University of Windsor
Scholarship at UWindsor

OSSA Conference Archive   OSSA 2

May 15th, 9:00 AM - May 17th, 5:00 PM

Persuasion Monologue

Chris Reed
University College London

Derek Long
University of Durham

Follow this and additional works at: http://scholar.uwindsor.ca/ossaarchive

Part of the Philosophy Commons

PERSUASIVE MONOLOGUE

Chris Reed
Department of Computer Science
University College London
Derek Long
Department of Computer Science
University of Durham
©1998, Chris Reed & Derek Long

Abstract:
The emphasis in most process-oriented models of argumentation is placed heavily upon analysis of dialogue. The current work puts forward an account which examines the argumentation involved in persuasive monologue, drawing upon commitment-based theories of dialogue. The various differences between monologue and dialogue are discussed, with particular reference to the possibility of designing a *monologue game* in which commitments are dynamically incurred and updated as the monologue is created. Finally, the computational advantages of adopting such an approach are explored in the context of an existing architecture for the generation of natural language arguments.

***

1. Introduction[1]

The rhetorical tradition since Aristotle has focused on the dialogical side to argumentation, assuming the interlocutors to be engaged *vis-à-vis*. More recent research in dialectics has placed further emphasis upon the study of the game of dialogue itself, yielding insights for both formal, logical investigations of argument, and so-called 'informal' accounts of fallacies.

One of the most prolific forms of persuasive argumentation, however, is the monologue: newspaper leaders, political addresses, academic writings, advertisements, and so on all share a number of key characteristics. Monologue allows no response from an audience: there is no opportunity for the speaker to assess progress, and no means of verifying assent to premises. It represents a one-shot deal, where errors can be neither detected nor corrected. It is a 'speaker-oriented' activity, relying solely on the speaker's knowledge, both of the domain of the argument, and also of the (presumed) *hearer* knowledge (of both the domain of the argument and the speaker, etc.).

Yet monologue clearly also shares many features with its multi-participatory counterpart: the structure of the argument form, the means of arranging components to maximum effect, the dominance of hearer beliefs in guiding the presentation of information. In fact, the similarities between monologue and dialogue make the distinction unclear and equivocal. Academic writing offers a good example: a paper clearly represents a discrete argument for a particular thesis, yet it both relies upon and criticises previous work, and in its turn is relied upon and criticised. Even more salient is the common scenario of a journal publishing an article followed by criticisms offered by peers, and then closing with a final riposte from the article's author. Although each piece is a distinct monologue, taken together they seem to form a complex dialogue—albeit one conducted on paper over an

extended period of time.

Perhaps, therefore, it would be appropriate to define monologue simply as a turn in dialogue: in the examples above, where audience response is not forthcoming, the dialogue simply happens to be one comprising only a single 'move'. Yet, as explained in more detail below, this view denigrates monologue—monologue *expects* no audience response, and is of necessity constructed on more assumptions regarding the hearer's beliefs.

O'Keefe (1977) distinguishes two meanings of the term 'argument' (essentially, two quite separate lexemes which happen to be homonymic): argument$_1$ and argument$_2$. The former refers to the structure created to support a case, the latter to the activity of disputation. Thus the distinction appears to be coextensive with that between monologue and dialogue, and indeed, O'Keefe seems to subscribe to this view considering his brief discussion of some linguistic ramifications (1977 p. 122): "The distinction here is evidenced in everyday talk by ... the difference between the sentences "I was arguing$_1$ that P" and "we were arguing$_2$ about Q."

However this is a rather summary conclusion, for monologue is not simply a turn in a dialogue, whereas O'Keefe implicitly suggests (p. 123) that an arguer$_2$ is perforce offering arguments$_1$. Further support for teasing apart the apparently similar distinctions of monologue/dialogue and argument$_1$/argument$_2$ comes from consideration of Brockreide's (1977) reply to O'Keefe, in which it is claimed that argument$_1$ focuses on *product*, and argument$_2$ on *process*. Phrased in these terms, this useful dichotomy seems to be orthogonal to that of monologue/dialogue, in that each can be viewed as either product or process. The distinction is of particular interest in considering monologue, where as O'Keefe implies, the usual stance is product-oriented (in contrast to the usual process-oriented analysis of dialogue). The fact that monologue too can be seen as a process is often overlooked. It is only when considering how to artificially generate monologue that the issue is forced: to create a monologue from a set of beliefs and goals is necessarily a process.

