

Human-Agent Communication

Ian Dickinson
Living the Vision
HP Laboratories Bristol
HPL-98-130
July, 1998

agent technology,
human-agent
interaction, HCI

Agent-based systems are potentially extremely powerful, but agent properties such as autonomy and proactivity raise significant challenges for usability. This paper argues that too little attention is being paid to human-agent communication from the user's perspective. To justify the aspirations of agent designers, fundamental usability questions must be addressed more systematically than is currently the case. The paper outlines some of the problems that are foreseen, and surveys some of the many research activities currently underway that will contribute to eventual solutions.

Human-Agent Communication

Ian Dickinson

*Hewlett-Packard Laboratories
Filton Rd, Stoke Gifford, Bristol, UK
ian_dickinson@hp.com*

Abstract

Agent-based systems are potentially extremely powerful, but agent properties such as autonomy and proactivity raise significant challenges for usability. This paper argues that too little attention is being paid to human-agent communication from the user's perspective. To justify the aspirations of agent designers, fundamental usability questions must be addressed more systematically than is currently the case. The paper outlines some of the problems that are foreseen, and surveys some of the many research activities currently underway that will contribute to eventual solutions.

Introduction

Agents, software and hardware, are currently one of the hot topics in computer science and software engineering. They are the focus of much innovative research, and a certain measure of speculative hyperbole, not all of which can be justified.

Our group at HP Labs in Bristol is investigating the application of agent technology to electronic commerce, specifically buying and selling over the Internet. Examples from this widely applicable domain will be used as motivators in this paper. Much of our research, like that of others in the field, is devoted to innovation in the underlying agent mechanisms and models: how they are constructed and how they interact with each other. Indeed, there is much crucial work to be done here. However, there is another equally large and challenging set of problems: those of human—agent interaction. Increasingly we are aware these problems must be addressed before agents can meet the expectations that many commentators are setting for them. Despite the many claims for the benefits of agents to users, relatively little work is in progress to assess and prove these claims.

The design of human-computer interaction (HCI) is hard to do well, and all too often is done badly

[1]. I argue that this general difficulty is made markedly more difficult when agents are introduced into the interface. Indeed, it is precisely the properties of agents that promise to make them useful that also make them problematic.

In this paper, I aim to introduce some of the additional HCI issues agents introduce, and touch on some themes in the research work already underway. I hope to motivate the further, systematic exploration of the whole of the human-agent interface, and perhaps the critical re-evaluation of some of the claims that have been made. This brief paper does not set out to exhaustively catalogue all of the potential issues, nor to survey all relevant work. A more comprehensive survey of relevant work can be found in [2].

The paper is structured as follows. I begin by summarising some of the reasons why agents will be valuable to their human users and designers, and what properties of agent will be required to obtain them. Then I go on to explore potential problems these properties bring. There follows an introduction to some of the ongoing work that will help to address such problems. I conclude with some thoughts on future directions.

Agents: why bother?

Agents present a compelling vision of future computational systems. Intelligent, aware, flexible software entities occur regularly in science fiction and “future-mongering” science books and articles (e.g. [3, 4, 5]).

Debate about exactly what an agent is continues (though a good response is advanced in [6]). Even so, we can identify two major sets of motivations for using agents: *metaphors for system design*, and *metaphors for system presentation*.

In the first case, agents, and in particular agent societies, help system designers and developers create scalable, robust distributed systems by

drawing inspiration from biological and social systems such as bee colonies and economies. Such approaches can provide powerful benefits to the software designer. They are, however, the subject of many other books and articles (e.g. [7]) and I will not discuss them here.

In the second case, agents can present an organising metaphor for a user's experience of interacting with a computer. This approach draws inspiration from social interactions between humans, to enrich the user's interaction with a non-human social partner. The value of an agent to its user (apart from any given domain functionality encoded by the agent) is likely to include at least some of the following components:

- Availability – the agent can act as the user's proxy in the task when the user is otherwise engaged;
- Speed – the agent can act and react in situations more quickly than the user could;
- Thoroughness – the agent can perform repetitive tasks without losing attention through boredom;
- Flexibility – the agent may be specifically designed to adapt to changing circumstances or user preferences;
- Attractiveness – some researchers propose that new interface modalities, such as anthropomorphic characters, may make computational systems more attractive to users put off by today's computer interfaces.

