Persuasive Argumentation in Human Computer Dialogue

Pierre Andrews* and **Marco De Boni**** and **Suresh Manandhar***

* Departement of Computer Science
University of York
York YO10 5DD
UK
{pandrews,suresh}@cs.york.ac.uk

** Unilever Corporate Research
Bedford MK44 1LQ
UK
Marco.De-Boni@unilever.com

Abstract

In the field of natural language dialogue, a new trend is exploring persuasive argumentation theories. Applying these theories to human-computer dialogue management could lead to a more comfortable experience for the user and give way to new applications.

In this paper, we study the different aspects of persuasive communication needed for health-care advising and how to implement them to produce efficient, computer directed persuasion.

Our opinion is that a persuasive dialogue will have to combine the current logical approach to persuasion with novel emotional cues to render the dialogue more comfortable to the user.

Keywords: *natural argumentation, rhetorics, dialogue, persuasion, health-care counselling, natural language processing*

Introduction

Dialogue is not new to computational linguistics but, up to now, has only been applied to very restricted domains. The computer is often either too task oriented or does not guarantee topic consistency during dialogue: few solutions have been found to this problem yet. This research intends to investigate if, by using the long studied theories of rhetoric in philosophy and linguistics, it will be possible to build a more human-like *automated* dialogue systems. Up to now, persuasive aspects of dialogue have been taken into account only in very specialised areas such as law (see Bench-Capon 2003a; Bench-Capon 2003b for example). The interest of this project is therefore to study how rhetoric and persuasion theories could be applied in some less restricted types of dialogue. We follow the current trend of computer assisted persuasion (see Fogg 2003; Guerini, Stock, & Zancanaro 2003 for instance) and will apply these theories in the field of health-care advising.

Developing applications in automatic health-care counselling allows us to rely on existing guidelines and resources (see Green 2003; Grasso 2003 for example) to ground persuasive dialogue development. Such guidelines have already been applied to automated counselling systems (Friedman *et al.* 1997) which will benefit from the new developments in human-computer natural language dialogue.

Human Computer Dialogue

From Eliza the psychologist (Weizenbaum 1966) to current automated phone reservation services (Seneff & Polifroni 2000; Mcglashan *et al.* 1992), human computer dialogue has been a research field in artificial intelligence for a long time.

Two approaches have mainly been used up to now: *chatbot systems*, which discuss without following a particular agenda with the user, and *task oriented* approaches, which follow a plan to achieve a specific task. Indeed, *task oriented* dialogues are designed to accomplish a task in the most efficient way, which is sometimes confusing for the users as they do no feel comfortable with the restrictions imposed on the dialogue (Shechtman & Horowitz 2003; Farzanfar *et al.* 2005). In the same way, *chatbot* dialogue can appear "unreal" because of the lack in discourse consistency of the bot, which "jumps" from one subject to the other or utters meaningless sentences. However, users seem to feel comfortable with chatbots and like to *play* with the systems for some time (Saarinen 2001).

Rhetoric

On the other hand, Aristotle (III century B.C.) started examining the art of producing persuasive discourse in his Rhetoric – see Barnes (1998) for a modern translation. Philosophers have, since then, been formulating a large number of different theories on rhetoric and argumentation, such as the popular New Rhetoric by Perelman & Olbrechts-Tyteca (1969).

Computational linguistics have rarely considered these theories. The few applications try to formalise human discourse as a logic reasoning task (Reed & Long 1998) and often forget the *emotional* part of rhetoric that has been emphasized by Aristotle's separation of persuasion in the three axes: *ethos* (ethical), *pathos* (emotional) and *logos* (logical).

However, natural argumentation is finding its way in the artificial dialogue field (Norman & Reed 2003), especially in legal argumentation (Bench-Capon 2003a; 2003b). By

doing so, it raises a number of new challenges, which, we believe, fit in the gap left between the chatbot approach and the task oriented technique.

We are currently conducting research on automatic promotion of physical activity through computer dialogue. This task implies the use of persuasive techniques to change the user behaviour. In the following sections we report the issues that have been identified for this task. A dedicated section discusses the different solutions that we believe can be applied to address these issues.

