) "Dilemmas in a General Theory of Planning", *Policy Sciences* 4, pp155-169
- (Robin, 1994) Robin, J. (1994) "Revision-Based Generation of Natural Language Summaries Providing Historical Background", Columbia University PhD Thesis, CUCS-034-94, New York
- (Rosner and Stede, 1992) Rosner, D. & Stede, M. (1992) "Customizing RST for the Automatic Production of Technical Manuals" in Dale, R., Hovy, E., Rosner, D., & Stock, O., (eds), *Aspect of Automated Natural Language Generation*, Springer-Verlag, Berlin, pp199-214
- (Rubinoff, 1992) Rubinoff, R. (1992) "Integrating Text Planning and Linguistic Choice by Annotating Linguistic Structures" in Dale, R., Hovy, E., Rosner, D., & Stock, O., (eds), *Aspect of Automated Natural Language Generation*, Springer-Verlag, Berlin, pp45-56
- (Sacerdoti, 1974) Sacerdoti, E. (1974) "Planning in a Hierarchy of Abstraction Spaces", *Artificial Intelligence* 5, pp115-135
- (Sacerdoti, 1975) Sacerdoti, E.D. (1975) "The Nonlinear Nature of Plans" in *Proceedings of the fourth International Joint Conference on Artificial Intelligene (IJCAI75)*, pp206-214
- (Sadock, 1977) Sadock, J. (1977) "Modus Brevis: The Truncated Argument" in *Papers from the 13th Regional Meeting of the Chicago Linguistics Society*, pp545-554
- (Sandell, 1977) Sandell, R. (1977) "Linguistic Style and Persuasion", Academic Press, London
- (Sanders *et al.*, 1992) Sanders, T.J.M., Spooren, W.P.M. & Noordman, L.G.M. (1992) "Toward a Taxonomy of Coherence Relations", *Discourse Processes* 15, pp1-35
- (Schegloff and Sacks, 1973) Schegloff, E.A. & Sacks, H. (1973) "Opening up closings", *Semiotica* 7, pp289-327
- (Scott and deSouza, 1990) Scott, D.R. & deSouza, C.S. (1990) "Getting the Message Across in RST-based Text Generation" in Dale, R., Mellish, C. & Zock, M., (eds), *Current Research in Natural Language Generation*, Academic Press, London, pp47-73
- (Scott and Paris, 1995) Scott, D. & Paris, C. (1995) "Identifying the mapping of semantics onto language:going beyond the text" in *Working Notes of the AAAI95 Fall Symposium on Empirical Methods in Discourse Interpretation and Generation*
- (Searle, 1969) Searle, J.R. (1969) "Speech Acts", Cambridge University Press, Cambridge
- (Shadbolt, 1992) Shadbolt, N.R. (1992) "Cooperative planning: A foundation for communicative negotiation", University of Nottingham Computer Science Dept Technical Report

- (Shiu *et al.*, 1996) Shiu, S., Luo, Z. & Garigliano, R. (1996) "Type Theoretic Semantics for SemNet" in Gabbay, D., Ohlbach, H.J., (eds), *Practical Reasoning: Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96* (Lecture Notes in AI Vol. 1085), Springer Verlag, pp582-595
- (Shoham, 1993) Shoham, Y. (1993) "Agent-oriented programming", *Artificial Intelligence* 60, pp51-92
- (Sider and Burger, 1992) Sider, J.S. & Burger, J.D. (1992) "Intention Structure and Extended Responses in a Portable Natural Language Interface", *User Modelling and User Adapted Interaction* 2, pp155-179
- (Sidner, 1979) Sidner, C.L. (1979) "Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse", AI Lab, MIT Technical Report 53
- (Sierra *et al.*, 1997) Sierra, C., Jennings, N.R., Noriega, P., Parsons, S. (1997) "A framework for argumentation-based negotiation" in *Proceedings of the Workshop on Agent Theories, Architectures and Languages (ATAL'97)*
- (Sierra *et al.*, 1997a) Sierra, C., Faratin, P., Jennings, N. (1997) "A Service-Oriented Negotiation Model Between Autonomous Agents" in *Proceedings of the Workshop on Modelling Autonomous Agents in a Multi Agent World (MAAMAW'97)*
