A Strategy for Generating Evaluative Arguments

Giuseppe Carenini
Intelligent Systems Program
University of Pittsburgh,
Pittsburgh, PA 15260, USA
carenini@cs.pitt.edu

Johanna D. Moore
The Human Communication Research Centre,
University of Edinburgh,
2 Buccleuch Place, Edinburgh EH8 9LW, UK.
jmoore@cogsci.ed.ac.uk

Abstract

We propose an argumentation strategy for generating evaluative arguments that can be applied in systems serving as personal assistants or advisors. By following guidelines from argumentation theory and by employing a quantitative model of the user's preferences, the strategy generates arguments that are tailored to the user, properly arranged and concise. Our proposal extends the scope of previous approaches both in terms of types of arguments generated, and in terms of compliance with principles from argumentation theory.

Introduction

Arguing involves an intentional communicative act that attempts to create, change or reinforce the beliefs and attitudes of another person. *Factual* and *causal* arguments attempt to affect beliefs (i.e. assessments that something is or is not the case), whereas *evaluative* arguments attempt to affect attitudes (i.e., evaluative tendencies typically phrased in terms of like and dislike or favor and disfavor).

With the ever growing use of the Web, an increasing number of systems that serve as personal assistants, advisors, or sales assistants are becoming available online[1]. These systems frequently need to generate evaluative arguments for domain entities. For instance, a real-estate assistant may need to compare two houses, arguing that one would be a better choice than the other for its user.

Argumentation theory (Mayberry and Golden 1996; Miller and Levine 1996; Corbett and Connors 1999) indicates that effective arguments should be constructed following three general principles. First, arguments should be constructed considering the dispositions of the audience towards the information presented. Second, sub-arguments supporting or opposing the main argument claim should be carefully arranged by considering their strength of support or opposition. Third, effective arguments should be concise, presenting only pertinent and cogent information.

In this paper, we propose an argumentation strategy for generating evaluative arguments that can be applied in systems serving as personal assistants or advisors. By following principles and guidelines from argumentation theory and by employing a quantitative model of the user's preference, our strategy generates evaluative arguments that are tailored to the user, properly arranged and concise.

Although a preliminary version of our argumentative strategy was cursorily described in a previous short paper (Carenini and Moore 1999), this paper includes several additional contributions. First, we discuss how the strategy is grounded in the argumentation literature. Then, we provide details on the measures of argument strength and importance used in selecting and ordering argument support. Next, we generalize the argumentative strategy and correct some errors in its preliminary version. Finally, we discuss how our strategy extends the scope of previous approaches to generating evaluative arguments in terms of coverage (i.e., types of arguments), and in terms of compliance with principles from argumentation theory. Because of space limitations, we only discuss previous work on generating evaluative arguments, rather than previous work on generating arguments in general.

---

[1] See for instance www.activebuyersguide.com

1 Guidelines from Argumentation Theory

An argumentation strategy specifies what content should be included in the argument and how it should be arranged. This comprises several decisions: what represents supporting (or opposing) evidence for the main claim, where to position the main claim of the argument, what supporting (or opposing) evidence to include and how to order it, and how to order supporting and opposing evidence with respect to each other.

Argumentation theory has developed guidelines specifying how these decisions can be effectively made (see (Mayberry and Golden 1996; Miller and Levine 1996; Corbett and Connors 1999; McGuire 1968) for details; see also (Marcu 1996) for an alternative discussion of some of the same guidelines).

(a) What represents supporting (or opposing) evidence for a claim – Guidelines for this decision vary depending on the argument type. Limiting our analysis to evaluative arguments, argumentation theory indicates that supporting and opposing evidence should be identified according to a model of the reader's values and preferences. For instance, the risk involved in a game can be used as evidence for why your reader should like the game, only if the reader likes risky situations.