2. Defining Monologue

Taking a process-oriented stance towards the analysis of monologue still leaves the problem of defining persuasive monologue, clearly delimiting it from dialogue. There are several potential misconceptions regarding the nature of monologue, and the following discussion covers three of the most plausible: (i) that monologue is a record of a chain of reasoning (or similarly, that it is a soliloquy); (ii) that monologue is a dialogue in which the speaker plays both roles; (iii) that monologue is a turn in dialogue.

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 of an audience (most usually to convince, but also to shed doubt, confound, impress, etc.), and to this end, makes careful use of assumed hearer beliefs. 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.

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). For a process of "having a dialogue with oneself" (or, "arguing$_2$ with myself about Q"— O'Keefe, 1977, p. 122) does not involve the characteristic dynamics of dialogue: most notably, the speaker is

never brought to retraction by her "opponent". The opponent position is often forced to retract by the speaker, and occasionally the speaker will preemptively concede ground, but the underlying asymmetry remains. It is generally accepted that to engage in dialogue requires both parties to have a certain amount of flexibility: retraction must not be too easy otherwise parties could never be held to anything. Equally, retraction must be possible, or there would be no hope of one party successfully changing the beliefs of another (Walton and Krabbe, 1995, p. 10). This lack of equal footing between the 'interlocutors'—that the proponent position never retracts—differentiates the pseudo-dialogue sometimes found in monologue from true colloquy. This differentiation is further aided by a number of surface differences, such as the absence in monologue of any of the conventional turn-taking and exchange structure characteristic of dialogue.

Finally, as touched upon above, monologue is not simply a turn in dialogue. Despite the frequency of this assumption in computational research (Fawcett and Davies (1992) for example), a number of factors preclude this analysis. Monologue makes far fewer assumptions about the 'state of play' (excluding hearer beliefs) than does a turn in dialogue: it is not constrained by what the opponent has previously uttered, 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 is 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 its inherently dynamic nature. As detailed below in section 5.1, this 'extra diligence' can be shown to have a computational foundation.

Having dismissed a number of candidates for definitions of monologue, it now remains to outline key features which characterise monologue rather more accurately. There seem to be two major components of monologue. Firstly, the intuitive 'case building' of presenting arguments in support of the thesis. Premises are supported by subarguments, 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 for supporting them with; (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).

Secondly, there is the more complex technique of presenting counterarguments to the thesis propounded, and then offering arguments which defeat those counterarguments. So although the monologue itself is more than simply a record of a dialogue between speaker and (speaker's approximation to) hearer, it can nevertheless include passages which appear to be dialogic (with the omission of proponent retraction). 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 (—to paraphrase, the objection is that humans are the only beings upon which God confers a soul, and the counter, that this impinges upon His omnipotence, inasmuch as He should be able to

confer a soul upon anything).

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 ... (p. 454)

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, 1997).

3. Commitment in Monologue

The role of commitment in argument analysis has attracted significant interest, particularly in accounts of fallacy identification (Hamblin, 1970; Walton & Krabbe, 1995). The intuition behind its use is that as a dialogue unfolds, the participants must remain consistent: later claims must not contradict earlier claims, and no claim should be at odds with "dark-side" commitments, deeply entrenched beliefs which are rarely subject to change (cf. Gärdenfors, 1992). Exchanges, characterised by dialogue game rules, challenge and verify the consistency of a participant's argument. The dynamic of dialogue is defined in terms of the retraction of commitments.

Given the success of commitment-based analyses of dialogue, the question arises as to whether such an approach might also be appropriate for the study of monologue. In particular, might the concept of commitment be productive in a computational system designed to generate persuasive monologue?