Emotions and Social Cues

Human computer dialogue systems – developed after extensive research – are still considered austere by the users that clearly act differently when engaged in computer dialogue than during human to human dialogue (Shechtman & Horowitz 2003). This is because they do not take into account aspects that appear to be perfectly natural to most people: *emotions* and *social cues*.

Users are still fully aware that they are speaking with a computer. They are therefore not encouraged to get involved in a real dialogue and either play with the system or rapidly accomplish their task. This limits the range of possible dialogue applications and does not suit the new needs formulated for such systems. For instance, in the field of health communication, Bickmore & Giorgino (2004) identify the following open issues for dialogue systems:

- continuity between multiple dialogues,
- language change over time (in the context of multiple dialogues),
- "patient-centered" communication,
- use of social cues and empathy to improve the relationship between the system and the user.

The task oriented approach to dialogue management is used in the majority of industrial-strength applications and comprises three main models:

- Finite State Machines (FSM) model the discourse structure with a finite number of states. These states usually encode the different stages of the task completion and are linked by transitions triggered by user inputs (see Pietquin & Dutoit 2003 for reference).
- Form-filling systems use a less restricted approach to obtain specific information from the user. The system models the different phases of the dialogue as form slots that the dialogue manager chooses to be filled in the current dialogue context. This model is used for example in reservation systems to help the user create a database query, retrieve information and submit a reservation (Seneff & Polifroni 2000; Mcglashan *et al.* 1992).
- Plan based approaches use planning techniques to construct the dialogue. The dialogue is modelled by a set of rules specifying preconditions and post conditions and dialogue actions that are chosen by a planner in function of the current user model, discourse history etc. This planning technique is used in complex tasks where the dialogue path cannot be anticipated in advance. Planning is used, for example, in tutorial dialogue (see Zinn, Moore, & Core 2002 for example) where the system helps students to acquire new concepts.

These systems work because the information they use is highly structured and domain dependent; for instance, in Ravenscroft & Pilkington (2000) the dialogue management is based on logical rules describing physical interactions. However, users are familiar with human interaction and expect more from the system: they tend to give a personality to the computer and think it can understand their own situation, needs and their current *emotional state*. Users want and try to be in control (De Angeli, Johnson, & Coventry 2001), even if that is not vital to the completion of the task. This will often be the case in health-care counselling, where the user needs the advisor to show empathy and be more "patient centered" (Farzanfar *et al.* 2005).

In addition, task oriented systems rarely perceive the social cues used by the users. In fact, the most basic ones will not make any difference between a grammatically correct, politely formulated sentence and a list of keywords. Therefore, they will not apprehend the emotional variations during the dialogue and the user will eventually feel some discomfort using such system. In an automatic advice system for physical activity, it is important to keep users motivated to interact with the system as it is rarely a vital issue for them to use it. Therefore, it is important to avoid this type of discomfort (see, for reference, the evaluation in Farzanfar *et al.* 2005).

Computer Assisted Dialogue in Health-care

Automatic dialogue with the patient is a topic embraced by the different actors of health-care (see Friedman *et al.* 1997; Slack 2000 for references) from the beginnings of human-computer dialogue research; for example, Eliza (Weizenbaum 1966) was an automatic psychiatrist. Automatic therapy and counselling is indeed an interesting prospect in health-care as it enables to lower medical costs and to reduce therapist time and the constraints on the patient. Therefore, automatic dialogue systems have been developed and evaluated extensively in this field.

For instance, the Telephone Linked Computer (TLC) systems (Smith *et al.* 1988) use an automatic dialogue interface, accessed with a telephone to provide various health-care services from nutrition advice to cognitive therapy. These systems provide pre-recorded answers that are uttered in response to the choice users enter using the phone keypad.

The TLC approach uses branching logic to control the interaction. The dialogue moves are encoded in the program, which chooses the right answer or question to give in function of the choices made by the user. This can be interpreted as a state-oriented approach, where each pre-recorded answer corresponds to a state and the multiple choices proposed to the user specify the transitions to the next states. In this, it is comparable to the standard approach to task-oriented dialogue that we have discussed earlier.