As an illustration, consider an agent tasked with purchasing a particular good in an electronic market. The agent can be "in" the market, monitoring activity and identifying, for example, good market prices, 24 hours a day, seven days a week. Reaction times to new events, such as the arrival of a new seller, can be measured in milliseconds. Similarly, the bargaining cycle, as agents negotiate the price for the good, will be very fast. The agent can perform price comparisons with all available sellers [8], without tiring, as a human shopper would do.

It is the properties of the agent that give rise to these values. In particular, Wooldridge and Jennings [9] suggest that agents can be

characterised as hardware or software entities that are (at least):

- Autonomous
- Reactive
- Intelligent
- Social

The agent's ability to act as a proxy for its user arises from these properties. An agent imbued with such properties could enter into negotiations, acting independently to help achieve the user's goals in an unpredictable environment, and communicate effectively back to the user. But it is also these properties, particularly autonomy, which raise some difficult issues for the agent interface designer.

To judge from the proceedings of agent technology conferences, the majority of researchers' efforts are directed at the underlying technology providing the agents with these properties. These technologies include agents' cognitive models and architectures, agent communication languages, infrastructure (including mobility) and security. Comparatively little systematic attention appears to be being paid to how users will perceive and react to these systems.

Why agent interfaces are difficult

Consider the following simple examples.

Example 1. User Sally wishes her agent to negotiate to buy 10000 minutes of network connection time on an open bandwidth trading exchange. Her priority is price, and she is willing to give up some, but not too much, quality of service (QoS) to get a good price. How does she express her (personal or corporate) view of "value for money", or even what constitutes unacceptably low quality of service?

Example 2. Continuing the previous example, Sally wishes to delegate to her agent the task of finding and negotiating a good deal, while she performs some other important tasks requiring her attention. How does she express the limits on the agent's ability to conduct negotiations, including the situations in which it can continue negotiating and those in which it should get confirmation before proceeding? How far can the agent go in the negotiation process, especially if a good deal

arises and Sally doesn't respond to requests for confirmation?

Example 3. Sally's agent has reported back that it has secured a good deal on the bandwidth purchase. But Sally is not satisfied with the result. How can Sally (a) find out whether this really was the best available, and (b) educate or punish the agent?

Example 4. During the dialogue in which Harry is instructing his agent to make a purchase, it detects that Harry has misconceptions about the structure of the market or the economic principles on which it runs. Should the agent take control of the dialogue and give Harry some coaching, and if so how?

Example 5. In a call service centre, responsibility for choosing the sequence of jobs to work on is reassigned from the well-trained, skilful engineers to a set of scheduling agents jointly responsible for overall call handling efficiency. The engineers are now required to consult their agent to be told their daily schedule.

The above examples are illustrative of some of the problems that may arise. These may be more generally described by the following groups of issues:

- Instructing agents to act and react
- Delegating task and authority
- Sharing context
- The user's conceptual model of the agent
- Dialogue issues

Although many of these groups are interrelated, and the distinctions between them sometimes rather fine, they help to identify the major research challenges.

Instructing agents to act and react

We want our autonomous agents to act on our behalf, so we must issue an instruction to them. Agents, as constructed by their designers, may be more or less general. For example consider a "CD buying agent" and a "small-ads selling agent". Intuitively, if we want the agent to appear useful and intelligent, it should have some model of the domain of discourse (e.g. "CD's" and "trading") to facilitate a useful dialogue with the user. This model may be explicit, or implicitly encoded into the agent's implementation and interface. If the

model is too weak, the agent provides too little support to be useful. *In extremis*, a programming language such as Java contains a complete model of agent operation; it's just up to the user to fill in the details! Conversely, too strong a model and the agent ceases to be a general enough to warrant the ascription of agenthood. The Jango¹ agent searches World Wide Web online stores to perform shopping comparisons on certain product categories. It is a useful tool, but the claims of its designers notwithstanding, is it perceived by its users as an agent or, because of its specificity, a good search engine?