(b) Positioning the main claim - Claims are often presented up front, usually for the sake of clarity. Placing the claim early helps readers follow the line of reasoning. However, delaying the claim until the end of the argument can be effective, particularly when readers are likely to find the claim objectionable or emotionally shattering.

(c) Selecting supporting (and opposing) evidence - Often an argument cannot mention all the available evidence, usually for the sake of brevity. Only strong evidence should be presented in detail, whereas weak evidence should be either briefly mentioned or omitted entirely.

(d) Arranging/Ordering supporting evidence – Typically the strongest support should be presented first, in order to get at least provisional agreement from the reader early on. If at all possible, at least one very effective piece of supporting evidence should be saved for the end of the argument, in order to leave the reader with a final impression of the argument's strength. This guideline proposed in (Mayberry and Golden 1996) is a compromise between the climax and the anti-climax approaches discussed in (McGuire 1968).

(e) Addressing and ordering the counterarguments (opposing evidence) – There are three options for this decision: not to mention any counterarguments, to acknowledge them without directly refuting them, to acknowledge them and directly refuting them. Weak counterarguments may be omitted. Stronger counterarguments should be briefly acknowledged, because that shows the reader that you are aware of the issue's complexity; and it also contributes to the impression that you are reasonable and broad-minded. You may need to refute a counterargument once you have acknowledged it, if the reader agrees with a position substantially different from yours. Counterarguments should be ordered to minimize their effectiveness: strong ones should be placed in the middle, weak ones upfront and at the end.

(f) Ordering supporting and opposing evidence – A preferred ordering between supporting and opposing evidence appears to depend on whether the reader is aware of the opposing evidence. If so, the preferred ordering is opposing before supporting, and the reverse otherwise.

Although these guidelines provide useful information on the types of content to include in an evaluative argument and how to arrange it, the design of a computational argumentative strategy based on these guidelines requires that the concepts mentioned in the guidelines be formalized in a coherent computational framework. This includes: explicitly representing the reader's values and preferences (used in guideline a); operationally defining the term "objectionable claim" (used in guideline b) through a measure of the discrepancy between the reader's initial position and the argument's main claim[2]; providing a measure of evidence strength (needed in guidelines c, d, and e); and

---

[2] An operational definition for "emotionally shattering" is outside the scope of this paper.

48

OBJECTIVES		COMPONENT VALUE FUNCTIONS
	ATTRIBUTES	

Value tree (House Value):
- Location 0.7
  - Neighborhood 0.4
  - Distance-from-park 0.6
- Size 0.3
  - #-of-room 0.8
  - Storage-space 0.2

Neighborhood component value function:

x1	v1(x1)
x1=n1	0
x1=n2	0.3
x1=n3	1

Distance-from-park component value function:

x2	v2(x2)
0=<x2<=5	1-(1/5 * x2)
x2>5	0

#-of-room · · · · · ·
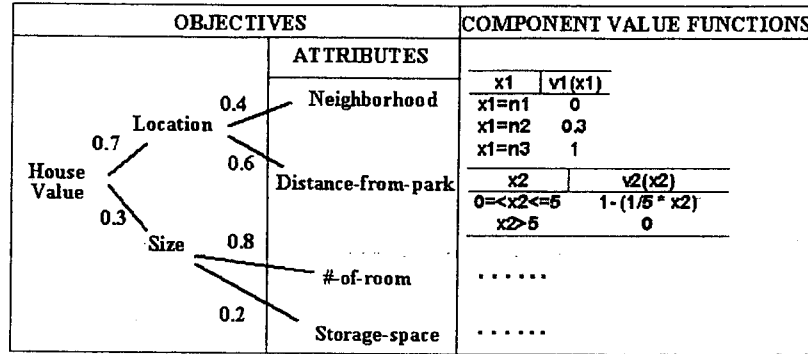
Storage-space · · · · · ·

Figure 1 Sample additive multiattribute value function (AMVF)

representing whether the reader is or is not aware of certain facts (needed in guideline f).