There are a number of criticisms which might be levelled at the adoption of a commitment-based framework for the synthesis of monologue. In the first place, the audience does not say anything—that is, no one other than the speaker makes any binding utterances, anything which belies or incurs commitments. Clearly, we can describe the commitments of the speaker on the basis of her monologue, but this seems to be a rather pointless task: these commitments are quite static, remaining unchanged from the start of the monologue to its conclusion.[2] Furthermore, the set of operations which are applied through the course of the monologue (i.e. the locution and commitment store rules) is greatly impoverished. The speaker makes assertions, arguments, and, more infrequently, concessions (as in dialogue, this latter alleviates the burden of proof—a concept still very much in evidence in monologue). But the dynamics of challenges, requests, extractors, confronters and retractions (to take a typical list, from Walton & Krabbe's (1995, p. 149) PPD game) and their associated commitment-store operations, are all absent.

Indeed, retraction plays a crucial role in persuasive dialogue (in that it is used to define the aim), and yet monologue does not involve retraction—at least, not of the proponent's claims (the speaker may 'retract' some opposing claims following presentation of opposing arguments in a 'pseudo-dialogue'). However, the importance of retraction is a key feature of dialogue, and not of commitment itself. That monologue should not possess this feature does not entail that monologue generation should not employ the concept of commitment. However, it is necessary to establish why the approach is warranted.

Modelling the speaker's commitments is a trivial task, and the pseudo-dialogue exchanges of counterargument and counter-counterargument similarly do not merit commitment modelling. The role that a commitment based approach might play in monologue is in the development of the *expected* beliefs of the hearer, in conjunction with the beliefs and utterances of the speaker. That is, the speaker assumes acquiescence on the part of the hearer, and on the basis of those concessions builds up her argument. Crucially, this approach depends upon viewing monologue as a process, during which the commitment-store of the speaker and the speaker's model of the hearer's commitment store are dynamically updated. Monologue represents a distinctly more continuous medium than discrete turn-based dialogue, so the update process must necessarily rely on a more imposed segmentation. Perhaps the most appropriate segmentation is between statements, i.e. with commitment update performed after utterance of any single premise or conclusion. Thus after each functional unit of the monologue, the speaker's commitment-store is modified to include the content of that utterance, and the speaker's model of hearer commitments updated to reflect hearer acquiescence, as appropriate. With such a process-oriented model employing dynamic commitment allocation, it may be appropriate to see monologue as a game, in just the same way as dialogue is frequently characterised as a game. The moves of the sole player are determined by the assumed hearer commitments and the beliefs assumed to be held by the hearer in regard to speaker commitments (the speaker's commitment store is a model of what she believes the hearer understands to be her commitments).

4. Towards the Game of Monologue

If monologue can be described in terms of a game, how might that game compare with more familiar dialogue games? Walton and Krabbe (1995) describe two classes of persuasive dialogue game: the permissive (PPD) and the rigorous (RPD). The former is the freer, with more scope for the creativity of both participants and with less constraints upon retraction; the latter is more restrictive with tighter controls on appropriate responses and the ability to retract. Importantly, RPD is asymmetric, distinguishing the proponent from the opponent—the former can make assertions and pose questions whilst the latter can only concede or challenge (this is in direct contrast to PPD, which is quite symmetrical with respect to the roles of the interlocutors). Furthermore, the close regulation of RPD means that relevance is enforced by its rules, whilst in the freer PPD, relevance has to be assumed on Gricean principles.

Monologue clearly shares a number of features with RPD—the asymmetry in the two games is immediately clear: in monologue, the role of opponent is reduced to a point at which no locution is permitted at all. There are also a number of other similarities. Retraction—rare and tightly controlled in RPD—is only permitted in monologue for the speaker's espousal of opposing views. In addition, the characteristically careful and explicit adherence to 'logical steps' in RPD is also frequently found in monologue—and for similar reasons. In monologue, the speaker attempts to lead the hearer faultlessly along a chain of inference—to do this requires that each step be clear and unequivocal (since any minor quibble, even if easily quashed, could disrupt the hearer's attention or make him sceptical of subsequent reasoning, as the speaker is not afforded the opportunity of identifying and addressing such quibbles). The same desire for unequivocal clarity necessitates the RPD game.