Sally wants to encode a task, and then assign responsibility for that task to an agent. She wants enough support from her agent to express detailed domain considerations like quality of service, without having to introduce to the system a new explanation of what that means. But she doesn't just want an opaque token, since the agent is being asked to negotiate about degree of QoS, and so must understand² it well enough to do that. The remaining choice seems to be that the designers of the agent anticipate all possible aspects of the dialogue, their interrelationships, and allow sufficient expressiveness in the dialogue for Sally to encode her task effectively. Such a demand is highly unrealistic. Note also that the problem is compounded by the realisation that Sally's agent will be negotiating with other agents about the task, so their models must be compatible.

To realise the benefit of flexibility, the agents should be able to react appropriately to new situations. This necessarily involves recognising the new situation, and then knowing what to do, or, more precisely, what the user wants done. That knowledge will again depend on input from the user, but aside from the demands upon the representational power of the user-agent dialogue, it seems unreasonable to ask the user to provide

¹ In 1997, Jango was bought by Excite. Previously a personal desktop shopping agent, the technology now drives an Excite product search service (see: <http://www.jango.com>). Even as an agent installed on the user's computer, the functionality was arguably closer in experience to a good search engine.

² Unfortunately, it is impossible to avoid using anthropocentric terms when discussing agent cognition, even though it is hard to conceive of a sense in which the agents we can build today really "understand" the symbols they are given.

information for all contingencies at the beginning of the task.

Delegating task and authority

The previous section looked at the difficulties around *describing* tasks so that they can be delegated to agents. The *act of delegating* itself can be problematic too. As often presented in the literature, delegating to agents requires that the agent accepts an (under-specified) goal from the user, thereby moving away from a metaphor of *direct manipulation* [10]. The agent will then, autonomously, sense its environment and act to achieve the user's goal.

This presents several problems. It asks the user to translate their conceptualisation of the task from an action oriented view to a goal oriented view, which, in a given task domain, may not be a natural thing to do, or at least do well. In [11], Milewski and Lewis report that delegation may not be an easy or natural behaviour for humans in general. One of the conclusions they draw is that powerful task monitoring and control mechanisms may be needed to reassure delegators and generate trust. Such mechanisms, of course, compromise the agent's autonomy.

Even more difficult to specify may be the degree to which the agent has the authority to act for the user, especially if the agent has been given the task of negotiating. The Adept system [12] uses a powerful task description language to scope authority in *service level agreements*. These however, are written at design time by programmers; it is not clear what opportunity the user has to vary agents' authority.

The question of authority becomes very important when the contracts being negotiated have legal force, such as the agreement between buyer and seller. If the user doesn't like the agent's deal after it has been struck, can they countermand the agreement (the *non-repudiation* problem)? If the party to an agreement is an agent, whose user then reneges on the deal (e.g. by not delivering the sold good), who is at fault? This is the question of *reciprocity*. The issue of how to deal with rogue agents (or users) perhaps strays away from the human-agent interface, towards the social and organisational context, but must still be answered.

Note that not all researchers agree that delegation is a valuable or achievable goal. See, for example, Shneiderman [13] for a counter-view. Erikson

[14] raises the point that, if the user does not understand exactly what a system did in response to a vague goal, it may be hard for them to predict how to correct the instruction in the case that they don't get the desired result.

Sharing context

One reason that both task description and delegation are hard relates to contextual or background information. Human beings draw upon a lot of contextual information to make their decisions, both to characterise a task, and to decide which actions to take to further a goal. Effective agents, making similar decisions on the user's behalf, must possess some of that context.

There are at least two parts to this contextual information. Part of it is the domain and task specific concepts that underlie the human-agent dialogue. This underlying model is often relatively stable, and could be encoded into the agent and its interface at design time. The other part of the context varies much more dynamically, and represents the influences on the user's decisions at a given time.

For example, some standard commodity items, such as electronic components, are *substitutable*. Certain Intel CPUs can be used in place of certain Cyrix CPUs, and vice versa. Sometimes, a purchasing agent would be acting correctly to recognise a substitution choice if it made for a better deal. The user asks "I need 350 Intel Pentiums by Thursday", to which the reply is "That can't be done, but these other chips could ship today". At other times, that solution may not be acceptable to the user because of the user's employer's product support policy. It doesn't sound too implausible, as presented above, that the user should review and authorise the agent's choice. But in a thriving electronic parts market, there could be hundreds or thousands of such decisions the agent could make. Even if the user retains the final say, the agent needs the user's context to at least rank the recommended solutions for presentation.