Recent TLC based systems include speech recognition, improving its ergonomics. However, their basic principle is still the same, the user is faced with a limited number of

choices that can either be typed on the keypad or spoken. This leaves the user with limited initiative in the dialogue. Although we have previously stated that this type of system is sometime frustrating, it seems that the benefits of using the telephone, its apparent anonymity and the reduction of the therapist time gave good results where the TLC system was used. For instance, in cognitive therapy for depression (Wright *et al.* 2005), computer-assisted patients reported similar improvement in their condition to those following a standard cognitive therapy.

However, in our research we focus on the complementary area of physical activity promotion and counseling. In this field, various evaluations gave fair results (Farzanfar *et al.* 2005; Glanz *et al.* 2003; Pinto *et al.* 2002) as most of the patients did not use the system on a long term or were intimidated by its design. Nevertheless, the results found by Farzanfar *et al.* (2005) are promising as some patients reported behaviour changes and perceived benefits, proving that such automated approach could be effective. The authors also clearly identify the remark on the system by to the patients who reported lower benefits:

- the system currently lacks empathy,
- the patients do not feel in control of the dialogue,
- monitoring might generate avoidance.

In our research, we intend to take a somehow different approach, trying to allow more freedom to the user and display social cues to show empathy. We intend to do this by using natural language techniques, allowing the user to format utterances more freely. In the next sections, we will review the possible approaches and the problem they raise.

Argumentation framework

The need for a complete planning system in persuasive dialogue has been pointed out recently by Gilbert *et al.* (2003). This study presents the complex steps to be followed to generate persuasive dialogue moves. The dialogue system would need to identify arguments in the user utterances, evaluate their correctness and build its utterances based on this processing.

The framework proposed by Gilbert *et al.* (2003) is still highly theoretical and relies on ongoing work in natural argumentation and computational linguistics. However, a number of working frameworks for planning and generating arguments have been developed in the natural language field (Reed & Grasso 2001; Grasso 2002) using rhetoric theories like the "New Rhetorics" (Perelman & Olbrechts-Tyteca 1969) or the "Rhetorical Structure Theory" (Mann 1999).

One major problem raised by research on argumentation based natural language dialogues – for example in the framework described in Gilbert *et al.* (2003) – is that of understanding the user's counter arguments. Detecting the argument scheme used, deciding whether or not it is a fallacy and verifying the veracity of each of the premises put forward require a large knowledge base and is currently difficult.

Persuasion and chatbots

In trying to change the users' behaviour and attitude, techniques of persuasion like argumentation and rhetoric have to be used, but it is currently difficult to see how to integrate them in computer based dialogue.

Indeed, persuasive communication textbooks (Stiff & Mongeau 2002) show that it is often needed to win trust and credibility by using social talk in order to be effective in the argumentation.

In some ways, chatbots are designed to display empathy and entertain the user. They are not tailored to any particular task and can chat freely with the user. This *freedom of speech* makes them appear as understanding emotions and using some social cues.

In this perspective, developing natural argumentation over a chatbot system would be tempting. Argumentative dialogue needs to be more open than traditional task oriented dialogue. It needs to leave some freedom to the user and to be able to use some real social cues to effectively deliver the argument.

However, *freedom of speech* is also the weakness of chatbots. They don't focus on any particular subject and change topic each time they cannot understand the user. Therefore, the users cannot achieve a well specified task with them and can be frustrated or disappointed by this lack of *discourse consistency*. This is also a barrier to achieve continuity in the discourse, as chatbot systems have short memories[1] and cannot adapt the discourse in regard to the topics already discussed.

In fact, argumentative dialogue is not just *chitchat* with the user. Argumentation always tries to eventually achieve a goal. In a persuasive argumentation, the goal is obviously to persuade the users of the speaker's point of view, to change his beliefs – and eventually his behaviour.