- (Sillince and Minors, 1992) Sillince, J.A.A. & Minors, R.H. (1992) "Argumentation, Self-Inconsistency, and Multidimensional Argument Strength", *Communication and Cognition* 25 (4), pp325-338
- (Simard, 1996) Simard, M. (1996) "Considerations on Parsing a Poorly Punctuated Text in French" in Jones, B., (eds), *Proceedings of the ACL/SIGPARSE Workshop on Punctuation in Computational Linguistics*, Santa Cruz, CA, pp67-72
- (Simons, 1976) Simons, H.A. (1976) "Persuasion: Understanding, Practice and Analysis", Addison-Wesley, Reading, MA
- (Slot, 1993) Slot, P. (1993) "How can you say that? Rhetorical questions in argumentative texts", IFOTT, Amsterdam
- (Smeltzer, 1996) Smeltzer, M.A. (1996) "Lying and Intersubjective Truth: A Communication Based Approach to Understanding Lying", *Argumentation* 10 (3), pp361-373
- (Smith and Peot, 1992) Smith, D.E. & Peot, M.A. (1992) "A Critical Look at Knoblock's Hierarchy Mechanism" in *Proceedings of the First International Conference on AI and Planning Systems (AIPS'92)*, pp307-308
- (Smith *et al.*, 1994) Smith, M.H., Garigliano, R. & Morgan, R.C. (1994) "Generation in the LOLITA system: An engineering approach" in *Proceedings of the 7th International Workshop on Natural Language Generation*, Kennebunkport, Maine
- (Smith *et al.*, 1994a) Smith, M.H., Garigliano, R., Morgan, R., Shiu, S. & Jarvis, S. (1994) "LOLITA: A Natural Language Engineered System", unpublished manuscript
- (Smith, 1980) Smith, R.G. (1980) "The Contract Net Protocol", *IEEE Transactions on Computers* C-29 (12), pp1104-1113
- (Snoeck Henkemans, 1995) Snoeck Henkemans, F. (1995) "'Anyway' and 'even' as indicators of argumentative structure" in Eemeren, F.H. van, Grootendorst, R., Blair, J.A. & Willard, C.A., (eds), *Proceedings of the 3rd ISSA Conference on Argumentation (ISSA'94)*, Volume III, SicSat, Amsterdam, pp183-191
- (Snoeck Henkemans, 1997) Snoeck Henkemans, A.F. (1997) "Verbal indicators of argumentation and explanation" in *Proceedings of the OSSA Conference on Argument and Rhetoric*, St. Catharines, Canada
- (Sparck Jones, 1991) Sparck Jones, K. (1991) "Tailoring output to the user: What does user modelling in generation mean?" in Paris, C.L., Swartout, W.R. & Mann, W.C., (eds), *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, Kluwer, pp201-225
- (Sperber and Wilson, 1986) Sperber, D. & Wilson, D. (1986) "Relevance: Communication and Cognition", Basil Blackwell, Oxford
- (Stede, 1993) Stede, M. (1993) "Lexical Choice Criteria in Language Generation" in *Proceedings of the 6th Meeting of the European Chapter of the Association for Computational Linguistics (EACL'93)*, pp454-459
- (Stefik, 1981) Stefik, M. (1981) "Planning with constraints (MOLGEN:Part 1)", *Artificial Intelligence* 16(2), pp111-139
- (Stefik, 1981a) Stefik, M. (1981) "Planning and meta-planning (MOLGEN: Part 2)", *Artificial Intelligence* 16(2), pp141-169
- (Strawson, 1964) Strawson, P. (1964) "Intention and convention in speech acts", *Philosophical Review* 73, pp439-460
- (Sussman, 1975) Sussman, G.J. (1975) "A Computer Model of Skill Acquisition", MIT Press, Cambridge, MA

- (Suthers, 1994) Suthers, D.D. (1994) "Strategies for Sequencing as a Planning Task" in Proceedings of the 7th International Workshop on Natural Language Generation, Kennebunkport, Maine, pp29-36
- (Sycara, 1989) Sycara, K.P. (1989) "Argumentation: Planning Other Agent's Plans" in Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI'89), Detroit, MI, pp517-523
- (Sycara, 1990) Sycara, K. (1990) "Persuasive Argumentation in Negotiation", Theory and Decision 28, pp203-242
- (Tate, 1975) Tate, A. (1975) "Interacting Goals and Their Use" in Proceedings of the fourth International Joint Conference on Artificial Intelligence (IJCAI'75), pp215-218
