

# The Intentional Interface

Peter Wallis

Centre for Policy Modelling  
Business School  
Manchester Metropolitan University  
Manchester, United Kingdom  
`pwallis@acm.org`

**Abstract.** The SERA project put “robot rabbits” in older peoples homes and recorded what happened. The challenge is now to use that data to develop better rabbits, but how? We are currently working on a methodology for distilling this data down into explanatory narratives, but in the mean time we are working on the idea that the essential nature of the SERA interface (and other conversational agents) is that it is *intentional* - it is an interface that sets out to have people ascribe beliefs and desires to it. According to Tomasello, this is not enough however. An intentional interface also needs to intend to help - it needs to be *cooperative*. What this means in detail is fleshed out in the context of an IVR system - a computer that answers the telephone.

## 1 Introduction

A year on from the SERA project - Social Engagement with Robots and Agents - this paper looks back on what we did, and attempts to put the lessons learned in a historical context. Our vision was to use a talking robot rabbit (an augmented Nabaztag) as long term “companion”. Obviously it is beyond us to create a perfect simulation of a human conversational partner, but was current technology able to capture the *essence* of what is needed? The answer was no, but the experience certainly prompted some thinking about that essence. This paper develops that thinking and in doing so, offers a “grand unified theory” of HCI based on Dennett’s *Intentional Stance* [1]. The theory is then used to develop an IVR system (Interactive Voice Response) that answers the telephone and, as such, is inevitably treated as a social actor.

## 2 SERA

The SERA project was funded under the FP7 Theme 2.2: Cognitive Systems, Interaction and Robots, and the aim was to put real robots in real people’s hallways and kitchens and record what happens. The work was done with the School of Health and Related Research at Sheffield (SchARR) which had extensive experience recruiting subjects from the broader community, and which was working with “smart homes” with a view to “life-style reassurance” in which

people living alone could be assured that, should something happen to them, help would be at hand. Building a state-of-the-art companionable robot is a project in itself [2] and so, rather than building a mobile autonomous robot, we decided to use a commercial off the shelf Nabaztag [3] that behaved as if it could sense it's environment, but which actually used the smart home sensors. Although the "robot" was not be mobile, it was able to sense its environment and was thus able to initiate action in a way that is expected of robots. It is in this sense of "robot" - autonomous action based on sensing the environment - that our nominally simple interface addressed the call. Figure 1 shows the setup in use.



**Fig. 1.** Sarah and Harvey.

That was the set-up, but another challenge was deciding what the rabbit should say. We settled on the popular scenario of an "exercise companion". Using the Trans Theoretical Model of behaviour change (TTM) [4] which places people in one of 5 stages, the system could introduce the advantages of being fit at the appropriate moment if the user was in the pre-contemplation stage, or identify progress if the user was in the maintenance stage and so on. The idea was to use "key-word spotting" speech technology and develop a system that was primarily "system initiative" with conversation being initiated by the following events:

- Keys off (participant going out)
- Keys on (participant returning home)
- PIR & first appearance in the morning
- PIR & last activity of the day has been done
- PIR & a new message/recommendation
- participant initiation - "Hey rabbit!"

If the keys came off when the subject had an entry in the diary for some exercise, then the rabbit could say things like “Going swimming? Have a good time” - which at least one subject found quite impressive even if she knew how it worked.

Before introducing the theory, the shared conclusions from the SERA project were, first, don’t try to use ASR in the wild - the Siri publicity (formal and viral) is a dream. That kitchen is not my kitchen: there are no kids practicing the tin whistle, no oil sizzling on the hob, no radio playing, no extractor fan, no traffic and no refrigerator humming away. Indeed she is perfect as well with a nice East Coast accent with no Yorkshire clipping or Australian vowels. And the recipe - no star annis, or dried paw paw; nothing out of the ordinary. Put a speech recognition system in a kitchen and, we discovered, word error rates are too low for even a handful of key phrases.