A recent project in our laboratory looked at the use of agents to manage telephone and other communications. Individual users could express preference rules to their agent, of the form "if I'm busy, route non-urgent incoming calls directly to voicemail". Here the context information changes very rapidly. Some sense of busyness can be captured from environmental sensing: is there

more than one person in the same location as the user (implying a meeting), or is the user actively using Microsoft Word? But ultimately, busyness is a property of the user's context determined *by* the user, and consequently difficult to obtain in a meaningful or convenient way.

The user's conceptual model of the agent

What do ordinary users think agents are? What do we, as researchers, want them to think? Various models or positions have been proposed for the latter question, some perhaps without due consideration to the user's perspective. Agents could be seen as general purpose aides, butlers, secretaries, and so, using metaphors of helper roles in human society. Agents can be specialists or generalists, single or aggregated, "owned" by the user or the system. More subtly perhaps, some agents have anthropomorphic appearance or use first-person dialogue ("I've found a great deal for you"), as though there is a putative *someone* the user is interacting with, whose words are being reported through the interface.

On the one hand, such interface trickery can be very useful: human beings have been evolving and developing communications behaviours for a long time. Tapping into the deep-rooted motivations underlying these behaviours may be a good way to open computer usage to people who don't like the way today's computer interfaces behave. Research by Nass and Reeves [15], shows that it's very easy to expose and exploit human social reactions to technological devices. King and Ohya [16] report that test subjects rated anthropomorphic faces more "agent-like", and a blinking face as significantly more intelligent than other visual forms.

On the other hand, raising users' expectations that they are interacting with a genuine social partner seems doomed to cause disappointment when the limits of the technology are encountered. And when that happens, will users reject the technology more strongly for having failed their elevated expectations? When Microsoft Word crashes, it is annoying. Will I feel worse when my trusted agent fails at its task?

Clearly, our users will have to have a different model of the agent than they do today of, say, a standard PC software package. Without such a model, it seems unlikely they will fully be able to exploit the potential values that agents offer.

It is tempting to try to characterise agent behaviour with some informal notion of "do-what-I-would-have-done" semantics, so that the agent becomes simplistically an extension of the user. Given the problems inherent in this view, it seems that the opportunity is there for us to develop new, more sophisticated models of the relationship between user and agent.

Dialogue issues

Human-computer dialogues today tend to fall into one of two categories. The largest category has the user owning the initiative in the dialogue: issuing commands and queries, and the computer responding. Alternatively, for example in tutoring systems, the computer controls the structure and direction of the dialogue.

Example 4, above, illustrates one way in which the dialogue between human and agent will exhibit *mixed-initiative* [17]. Sometimes the user has control of the exchange, and sometimes the computer. Change of initiative is inevitable when agents are acting autonomously. One potential problem here is that a sense of control is long established as a principle for good HCI design.

Beyond the questions of initiative, other properties of human-agent dialogues need to be addressed. Naturalistic styles, such as parsing typed or spoken natural language help to support the impression (or illusion) of interacting with an intelligent, social partner. As long as the dialogue is presented in a non-naturalistic style, users may not be persuadable that they are interacting with an agent *per se*. Whether or not this is a good thing has yet to be established. For example, MIT's Kasbah bargaining agent system [18] asks the user to select which kind of agent (simple or customisable) they want to sell a good on their behalf. Once selected, however, the agent is instructed through a standard form-filling interface. The sense of a social partner or personality is very limited.

Another interesting challenge is the use of emotional tone in dialogue. The KimSac project [19] provides agents in an information kiosk to advise potential claimants whether they are eligible to claim state benefits. Given the likely emotional state of at least some potential claimants, should the agents adopt a different tone

when delivering bad news, e.g. that the user is not eligible³?