2 From Guidelines to the Argumentation Strategy

We assume that the reader's values and preferences are represented as an additive multiattribute value function (AMVF), a conceptualization based on multiattribute utility theory (MAUT)(Clemen 1996). Besides being widely used in decision theory (where they were originally developed), conceptualizations based on MAUT have recently become a common choice in the field of user modelling (Jameson, Schafer et al. 1995). Similar models are also used in Psychology, in the study of consumer behaviour (Solomon 1998).

2.1 Background on AMVF

An AMVF is a model of a person's values and preferences with respect to entities in a certain class. It comprises a *value tree* and a set of *component value functions*, one for each attribute of the entity. A value tree is a decomposition of the value of an entity into a hierarchy of aspects of the entity[3], in which the leaves correspond to the entity primitive attributes (see Figure 1 for a simple value tree in the real estate domain). The arcs of the tree are weighted to represent the importance of the value of an objective in contributing to the value

---

[3] In decision theory these aspects are called *objectives*. For consistency with previous work, we will follow this terminology in the remainder of the paper.

of its parent in the tree (e.g., in Figure 1 *location* is more than twice as important as *size* in determining the value of a house). Note that the sum of the weights at each level is equal to 1. A component value function for an attribute expresses the preferability of each attribute value as a number in the [0,1] interval. For instance, in Figure 1, neighborhood n2 has preferability 0.3, and a distance-from-park of 1 mile has preferability $(1 - (1/5* 1))=0.8$.

Formally, an AMVF predicts the value $v(e)$ of an entity e as follows:

$v(e) = v(x_1, ..., x_n) = \Sigma w_i v_i(x_i)$, where

- $(x_1, ..., x_n)$ is the vector of attribute values for an entity e
- $\forall$ attribute i, v_i is the component value function, which maps the least preferable x_i to 0, the most preferable to 1, and the other x_i to values in [0,1]
- w_i is the weight for attribute i, with $0 \le w_i \le 1$ and $\Sigma w_i = 1$
- w_i is equal to the product of all the weights from the *root* of the value tree to the attribute i

A function $v_o(e)$ can also be defined for each objective. When applied to an entity, this function returns the value of the entity with respect to that objective. For instance, assuming the value tree shown in Figure 1, we have:

$v_{Location}(e) =$
$= (0.4 * v_{Neighborhood}(e)) + (0.6 * v_{Dist-from-park}(e))$

Thus, given someone's AMVF, it is possible to compute how valuable an entity is to that

individual. Furthermore, it is possible to compute how valuable any objective (i.e., any aspect of that entity) is for that person. All of these values are expressed as a number in the interval [0,1].

2.2 Computational Definition of Concepts Mentioned in Guidelines

Presenting an evaluative argument is an attempt to persuade the reader that a value judgment applies to a subject. The value judgement, also called the *argumentative intent*, can either be positive (in favour of the subject), or negative (against the subject)[4]. The *subject* can be a single entity (e.g., "This book is very good"), the difference between two entities (e.g., "City-a is somewhat better than city-b"), or any other form of comparison among entities in a set (e.g., "This city is the best in North America").

Guideline (a) - Given the reader's AMVF, it is straightforward to establish what represent supporting or opposing evidence for an argument with a given argumentative intent and a given subject. In fact, if the argumentative intent is positive, objectives for which the subject has positive value can be used as supporting evidence, whereas objectives for which the subject has a negative value can be used as opposing evidence (the opposite holds when the argumentative intent is negative). The value of different subjects is measured as follows. If the subject is a single entity e, the value of the subject for an objective o is $v_o(e)$, and it is positive when it is greater than 0.5, the midpoint of [0,1] (negative otherwise). In contrast, if the subject is a comparison between two entities (e.g., $v(e_1) > v(e_2)$), the value of the subject for an objective o is $[v_o(e_1) - v_o(e_2)]$, and it is positive when it is greater than 0 (negative otherwise).

