Speaker-hearer beliefs for discourse planning

Massimiliano Garagnani
Department of Computing
The Open University
Walton Hall, Milton Keynes, MK7 6AA, U.K.

Abstract *Using sophisticated and expressive belief languages for generating complex discourse plans requires a clear distinction between the belief system and the planning machinery which makes use of it. The main results presented in this paper consist of: 1) the formalisation of an expressive 'speaker-hearer' belief system which allows the representation of a structure of belief justification and grounding, uncertain and contradictory beliefs and communicative intentions; 2) a set of discourse planning operators — based on the belief system — which have been used by the 'IPP' planner to generate 'persuasive' discourse plans.*

Keywords: Belief System, Planning, Natural Language.

1 Introduction

One of the main areas of research in the field of Natural Language (NL) processing focuses on the problem of automatic generation of *discourses*. A 'discourse' can be defined as a linearised structure of semantically related statements generated by a source ('*speaker*', S) and conveyed to an audience ('*hearer*', H) in order to achieve a specific '*communicative intention*' (or goal) [13] [16]. The communicative intentions constitute the 'purposes' which lie behind the generation of the discourse. For example, the speaker might want to 'explain' a concept, or 'persuade' the hearer about the validity of a statement.

Because of the inherently hierarchical and decompositional nature of the communicative goals, the AI *planning* technique has been successfully used by many as a method to tackle this problem (e.g. [3] [14] [16] [19] [20] [27]). This work adopts the same approach, and sees the problem of discourse generation as an application domain for the theory of planning, in which the output of the planner will consist of a *discourse plan*.

The type of communication considered has been limited to that of *sincerely*-generated *persuasive* discourses, i.e. mono-directional flows of truthful statements produced with the aim of convincing the audience about the validity (or fallacy) of a certain proposition.

Planning a communication that achieves a *change* of the hearer's mental attitude towards a specific fact requires the speaker to hypothesise and maintain an adequate *model* of the hearer's *beliefs* and to update it according to the effects that the speaker's utterances (are expected to) have on the hearer's doxastic state. Indeed, in the theories of discourse and communication the need for the speaker to maintain an internal representation of the hearer's knowledge and beliefs has been widely acknowledged (e.g. [1] [14] [26] [20]).

However, all of the discourse planning systems described in the recent research literature lack a clear separation between the *model of beliefs* adopted and the *planning system* which uses it. They do not give a precise formalisation of an expressive language of beliefs and seem to neglect the issues of consistency, belief grounding and logical omniscience, normally associated with the introduction of a system of beliefs. Moreover, the representation of the hearer model is often not clearly distinguished from that of the speaker, and suffers from the same representational shortcomings.

A clear definition of a formal speaker-hearer belief system for discourse plan construction is needed for the realisation of complex, real-world persuasive discourses.

This paper presents the formalisation of such a system and describes a set of '*discourse operators*' defined accordingly, which have been used by 'IPP' [17] (a Graphplan-based planner [2]) for the automatic generation of persuasive discourse plans.[1]

2 The belief system

According to the classification given by Konolige in [18], the belief system developed here belongs to the class of 'syntactic' — or interpreted symbolic structures — models. In particular, following Ellis [7], a belief system is, in its simplest form, a set of assignments of the values *True*, *False* and X to the sentences of a language[2]. Similarly, in this work a belief system is computationally defined as a calculable function (called 'Belief Prover', or BP function) which, given the current *state* — i.e. the set of beliefs currently held — assigns either *True* or *False* to each sentence of a language Lo (see Fig. 1).

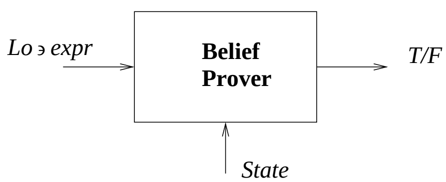

Figure 1: A belief system.