Second, managing attention is a big issue when the interface is sensing the environment, is always on, and can be proactive. The classic HCI interface is passive (see below) and needs to be “poked” before it behaves. Our existing model for a proactive interface is the alarm that *demand*s attention. In between is the telephone that could be demanding when its location was fixed but mobile phones would, ideally, be more “socially aware” of the context. The SERA rabbits used a PIR security sensor to detect with people were near but is she in a hurry? Is she making an omlette? [5] or is she just after a glass of water in the middle of the night?

Third, people have “idiosyncratic” behaviour. Where one person gets cross and yells, another laughs while another roles her eyes. Another subject may frown or not respond at all. Naturally the notion of a response to “the same” event is also problematic without a framing theory but this issue is well known; what stood out for us all was the huge range of responses across socio-economic backgrounds.

Finally, we can say that there is no consensus on what to do with the data we collected. We could all publish papers, but how can the data be used to advance the state-of-the-art? There has been some work on a better methodology for looking at the data [6], but in this paper a theory of HCI is introduced which may help “frame” the questions one might ask of the data and so help identify the issues and suggest improvements. The proposed theory is based on how people view artifacts around them.

### 3 How the mind works<sup>1</sup>

The proposal is that the SERA rabbits were not simply a conventional human-computer interface with a speech recognition front end, but were instead an attempt at an *intentional interface*. This is not to say that the SERA interface was unique – many have attempted similar things – the point is to introduce a class for interface for which SERA is an example. In order to compare and contrast, the observation is that we can classify human-computer interfaces based

---

<sup>1</sup> Thanks to Steven Pinker for the title of this section and the next.

on how the user goes about understanding the computer, and that interesting distinctions can be drawn by looking at Dennett's position on intentional systems.

Dennett, argues that the study of minds is different to the study of brains, and that the wide spread use of "folk psychology" in the Social Sciences is perfectly valid as science. For realists there is little doubt that minds reside in the hardware of brains, but studying brains is not necessarily going to provide explanations for why things are the way they are. As a scientist one might have a theory of id, ego and super-ego, or as a mathematician one might have an elegant Bayesian model of how brains work that is meant to explain things, but Dennett's line is that the psychology we use in our everyday lives is equally valid as a scientific theory *and more efficient* .

Dennett argues that humans use three different approaches, or stances, when trying to predict the behaviour of something. When a system is fairly simple - balls on a level table perhaps - then we can use a causal model to predict future events. Tapping the white ball in a particular way will cause it to role over to the red ball and knock it into the centre pocket. Taking this **physical stance**, people use their knowledge of hundreds (if not thousands) of highly reliable "facts" about the way things behave to assemble chains of causal events to predict the future. Dennett was writing at the time of good old fashioned AI and so the nature of these facts, as we now know from the work in computer science, is problematic and (apparently) based on situated action. But possible enumerations and classification of the base facts is not the point; the point is we can and do reason causally. Taken to its extreme, this is the idea of a clockwork, deterministic universe and that ultimately "there is only physics".

Another way we humans predict the future is by knowing what something is designed to do. Pressing the brake pedal when driving, one does not reason about hydraulic fluid, but simply knows what that pedal is *meant* to do. An alarm clock is too complex to follow the internal workings in a causal sense but, knowing what it is designed to do, one can set it in the evening and predict that it will wake you in the morning. This is of course where classic HCI is based with advice on how to create good interfaces being things like making sure that the system works as designed, and that the user has a clear idea of the function of the design (e.g. Interaction Design: beyond human-computer interaction (2ed) [7]). When we use this **design stance** , note how it licences the notion of something being "broken".