- (Tate, 1977) Tate, A. (1977) "Generating Project Networks" in Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'77)
- (Toulmin, 1958) Toulmin, S. E. (1958) "The Uses of Argument", Cambridge University Press, Cambridge
- (Tracy, 1990) Tracy, K. (1990) "Multiple Goals in Discourse: An Overview of Issues", Journal of Language and Social Psychology 9 (1-2), pp1-13
- (Triandis, 1971) Triandis, H.C. (1971) "Attitude and Attitude Change", John Wiley, New York
- (Turing, 1950) Turing, A.M. (1950) "Computing machinery and intelligence", Mind 59, pp433-460
- (Turner, 1994) Turner, E.H. (1994) "Exploiting Problem Solving to Select Information to Include in Dialogues between Cooperating Agents" in Proceedings of the 16th Conference of the Cognitive Science Society, Atlanta, GA, pp882-886
- (Verheij, 1996) Verheij, B. (1996) "Two approaches to dialectical argumentation: Admissible sets and argumentation stages" in Proceedings of the FAPR'96 Workshop on Computational Dialectics, Bonn
- (Vorobej, 1997) Vorobej, M. (1997) "What Exactly is a Persuasive Monologue?" in Proceedings of the OSSA Conference on Argument and Rhetoric, (to appear), St. Catharines, Canada
- (Vreeswijk, 1992) Vreeswijk, G. (1992) "Reasoning with Defeasible Arguments" in Wagner, G., Pearce, D., (eds), Proceedings of the European Workshop on Logics in AI, Springer Verlag, Berlin, pp189-211
- (Vreeswijk, 1995) Vreeswijk, G.A.W. (1995) "Open Protocol in Multi-Agent Systems", Technical report CS 95-01, Dept. of CS, University of Limburg
- (Wahlster and Kobsa, 1989) Wahlster, W. & Kobsa, A. (1989) "User Models in Dialog Systems" in Kobsa, A., Wahlster, W., (eds), User Models in Dialog Systems, Springer-Verlag, Berlin, pp4-33
- (Wahlster, 1993) Wahlster, W. (1993) "VERBMOBIL: A translation of face-to-face dialogs" in Proceedings of the MT Summit IV, Kobe
- (Waldron *et al.*, 1990) Waldron, V.R., Cegala, D.J., Sharkey, W.F. & Teboul, B. (1990) "Cognitive and Tactical Dimensions of Conversation Goal Management", Journal of Language and Social Psychology 9 (1-2), pp101-118
- (Walker, 1992) Walker, M.A. (1992) "Redundancy in Collaborative Dialogue" in Proceedings of the Conference on Computational Linguistics (COLING'92), Nantes, pp345-351
- (Walker, 1996) Walker, M.A. (1996) "The effect of resource limits and task complexity on collaborative planning in dialogue", Artificial Intelligence 85, pp181-243
- (Walton and Krabbe, 1995) Walton, D.N. & Krabbe, E.C.W (1995) "Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning", State University of New York Press, New York
- (Walton, 1992) Walton, D.N. (1992) "The Place of Emotion in Argument", Penn State Press
- (Walton, 1995) Walton, D. (1995) "Appeal to Pity: A Case Study of the Argumentum Ad Misericordiam", Argumentation 9, pp769-784
- (Walton, 1996) Walton, D.N. (1996) "Arguments from Ignorance", Penn State Press
- (Wellman, 1990) Wellman, M.P. (1990) "Formulation of tradeoffs in planning under uncertainty", Pitman, London
- (Whately, 1855) Whately, R. (1855) "Rhetoric", Richard Griffin & Co., London
- (Wilkins, 1988) Wilkins, D. (1988) "Practical Planning: Extending the Classical AI Paradigm", Addison-Wesley
- (Willard, 1976) Willard, C. A. (1976) "On the utility of descriptive diagrams for the analysis and criticism of arguments", Communication Monographs 43, pp308-319
- (Willard, 1989) Willard, C.A. (1989) "A theory of argumentation", University of Alabama Press, Tuscaloosa
- (Williams, 1893) Williams, W. (1893) "Composition and Rhetoric", Heath & Co., Boston
- (Wilson, 1980) Wilson, B.A. (1980) "The Anatomy of Argument", University Press of America, Washington, D.C.