According to the previous definition, a belief system can be also seen as a 'black-box' whose behaviour is unequivocally described by a *formal theory*. In fact, given the set Lo of well-formed formulas (*wff*s) which constitute the language of the system, the current *state* of the belief system can be thought of as the set of *axioms* of the theory. A certain query $expr \in Lo$ will be evaluated *True* by the BP function *if and only if* '*expr*' is a *theorem* of the theory, where the theorems of a formal theory consist of all the axioms plus all the *wff*s which can be *derived* from the axioms through the use of a set of *inference rules.*

The two following subsections introduce and discuss, respectively, the language of beliefs Lo adopted, and the inference-rule based BP function of the system implemented.

2.1 The Language

The BNF formalisation of the language Lo developed for the belief system is given below:

$$
\begin{aligned}
\langle Lo\rangle &::= \langle Mo\rangle(\langle A\rangle) \mid \mathrm{NOT}(\langle Mo\rangle(\langle A\rangle))\\
\langle Mo\rangle &::= \mathrm{HBEL} \mid \mathrm{HUND} \mid \mathrm{HUNK}\\
&\quad\;\; \mid \mathrm{SBEL} \mid \mathrm{SUND}\\
\langle A\rangle &::= \langle E\rangle \mid \mathrm{NOT}(\langle E\rangle)\\
\langle E\rangle &::= \mathrm{SR}\,(\langle Es\rangle) \mid \mathrm{HR}\,(\langle Es\rangle) \mid \langle Es\rangle\\
\langle Es\rangle &::= \mathrm{SUP}(\langle En\rangle, \langle En\rangle) \mid \langle Ei\rangle\\
\langle En\rangle &::= \langle Es\rangle \mid \mathrm{NOT}(\langle Es\rangle)\\
\langle Ei\rangle &::= \mathrm{E}_0 \mid \mathrm{E}_1 \mid \mathrm{E}_2 \mid \mathrm{E}_3 \mid \ \ldots \ \mid \mathrm{E}_n \mid \ \ldots
\end{aligned}
$$

Examples of correct expressions are $\mathrm{NOT(HBEL(E_2))}$, $\mathrm{HUND(SUP(E_0,NOT(E_5)))}$, $\mathrm{SBEL(NOT(HR(E_8)))}$. Examples of expressions which cannot be generated by $\langle Lo\rangle$ are $\mathrm{HBEL(SBEL(E_1))}$, $\mathrm{HR(NOT(E_4))}$.

The first three modal operators of $\langle Mo\rangle$ should be read as "Hearer BELieves", "Hearer is UNDecided" and "Hearer UNKnows". The attitude of undecision 'HUND(e)' towards the proposition (or '*event*') 'e' expresses the hearer *uncertainty* in believing one of the two events 'e' and 'NOT(e)'. 'HUNK(e)' means that (the speaker believes that) the concept 'e' has *never been formulated* in the hearer's mind.[3] A similar but somewhat reciprocal notion (namely, 'KNOW-ABOUT'), expressing 'acquaintance' with a concept, has been adopted by Moore and Paris in [20].

While the last two expressions of $\langle Mo\rangle$ are the 'Speaker version' of the first two, the modal operator SUNK( ) has not been introduced in the syntax. This is because whenever the belief SUNK(e) is conceived by the speaker, it becomes

[1] The actual *integration* of these two elements — namely, the belief system and the operators — into a single discourse planning system constitutes an entirely different problem in itself, and has been treated by the author elsewhere (see [11] [12]).

[2] The value X indicates the absence of any clear attitude towards a sentence.