The **intentional stance** is what we use when a system is too complex to predict with the physical stance, and the purpose of system - what it is designed for - is inaccessible to us. We humans have a strong tendency to assume that something capable of autonomous action will do what it believes is in its interests. That is, that the system will have desires, and that it can plan its actions to achieve (some of) those desires given its beliefs about the current state of the world. This tendency is very strong in us. Seeing two children tugging at a teddy bear, the casual observer will assume they both *want* it. When playing chess against a computer, I do not reason about the causal behaviour of registers and

electricity, but rather predict the future by reasoning along the lines of it *wanting* to take my bishop. The consequences of a rational agent wanting something do not need to be spelt out for us; we just know. We are also likely to explain things that are not rational action with this model and Dennett gives a lovely example of someone explaining that electricity normally *wants* to take the shortest path but sometimes it “gets confused”.

### 3.1 The Human-Computer Interface from the stances

Current HCI best practice can be critiqued as using a “tool” metaphor in which the computer is wielded by the user to achieve his or her goals. This is fine as far as it goes and has the advantage that as long as the tool does what it is *designed* to do, the user is responsible for outcomes. Hit your thumb with a hammer and there is only yourself to blame. Using such a metaphor, the guidance on HCI design is about making the design clear, and the consequences of an action explicit and immediate [7]. In retrospect this is Dennett’s design stance. The human is expected to understand what the interface is designed to do, and then wield it appropriately.

Extending the metaphor, the sexy human-computer interfaces are those based on the physical stance. On the surface there is a class of interface that exploits the “facts” we have about the physical world with desk tops as a place to “put” things temporarily, folders that “contain” stuff, and recycle bins for the things we don’t want any more. Today’s touch screens allow things to be “flicked” and multi touch screens allow things to be “stretched” in a way that tend to obey our facts about the (physical) world. At a deeper level, many modern interfaces - especially those designed for new markets such as children - not only allow, but actively encourage exploration. In effect they encourage the user to discover things about causality in the virtual world that mirror the “hundreds or thousands of facts” we know that support the physical stance in the physical world.

This exploration process is clearly what Suchman points to in her classic work on situated action and the photocopier [8].

The proposal is that the *essence* of our rabbit interface is that (it looks as if) it behaves in accordance with our intentional stance. At first blush the distinguishing feature of the SERA rabbits was the speech recognition. On reflection the distinguishing feature was that the PIR meant that our rabbits were proactive about initiating a conversation. In accordance with the call, our engineering aim was indeed to sense and react to the environment as a robot is expected to do and this meant that it is hard not to think of the rabbit as *wanting* to do things. From the user’s perspective, a rabbit has its own agenda and the user slips very easily into taking an intentional stance. Once a conversation was started - as was clear from the video evidence - the rabbits were never good at negotiating *shared* goals. The problem was that the system did not take an intentional stance on the functioning of its user and was thus not able to negotiate a shared intention.

It turns out that the intentional nature of human-human communication is well recognised in linguistics proper. What our rabbits need is a better approach to dialogue management.

## 4 The language instinct

Computer science as applied to natural language moved out of the arm chair in 1989/90 and that research community generally accept that data driven research is the way forward. The critical mass however use statistical models and, like the behaviourists of old, abhor any notion involving “mental attitudes”. In the last 10 years this has been applied to dialogue and so, the argument goes, we do not need to study how language works because (given enough data) machine learning techniques will enable computers to simulate conversational behaviour without theory. Partially Observable Markov Decision Processes (POMDP) have been applied to the dialogue problem [9] and the claim is that the goal of the user - the conversational partner’s intent - can be treated as a “hidden state” in a POMDP.

This is a noble aim but in practice human intervention is required to make these systems work. In practice ML techniques are not trained on raw speech or text, but rather on tag sequences where the tags are from a set of dialogue acts or DAs. There is no consensus on what should go into these sets of tags and in general each annotation scheme adapts an accepted set of dialogue acts to the particular application domain. From a linguistics perspective the methodology is sequence analysis [10] which has been unfavourably critiqued by Levinson [11] (page 289), and which in practice produces results with questionable repeatability [12]. Indeed Eduard Hovy at ISI has for some time been pointing out just how much theory is embedded in the choice of DAs and argues for more public discussion of underlying theory [13].

