Autonomous agent as helper – Helpful or Annoying?

Paul Rudman and Mary Zajicek
Oxford Brookes University
prudman, mzajicek @brookes.ac.uk

Abstract

This paper describes an investigation into the feasibility of agent-provided assistance for two specific situations, (1) when a previously-interrupted task is resumed, offer to open previously-used contextual documents and (2) when a non-optimal task is begun, suggest a more suitable alternative task. A paper-based task places participants in the situations described. The concern is to maximise the balance between helpfulness and annoyance.

The results are discussed in terms of timing of interruption and social effects. Overall, the agent needs to take account of the human's likely feelings towards any intervention; interventions must be both useful and perceived to be useful.

1. Introduction

The last decade has seen the advent of the autonomous agent as helper. The principle is sound—have the agent monitor the user and make suggestions when required. Simple implementations, such as "auto-complete" have proved successful. However, it soon became clear that this approach was not without its pitfalls. Offering help to the user means interacting with that user. For the user, this interaction takes time and effort away from the actual task at hand. As the complexity of the agent increased, and it attempted to be more "perceptive", so too did the opportunity to misjudge the user's needs, creating interruptions with no helpful purpose.

Perhaps the most notorious problematic implementation of an autonomous agent as helper was the Microsoft Office "Clippit" office assistant. Here, while the agent had access to an extensive database of how help information, it had relatively little ability to assess the user's current activity—their tasks, intentions and goals—or their immediate needs—if any. For example, typing an address followed by "Dear" on the word processor would cause the agent to appear on-screen as an animated character, with the text "It looks like you're writing a letter. Would you like help?". The result was a system that was generally disliked [1].

Early implementations of the autonomous agent as helper, then, have suffered from an inability to assess the user's activity and from an inaccurate user model with which to infer needs from the user's activity. The fundamental problem is that any form of intervention will be, by definition, an interruption to the task at hand—attending to the intervention will divert resources away from the original task. It is, of course, the agent's intention (or rather, the designers of that agent) that the intervention will be of sufficient assistance to outweigh the disadvantage of also being an interruption. However, what the agent intends and what the user perceives may not be the same thing.

The investigation described here is a part of the European Union's Sixth Framework Programme for Technology-enhanced Learning. The "AtGentive" project as a whole looks at the use of embodied agents to assist adult and child learners in allocating their attention, both at the general level—planning and the selection of tasks—and the specific—maintaining attention on a task. The project looks at a number of potential interventions. Reported here is an early part of an investigation into the feasibility of agent-provided assistance. This investigation considers two forms of intervention intended to assist adults in directing their attention:

1. when a previously-interrupted task is resumed, offer to open previously-used contextual documents
2. when a non-optimal task is begun, suggest a more suitable alternative task

Our concern here, however, is to consider the more general effects of interrupting users in order to offer help and in particular to consider the balance between helpfulness and annoyance.

Proceedings of the IEEE/WIC/ACM International
Conference on Intelligent Agent Technology (IAT'06)
0-7695-2748-5/06 $20.00 © 2006 IEEE

2. Related work

Agents are often designed—and used—with the idea of a human PA (Personal Assistant) in mind who "looks over the shoulder" [2] of the user, looking for opportunities to render assistance. Indeed, the idea of having a knowledgeable and (apparently) intelligent entity metaphorically sitting beside the user, offering assistance where appropriate, is very easy to anthropomorphise. It is a tendency that has contributed to the use of artificial animated characters as part of the agent's interface, despite the relatively sparse empirical evidence for their effectiveness [3].

The process of interacting with computers is both logical and emotional, and this affective aspect needs to be considered [4] as it can bring problems of its own. The generation of such problems have been characterised as "social" effects [5]—the user treats the agent as they would a person, responding to social ineptness with characteristic negative emotions, such as annoyance, dislike and a preference to avoid future encounters with the transgressor. Such effects have contributed to the demise of earlier "helpful" agents (Microsoft's "Clippit" system being one high-profile example [6]).

One important approach to avoiding these problems is to record the agent's interactions with the user and from this record try to predict the most appropriate interventions. For example, Schiaffino and Amandi [7] produced a system for learning to assess the extent of a user's need for help. A decision is made between preferring only a warning about a situation or being amenable to the offer of suggestions; an important third possibility is that the user may wish the agent to remain silent. Further work [8] has attempted to define "eight major issues interface agents have to achieve". These focus on issues such as the agent—type and behaviour—the user—requirements and tolerance—and control of the agent through feedback.