Indeed, the desired system will need to foresee the dialogue moves to accomplish in order to reach a point in the dialogue where the user seems convinced. In the following section, we report results that show how chatbot systems are not quite suited for building a natural argumentation system with such characteristic.

Case Study

We investigated a state-of-the-art chatbot system to evaluate the possibility of a better control of the dialogue path in such system to improve the continuity of the dialogue. This section reports the results of this evaluation and our conclusions. The aim of our study was to find if it was possible to add goal management in the apparently unrestricted pattern matching system.

Pattern matching chatbots are designed to take an approach where there is no need to explicitly encode states and transitions to construct a dialogue. The system is based on a set of patterns that match the user input and are linked to a template to generate an answer. For example:

```
I AM * YEARS OLD
=> Only *?  You are quite mature.
```

[1] usually, one or two utterances.

will match an utterance like:

"I am twenty years old"

and reply:

"Only twenty? You are quite mature"

Knowledge of how discourse is structured and on the social rules of conversation are implicitly encoded in the 40'000 patterns of the system. We wanted to evaluate if these patterns contained implicit state and transitions that embed this knowledge. Indeed, the patterns are also often associated with the last utterance the bot would have said before. Therefore, connecting each pattern with the patterns that match the previous utterance specified creates a graph like structure equivalent to a state machine where each state represents the last user utterance and each transition the bot answer.

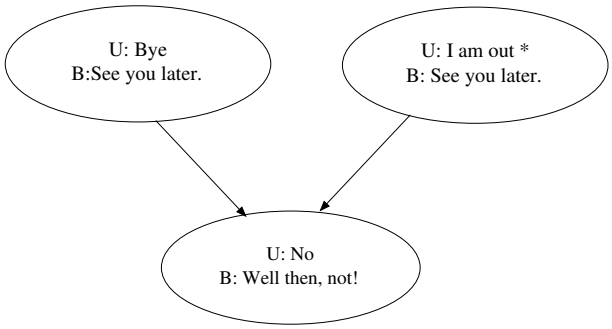

Figure 1: Sample of the chatbot internal state transitions. (U: user; B: bot)

Figure 1 shows a state transition graph that we found in the patterns and that could apply in a dialogue of the type:

– ...
– (User) I am out of here,
– (Bot) See you later.
– (User) No
– (Bot) Well then, not!
– ...

This is a small example of the whole set of patterns we decided to study, trying to find precious patterns sets implicitly encoding important conversation rules for given situations.

However, studying the possible state transitions between all the possible pattern/answer pairs in the system knowledge base gave no interesting results. The pattern matching system showed no internal structured state. State transitions were either:

- in small (2 or 3 states) clusters (see Figure 1 for example) that provided no information on any global strategy,
- in large, strongly interconnected clusters that displayed no logic in the dialogue "path" the bot could follow.

There is no isolated set of patterns representing the heterogeneous discourse knowledge implicitly encoded in the chatbot that gives it this apparent comfort in the conversation.

These observations confirmed the assumption that such chatbot systems could hardly be restricted to achieve explicit goals by constraining the dialogue to different isolated set of patterns in function of the current dialogue phase.

Discussion

Initiative in Persuasion

If we observe persuasive dialogues, it appears that the conversation is not symmetric (see for example Plato's "Apology of Socrates"): the main speaker tries to convince its interlocutor(s) by leading them to a persuasive goal. Therefore we believe that a framework for persuasive dialogue should give such initiative to the dialogue system. The system interlocutor(s) will then rarely make unrestricted moves in the process – as long as the argumentation is effective – because of the social rules of conversations. We hypothesise that this approach could lead to a simpler planning scheme.

Moreover, our position is that if the system's arguments are presented in simple utterances and formulated in order to give the initiative to the system, the user will respond with simple answers. The system will only have to detect the level of agreement or disagreement to the proposed argument and the emotional impact of the last argument.