[3] This attitude should not be confused with a situation of *incomplete* knowledge.

immediately *False* by definition, as the event '*e*' is part of the thought and loses its state of unknown (read 'non-formulated') concept.[4]

The "SUPport" relation SUP( ) allows the representation of a structure of belief *justification*. In other words, the expression $\text{SUP}(e, e')$ indicates that the belief in the event e is a sufficient reason to believe e'. The expressions 'HR' and 'SR' mean, respectively, "Hearer Real" and "Speaker Real", and have been introduced to allow the representation of *grounded* beliefs, i.e. beliefs whose truth is, from the agent's point of view, unquestionable.[5]

Looking at *Lo* as a language for discourse generation, it is possible to see that *Lo* is expressive enough to capture the relevant information and concepts which will be needed during the process of construction of a *sincere*[6] discourse. These include the *functional relations* which can hold between the different parts of a discourse (as accounted by the 'standard' analysis of arguments [10] and by the Toulmin schema [25] — see Fig. 2), the *intentional structure* of the discourse (i.e. the communicative goals, like "persuade the Hearer to believe that '*e*' ") and the (defeasible) relations of '*support*' and '*attack*' between the different sub-parts of an argument (cf. [6] [22]). Figure 2 shows how the five basic functional relations can be expressed using relations of support.

Looking at *Lo* as a language for belief systems, the distinction between speaker and hearer models results clear. This separation and the presence of expressions for belief grounding and justification allow the representation of belief structures that satisfy both the 'foundations' and 'coherence' paradigms[7] [4]. Moreover, the presence of *uncertainty* in the model is allowed and dealt with through the adoption of the 'HUND' and 'SUND' expressions.

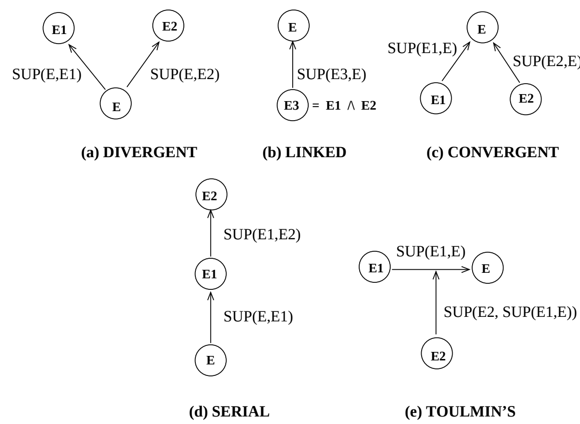

Figure 2: Standard and Toulmin's argument structures realized with 'SUP'.

2.2 *BP* function and Inference rules

Given a language *Lo*, a set R of inference rules and a (belief) state $S \subseteq Lo$, the BP function can be formally defined as follows:

$$BP(x) = \begin{cases} True & \text{if } x \in C_R(S) \\ False & \text{otherwise} \end{cases}$$

where $x \in Lo$ and $C_R(S)$ is the *deductive closure*[8] of S on R.

The entire set R of inference rules actually implemented in the system contains more than thirty rules. For reasons of space, only the most representative ones have been reported here:

[4]This situation presents an interesting symmetry with the notion of Cartesian basic belief "*I think, therefore I am*" (or '*Cogito ergo sum*'), which becomes immediately *True* (for the thinker) as soon as it is conceived [23].

[5]In this work, a belief is considered 'grounded' when the event object of the belief has been *witnessed* by the believer. For example, the expression $\text{HR}(\text{E}_3)$ indicates that the happening of the event E_3 has been perceived *directly* by the Hearer through *sensorial experience*.

[6]It is possible to show that only if the possibility of *deception* is considered by the system, will expressions containing two levels of belief nesting – like HBEL(SBEL(...)) – be required.

[7]According to the 'foundations' approach, if an agent believes an event '*e*', then '*e*' must either be *justified* by the presence of other beliefs, or constitute a *grounded* belief. However, when planning a discourse, a speaker's belief may be left 'unjustified' if it is also believed by the audience (as suggested by Cohen and Perrault in [3]). This is in accordance with the 'coherence' approach, which forbids *contradiction* but does not require justification or grounding of the beliefs.