Much of Linguistics and the social scientists however do data driven research but take the line that mental attitudes are causal in human affairs and, as Dennett argues, that a valid science can be based on such concepts. The argument is made beautifully by ten Have [14] but the remainder of this paper is based on the hypothesis by Michael Tomasello [15] that human communication is not only intentional in nature but also a fundamentally cooperative process. Rather than the language instinct being some hard wired ability to recognise mathematical patterns [16], it is the hard wired ability to recognise the intention of others, and the propensity to cooperate in the communicative process.

As is often the case there are notable exceptions - classically Grosz and Sidner [17] talk of attention and intention, and the people working on Max, an embodied conversational agent that has been deployed in the wilds of a museum [18], have talked about attention and mixed initiative at the “discourse level,” and in this paper we use a model of intention recognition that has been used in military simulation based on the pre-compiled plans of a BDI architecture [19]. This is discussed further in the next section but first a brief discussion of cooperation may be required.

The need for cooperation is made clear when we take a closer look at what linguists have said about the mechanism of language. Conversation Analysis [20, 21, 14] (CA) is a methodology that enables researchers to notice the detail of language in use and the approach has certainly been prolific. Seedhouse [22] summarises the findings of CA as follows. At any point in a conversation, an utterance will go **seen but unnoticed** in that it is (one of a small set of) expected response, it will go **noticed but accounted for** where it wasn't the expected response but the recipient could figure out why it was said, or the utterance will **risk sanction**. When talking to computers, the sanction is swearing [23] or users not wanting "to use the system on a regular basis" [24]. The point is that the "accounting for" requires us to work hard at recognizing the intent of the speaker. Consider the text book example from Eggins and Slade with which they introduce the notion of sequential relevance:

**A:** What's that floating in the wine?

**B:** There aren't any other solutions.

You will try very hard to find a way of interpreting B's turn as somehow an answer to A's question, even though there is no obvious link between them, apart from their appearance in sequence. Perhaps you will have decided that B took a common solution to a resistant wine cork and poked it through into the bottle, and it was floating in the wine. Whatever explanation you came up with, it is unlikely that you looked at the example and simply said 'it doesn't make sense', so strong is the implication that adjacent turns relate to each other [25].

The appearance of an utterance immediately after another in an interaction to which the partners are committed (that is, a conversation) causes the hearer to work hard at recognising the intent of the speaker. The social pressure on doing this and cooperating in general is captured by Tomasello. Quoting at length:

Thus, from the production side, we humans must communicate with others or we will be thought pathological; we must request only things that are reasonable or we will be thought rude; and we must attempt to inform and share things with others in ways that are relevant and appropriate or we will be thought socially weird and will have no friends. From the comprehension side, we again must participate, or we will be thought pathological; and we must help, accept offered help and information, and share feelings with others, or we risk social estrangement. The simple fact is that, as in many domains of human social life, mutual expectations, when put into the public arena, turn into policable social norms and obligations. The evolutionary bases of this normative dimension of human communication in terms of public reputation, will be ... [15]

Looking again at Eggins and Slade's example, apes it seems are not hard-wired, preprogrammed and/or socialized into putting the effort in and would simply say "it doesn't make sense" and move on. What our computers need as

social actors is the ability to *account for* the communicative acts of its human companions and to do that requires intention recognition and a willingness to put in the effort.

## 5 Practical intention recognition

Intention recognition and pro-active cooperation are core to human communication of all kinds but it is not enough to say this or even prove it. If researchers in an engineering faculty are going to embrace it, there needs to be a means of implementing it, and the rest of this paper shows how this can be done for limited, but useful, cases. As is often the case with non-incremental development [26] a holistic solution can result in problems cancelling each other out. The challenge of intention recognition and the challenge of proactively helping can be beneficially addressed together by working from a pool of pre-compiled partial plans as used in BDI agent architectures [27–29]. The limited domain used to demonstrate the process is the very applied task of accessing information in a relational database via the telephone.