- (Woods and Walton, 1989) Woods, J. & Walton, D.N. (1989) "Fallacies: Selected Papers 1972-1982", Foris, Dordrecht
- (Wooldridge and Jennings, 1995) Wooldridge, M. & Jennings, N.R. (1995) "Intelligent Agents: Theory and Practice", *The Knowledge Engineering Review* 10 (2), pp115-152
- (Yang *et al.*, 1996) Yang, Q., Tenenbergs, J. & Woods, S. (1996) "On the Implementation and Evaluation of AbTweak", *Computational Intelligence* 12
- (Young and Moore, 1994) Young, R.M. & Moore, J.D. (1994) "DPOCL: A principled approach to discourse planning" in *Proceedings of the 7th International Workshop on Natural Language Generation*, Kennebunkport, Maine, pp13-20
- (Young and Moore, 1994a) Young, R.M. & Moore, J.D. (1994) "Does Discourse Planning Require a Special-Purpose Planner?" in *Working Notes of the AAAI'94 Workshop on Planning for InterAgent Communication*
- (Young *et al.*, 1994) Young, R.M., Moore, J.D. & Pollack, M.E. (1994) "Towards a Principled Representation of Discourse Plans" in *Proceedings of the 15th Conference of the Cognitive Science Society*
- (Young *et al.*, 1994a) Young, R.M., Pollack, M.E. & Moore, J.D. (1994) "Decomposition and Causality in Partial Order Planning" in *Proceedings of the Second International Conference on AI and Planning Systems (AIPS'94)*, Chicago, IL
- (Young, 1996) Young, R.M. (1996) "A Developer's Guide to the Longbow Discourse Planning System", University of Pittsburgh, ISP TR-96-2
- (Zlotkin and Rosenschein, 1990) Zlotkin, G. & Rosenschein, S. (1990) "Negotiation and Conflict Resolution in Non-Cooperative Domains" in *Proceedings of the 9th National Conference on Artificial Intelligence (AAAI'90)*, AAAI, pp100-105
- (Zlotkin and Rosenschein, 1994) Zlotkin, G. & Rosenschein, J. (1994) "Rules of Encounter", MIT Press, Cambridge, MA
- (Zlotkin and Rosenschein, 1996) Zlotkin, G. & Rosenschein, J.S. (1996) "Compromise in negotiation: exploiting worth functions over states", *Artificial Intelligence* 84, pp151-176
- (Zukerman and McConachy, 1995) Zukerman, I. & McConachy, R. (1995) "Generating Discourse across Several User Models: Maximizing Belief while Avoiding Boredom and Overload" in *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, pp1251-1257
- (Zukerman *et al.*, 1996) Zukerman, I., Korb, K. & McConachy, R. (1996) "Preambulations on the Way to an Architecture for a Nice Argument Generator" in *Proceedings of the ECAI'96 Workshop on Planning and Natural Language Generation*, Budapest, pp31-36