However, monologue also manifests a number of characteristics which are much more in line with PPD. The freedom for creatively constructing lines of argumentation—rather than being rigorously (and even grammatically) tied to previous utterances—is crucial for ensuring that monologue can be as rich as dialogue. Similarly, the ability to question or challenge the assertions of opposing views is also very loosely regulated.

With similarities with both PPD and RPD, it is hardly surprising, therefore, that monologue also possesses aspects which straddle the two extremes on offer. Relevance in monologue, for example seems to be part PPD-like and part RPD-like. The freedom in PPD means that relevance can only be assumed though Gricean principles of co-operation. Similarly in monologue, for the major shifts of focus between large subarguments (eg. from one section of a paper to another), relevance must be based on cooperation. But equally, at a lower level the RPD rule-based relevance is appropriate for ensuring continuity from one sentence to the next. PPD and RPD are also at odds with regards to the importance of the 'pay-off'—in RPD, winning or losing is everything, whilst in PPD other factors (maintaining face, etc.) also play a role. In monologue, the pay-off is not all important (persuasion is not to be accomplished at all costs), but it is rather more important than in many examples of dialogue, since the speaker's investment is usually much greater (constructing a monologue is generally a much more taxing task than entering into a dialogue).

Viewing monologue—from a process-oriented viewpoint—as a commitment-based game can offer insights into how natural language monologue might be generated artificially.

5. Natural Language Generation

5.1 Monologue, Dialogue and Uncertainty

As discussed in (Reed and Long, 1997), the automatic synthesis of arguments in natural language is attracting increasing interest in artificial intelligence, due to the wide range of applications in which it is of use—expert system justifications, decision support systems, tailored medical information, in fact, any situation in which nontrivial reasoning needs to be communicated to a human user. The decision of whether to present such reasoning in monologic or dialogic form is one which has important ramifications for the underlying computational model, for the differences between monologue and dialogue are directly reflected in the planning processes that contribute towards their construction. These differences can be traced to the distinction between planning strategies which are commonly referred to as on-line and off-line. In the former, a planner is a component of a situated agent, interacting with its environment and interleaving its planning with execution, while in the latter the planner might well be envisaged as a component of an agent capable of executing the plans produced, but the planning process itself is completed prior to any execution of the resulting plans.

A planner can be seen as a system which takes a description of a state of an environment (the initial state), the actions which can be executed in that environment and a goal (a partial characterisation of a state) and that produces a sequence of the given actions which, it is hoped, can be executed from the initial state to result in a state satisfying the goal. Before considering the planning process in its application to the construction of utterances, either in a dialogue or a monologue, firstly consider the implications of exploiting either on- or off-line planning strategies.

A well-known problem confronting situated agents is the inherent uncertainty in the behaviour of most environments, making planning a potentially wasteful use of resources. Uncertainty can arise in at least two ways: actions might not have deterministic effects in the environment and the environment might include other autonomous agents. In neither case can the state of the world following execution of an action be accurately predicted. Planning involves constructing a sequence of actions, with a dependency between earlier actions and later actions, so that later actions will be executable by virtue of the successful execution of earlier actions. If actions have unpredictable effects or the environment can change in ways that can neither be predicted nor