Guidelines (b) - Since argumentative intent is a value judgment, we can reasonably assume that instead of being simply positive or negative, it may be specified more precisely as a number in the interval [0,1] (or as a specification that can be normalized in this interval). Then, the term

---

[4] Arguments can also be neutral. However, in this paper we do not discuss arguments with a neutral argumentative intent.

"objectionable claim" can be operationally defined. If we introduce a *measure-of-discrepancy*(MD) as the absolute value of the difference between the argumentative intent and the reader's expected value of the subject before the argument is presented (based on her AMVF), a claim becomes more and more "objectionable" for a reader as MD moves from 0 to 1.

Guidelines (c) (d) (e) - The strength of the evidence in support of (or opposition to) the main argument claim is critical in selecting and organizing the argument content. To define a measure of the strength of support (or opposition), we adopt and extend previous work on explaining decision theoretic advice based on an AMVF. (Klein 1994) presents explanation strategies (not based on argumentation theory) to justify the preference of one alternative from a pair. In these strategies, the *compellingness* of an objective measures the objective's strength in determining the overall value difference between the two alternatives, other things being equal. And an objective is *notably-compelling?* (i.e., worth mentioning) if it is an outlier in a population of objectives with respect to compellingness. The formal definitions are:

$compellingness(o,a_1, a_2, refo) =$
$= w(o,refo)[v_o(a_1) - v_o(a_2)]$, where
- o is an objective, a_1 and a_2 are alternatives, *refo* is an ancestor of o in the value tree
- $w(o,refo)$ is the product of the weights of all the links from o to *refo*
- v_o is the component value function for leaf objectives (i.e., attributes), and it is the recursive evaluation over *children(o)* for nonleaf objectives

$notably\text{-}compelling?(o,opop,a_1, a_2, refo) \equiv$
$|compellingness(o,a_1, a_2, refo)| > \mu_x + k\sigma_x$, where
- o, a_1, a_2 and *refo* are defined as in the previous Def; *opop* is an objective population (e.g., *siblings(o)*), and $|opop| > 2$
- $p \in opop$; $x \in X = |compellingness(p,a_1, a_2, refo)|$
- μ_x is the mean of X, σ_x is the standard deviation and k is a user-defined constant

We have defined similar measures for arguing the value of a single entity and we named them *s-compellingness* and *s-notably-compelling?*.

An objective can be s-compelling either because of its strength or because of its weakness in contributing to the value of an alternative. So, if m_1 measures how much the value of an objective contributes to the overall value difference of an alternative from the worst possible case[5] and m_2 measures how much the value of an objective contributes to the overall value difference of the alternative from the best possible case, we define s-compellingness as the greatest of the two quantities m_1 and m_2. Following the terminology introduced in the two previous Equations we have:

$s\text{-}compellingness(o,a,refo) =$

$= w(o,refo)[\max[v_o(a) - 0]; [1 - v_o(a)]]$

We give to *s-notably-compelling?* a definition analogous to the one for *notably-compelling?*

$s\text{-}notably\text{-}compelling?(o,opop,a,refo) \equiv$
$| s\text{-}compellingness(o,a,refo) | > \mu_x + k\sigma_x,$

Guideline (f) - An AMVF does not represent whether the reader is or is not aware of certain facts. We assume this information is represented separately.

2.3 The Argumentation Strategy

We have applied the formal definitions described in the previous section to develop the argumentative strategy shown in Figure 2. The strategy is designed for generating honest and balanced arguments, which present an evaluation of the subject equivalent to the one you would expect the reader to hold according to her model of preferences (i.e., the argumentative intent is equal to the expected value, so MD=0)[6]. We now examine the strategy in detail, after introducing necessary terminology. The *subject*

is either a single entity or a pair of entities in the domain of interest. *Root* can be any objective in the value tree for the evaluation (e.g., the overall value of a house, its location, its amenities). *ArgInt* is the argumentative intent of the argument, a number in [0,1]. The constant k, part of the definitions of *notably-compelling?* and *s-notably-compelling?*, determines the degree of conciseness of the argument. The Express-Value function, used at the end of the strategy, indicates that the objective applied to the subject must be realized in natural language with a certain argumentative intent.