For the next project the aim is to demonstrate an intentional interface with an IVR system and the scenario under consideration is the classic directory assistance application. With these systems a caller can ring the institution and talk to a computer which puts them through to the required individual. The classic approach would hold the relevant information in a relational database and, in the spirit of Meaning-Text Theory [30] would focus on the information. Consider:

**M/C** Welcome to University of Sheffield directory assistance. Who do you wish to contact?

**USR** Mark Hepple please.

An ASR module would produce the text from the voice signal, a parser normalize the grammar, and a language understanding module might map that into a canonical representation of the meaning. In the case of database access, the canonical form might take the shape of the SQL query:

```
SELECT ALL FROM phonebook WHERE "familyName='Hepple' &
givenName='mark'"
```

For the University of Sheffield phone book, such a query returns:

```
givenName familyName dept extn
Mark        Hepple      DCS 21829
```

From this result a text generation system that uses some form of pronoun rewriting could say “His number is 21829.”

This is great when things go well. Trouble occurs however when the user’s query does not return a single row.

## 5.1 Trouble in text

If the caller asks for Mark Hawley in Computer Science, the resulting query returns no rows as Hawley is not in the Department of Computer Science. What should the system do?

Using a classic HCI approach the aim would be to make it clear to the user that he or she is using a relational database and to point out that it is the user's query that is resulting in unhelpful output. For many of the DARPA Communicator systems a common solution to no result (no rows in the resulting table) was to remind the caller that he or she could change their query by adjusting the parameters. The user is using a tool, and it is the user's responsibility to use it as *designed*.

At the other extreme, if the caller asks for Mark in Computer Science, the resulting query returns:

<b>givenName</b>	<b>familyName</b>	<b>dept</b>	<b>extn</b>
Mark	Hepple	DCS	21829
Mark	Stevenson	DCS	21921
Mark	Ellerby	DCS	21856

But which Mark does he or she mean? When the user's query matches multiple rows a graphical user interface can present all the rows and this is sometimes attempted with IVR systems. Once again in the spirit of the Communicator systems the system might say:

**M/C:** "There are 16 people with that name,  
the first is Mark Hepple in Computer  
Science. Is that who you are after?"  
**Usr:** No  
**M/C:** The second is Mark Stevenson in Com-  
puter Science. Is that the person you  
are after?  
**Usr:** No  
...

The computer as tool metaphor may work, but can an intentional approach provide an alternative?

## 5.2 An Intentional interface

It seems a computer behaving as a social actor needs not only to be right, but also seen to be helpful, and the challenge in the first instance is to come up with helping strategies that the system can introduce. Introducing a strategy requires mixed initiative, not just at the information level, but at the level of intent. The following discussion shows what this means and does it in terms of conversational strategies implemented as plans in a BDI architecture.

The Belief, Desire and Intention architecture was introduced by the software agents community as a means of balancing reactive and deliberative behaviour in a constantly changing environment. The approach does not *do* planning in the traditional AI sense, but rather manages commitment to plans. The usual

BDI approach is to work from a library of pre-compiled plans and “intention recognition” can be implemented (in a limited sense) as a variant of plan choice.

In the case of 2 or perhaps 3 rows, the HCI approach of presenting the list can be used:

**Usr:** Mark in Computer Science  
**M/C:** Mark Hepple or Mark Stevenson?  
**Usr:** Stevenson  
**M/C:** Mark Stevenson is on 219...

This strategy is good as far as it goes, but the machine’s question assumes the user knows. Thus, this helping strategy might fail if the user is unsure. For a BDI architecture this is not a problem — the architecture was introduced to handle plan failure — and the system simply looks for another plan. The success or failure of this plan will of course depend on what the user says next. Failure however is not bad; what is important is that the system is seen to be trying to help. Consider:

**Usr:** Mark in Computer Science  
**M/C:** Mark Hepple or Mark Stevenson?  
**Usr:** Err I was talking with Mark about doing a masters course  
**M/C:** Mark Hepple is the masters coordinator.  
**M/C:** Mark Hepple is on 219...