One problem with an agent relying on being able to learn the user's preferences is the time taken for an agent to learn what is required of it. During this time, there is the opportunity for "social effects" to take place—the user may "learn" that the agent is not well adapted to his or her needs, and discontinue its use before it has adapted to that user's preferences.

Adamczyk and Bailey [9] looked at "the effects of interruption at different moments within task execution". Their work showed that the negative impact of interruptions can be significantly mitigated by careful timing of interruptions—delaying them as necessary until the user is less busy, and therefore less affected by interruptions. Bailey and Konstan [10] looked further at the timing of interruptions, finding that users were more affected by interruptions when deeply engaged in the task. Conversely, they were less affected if the interruptions came at task boundaries.

While perhaps less effective at predicting the user's needs in the longer term than learning algorithms, this approach provides the opportunity to enhance the human-agent interaction from the start. The work described here looks at the question of maximising the helpfulness : annoyance ratio for the two specific circumstances described above, with particular emphasis on the timing of interruptions.

3. Method

The investigation described here placed participants in situations where each of these two interventions occurred. Afterwards, they filled in a questionnaire and took part in a short interview, both to elicit their feelings and opinions about the interventions. The purpose of the investigation was to look for potential user-related problems with these specific interventions, so that such problems can be circumvented or minimised as far as possible in any future agent implementation.

The investigation was conducted using low fidelity prototyping tools. The intention was to ensure that any problems found by the "users" were not created by the software interface used by the investigation, rather than the task situation itself. Therefore participants were given a "pen and paper" task, during which they would be put in each of the two situations under investigation.

The domain of herbs and their purported medicinal values was chosen, as this comprises a large amount of well documented and inter-related information, and has been used previously by one of the researchers [11].

3.1. Materials

The chosen domain comprised general information on 43 herbs, including their leaf and flower descriptions and reputed medicinal uses. This information was taken from a combination of two encyclopaedias of herbs [12, 13].

- 43 herb index cards. (The books contain one page per herb; each page may be reduced to an A5 sized card.)
- three A4 sheets of "A-Z of herbal treatments" listing a selection of ailments and herbs traditionally used to treat them
- an A6 piece of coloured paper
- a set of question—one per A4 question sheet (one sheet has a choice of two questions)

3.2. Participants

13 participants were enlisted from within Oxford Brookes University, comprising administrative staff and Computer Science PhD students.

3.3. Procedure

The participants took part individually. The participant was initially given a set of herb cards, the index sheets and a pen. The investigation proceeded by giving out one question sheet at a time, interrupting the participant where necessary as described below.

3.3.1. Question one. The participant was given the A4 sheet containing question one (with a "box" in which to write the answer). The purpose of this question was to start the trial and to help the participant familiarise them self with the documents. The question was:

> Find a herb that can help with headaches. Briefly describe its flowers (so that you can identify it in the herb garden).

The participant was allowed up to five minutes to complete the question.

3.3.2. Question two (part 1). This question investigated the first potential intervention point: restore task context (on resuming a task). The participant was given the A4 sheet containing question two. The participant was interrupted part-way through the question, with a record being kept of some of the contextual information. Later, the participant was asked to continue with the question:

> You are thinking ahead to winter. Choose one or more herbs that may be useful for winter ailments. Briefly describe their leaves (so that you can identify them in the herb garden).

As soon as the participant began either to look at the herb cards (not the index) or to write on the paper the task was swiftly ended. The participant was then given a small (A6) piece of coloured paper (coloured so that it may be easily referred to during the interview) and asked to write down any herbs they currently had in mind.

> "I would like to interrupt you at this point please and ask you to stop what you are doing. Could you please write on this paper the names of any herbs that you currently have in mind for this question"

The question sheet and the coloured paper were then collected.

3.3.3. Question three. This question investigated the second potential intervention point: propose an alternative task (more appropriate for the time available). The participant was given the A4 sheet containing question three. Two alternative questions were shown as follows:

> **Please answer ONE of the following questions:**
> Your friend has an elderly aunt. Choose one or more herbs that may be helpful to her. Briefly describe their leaves (so that you can identify them in the herb garden).
> **or**
> You are going on a long journey. Choose one or more herbs to take with you that may be helpful on the journey. Briefly describe their leaves (so that you can identify them in the herb garden before you go).