[8]The deductive closure of a set of propositions $S \subseteq Lo$ on a set of inference rules R is defined by Konolige [18] as $C_R(S) = \{\omega \in Lo \mid S \vdash_R \omega\}$ where $S \vdash_R \omega$ means that ω can be derived from S using only rules in R.

i_1) $e \vdash$ SBEL(e)
i_2) SR(e) $\vdash e$
i_3) HR(e) $\vdash e$
i_4) HBEL(SR(e)) $\vdash$ HBEL(e)
i_5) $\neg$ SR(e) $\vdash$ NOT(SR(e))
i_6) $\neg$ (e $\vee$ NOT(e)) $\vdash$ SUND(e)
i_7) $\neg$ (HBEL(e)$\vee$ HBEL(NOT(e))$\vee$ HUND(e)) $\vdash$ HUNK(e)
i_8) HUND(e) $\vdash$ HUND(NOT(e))
i_9) HUNK(e) $\vdash$ HUNK(NOT(e))

For each rule, the variable 'e' can match any expression of $\langle A \rangle$ which produces a rule containing *wff*s of $\langle Lo \rangle$. The symbols '$\neg$' and '$\vee$' in rules i_5)-i_7) should be interpreted as classical logic operators. For example, the premiss of i_6) is verified iff neither the belief 'e' nor the belief 'NOT(e)' is evaluated *True* by the *BP* function.

In order to be usable in frameworks with limited resources and avoid the problem of *logical omniscience* [15], the *BP* function implemented needed to be able to deduce new beliefs from existent ones according to the given rules and, at the same time, limit such process of deduction to the minimum amount necessary. This has been achieved by: 1) using a *lazy* approach to calculate the deductive closure of the current belief state; 2) defining the 'SUP' relation as a *defeasible* [22] concept of 'support' instead of strict *logical implication*.

On the one hand, the 'lazyness' of the process of calculation of the deductive closure has been realized by limiting the application of the inference rules to the set of propositions contained in the current '*discourse context*' [21]. In other words, the system uses the given discourse communicative goal(s) to select, among the set of expressions contained in $\langle A \rangle$, the *finite* list of events which are *relevant* for the construction of the current discourse, and calculates a (finite) deductive closure of the state by limiting the *instantiation* of the inference rules exclusively to such relevant events.

On the other hand, the support relation is seen as 'defeasible' in that even if BEL(e) and BEL(SUP(e, e')) hold, it is *not necessarily* true that BEL(e'). This possibility is allowed by the presence of *uncertainty* in the model.[9] Therefore, the beliefs of the system are not (necessarily) closed under *logical implication*, and the doxastic version[10] of the normal modal logic axiom K) $\models \Box(\phi \Rightarrow \psi) \Rightarrow (\Box\phi \Rightarrow \Box\psi)$ (see [24]) does not hold. Because of this feature, the model can formalise situations containing *contradictory* beliefs and yet be *locally coherent* and avoid the widespreading of inconsistency across the rest of the system, as suggested by Doyle in [5].

3 The discourse operators

This section describes the discourse planning operator-schemes defined using the language *Lo*. The operators developed and the *BP* function have been used by the IPP planner [17] to produce plans for truthful, persuasive discourses. An example of discourse plan has been reported at the end of the section.

Because of the requirement of *sincerity*, the speaker's strategy of persuasion cannot make use of deceptive techniques or rhetorical fallacies. Hence, the fundamental method adopted for convincing the hearer about the validity of a fact relies upon the '*Modus Ponens*' (MP) rule: $a \wedge (a \Rightarrow b) \vdash b$. More precisely, given a situation in which HUND(b) holds, the speaker will first (try to) achieve the communicative goals HBEL(a) and HBEL(SUP(a, b)), and then lead the hearer to the conclusion HBEL(b).

The '*abstract*' operator[11] "Persuade" reported below encodes the application of such strategy:

[9]In fact, suppose that all of the following events are believed true: { e', e'', SUP(e', e), SUP(e'',NOT(e)) }. The only coherent attitude which can be adopted towards 'e' is one of *undecision*. But this means that BEL(e') and BEL(SUP(e', e)), and yet UND(e).

[10]Obtained by replacing '$\Box$' with 'BELieve'.