Which is a successful outcome based on the system having a strategy in the plan library for callers looking for information on the masters programme. Critically however it is socially acceptable (i.e. does not risk sanction) for the user’s plan to fail:

**Usr:** Mark in Computer Science  
**M/C:** Mark Hepple or Mark Stevenson?  
**Usr:** Err I was talking with Mark about doing a masters course  
**M/C:** Right.  
...

The point here is the “unfolding” of the conversation and, like a game of football, plan failure is routine. What matters is that the system is seen to be trying so it does not “risk social estrangement ... and have no friends”.

If the user’s query returns no rows, it is the system that knows what it has and the machine can push information:

**Usr:** Mark Hawley in Computer Science  
 please  
**M/C:** Err no Mark Hawley in Computer Sci-  
 ence. (1 second)  
**M/C:** There is a Mark Hawley in Health? (1  
 second)  
**M/C:** I can give you the number for the De-  
 partmental Secretary in Computer Sci-  
 ence?  
**Usr:** Mark Hawley please  
**M/C:** Professor Mark Hawley in the School  
 of Health and Related Research is on  
 219..."

Once again the point is the "unfolding" of conversation and a socially ept inten-  
 tional interface has a responsibility to help.

The fourth case is where there are many rows in the table - 0,1,2,many - and  
 when this happens there is often a misunderstanding. Consider someone who  
 thinks he is calling the number for the Department of Computer Science and  
 says:

**M/C:** Good morning how can I help?  
**Usr:** Mark please.  
**M/C:** Err you have called directory assistance  
 for the University of Sheffield.  
**M/C:** I'm sorry, who are you after?

Putting on one's CA hat, the "work done" but the machine's response is to  
 appeal to the caller's sense of fairness. As Tomasello says, people have a sense  
 of fairness and the strategy here is for the system to explain what its job is,  
 suggesting that it is unfair to expect it to be able to help in this case.

Intention recognition is hard for a machine but we can get some way there  
 by working from a fixed set of plans. At this stage the above conversational  
 strategies have been implemented but the system has not been evaluated in an  
 operational setting at this stage. The point of this paper however has been to  
 introduce an alternate model for HCI, and to demonstrate that it is not just  
 hand waving - Tomasello's claims are concrete and implementable.

## 6 Conclusion

ICT is amazingly versatile, enabling us to create the information systems we  
 want, with the interfaces we want. Without limitations, the designer is ultimately  
 responsible for *any* problems. It is very tempting in these circumstances for  
 us to favour interfaces that exploit the user's *design stance* which shifts some  
 responsibility to the user - the user *ought* to RTFM (read the manual) and then  
 wield the tool as we designed it to be wielded.

The sexy new interfaces - be it 2010 or 1985 - exploit the user's *physical  
 stance* in which our understanding of cause and effect in the physical world is  
 mapped onto virtual events.

The claim being made is that the *essence* of “human-like” interfaces — from embodied conversational agents to robot companions through chat-bots to speech interfaces and IVR systems — is that the user takes an *intentional stance*. Although making these systems more like humans is interesting in its own right — adding micro movements to ECA, emotion or persona to chat-bots — the feature of human communication that provides an opportunity for HCI is the intentional nature of the human interface. This is not enough however because, according to Tomasello, a social actor in human society also needs to be *cooperative*.

Such claims might be seen as too abstract, but the paper gives an interpretation of these principles in the context of an IVR system providing directory assistance.