As soon as the participant looked at any one specific herb card or wrote a few words on the paper, the researcher gave the impression of reading the herb name or written words. The participant was then interrupted to be told that the "other" question was quicker.

> "Just to interrupt a moment, you may find the other question quicker in the time available, but it's up to you. Sorry to interrupt."

The participant was allowed up to five more minutes to complete the question.

3.3.4. Question two (part 2). The participant was asked to resume their answering of question two. A new question sheet was given out for this purpose, identical to the one given earlier (but with nothing the participant may have written on the earlier sheet).

The participant was given a short time to consider the problem and begin the task. This was until they did anything other than look at the question (i.e. pick up or look at any herb card, the index, or show signs of beginning to write).

The researcher then interrupted the participant by handing him/her the coloured paper they wrote on earlier, stating that "You may find this helpful".

3.3.5**. Questionnaire.** After the main part of the trial, the participants were asked to fill in a questionnaire and take part in a short interview. The questionnaire comprised 14 statements for each of which the participants were asked to rate their agreement with the statement using a seven point Likert scale (from "Completely Disagree" (1) through "Neither Agree nor Disagree" (4) to "Completely Agree" (7)). The

questionnaire used was part of the process of creating and validating a questionnaire for measuring the success of future agent interventions. For the purposes of this paper, only relevant questions are reported.

4. Results

We present a qualitative analysis of this initial study, based on selected Likert-scale questions and the post-test interviews. The Likert scale has been condensed here for reporting purposes to Disagree (1-3), Neither (4) and Agree (5-7).

It should be borne in mind that the sample size of 13 makes the results relevant as pointers to likely problems with agent interventions, rather than representative proportions of the population.

4.1. Intervention one: restore task context (on resuming a task)

The statement "I used the information about herbs I had previously been thinking about" yielded nine participants who agreed that they had used the information offered and three who disagreed (one selected Neither and is discounted here).

Of the nine participants who used the information, seven agreed with the statement "I was satisfied with being told which herbs I had previously been thinking about". Answers to their interview question "How did you feel about being told which cards you had previously been using?" were fairly plain, such as "Helpful", "Happy", "Didn't mind" etc. One person disagreed, but in the interview (s)he stated that (s)he had intended to restart the question in a different manner, but on receiving the information (the coloured paper) (s)he decided to use it ("Because you passed it to me, it changed what I was going to do"). One person neither agreed nor disagreed; in the interview that person said (s)he initially ignored it, but later used it to "cross-check" the answer.

Of the three who did not use the information, two did not agree that they were satisfied with the information. The third person stated neither—in the interview (s)he said that (s)he simply ignored the coloured paper.

Also of the nine who used the information, seven disagreed with the statement "I wish I had not been told which herbs I had previously been thinking about". The same person who had decided to use because it was offered agreed that (s)he wished she had not been told. One person stated neither agree nor disagree ((s)he hardly used the information).

Of the three who did not use the information, one agreed that (s)he wished (s)he had not been told. One person disagreed—that person was answering the question differently and ignored the coloured paper. The same person as before stated neither. These findings are summarised in Table 1 below.

	Used (9)	Did not use (3)
Satisfied-agree	7	0
Satisfied-neither	1	1
Satisfied-disagree	1	2
Wish not told-agree	1	1
Wish not told-neither	1	1
Wish not told-disagree	7	1

Table 1 - Summary of findings for Intervention one

4.2. Intervention two: propose alternative task (more time-appropriate)

Seven people agreed to the statement "I wish I had not been given the suggestion about changing questions". Three participants gave negative answers ("Cross", "Disconcerting" and "Frustrating") to the interview question of "How did you feel about being told that a certain question was quicker?". The other four participants all stated that they ignored the suggestion (no affect) as they had already started the question. Only one of the seven actually changed questions as a result of the suggestion (the person who found it disconcerting). This was also the only person of this seven who stated that they believed the information to be accurate.

Three people neither agreed nor disagreed with the statement. One person had not decided on a question at the point where the suggestion was given, so could not say how it affected his/her decision (but stated that they would have been confused had the suggestion come after making a decision). One person just ignored the suggestion (no affect). The other person began thinking out-loud, eventually convincing him/herself to change questions. These two stated that they believed the information to be accurate.