In the first part of the strategy, depending on the nature of the subject, an appropriate measure of evidence strength is assigned, along with the appropriate predicate that determines whether a piece of evidence is worth mentioning. After that, only evidence that is worth mentioning is assigned as supporting or opposing evidence by comparing its value to the argument intent. In the second part, ordering constraints from argumentation theory are applied[7]. Notice that we assume a predicate *Aware* that is true when the user is aware of a certain fact, false otherwise. Finally, in the third part of the strategy, the argument claim is expressed in natural language. The opposing evidence (i.e., *ContrastingSubObjectives*), that must be considered, but not in detail, is also expressed in natural language. In contrast, supporting evidence is presented in detail, by recursively calling the strategy on each supporting piece of evidence.

2.4 Implementation and Application

The argumentation strategy has been implemented as a set of plan operators. Using these operators the Longbow discourse planner (Young and Moore 1994) selects and arranges the content of the argument. We have applied our strategy in a system that serves as a real-estate personal assistant (Carenini 2000a). The system presents information about houses available on the market in graphical format. The user explores this information by means of interactive techniques, and can request a natural

---

[5] a_{worst} is an alternative such that $\forall o\ v_o(a_{worst})=0$, whereas a_{best} is an alternative such that $\forall o\ v_o(a_{best})=1$

[6] An alternative strategy, for generating arguments whose argumentative intent was greater (or lower) than the expected value, could also be defined in our framework. However, this strategy should boost the evaluation of supporting evidence and include only weak counterarguments, or hide them overall (the opposite if the target value was lower than the expected value)

---

[7] The steps in the strategy are marked with the guideline they are based on.

51

Argue(subject, Root, ArgInt, k)
;; **assignments and content selection**
    If subject = single-entity = e **then** $SV_{o_i} = v_{o_i}(e)$

                                Measure-of-strength = s-compellingness
                                Worth-mention? = s-notably-compelling?

    Else If subject = e_1, e_2 **then** $SV_{o_i} = |v_{o_i}(e_1) - v_{o_i}(e_2)|$

                         Measure-of-strength = compellingness
                         Worth-mention? = notably-compelling?

Eliminate all objectives $o_i | \neg$ Worth-mention? (o_i, siblings(o_i), subject, Root)   ;guideline(c)
AllEvidence ← *children(Root)*
AllInFavor ← all $o \,|\, o \in$ *AllEvidence* ∧ *(SV$_o$* ≈ *ArgInt)*   ;guideline(a)
SecondBestObjInFavor ← second most compelling objective $o \,|\, o \in$ *AllInFavor*
RemainingObjectivesInFavor ← *AllInFavor - SecondBestObjInFavor*
ContrastingObjectives ← *AllEvidence - AllInFavor*   ;guideline(a)
;; *ordering the selected content*

AddOrdering(*Root* ≺ *AllEvidence*) ;; we assume MD=0, so claim is not objectionable ;guideline(b)
If Aware(User, *ContrastingObjectives*) **then**   ;guideline(f)
        AddOrdering(*ContrastingObjectives* ≺ *AllInFavor*)
Else AddOrdering(*ContrastingObjectives* ≻ *AllInFavor*);
AddOrdering(*RemainingObjectivesInFavor* ≺ *SecondBestObjInFavor*)   ;guideline(d)
Sort(*RemainingObjectivesInFavor;* decreasing order according to Measure-of-strength) ;guideline(d)
Sort(*ContrastingObjectives;* strong ones in the middle, weak ones upfront and at the end) ;guideline(e)