Finally, it is very likely that acceptable and effective dialogue structures are strongly influenced by culture. Thus the problem of internationalising agent software goes beyond the already difficult problems of language, character sets and layout: the very basis of the social interface may not be easily applied in different cultures [20].

Summary of Issues

Through properties such as autonomy and sociability, and our expectations of what value these can offer the user, new problems, indeed new classes of problem, arise for the agent interface designer. Doubtless agents are already being created that don't address these problems, and yet provide novel, useful functionality. I argue that we will not reap the full potential of agent technology without addressing these issues systematically. It is clear that the issues we are facing cannot be solved by computer scientists alone. Instead, a multi-disciplinary approach including contributions from psychology, ethnography, linguistics, game theory, graphic design and probably several others will be needed.

The next section briefly introduces a sampling of research work already underway that will contribute to a resolution of the issues in the agent interface. The work reported here primarily addresses individual technologies or approaches in the human-agent interface.

A brief survey of work in progress

Mixed initiative dialogue was the subject of an AAAI Spring Symposium in 1997 [21]. Some 31 papers were presented on various aspects of this topic. Several papers wrote about implemented or under-development prototypes of systems that showed flexibility in dialogue initiative. Other papers explored the dimensions of the problem, highlighting different aspects of the single term "initiative". In particular, [17] highlights control of the *choice of task*, *choice of speaker* and *choice of outcome* as the main components of initiative.

³ This idea is from a personal conversation with Prof. E. Mamdani of Imperial College, who works on KimSac.

Other papers analysed the cues in natural dialogue that indicate when initiative is being transferred.

Socially Intelligent Agents was the title of another interesting AAAI Fall Symposium in 1997 [22]. This diverse collection of papers includes contributions on agent personalities and emotions, analyses of social behaviour, and agent societies. Here, social intelligence is explored both as a principle of agent interface design, and as a principle of inter-agent interaction. In [23], Dryer examines the use of agent technology to provide an assistant to users of IBM's OS/2 operating system, exploring, in particular, the personality dimension of such agents. A strength of this work is that laboratory experiments were employed to test users' reactions to different personalities. Other commentators have noted that more such work is needed (e.g. [13]). In [24], Dautenhahn surveys the elements of social intelligence and motivations for social agents. Cañamero and Van de Welde [25] discuss agents using emotional states to ground agent attitudes in dialogue. They refer to a physiological model of emotions developed to give a practical basis to developing agents with attitude. Castelfranchi et al [26] motivate and examine agent personalities, identifying certain traits and attitudes that can be used to explain or construct an artificial personality.

Many researchers are exploring the possibilities (and difficulties) of anthropomorphic interfaces. Interfaces that present some aspect of the system through a graphically rendered human face form a subset of so-called *character-based interfaces*⁴. Some early character-based interfaces are available in products (e.g. Microsoft "Bob" [27]), but these characters are not complex or complete enough to make plausible agents. Indeed, anecdotal evidence in our lab suggests that the majority of (though not all) users find the Microsoft Office '97 assistant, another animated character, irritating and of little value. Sproull et al [28] report a detailed study of the reactions of a group of test subjects to an animated face display. Differences in behaviour were observed, for example the subjects were more aroused and presented themselves in a more positive light than

⁴ Those of us who started our computing careers struggling with "character interfaces" on glass teletypes attached to mainframes may find this designation slightly ironic.

subjects seeing the same information in text form. The authors conclude that more studies should be conducted to further analyse the factors affecting the responses they observed.

Several researchers have proposed machine learning as a solution to the problems of instructing an agent without imposing too great a burden on the user. For example in [29], Maes suggests four ways in which a learning agent can acquire the knowledge it needs: by “looking over the shoulder” of its user, direct and indirect feedback on the agent’s actions, learning from examples provided by the user and expertise sharing with other agents. Such approaches seem promising for capturing the user’s context and other background information, but in my view do not yet address the problem of fully characterising a task that a user wishes to delegate to an agent. In Maes’ work, agents can also anticipate user needs and proactively suggest, or take, an action. However, given users’ needs to understand and control what the agent is doing, the supplied examples (e-mail and news filtering) suggest that the net gain to the user may be small once the agent has communicated its suggested intent for review. This is then another instance requiring detailed study of actual user reactions to understand whether the anticipated benefit is realisable.