entirely controlled, the dependencies cannot be relied on and later components of the plan might well be unexecutable (and even entirely inappropriate) as a consequence of unforeseen circumstances. For this reason, planning in uncertain environments is better restricted so that only local elements of the plan are constructed in detail, with more distant elements left unplanned, at least at a detailed level (the role of abstracted plan descriptions will be explored briefly later in this paper). Indeed, it is just this unpredictability that has motivated work in the construction of situated agents which abandon all efforts at planning and rely entirely on a reactive interaction with their environment. Purely reactive behaviour is unsuitable in environments in which the goals being sought can be a significant number of actions away from the initial state and in which the states through which an appropriate sequence of actions might progress towards the goal carry no intrinsic indication of convergence towards a goal state. In uncertain environments on-line planning is a sensible strategy to adopt, interleaving planning effort with execution, producing a detailed plan for only a few steps before executing them and monitoring their effects. The degree of uncertainty in the effects of actions and in the behaviour of the environment will influence how long such planned action sequences should be to achieve an efficient use of resources. However, each sequence of actions which is planned and then tentatively executed does not, alone, represent a complete plan to achieve the overall goal towards which the planner is working—it is only a component in such a plan and achieves local goals on the route towards that overall goal. Additional important influences on the length of the action sequences which can be usefully planned by a situated agent are the relative rates of activity within the planner itself, the actions of the agent and the changes wrought on the environment by other agents. In particular, the planner must face a bound on the time spent planning a sequence of actions and the related execution time for the planned sequence if the plan is intended to anticipate activities of other agents in the environment. Further elements that arise in the context of planning within situated agents are the problems of recognising plan failure and of replanning through the reuse of components of failed plans. The first of these elements demands some ability, on the part of the agent, to sense and interpret aspects of its environment.

In contrast, off-line planning is suitable when the planner confronts a certain and entirely predictable environment. In these circumstances, an entire sequence of actions can be planned from the initial state to the goal state prior to execution. Resource consumption during planning is not a critical issue, although plans must obviously be produced within reasonable periods of time. The plans which are produced might, of course, be required to conform to bounds on time of execution in order to achieve their intended effects.

Turning, now, to the characteristics of the problems of planning persuasive monologue and dialogue, the following observations may be noted. In dialogue the speaker is situated. The environment in which the actions—utterances—take place is described by the states of the beliefs of the participants. The utterances of the speaker can change the environment by changing the beliefs of both speaker and hearer. Unfortunately, the changes these actions bring about are non-deterministic—an assertion might be accepted and adopted, merely observed, or perhaps even misinterpreted. In addition, the hearer can change the environment autonomously, both by privately revising his own beliefs and by making utterance-actions of his own. The access the speaker has to the state of the environment is limited, since the beliefs of the hearer are not visible but can only be inferred through observation of the reactions of the hearer to the speaker's actions. There are resource bounds set on the interactions possible in dialogue: participants are governed by unwritten rules of social interaction which constrain the lengths of their turns. Finally, the important shared resource, the focus of the dialogue, cannot be exclusively controlled by the speaker and the speaker will frequently be forced to replan elements of her discourse to turn focus back towards issues relevant to achieving her goals.

In monologue, the speaker typically does not have access to a particular hearer and cannot observe the reaction of a hearer to the monologue. Instead, 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.

These observations illustrate that dialogue and monologue represent planning and execution tasks set against different environments. The difference in the environments suggests different planning strategies are appropriate for each of the tasks.

This section concludes with some comments on the role of abstraction-based planning in the construction of dialogue. When a situated agent confronts the task of planning a sequence of actions to achieve a goal, as has already been discussed, the length of the planned sequence must be balanced against the expectation of failure in execution of long sequences. However, a criticism levelled at purely reactive behaviour is that it is unsuitable in environments where locally possible actions do not have obviously goal-directed effects, so that the effects of sequences of actions must be foreseen in order to perceive the utility of a locally possible action. It would appear that without a clear goal-directed purpose, longer sequences of actions fall prey to the same criticism. This can be pre-empted by the use of abstraction-based planning, in which coarse-grained abstract plans are formed which achieve the goal, from the initial state for a planning problem, with relatively few abstracted actions. These abstracted actions can then be refined towards detailed plans, acting as a scaffolding within which the details can be developed. This abstract plan ensures that the detailed local plan sequences are goal-directed, since they sit within the context of a framework which directs their continuation. At the same time, abstracting out details of execution, an abstract plan can prove more robust in the face of local plan failure, so that failed plans can be detected and restructured during execution, while maintaining the abstract plan structure which guides the process. Abstraction-based planning has, in fact, many other advantages to offer in the construction of discourse plans, both for monologue and for dialogue, allowing a hierarchical development of discourse from logical structure through textual structure to lexical structure of individual utterances.