This approach to interaction can be implemented for example using the dialogue game theory (Levin & Moore 1977). In this theory, dialogue is modelled as a game between the two interlocutors regulated by rules. At every stage of the game, an interlocutor can only choose between a limited number of dialogue moves. This approach eases the interpretation of the user input: user's utterance will match one of the dialogue moves specified by the dialogue game or, if it is not the case, then the system has either misunderstood or the user is not following the normal conversational rules. An example of this approach to argumentation is the work by Ravenscroft (2000) that uses the dialogue game theory to plan "debating style" moves for knowledge refinement in the context of computer assisted education.

Social Talk

It is also important that – in the framework chosen and in the resulting planning strategies – the system keeps the user comfortable. The user will need to feel familiar with the system to be receptive to its arguments. As it has been emphasised in section "Emotions and Social Cues", taking into account the emotional state of the user is important to formulate effective arguments. The choice of the argumentation schema has to be directed by the user sensibility, as a result the persuasive moves have to be strongly bound to the user emotional state and will not be the same in all the contexts.

In that perspective, even if chatbot systems do not appear to be the good starting point for dialogue planning – as their strategies only span on a few utterances – the framework developed could find useful inspiration from the dialogue moves used by chatbots.

A planned approach to social language modeling has been proposed by Cassell & Bickmore (2002) who use small-talk to build the user trust. The authors use an activation network to plan dialogue moves, choosing either between social talk moves or task resolution moves according to the user model. However, one could see a drawback in this system in the fact

that it mixes small-talk and task planning which sometimes disadvantages the portability of the system (Allen, Ferguson, & Stent 2001).

Proposed Approach

Indeed, the assumption we are making in this research is that the management of the small-talk in the dialogue does not need to be so complex. Separating the dialogue management in layers of processing as proposed by Lemon, Cavedon, & Kelly (2003) or Zinn, Moore, & Core (2002) will allow the independent management of user trust and comfort in the dialogue and simplify the planning of the argumentation.

It is in that perspective that we believe a chatbot approach could be mixed with planning techniques:

- the chatbot, shallow discourse parsing approach is, in our opinion, suited to manage surface discourse trait needed for small-talk management,
- the planning techniques are needed to keep the consistency of the dialogue and guarantee the accomplishment of persuasive goals.

Therefore, we propose the following layered approach:

1. A *planning* layer, directed by dialogue game rules, decides the current dialogue phase but is not responsible for the actual dialogue realisation.
2. An *interaction* layer receives constraints from the previous one and uses chatbot techniques to construct the dialogue with the user.

With such architecture, the dialogue management clearly separates the two distinct components of the dialogue. This makes the planning of the persuasive argumentation more direct as it is not mixed with the dialogue's surface traits management. In addition, concentrating on two independent layers eases the development of the emotional and social cues handling which allows the user to feel comfortable and interested in the dialogue.

Conclusion

Up to now, the steps followed by our project have shown that the realisation of a persuasive argumentation system for health-care would need a broad range of techniques in human computer dialogue, eventually taking inspiration from both chatbot systems and purely task oriented approaches.

Moreover, the "emotional" part of argumentation – rarely emphasised by current road map in natural argumentation work – puts an accent on the type of planning that will be needed in this project:

- planning persuasive argumentation by giving the position of the expert – i.e. the leader of the conversation – to the computer will eventually ease the achievement of its goals,
- to complement this approach, the dialogue manager needs to employ social cues to be effective and show empathy in order to offer human-like health-care counselling.

Social cues are needed in the dialogue for the user to feel comfortable and we believe using a chatbot approach will be sufficient. In addition, if this processing is separated from persuasive goal planning, in a multi-layered dialogue management, it will benefit the planning system by giving a better control of the dialogue.

Indeed, the dialogue is divided in two distinct parts:

1. the persuasion discourse which spans over the whole dialogue and forms the main content of the dialogue. The persuasion techniques used for this component of the dialogue are dependent on the application field,
2. the surface traits of the dialogue that are common features of any human-like dialogue. These parts of the dialogue are therefore not dependent on the application but are necessary to provide a comfortable experience to the user.

Hence, separating the dialogue management in two layers will improve the portability and extensibility of the system.