;; **steps for expressing or further argue the content**
Express-Value(subject, Root, ArgInt)
For all $o \in$ *AllInFavor* , **If** ¬leaf(o) **then** Argue(subject, o, SV_o, k)
                        Else Express-Value(subject, o, SV_o)
For all $o \in$ *ContrastingObjectives,* Express-Value(subject, o, SV_o)   ;guideline(e)

Legend: *(a ≺ b)* ↔ a preceeds b

        $(v_1 ≈ v_2)$ ↔ v_1 and v_2 are both positive or negative values

        (see Section 0 for what this means for different subjects)

Figure 2 The Argumentation strategy

language evaluation of any house just by dragging the graphical representation of the house to a query button. The evaluative arguments generated by the system are concise, properly arranged and tailored to the user's preferences[8]. For sample arguments generated by our strategy see (Carenini 2000b) in this proceedings.

3 Previous Work

Although considerable research has been devoted to study the generation of evaluative arguments, all approaches proposed so far are limited in the type of evaluative arguments generated, and in the extent to which they comply with guidelines from argumentation literature.

(Elhadad 1992) investigated a general computational framework that covers all aspects of generating evaluative arguments of single entities, from content selection and structuring to fine-grained realization decisions. However, his work concentrates on the linguistic aspects. His approach to content selection and structuring does not provide a measure of evidence strength, which is necessary to implement several of the guidelines from argumentation literature we have examined.

Other studies have focused more on the process of content selection and structuring. However, with respect to our proposal, they still suffer from some limitations. (Morik 1989) describes a system that uses a measure of evidence strength to tailor evaluations of hotel rooms to its users. However, her system adopts a qualitative measure of evidence strength (an ordinal scale that appears to range from *very-important* to *not-important*). This limits the ability of the system to select and arrange argument evidence, because qualitative measures only support approximate comparisons and are notoriously difficult to combine (e.g., how many "somewhat-important" pieces of evidence are equivalent to an "important" piece of evidence?).

(Elzer, Chu-Carroll et al. 1994; Chu-Carroll and Carberry 1998) studied the generation of evaluative arguments in the context of collaborative planning dialogues. Although they also adopt a qualitative measure of evidence strength, when an evaluation is needed this measure is mapped into numerical values so that preferences can be compared and combined more effectively. However, with respect to our approach, this work makes two strong simplifying assumptions. It only considers the decomposition of the preference for an entity into preferences for its primitive attributes (not considering that complex preferences frequently have a hierarchical structure). Additionally, it assumes that the same dialogue turn cannot provide both supporting and opposing evidence.

(Kolln 1995) proposes a framework for generating evaluative arguments which is based on a quantitative measure of evidence strength. Evidence strength is computed on a fuzzy hierarchical representation of user preferences. Although this fuzzy representation may represent a viable alternative to the AMVF we have discussed in this paper, Kolln's proposal is rather sketchy in describing how his measure of strength can be used to select and arrange the argument content.

Finally, (Klein 1994) is the previous work most relevant to our proposal. Klein developed a framework for generating explanations to justify the preference of an entity out of a pair. These strategies were not based on argumentation theory. As described in Section 2.2, from this work, we have adapted a measure of evidence strength (i.e., *compellingness*), and a measure that defines when a piece of evidence is worth mentioning (i.e., *notably-compelling?*).

Conclusions and Future Work

In this paper, we propose an argumentation strategy that extends previous research on generating evaluative arguments in two ways. Our strategy covers the generation of evaluations of a single entity, as well as comparisons between two entities. Furthermore, our strategy generates arguments, which are concise, properly arranged and tailored to a hierarchical model of user's preferences, by

---

[8] The generation of fluent English also required the development of microplanning and realization components. For lack of space, we do not discuss them in this paper.

following a comprehensive set of guidelines from argumentation theory.