5.2 Generating Monologue

The majority of recent work in natural language generation (NLG) (surveyed by Hovy, 1993) concentrates upon the issue of generating text which successfully conveys the requisite information in a coherent, understandable form. As a consequence, the emphasis is placed heavily upon monologue, with architectures drawing upon theories of sentence and paragraph structure developed in computational linguistics (such as Mann and Thompson's (1987) Rhetorical Structure Theory). Some systems purport to handle full, 'end-to-end' dialogue (i.e. to parse and understand the user's utterances, to design an appropriate response and then to plan the realisation of that response), and a few do, such as Allen's (1996) TRAINS project. However, without exception these systems are only capable of short exchanges composed of short utterances. Such oversimplification is essential if the highly complex problems of computational dialogue are to be assailable: quite apart from the task of language understanding (from syntactic parsing to pragmatic analysis) there are issues

regarding recognition of the user's plan, revision of the system model of hearer beliefs, dealing with the various forms of communication failure (disagreement, miscomprehension, non-comprehension, etc.), following rules of exchange, and so on.

By concentrating on monologue, it is possible to focus on the issue of building arguments in natural language without incurring the additional, overwhelming problems of full dialogue. This is the approach taken in (Reed *et al.*, 1996), in which argument construction proceeds by linking pieces of information with operators which represent various types of reasoning. Ordering is then imposed to ensure coherency, and the result then undergoes postprocessing to introduce appropriate lexicalisation (such as 'clue words' linking premises and conclusions in various arrangements), and to reorder components according to rhetorical maxims to improve persuasive effect.

There are two key problems which face automated generation of monologue: the selection of information to include and the subsequent ordering of that information to ensure the resulting text is both coherent and persuasive. The content selection task is composed of two consecutive phases: what information to include in the structure of the argument, and then which items to express in the final text (—compare this deep structure to surface structure synthesis with the reverse process, common to many analysis techniques, of reconstruction (Eemeren *et al.*, 1993)). It is important to distinguish structural content from that expressed in order to permit enthymematic contraction and other devices for avoiding redundancy: the primary phase of content selection determines the information which needs to be made *salient* to the hearer, the secondary how that salience is to be achieved. For example, in a Modus Ponens step, A, A ->B/ B, the second premise is rarely uttered explicitly, though clearly it is assumed to be salient to the hearer. There are a number of heuristics which can be brought to bear upon the question of how to achieve saliency, based upon studies of a corpus of arguments. The primary content selection task, however, is more difficult. In section 2, above, it was shown that arguments are ultimately supported by basic premises which either have no further support in the speaker's beliefs; thought by the speaker to be believed by the hearer; or thought by the speaker to be accepted by the hearer without further argumentation. It is this final condition which is problematic for a computational account, as there is no obvious means of determining whether a hearer would need further argumentation in order to believe any particular proposition. The problem of determining an appropriate level of detail is one which confronts all NLG work, and not just argument synthesis.

The other central issue in generating monologue is ordering. Constraining order to ensure coherency is a complex computational problem, but one which is well defined and well understood. Maximising persuasive effect, on the other hand, is an ill structured problem with no consensus on rules or approach (even between rhetoricians), and no clear rubric for evaluation. Although there a number of heuristics presented in (Reed *et al.*, 1996) which have been gleaned from rhetoric texts and corpus study, it is unclear that the list is complete or even sufficient for improving persuasiveness.

It is proposed that the framework set out in (Reed *et al.*, 1996) could benefit from the use of commitment modelling and the notion of a 'monologue game' in tackling the challenge posed by the significant problems outlined above. Determining an appropriate level of detail on the basis of what the hearer will accept without further argumentation can be considered in terms of the monologue game as whether the hearer would happily make concession to a proposition. In turn, this question can be addressed through consideration of potential hearer counterarguments (based, of course, upon the speaker's model of the hearer's beliefs). If the speaker cannot envisage the hearer forming substantial counterarguments, she may assume that he will concede. If however, his counterarguments are sufficiently weighty (as they would need to be to prevent the hearer from

conceding—in the absence of undue scepticism or bias), the speaker has three choices: (i) to present further supporting argumentation for the premise (assuming such support exists, and the hearer's counterargument is not strong enough to cancel the effect of any such argumentation); (ii) to present the counterargument and defeat it with a counter-counterargument (again, assuming such a counter exists); or (iii) to avoid using the proposition altogether (assuming the argument can be (re)constructed appropriately).