References

Allen, J.; Ferguson, G.; and Stent, A. 2001. An architecture for more realistic conversational systems. In *IUI '01: Proceedings of the 6th international conference on Intelligent user interfaces*, 1–8. New York, NY, USA: ACM Press.

Barnes, J., ed. 1998. *The Complete Works of Aristotle: Revised Oxford Translation*. Princeton University Press.

Bench-Capon, T. J. M. 2003a. Persuasion in practical argument using value-based argumentation frameworks. *Journal of Logic and Computation* 13(3):429–448.

Bench-Capon, T. J. M. 2003b. Try to see it my way: Modelling persuasion in legal discourse. *Artificial Intelligence and Law* 11(4):271–287.

Bickmore, T., and Giorgino, T. 2004. Some novel aspects of health communication from a dialogue systems perspective. In *AAAI Fall Symposium on Dialogue Systems for Health Communication*.

Cassell, J., and Bickmore, T. 2002. Negotiated collusion: Modeling social language and its relationship effects in intelligent agents. *(in press) User Modeling and Adaptive Interfaces*.

De Angeli, A. D.; Johnson, G. I.; and Coventry, L. 2001. The unfriendly user: Exploring social reactions to chatterbots. In Helander, M. G.; Kalid, H. M.; and Po, M. M., eds., *Proc. Int. Conf. Affective Human Factor Design*, 467–474. Asean Academic Press.

Farzanfar, R.; Frishkopf, S.; Migneault, J.; and Friedman, R. 2005. Telephone-linked care for physical activity: a qualitative evaluation of the use patterns of an information technology program for patients. *J. of Biomedical Informatics* 38(3):220–228.

Fogg, B. J. 2003. *Persuasive technology: using computers to change what we think and do*. Morgan Kaufmann Publishers.

Friedman, R. H.; Stollerman, J. E.; Mahoney, D. M.; and Rozenblyum, L. 1997. The virtual visit: Using telecommunications technology to take care of patients. *Journal of the American Medical Informatics Association* 4(6):413–425.

Gilbert, M. A.; Grasso, F.; Groarke, L.; Gurr, C.; and Gerlofs, J. M. 2003. The persuasion machine. In Reed, C., and Norman, T. J., eds., *Argumentation Machines: New Frontiers in Argument and Computation*. Springer. chapter 5.

Glanz, K.; Shigaki, D.; Farzanfar, R.; Pinto, B.; Kaplan, B.; and Friedman, R. H. 2003. Participant reactions to a computerized telephone system for nutrition and exercise counseling. *Patient Education and Counseling* 49(2):157–163.

Grasso, F. 2002. Towards a framework for rhetorical argumentation. In Bos, J.; Foster, M.; and Matheson, C., eds., *EDILOG 02: Proceedings of the 6th workshop on the semantics and pragmatics of dialogue*, 53–60.

Grasso, F. 2003. Rhetorical coding of health promotion dialogues. In Barahona, E., ed., *Artificial Intelligence in Medicine - 9th Conference on Artificial Intelligence in Medicine in Europe, AIME*, volume 2780 of *LNAI*, 179–188. Springer-Verlag.

Green, N. 2003. Towards an empirical model of argumentation in medical genetics. In Reed, C.; Grasso, F.; and Carenini, G., eds., *Proceedings of the IJCAI Workshop on Computational Models of Natural Argument*, 39–44.

Guerini, M.; Stock, O.; and Zancanaro, M. 2003. Persuasion models for intelligent interfaces. In Reed, C.; Grasso, F.; and Carenini, G., eds., *Proceedings of the IJCAI Workshop on Computational Models of Natural Argument*.

Lemon, O.; Cavedon, L.; and Kelly, B. 2003. Managing dialogue interaction: A multi-layered approach. In *Proceedings of the 4th SIGdial Workshop on Discourse and Dialogue*.

Levin, J. A., and Moore, J. A. 1977. Dialogue-games: metacommunication structures for natural language interaction. *Cognitive Science* 1(4):395–420.

Mann, B. 1999. An introduction to rhetorical structure theory (rst). web page.