[11]An abstract operator consists of an external '*shell*' and of a '*body*' (cf. [9]). The shell specifies the *filters* — conditions which must be true at the moment of the operator application — and the *effects* (lists of propositions added to and removed from the current state) that its execution will produce. The body contains the list of goals whose achievement enables the application of the operator's effects.

Persuade(?a,?b)
Filter: SBEL(SUP(a,b)), SBEL(a), SBEL(b), HUND(b)
Body: HBEL(a), HBEL(SUP(a, b))
Add: HBEL(b)
Delete: HUND(b)

Notice that the conditions in the filter guarantee the *sincerity* of the speaker during this step of persuasion. If required, the proposition HUND(b) can be moved into the 'body' of the operator, allowing the planner to consider it as a goal and achieve it through the application of a different operator.

The alternative method which can be adopted for convincing the hearer (H) about the validity of an event 'a' consists of asserting that 'a' is a Speaker-Real event — SR(a), i.e., that the happening of 'a' has been witnessed directly by the speaker. Because of the sincerity of S — of which H *is assumed to be aware* — H is then expected to believe that SR(a), and, as a consequence of the inference rule i_4), that 'a' is true:

Assert-Real(?a)
Filter: SBEL(SR(a)), HUNK(SR(a))
Body:
Add: HBEL(SR(a))
Delete: HUNK(SR(a))

An abstract operator with empty body like "Assert-Real" is called '*primitive*'. The effects of a primitive operator are a direct consequence of the filters being true and do not require the achievement of a separate list of goals.

Another (primitive) discourse operator which has been implemented is shown below. Its role simply consists of asserting a proposition 'a'. When "Assert(a)" is executed, the listener is forced to 'conceive' the event 'a' itself, abandon the current attitude of *unawareness* – HUNK(a) – and adopt one of '*uncertainty*' – HUND(a):

Assert(?a)
Filter: SBEL(a), HUNK(a)
Body:
Add: HUND(a)
Delete: HUNK(a)

The three remaining operators developed implement three different techniques of belief *undermining*. The first two are based on the idea of *undercutting* argument [22]. They effect the 'spoiling' of the hearer's belief in 'a' by making H doubt the validity of the event 'b' which *supports* it, or, alternatively, the validity of the support relation SUP(b, a) itself. The operator defined for the first case is given below[12]:

Spoil-Arg(?a,?b)
Filter: HBEL(a), HBEL(SUP(b, a)), NOT(SBEL(a)), NOT(SBEL(b)) NOT(HBEL(HR(a)))
Body: NOT(HBEL(b))
Add: HUND(a)
Delete: HBEL(a)

The undermining of the supporting event 'b' is achieved in the body by the goal NOT(HBEL(b)). This goal, in turn, may be achieved by the application of another undermining action. Notice that beside the conditions which guarantee sincerity, the filter contains also the condition NOT(HBEL(HR(a))). This requirement prevents the operator from being applied with the aim of undermining a Hearer-Real event. In fact, due to the unchangeable character of beliefs which are grounded on the basis of personal *experience*, such persuasion attempt is expected to fail.

The third and last undermining operator is based on the concept of *rebuttal* of an argument [22], and consists of attacking the hearer's belief in an event 'b' by supplying the audience with supporting evidence for the contrary event, NOT(b):

Attack(?a,?b)
Filter: HBEL(b), NOT(HBEL(HR(b))), SBEL(a), NOT(SBEL(b)), SBEL(SUP(a,NOT(b)))
Body: HBEL(a), HBEL(SUP(a,NOT(b)))
Add: HUND(b)
Delete: HBEL(b)

The above operator-scheme 'attacks' the belief 'b' through the achievement of the (communicative) goals HBEL(SUP(a,NOT(b))) and

[12] The operator defined for the alternative strategy — i.e. spoiling the support — is identical to this one, except for the body, which requires the achievement of the goal NOT(HBEL(SUP(b, a))).