Consideration of the monologue game can also assist in effecting appropriate ordering. Coherency ordering in (Reed *et al.*, 1996) is achieved through the use of a context parameter in the goals expressing propositional salience, and this parameter fulfils an almost identical role to the notion of 'local thesis' defined for RPD. In addition, the locution and commitment rules of the monologue game provide a clear framework within which to set the persuasion ordering constraints, and also offer scope for a new set of such constraints—that the hearer should not be forced to make too many consecutive concessions without any supporting argumentation, for example.

The employment of a commitment-based monologue game can thus offer a clear, structured means of tackling a number of conceptual and computational problems in the generation of natural language arguments.

6. Conclusions

The adoption of a process-oriented view towards the structure of monologue facilitates the design of a commitment-based model of monologue, drawing upon similarities and differences between the processes of monologue and dialogue. Such a model offers two useful insights. In the first place, monologue is more than just a record of speaker reasoning, an internalised dialogue, or a turn in dialogue: the notion of commitment assists in more accurately characterising the nature of monologue. Secondly, the concept of the monologue game, licensed by the underlying commitment-based theory, is of use in tackling a number of major issues facing the automatic generation of monologue, and in particular, in dealing with content selection, a fundamental issue in NLG as a whole.

Notes

1. There are a number of conventions adopted for this paper. The author of a monologue, whether it is written or spoken, is referred to as the 'speaker', and is pronominalised using the feminine 'her' when alternative phrasing would be cumbersome. The target recipients of the monologue—whether there are one or more—are referred to synonymously with 'audience' and 'hearer' (i.e. an audience comprising several individuals is assumed to be modelled by a single prototypical set of beliefs or commitments), and pronominalised using 'him'. The gender of the pronouns is not intended to convey aspects of any particular scenario: both should be read as neuter.

2. It is assumed, as with all theories based upon the employment of commitment-stores, that the model is somewhat idealised, in that the participants (in monologue, just the speaker) do not forget their commitments. This idealisation is even more reasonable in monologue than in dialogue, as the former is often produced as the result of more careful preparation (eg. in writing an academic paper or delivering a political address).

References

(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), pp. 7-48.

(Brockreide, 1977) Brockreide, W. (1977) "Characteristics of Argument and Arguing", *The Journal of the American Forensic Association* **13** pp. 129.

(Eemeren *et al.*, 1993) Eemeren, F.H. van, Grootendorst, R., Jackson, S. & Jacobs, S. (1993) *Reconstructing Argumentative Discourse*, University of Alabama Press, Tuscaloosa.

(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), *Aspects of Automated Natural Language Generation*, Springer Verlag, pp. 151-166.

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

(Hamblin, 1970) Hamblin, C.L. (1970) *Fallacies*, Methuen, London (Hovy, 1993) Hovy, E.H. (1993) "Automated discourse generation using discourse structure relations", *Artificial Intelligence* **63**, pp. 341-385.

(Mann and Thompson, 1987) Mann, W.C. & Thompson, S.A. (1987) "Rhetorical Structure Theory" in G. Kempen (ed) *Natural Language Generation: New Results in Artificial Intelligence, Psychology and Linguistics,* Martinus Nijhoff, pp. 85-95.

(O'Keefe, 1977) O'Keefe, D.J. (1977) "Two concepts of argument", *The Journal of the American Forensic Association* **13** pp. 121-128.

(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)* Springer Verlag, Berlin, pp. 555-566.

(Reed and Long, 1997) Reed, C.A. & Long, D.P. (1997, to appear) "Multiple Subarguments in Logic, Argumentation, Rhetoric and Text Generation" in *Proceedings of the International Joint Conference on Quantitative and Qualitative Practical Reasoning*, Springer Verlag, Berlin.

(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.