Three people disagreed with the statement that they wished they had not been given the suggestion. One person was happy to change questions. One person "didn't like it to start with" but later decided it was helpful and changed their question. The third person stated that they were "Confused" and did not change. All three stated that they believed the information to be accurate.

In total then, four of the participants were annoyed by the suggestion, with only one person happy to accept the advice and a further two spending some time

contemplating the change before making it. Six people believed the researcher (that the "other" question was quicker), four of whom changed their question. The other six did not believe and did not change. (One person is discounted as (s)he was given the suggestion before having made a decision). These findings are summarised in Table 2 below.

	Agreed (7)	Neither (3)	Disagreed (3)
Annoyance	3	0	1
Happy	0	0	1
Changed	1	1	1
Undecided	0	1	0
Believed	1	2	3

Table 2 - Summary of findings for Intervention two

5. Discussion

5.1. Intervention one: restore task context (on resuming a task)

The intention of this intervention was to help the participant restore the task's context by reminding him/her of information used previously in this context—in this case herb names.

Participants generally appreciated being given assistance in restarting this question, finding it helpful and being pleased that someone was trying to help. An interesting question that further research could investigate is whether this assistance would be appreciated to the same extent with repeated interventions of this nature over time.

Of particular interest are the four participants who found this information unhelpful or annoying, and the person who used the information only doing so because (s)he felt (s)he ought to.

It seems that there was a separate group with an identifiable reason for not wanting the "assistance". These participants, on seeing the question a second time (or possibly in the intervening time) decided to take a different approach to their answer. Thus, when the information on the coloured paper was made available to them it was no longer relevant. How they dealt with this varied, from disinterest to annoyance. It is possible that the offer of information was taken as a suggestion that their new direction was somehow being called into question.

Thus, timing is not a major issue for this intervention (although the sooner the information is given the more useful it is likely to be). Of greater relevance is the continuing usefulness of the stored information. Any agent offering this intervention needs to consider that the contextual information may be out of date, and that offering it may well cause a negative emotional response.

5.2. Intervention two: propose alternative task (more time-appropriate)

The intention for this intervention was to help the participant to be more productive overall by suggesting they change from a recently-begun task to one that will be more effective—in this case quicker.

Telling participants that the alternative question was (supposedly) quicker than the one they were attempting was generally seen as annoying and/or unhelpful. The reason given was consistent: once they had begun a task they were committed to that course of action. They did not want to change and restart in a different direction. Even when they believed that the advice was correct, and it would indeed be quicker at that point to change questions as suggested, they did not necessarily take the advice.

The investigation was deliberately set up to offer the suggestion very shortly after this decision. The reason for this was to move away from offering the suggestion at the same time as the question itself. Clearly, stating on the question sheet "Please note that question n is quicker to answer" would be different to receiving advice based upon the user's actions. We were interested in the latter, since it is people's actions on beginning a task that will trigger any future agent to intervene—it cannot be party to people's internal considerations as to which task to attempt next. The investigation therefore attempted to offer the suggestion at the point where their actions reveal their decision.

In one case this strategy failed, with the suggestion being given inadvertently before the participant had decided which question to answer. As expected, the person involved was not unhappy with the suggestion (but stated that (s)he would have felt confused if the suggestion had come after making a decision).

As one would expect, all participants who changed their question considered the information (that the "other" question would be quicker) to be accurate (i.e. they believed the researcher). What is interesting is that most (six out of eight) people who did *not* change their question said that they did not believe that the other question would be quicker. There was no evidence in the materials provided that either question would be quicker to complete and no obvious reason for this

being the case. Indeed, participants were told the "other" question would be quicker whichever question they attempted. (The need for this deception was explained to them after the trial.) Any decision to ignore the suggestion then was a subjective decision. It seems, in fact, to be a matter of belief—or otherwise—in the information given.

Thus, it seems that believability and related issues (trust, and likeability for example) are important attributes for any agent in this situation. This is especially relevant where an agent may be perceived as a "dumb" program, and may have greater difficulty than a human in being accepted.

In addition, giving details of why the suggestion has been made may help make the intervention less annoying, as it would assist the person to justify any change of decision, something that seemed important in this study.