HBEL(a). Given that HBEL(b), the achievement of these goals is expected — at least — to produce in the hearer an attitude of undecision towards 'b'.

The operators described above have been used by the IPP planner to produce plans for sincere persuasive discourses[13]. An example of discourse planning problem (represented as an initial belief state and a set of communicative goals) which the system has successfully solved is represented graphically in Figure 3. In this example, the only assigned goal was HUND(E1):

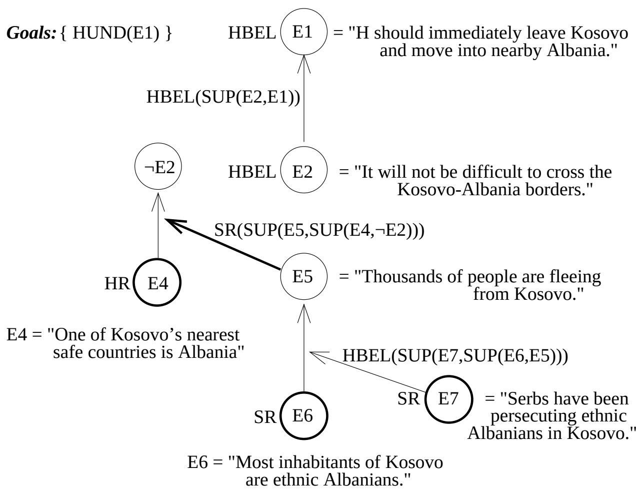

Figure 3: Example of communicative goal and initial belief state.

The elements in bold graphics represent S's or H's *Real* events; '¬ E2' is a short for 'NOT(E2)'. The plan found by IPP is reported below[14]:

```
0: Assert(E5), Assert-Real(E6),
   Assert-Real(SUP(E5,SUP(E4,NOT(E2)))),
   Assert-Real(E7), Assert(SUP(E6,E5)),
   Assert(SUP(E4,NOT(E2)))
1: Persuade(E7,SUP(E6,E5))
2: Persuade(E6,E5)
3: Persuade(E5, SUP(E4, NOT(E2)))
4: Attack(E4,E2)
5: Spoil-arg(E1,E2)
```

[13]The syntax of IPP does not accept abstract operators with filters and body. These two lists have been actually merged into a single '*precondition*' list, like in the STRIPS [8] formalism.

[14]The specific NL interpretation given to the events of this example is not taken into account by the planning process, at this level of abstraction.

In step 0, several facts are asserted (in 'parallel') in order to create the appropriate discourse *context*. These include the relevant *real* events situated at the 'basis' of the belief structure — representing the actual *premisses* of the discourse — and the events of which the hearer is currently unaware. In the remaining steps, the plan achieves the goal of 'shaking' the H's belief in E1 by *attacking* its support, E2. This, in turn, is done by providing H with supporting evidence for NOT(E2) through a chain of consecutive steps of persuasion.[15]

4 Conclusions

In summary, this work has argued that using sophisticated and expressive belief languages for generating complex persuasive discourse plans requires a clear distinction between the belief system and the discourse planning machinery which makes use of it. The main results presented consist of: *i*) the formalisation of an expressive speaker-hearer belief system, identified by a set of inference rules and a language *Lo*; *ii*) a set of abstract discourse planning operators (defined using *Lo*) which encode different strategies of belief persuasion and undermining.

The belief system implemented can be also integrated into planners which are more suitable than IPP for generating discourses, such as *causal link* planners [27] or *hierarchical* planners, which allow the direct adoption of abstract operators (e.g. [9]).

Therefore, the results presented constitute a contribution towards the identification of a framework for the development of expressive belief systems for discourse planning and, more in general, for the implementation of new AI systems with sophisticated NL processing capabilities.

References

[1] Asher, N., Koons, R. (1993) "The revision of Beliefs and Intentions in a Changing World", in *Proceedings of the AAAI Spring Symposium*

[15]This plan was found by IPP — running on a SPARC workstation — in less than 3 seconds of CPU time.