Successful delegation to an autonomous agent will require trust on behalf of the user. Hall [30] reports an approach to enhancing an agent’s trustworthiness through an *incremental validation*. This approach comprises four elements: defining task behaviour, proof that task behaviour is in accordance with user desires, guarded execution of incompletely trusted behaviours and incremental trust growth based on scenario generalisation. The scale of the problem is perhaps indicated by the scope of the proposed mechanism to address it. Sycara et al [31] report on the impact of (experimentally determined) errors made by an agent on the trust placed in that agent by a user in a high-pressure task environment.

Milewski and Lewis [11] investigate delegation to agents as a general concern, reviewing social science literature on delegation as well as agent-based approaches. They make a number of design recommendations for agent systems, including some initial predictors of when delegation may be a poor model. For example, in non-stable environments, people tend to underestimate the

effect of the environment, and attribute failures to the agent (the *fundamental attribution error*). A positive suggestion is to keep agents small and well circumscribed in their abilities, to aid users to understand and predict the agent’s behaviour. Perhaps the most interesting question here is whether “delegation”, as we know it from the human context, is really the right metaphor when we consider agents.

Conclusions

The original motivation for this work arose from discussions within our team where we found it relatively easy to construct scenarios in which the potential benefits of autonomous agents were outweighed by apparent problems in the human-agent interface. Since then, it has become clear that other researchers and observers are similarly concerned, or are working on solutions to aspects of the puzzle. A recent ACTS report concluded:

“There are many more such issues which need urgent discussion and resolution. The technology push to the use of agents and brokers will not stop, no matter what technical obstacles to performance may exist. In the history of man, there has never been a single case of technology push being resisted by concerns for the human implications.” [32]

Two things seem imperative for agent based research as it matures. First, the impact of agent technologies on their human users needs much more attention. Second, the claims of radical improvements in the effectiveness of computer systems based on agents are plausible, but they need verification and refinement by user trials and laboratory experiments. Such work must be treated with equal priority to the excellent work underway in the underlying technologies.

Work in HP Laboratories

Our own response to this challenge has been to commence a research programme into the reactions of human users to agent interfaces. Our approach will be experimental, presenting groups of users of differing backgrounds with a variety of agent and non-agent interfaces under controlled task conditions. Of particular interest to us is the mental or conceptual model that users form to explain and predict the actions of intelligent autonomous agents, especially ones performing

important personal or business tasks on the user's behalf. A principal goal of this work will be to produce a more compelling statement of the value of agent interfaces than the informal anecdotal evidence that characterises many reports.

Acknowledgements

I would like to thank Janet Bruten for reviewing this paper, and Sharon Dickinson, Julie Parker and Maarten van Tol for helpful comments.