Several issues require further investigation. First, we plan to generalize our approach to more complex models of user preferences. Second, although our strategy is based on insights from argumentation theory, the ultimate arbiter for effectiveness is empirical evaluation. Therefore, we have developed an evaluation environment to verify whether arguments generated by our strategy actually affect user attitudes in the intended direction (Carenini 2000b). A third area for future work is the exploration of techniques to improve the coherence of arguments generated by our strategy. In the short term, we intend to integrate the ordering heuristics suggested in (Reed and Long 1997). In the long term, by modelling user attention and retention, we intend to enable our strategy to assess in a principled way when repeating the same information can strengthen argument force. Finally, we plan to extend our strategy to evaluative arguments for comparisons between mixtures of entities and set of entities.

Acknowledgements

References

Carenini, G. (2000a). *Evaluating Multimedia Interactive Arguments in the Context of Data Exploration Tasks*. PhD Thesis, Intelligent System Program, University of Pittsburgh.

Carenini, G. (2000b). *A Framework to Evaluate Evaluative Arguments*. Int. Conference on Natural Language Generation, Mitzpe Ramon, Israel.

Carenini, G. and J. Moore (1999). *Tailoring Evaluative Arguments to User's Preferences*. User Modelling, Banff, Canada : 299-301.

Chu-Carroll, J. and S. Carberry (1998). *Collaborative Response Generation in Planning Dialogues*. Computational Linguistics 24(2): 355-400.

Clemen, R. T. (1996). *Making Hard Decisions: an introduction to decision analysis*. Duxbury Press

Corbett, E. P. J. and R. J. Connors (1999). *Classical Rhetoric for the Modern Student*, Oxford University Press.

Elhadad, M. (1992). *Using Argumentation to Control Lexical Choice: A Functional Unification Implementation*. PhD Thesis, CS. Columbia. NY.

Elzer, S., J. Chu-Carroll, et al. (1994). *Recognizing and Utilizing User Preferences in Collaborative Consultation Dialogues*. Proceedings of Fourth Int. Conf. of User Modeling. Hyannis, MA: 19-24.

Jameson, A., R. Schafer, et al. (1995). *Adaptive provision of Evaluation-Oriented Information: Tasks and techniques*. Proc. of 14th IJCAI. Montreal, Canada.

Klein, D. (1994). *Decision Analytic Intelligent Systems: Automated Explanation and Knowledge Acquisition*, Lawrence Erlbaum Associates.

Kolln, M. E. (1995). *Employing User Attitudes in Text Planning*. 5th European Workshop on Natural Language Generation, Leiden, The Netherlands.

Marcu, D. (1996). *The Conceptual and Linguistic Facets of Persuasive Arguments*. ECAI workshop - Gaps and Bridges: New Directions in Planning and Natural Language Generation.

Mayberry, K. J. and R. E. Golden (1996). *For Argument's Sake: A Guide to Writing Effective Arguments*, Harper Collins, College Publisher.

McGuire, W. J. (1968). *The Nature of Attitudes and Attitudes Change*. The Handbook of Social Psychology. G. Lindzey and E. Aronson, Addison-Wesley. 3: 136-314.

Miller, M. D. and T. R. Levine (1996). *Persuasion. An Integrated Approach to Communication Theory and Research*. M. B. Salwen and D. W. Stack. Mahwah, New Jersey: 261-276.

Morik, K. (1989). *User Models and Conversational Settings: Modeling the User's Wants*. User Models in Dialog Systems. A. Kobsa and W. Wahlster, Springer-Verlag: 364-385.

Reed, C. and D. Long (1997). *Content Ordering in the Generation of Persuasive Discourse*. Proc. of the 15th IJCAI, Nagoya, Japan.

Solomon, M. R. (1998). *Consumer Behavior: Buying, Having, and Being*. Prentice Hall.

Young, M. R. and J. D. Moore (1994). *Does Discourse Planning Require a Special-Purpose Planner?* Proc. of the AAAI-94 Workshop on planning for Interagent Communication. Seattle, WA.