Series: Reasoning about Mental States: Formal Theories and Applications.

[2] Blum, A.L., Furst, M.L. (1997) "Fast Planning Through Planning Graph Analysis", *Artificial Intelligence*, 90:281–300.

[3] Cohen, P.R., Perrault, C.R. (1979) "Elements of a Plan-Based Theory of Speech Acts", *Cognitive Science*, 3:177–212.

[4] Doyle, J. (1992) "Reason Maintenance and Belief revision: Foundations vs. Coherence Theories", in Gärdenfors, P., (ed), *Belief Revision*, Cambridge University Press, Cambridge, pp. 29–51.

[5] Doyle, J. (1994) "Reasoned assumptions and rational psychology", *Fundamenta Informaticae*, 20(1–3):35–73.

[6] Dung, P.M. (1995) "On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and *n*-person games", *Artificial Intelligence*, 77:321–357.

[7] Ellis, B. (1979) *Rational Belief Systems*, Blackwell, Oxford.

[8] Fikes, R.E., Nilsson, N.J. (1971) "STRIPS: A New Approach to the Application of Theorem Proving to Probelm Solving", *Artificial Intelligence*, 2:189–208.

[9] Fox, M., Long, D. (1995) "Hierarchical planning using abstraction", *IEEE Process-Control Theory and Applications*, 142(3):197–210.

[10] Freeman, J.B. (1991) *Dialectics and the Macrostructure of Arguments*, Foris, Dordrecht.

[11] Garagnani, M. (1999) "A Sound Linear Algorithm for Pre-processing planning problems with Language Axioms", *Proceedings of PLAN-SIG '99*, Manchester, UK, pp. 40–53.

[12] Garagnani, M. (1999) "Belief Systems for Persuasive Discourse Planning", unpublished Ph.D. thesis, Department of Computer Science, Durham University, UK.

[13] Grosz, B.J., Sidner, C.L. (1986) "Attention, Intentions and the Structure of Discourse", *Journal of Computational Linguistics*, 12(3):175–204.

[14] Grosz, B.J., Sidner, C.L. (1990) "Plans for Discourse", in Cohen P., Morgan J., Pollack M.E., eds. *Intentions in Communication*, MIT Press, Cambridge, MA, pp. 417–444.

[15] Hintikka, J. (1962) *Knowledge and Belief*, Cornell University Press, Ithaca, New York.

[16] Hovy, E.H. (1993) "Automated discourse generation using discourse structure relations", *Artificial Intelligence*, 63:341–385.

[17] Koehler, J., Nebel, B., Hoffmann, J., Dimopoulos, Y. (1997) "Extending Planning Graphs to an ADL Subset", *Proceedings of the 4th European Conference on Planning (ECP '97).*

[18] Konolige, K. (1986) *A Deduction Modelof Belief*, Pitman Publishing, London.

[19] Maybury, M.T. (1992) "Communicative acts for explanation generation", *International Journal of Man-Machine Studies*, 37(2):135–172.

[20] Moore, J.D., Paris, C.L. (1993) "Planning Text for Advisory Dialogues: Capturing Intentional and Rhetorical Information", *Computational Linguistics*, 19(4):651–695.

[21] Morishima, Y. (1999) "Effects of Discourse Context on Inference Computation during Comprehension", *Lecture notes in computer science*, 1688.

[22] Pollock, J.L. (1994) "Justification and defeat", *Artificial Intelligence*, 67:377–407.

[23] Scruton, R. (1995) "A short history of modern philosophy", Routlege, London.

[24] Spencer-Smith, R. (1991) "Modal Logic", *Artificial Intelligence Review*, 5:5–34.

[25] Toulmin, S.E. (1958) *The Uses of Argument*, Cambridge University Press, Cambridge, England.

[26] Walker, M.A., Rambow, O. (1994) "The role of Cognitive Modelling in Achieving Communicative Intentions", in *Proceedings of the 7th International Workshop on Natural Language Generation*, Kennebunkport, Maine.

[27] 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.