As regards the timing of this intervention, there seems to be a critical point at which the effect of the suggestion changes; it is the point at which the participant makes a decision as to which question to answer. This is problematic to any agent that intends to offer advice based on the user's decision to start a task. If, for example, an agent waits until a document has opened and the first line has been read (which is likely as the agent software may take a moment to react) the user will feel that they have "begun" the task, and may not appreciate being told, for example, that reading an important email now would be better than starting this long document. Strategies that put the user more in control, such as offering to assist but without saying immediately what the assistance is, may be worth employing in this type of situation.

6. Conclusion

In both interventions trialled in this investigation, the source of negative feelings was similar. Where a participant's viewpoint was in some way called into question, without there appearing to be sufficient reason, the result was negative feelings, such as annoyance and frustration.

With intervention one (restore task context) timing was not a major issue (although the sooner the information is given the more useful it is). What is important is the possibility that the person may have changed the manner in which they intend to approach the task, thereby rendering the contextual information out of date. Offering this information may generate a negative emotional response, possibly based on the person not wanting to have the new approach undermined, rather than the information simply being unhelpful.

Results from intervention two (propose alternative task) show that suggesting to a person that they switch tasks, having just begun a task, will be difficult to achieve consistently without a negative emotional response. Timing is critical in that the suggestion is much more acceptable before the person feels they are committed to that task. However, this may be difficult to achieve in practice. An alternative strategy would be to maximise the believability of the person/agent making the suggestion as this appears to influence the person (related factors, such as trust and likeability may also be relevant).

It is clear, then, that in any human-agent interaction the agent needs to take account of the human's likely feelings towards any intervention. Simply giving information that "should" be helpful, in terms of task efficiency, speed, etc., is not sufficient. This study suggests that a software agent, at least in these situations, should be likable and offer advice that is timely and believable. Above all, it needs to take into account the possibility that the user may know best and not make suggestions that may be to the contrary without backing them up—something that humans have known for thousands of years.

7. References

[1] H. Schaumburg, "Computers As Tools Or As Social Actors? — The Users' Perspective On Anthropomorphic Agents," *International Journal of Cooperative Information Systems*, vol. 10, pp. 217-234, 2001.

[2] P. Maes, "Agents that Reduce Work and Information Overload," *Communications of the ACM*, vol. 37, pp. 30-40, 1987.

[3] D. M. Dehn and S. v. Mulken, "The impact of animated interface agents: a review of empirical research," *International Journal of Human-Computer Studies*, vol. 52, pp. 1-22, 2000.

[4] M. Muller, "Multiple Paradigms in Affective Computing," *Interacting with Computers*, vol. 16, pp. 759-768, 2004.

[5] C. Nass, J. Steuer, and E. R. Tauber, "Computers are Social Actors," presented at CHI 94, Boston, MA, 1994.

[6] B. Trott (CNN), "Microsoft's paper-clip assistant killed in Denver," Accessed 29 Sep 2006 http://www.cnn.com/TECH/computing/9810/16/clipdeath.idg.

[7] S. Schiaffino and A. Amandi, "Personalizing user-agent interaction," *Knowledge-Based Systems*, In Press, Corrected Proof, 2005.

[8] S. Schiaffino and A. Amandi, "User - interface agent interaction: personalization issues," *International Journal of Human-Computer Studies*, 60, pp. 129-148, 2004.

[9] P. D. Adamczyk and B. P. Bailey, "If not now, when?: the effects of interruption at different moments within task execution," presented at Proceedings of the SIGCHI conference on Human factors in computing systems, Vienna, Austria, 2004.

[10] B. P. Bailey and J. A. Konstan, "On the need for attention-aware systems: Measuring effects of interruption on task performance, error rate, and affective state," *Computers in Human Behavior*, vol. 22, pp. 685-708, 2006.

[11] P. D. Rudman, "Investigating Domain Information as Dynamic Support for the Learner during Spoken Conversations," Unpublished PhD thesis. Birmingham UK: University of Birmingham, 2005.

[12] L. Bremness, "DK Pocket Encyclopaedia, Herbs," 10 ed. London: Dorling Kindersley, 1990, pp. 240.

[13] A. Bonar, *Herbs: A Complete Guide to their Cultivation and Use*. London: Tiger Books International, 1992.