References

1. Norman, D.: "Why interfaces don't work" in *The Art of Human-Computer Interface Design*" Laurel D (ed.), Addison-Wesley 1990.
2. Miceli M., Cesta A. & Rizzo P "Distributed artificial intelligence from a socio-cognitive standpoint: looking at reasons for interaction" *Artificial Intelligence and Society* no. 9, 1995 pp 287 - 320
3. Gibson, W. "Neuromancer" Harper-Collins, 1995
4. Negroponte, N., "Being digital" Vintage Books 1996
5. Gates, B "The Road Ahead" Penguin 1995
6. Franklin S. & Graesser A: "Is it an agent or just a program?: a taxonomy for autonomous agents" in *Intelligent Agents III* (eds. Muller JP, Wooldridge M & Jennings N) *Lecture Notes in Artificial Intelligence*, Volume 1193 Springer Verlag, 1996, pp 21 – 35.
7. Jennings, N. & Wooldridge M: "Agent Technology: Foundations, Applications, and Markets", Springer 1998.
8. Zacharia, G., Moukas, A., Guttman R. & Maes P. "An agent system for comparative shopping at the point of sale" *to appear in Proceedings of the European Conference on Multimedia and e-commerce*, Bordeaux France, September 1998. Also available from <http://ecommerce.media.mit.edu/>
9. Wooldridge MJ & Jennings, NR: "Intelligent Agents: Theory and Practice", *The Knowledge Engineering Review*, Vol. 10 (2), pp. 115-152, 1995 Also available from: <http://gryphon.elec.qmw.ac.uk/dai/pubs/#list>
10. Negroponte, N: "Agents: from Direct Manipulation to Delegation" in J. Bradshaw (ed) *Software Agents*, AAAI Press / The MIT Press, 1997, pp57 – 66.
11. Milewski AE & Lewis S: "Delegating to software agents" *International Journal of Man-Machine Studies* no 46, 1997, pp485 – 500.
12. Norman TJ., Jennings N., Faratin P. & Mamdani EH: "Designing and implementing a multi-agent architecture for business process management", in *Intelligent Agents III* (eds. J. P. Mueller, M. J. Wooldridge and N. R. Jennings) *Lecture Notes in Artificial Intelligence*, Volume 1193, Springer Verlag, 1996 pp261 – 275. Also available from: <http://gryphon.elec.qmw.ac.uk/dai/pubs/#list>
13. Shneiderman, B: "Direct manipulation versus agents: paths to predictable, controllable and comprehensible interfaces" in J. Bradshaw (ed) *Software Agents*, AAAI Press / The MIT Press, 1997, pp97 – 106.
14. Erikson T: "Designing agents as if people mattered", in J. Bradshaw (ed) *Software Agents*, AAAI Press / The MIT Press, 1997, pp79 – 96.
15. Nass C., Reeves, B: "The Media Equation", *CSLI Lecture Notes*, 1996
16. King, W.J. & Ohya J: "The representation of agents: a study of phenomena in virtual environments", *IEEE International Workshop on Robot and Human Communication*, 1995, pp199 – 204.
17. Novick DG., & Sutton S: "Mixed initiative dialogue" in [21] pp 114 – 116.
18. Project Kasbah. See <http://ecommerce.media.mit.edu/Kasbah/index.html>

19. Project KimSac: “Kiosk-based Integrated Multimedia Service Access for Citizens” see <http://www.broadcom.ie/Kimsac/>
20. O’Neill-Brown P: “Setting the stage for the culturally adaptive agent”, in [22], pp 93 – 97.
21. Haller, S & McRoy S (co-chairs) “Computational models for mixed initiative interaction. Proceedings of the 1997 AAAI Symposium”. Technical Report SS-97-04, AAAI Press, 1997
22. Dautenhahn, K (chair), Masthoff J & Numaoka C (co-chairs): “Socially intelligent agents. Papers from the 1997 AAAI Fall Symposium”. Technical report FS-97-02, AAAI Press, 1997.
23. Dryer, DC: “Ghosts in the machine: personalities for socially adroit software agents”, in [22], pp 31 – 36.
24. Dautenhahn K: “Ants don’t have friends – thoughts on socially intelligent agents”, in [22] pp22 – 27.
25. Cañamero, D & Van der Welde, W: “Socially emotional: using emotions to ground social interaction”, in [22], pp10 – 15.
26. Castelfranchi C, de Rosis F & Falcone R: “Social attitudes and personalities in agents”, in [22] pp16 – 21.
27. Winograd T: “Microsoft Bob” in Winograd T (ed) *Bringing design to software*, Addison-Wesley 1996, pp 146 – 150.
28. Sproull L, Subramani M, Kiesler S, Walker JH & Waters K: “When the interface is a face”, *Human-Computer Interaction*, volume 11 no. 2, pp 97 – 124, 1996.
29. Maes, P. “Agents that reduce work and information overload” *Communications of the ACM* vol. 37 no. 7, pp 30 – 40, July 1994
30. Hall, RJ. “Trusting your assistant” *Proceedings of KBSE ’96*, pp 42 – 51
31. Sycara, KP., Lewis M, Lenox T, Roberts L. “Calibrating trust to integrate intelligent agents into human teams” *Proceedings of the 31st Annual Hawaii International Conference on System Sciences*, pp 263 – 268, 1998
32. Clarke A., Carter C., Kaasinen E., Kantola K., Chester J., Concejero P., Buss R., Muehlbach L., & Ruschin D. “Further HF Issues in ACTS” available from <http://atwww.hhi.de/draft/USINACTSnorm/start2n.htm>