## References

1. Dennett, D.C.: The Intentional Stance. The MIT Press, Cambridge, MA (1987)
2. : The companions project (2007) <http://www.companions-project.org/>.
3. : Nabaztag (2010) [http://www.violet.net/\\_nabaztag-the-first-rabbit-connected-to-the-internet.html](http://www.violet.net/_nabaztag-the-first-rabbit-connected-to-the-internet.html).
4. Prochaska, J., Velicer, W.: The transtheoretical model of behaviour change. *American Journal of Health Promotion* **12** (1997) 38–48 TTM or ttm.
5. Wallis, P.: A robot in the kitchen. In: *ACL Workshop WS12: Companionable Dialogue Systems*, Uppsala (2010)
6. Wallis, P.: From data to design. *Applied Artificial Intelligence* **25** (June 2011) 530–548
7. Sharp, H., Rogers, Y., Preece, J.: *Interaction Design: beyond human-computer interaction* (2ed). John Wiley and Sons, Chichester, UK (2007)
8. Suchman, L.A.: *Plans and situated actions - the problem of human-machine communication. Learning in doing: social,cognitive,and computational perspectives*. Cambridge University Press (1987)
9. Young, S.J.: *Spoken dialogue management using partially observable markov decision processes* (2007) EPSRC Reference: EP/F013930/1.
10. Bakeman, R., Gottman, J.M.: *Observing Interaction: An Introduction to Sequential Analysis*. Cambridge University Press (1997)
11. Levinson, S.C.: *Pragmatics*. Cambridge University Press (2000) discussion of discourse analysis and mark up is page 289.
12. Carletta, J., Isard, A., Isard, S., Kowtko, J.C., Doherty-Sneddon, G., Anderson, A.H.: The reliability of a dialogue structure coding scheme. *Computational Linguistics* **23**(1) (1997) 13–31
13. Hovy, E.: *Injecting linguistics into nlp by annotation* (July 2010) Invited talk, *ACL Workshop 6, NLP and Linguistics: Finding the Common Ground*.
14. ten Have, P.: *Doing Conversation Analysis: A Practical Guide (Introducing Qualitative Methods)*. SAGE Publications (1999)
15. Tomasello, M.: *Origins of Human Communication*. The MIT Press, Cambridge, Massachusetts (2008)
16. Pinker, S.: *The Language Instinct*. Penguin Books, London (1994)
17. Grosz, B., Sidner, C.: Attention, intention, and the structure of discourse. *Computational Linguistics* **12**(3) (1986) 175–204

18. Kopp, S., Gesellensetter, L., Kramer, N., Wachsmuth, I.: A conversational agent as museum guide - design and evaluation of a real-world application. In: 5th International working conference on Intelligent Virtual Characters. (2005) <http://iva05.unipi.gr/index.html>.
19. Heinze, C.: Modelling intention recognition for intelligent agent systems (November 2004)
20. Sacks, H., Schegloff, E., Jefferson, G.: A simplest systematics for the organisation of turntaking in conversation. *Language* **50**(4) (1974) 696–735
21. Hutchby, I., Wooffitt, R.: *Conversation Analysis: principles, practices, and applications*. Polity Press (1998)
22. Seedhouse, P.: *The Interactional Architecture of the Language Classroom: A Conversation Analysis Perspective*. Blackwell (September 2004)
23. Wallis, P.: Robust normative systems: What happens when a normative system fails? In Antonella de Angeli, S.B., Wallis, P., eds.: *Abuse: the darker side of human-computer interaction*, Rome (September 2005)
24. Wallis, P.: Revisiting the DARPA communicator data using Conversation Analysis. *Interaction Studies* **9**(3) (October 2008)
25. Eggins, S., Slade, D.: *Analysing Casual Conversation*. Cassell, Wellington House, 125 Strand, London (1997)
26. Constant, E.W.: *The Origins of the turbojet revolution*. The John Hopkins Press Ltd, London (1980)
27. Bratman, M.E., Israel, D.J., Pollack, M.E.: Plans and resource-bound practical reasoning. *Computational Intelligence* **4** (1988) 349–355
28. Rao, A., Georgeff, M.: *BDI agents: from theory to practice*. Technical Report TR-56, Australian Artificial Intelligence Institute, Melbourne, Australia (1995)
29. Wooldridge, M.: *Reasoning about Rational Agents*. The MIT Press, Cambridge, MA (2000)
30. Mel'cuk, I.: Meaning-text models: a recent trend in soviet linguistics. *Annual Review of Anthropology* **10** (1981) 27–62