Mcglashan, S.; Fraser, N.; Gilbert, N.; Bilange, E.; Heisterkamp, P.; and Youd, N. 1992. Dialogue management for telephone information systems.

Norman, T. J., and Reed, C. 2003. *Argumentation Machines : New Frontiers in Argument and Computation (Argumentation Library)*. Springer.

Perelman, C., and Olbrechts-Tyteca, L. 1969. *The New Rhetoric: A Treatise on Argumentation*. University of Notre Dame Press.

Pietquin, O., and Dutoit, T. 2003. Aided design of finite-state dialogue management systems. In *Proc. of the IEEE International Conference on Multimedia & Expo (ICME 2003)*, volume 3, 545–548.

Pinto, B. M.; Friedman, R.; Marcus, B. H.; Kelley, H.; Tennstedt, S.; and Gillman, M. W. 2002. Effects of a computer-based, telephone-counseling system on physical activity. *American Journal of Preventive Medicine* 23(2):113–120.

Ravenscroft, A., and Pilkington, R. M. 2000. Investigation by design: Developing dialogue models to support reasoning and conceptual change. *International Journal of Artificial Intelligence in Education: Special Issue on Analysing Educational Dialogue Interaction: From Analysis to Models that Support Learning* 11(1):273–298.

Ravenscroft, A. 2000. Designing argumentation for conceptual development. *Computers & Education* 34(3-4):241–255.

Reed, C., and Grasso, F. 2001. Computational models of natural language argument. In Alexandrov, V. N.; Dongarra, J.; Juliano, B. A.; and Renner, R. S. a., eds., *Proceedings of the International Conference on Computational Sciences-Part I.*, volume 2073 of *Lecture Notes in Computer Science*, 999–1008. London, UK: Springer-Verlag.

Reed, C. A., and Long, D. P. 1998. Generating the structure of argument. In Hoeppner, W., ed., *Proceedings of the 17th International Conference on Computational Linguistics and 36th Annual Meeting of the Association for Computational Linguistics (COLING-ACL98)*, 1091–1097.

Saarinen, L. 2001. Chatterbots: crash test dummies of communication. Master's thesis, Te University of Art and Design Helsinki UIAH.

Seneff, S., and Polifroni, J. 2000. Dialogue management in the mercury flight reservation system. In *Proc. Dialogue Workshop, ANLP-NAACL*.

Shechtman, N., and Horowitz, L. M. 2003. Media inequality in conversation: how people behave differently when interacting with computers and people. In Cockton, G., and Korhonen, P., eds., *Proceedings of the SIGCHI conference on Human factors in computing systems*, 281–288. Ft. Lauderdale, Florida, USA: ACM Press.

Slack, W. 2000. Patient-computer dialogue: a review. In van Bemmel, J., and Mccray, A., eds., *Yearbook of medical informatics*. Stuttgart, Germany: Schattauer. 71–78.

Smith, M. B.; Burke, K. E.; Torgerson, J. S.; Stollerman, J. E.; Kern, D. C.; Hardy, W. L.; and Friedman, R. H. 1988. Logical and efficient conversation between patients and the telephone linked computer system. In *Proceeding of the 12th Annual Symposium on Computer Applications in Medical Care*, 463–469.

Stiff, J. B., and Mongeau, P. A. 2002. *Persuasive Communication, Second Edition*. The Guilford Press, second edition.

Weizenbaum, J. 1966. Eliza-a computer program for the study of natural language communication between man and machine. *Commun. ACM* 9(1):36–45.

Wright, J. H.; Wright, A. S.; Albano, A. M.; Basco, M. R.; Goldsmith, J. L.; Raffield, T.; and Otto, M. W. 2005. Computer-assisted cognitive therapy for depression: Maintaining efficacy while reducing therapist time. *Am J Psychiatry* 162(6):1158–1164.

Zinn, C.; Moore, J. D.; and Core, M. G. 2002. A 3-tier planning architecture for managing tutorial dialogue. In *ITS '02: Proceedings of the 6th International Conference on Intelligent Tutoring Systems*, 574–584. London, UK: Springer-Verlag.