An empirical study in agent-based interface issues

Marcelo Armentano*, Daniela Godoy* and Analía Amandi

ISISTAN Research Institute, UNICEN University
Campus Universitario (CP 7000), Tandil, Bs. As., Argentina
*Also CONICET
{marmenta,dgodoy,amandi}@exa.unicen.edu.ar

Abstract Interface agents are computer programs that provide assistance to a user dealing with computer based applications. The introduction of agents to user interfaces caused the exploration of new metaphors to enhance user ability to directly manipulate interfaces. In this regard, mixed-initiative interaction refers to a flexible interaction strategy in which agents contribute with users by providing suitable information at the most appropriate time. Mixed-initiative approaches promise to dramatically enhance human-computer interaction by allowing agents to reassemble human assistants. In this paper we report a study on how the interaction metaphor can affect the user perception of agent capabilities and in turn the final agent success.

1 Introduction

Interface agents, also known as personal agents, are autonomous software entities that provide assistance to users dealing with computer-based applications, such as information filtering, meeting scheduling, entertainment selection and so forth. These agents act as human assistants, collaborating with the user in the same work environment and becoming more efficient as they learn about user interests, habits and preferences.

From the human-computer interaction perspective, the emergence of interface agents changed the traditional interaction metaphor of direct manipulation to a complementary style of interaction, which has been referred to as indirect manipulation or mixed-initiative interaction. Instead of user-initiated interaction via commands and/or direct manipulation, the user is engaged in a cooperative process in which both human and software agents initiate communication, monitor events and perform tasks [12].

Although an agent-based interface should support completing cooperating tasks based on the user delegation of tasks to agents, users may be unaware of the tasks agents are performing in order to make it possible for the overall system to provide its functionality. The same software can be presented to users with a direct manipulation interface where an agent is acting in background to fulfill agent-related application requirements or with a mixed-initiative interface where synthetic, personified or even animated agents interact with users.

Having an agent operating directly in the user interface rather than as a background process increases the extent to which the user perceives the software

as acting like an assistant [11]. This paper discusses an empirical study which was undertaken to investigate how the perception of agents in a mixed-initiative interaction interface influences users and how this can impact on the learning process. In pursuing this goal, we carried out an experiment with real users interacting with a personal agent by means of interfaces corresponding to direct manipulation and mixed initiative interaction metaphors.

This work is organized as follows. Section 2 discusses adaptive systems and their relation with agent-based interfaces and human-computer interaction issues. Section 3 describes the study we performed and summarizes the achieved results. Concluding remarks and future lines of research are given in Section 4.

2 Personal Agents as Adaptive Systems

Adaptive systems are characterized by the presence of explicit user models or profiles representing user knowledge, goals, interests and other features that enable the system to distinguish among various users [3]. An adaptive system collects data for the user model from various sources including implicit observation of user interactions and/or explicit asking users to provide feedback. Finally, user models are used to provide an adaptation effect, that is, to tailor interaction to different users in the same context.

Personal agents are adaptive systems in which the user model is used to make suggestions, correct misconceptions, and generally guide agent actions [12]. Even when these agents may serve a variety of purposes, all of them provide proactive support to users and operate to some degree autonomously. In consequence, agent interfaces are inherently delegative since users turn things over to their agents to be done rather than do them by themselves. Typically, agents assist users by making suggestions in a non-invasive fashion, whereas decisions are still under the control of users.

In a graphical user interface, two main approaches can be distinguished in relation with the presentation of agents. The first consists in adapting the interface of the base application and controlling the agent using direct manipulation. The second focuses on making the agent explicit with a representative figure, which can enhance the level of engagement of the user in a mixed initiative style of interaction.

In the first approach, the effects of adaptation can be visualized in a variety of ways. For example, consider an adaptive Web system [2], in which adaptation can be materialized by three main techniques:

- Adaptive content selection: when the user searches for relevant information, the system can adaptively select and prioritize the most relevant items, as in *Letizia* [10] and *PersonalSearcher* [7].
- Adaptive navigation support: when the user navigates from one item to another, the system can manipulate the links to provide adaptive navigation support (agents like *Syskill&Webert* [16] and *WebWatcher* [1] exemplify this case).

– Adaptive presentation: when the user gets to a particular page, the system can present its content adaptively, as in *NewsAgent* [4].

All referenced systems have the common feature of adapting themselves according to the user model. Presentation and interaction mechanisms in these agents are always interweaved with the base application and are mainly limited to highlighting links and presenting recommendation in a separate window.

In the second approach, the user both retains full control over the direct manipulable interface and interacts with agents acting autonomously to perform tasks on behalf of users. Although agents are intended to execute tasks in a relatively autonomous fashion, users can customize agent functionality as well as assign, suspend, resume and cancel delegated tasks.

When agents show themselves with an external figure, they are called synthetic, personified or anthropomorphic agents. This representation of agents create engaging environments for users by introducing lifelike characters into the interface. Furthermore, they provide the illusion of autonomous, animistic entities with human-like capacities. It is worth noticing that the mere fact of adding a face to the agent does not increase its acceptance. For example, personified agents like Microsoft's paper-clip assistant[1], that constantly interrupt users irrespective of their focus of attention have shown being inefficient [18] and users prefer not to use them. Examples of personified agents can be found in [13,5,15].

Personified agents have the ability of receiving orders from human users and answering in a personalized way. Moreover, these agents can pro-actively realize when they could help their users and either suggest an action or directly act on behalf of users in a given situation. These systems respond to a mixed-initiative interaction metaphor and consider the pair user-agent as a relation in which the control shifts between the two according to the situation, the shared knowledge and the user model [6]. The main idea of mixed initiative agents (either personified or not) is to develop an active, cooperative, and adaptive system, in which users retain control over agents. In many cases, the acceptance of the system strictly depends on the effective interweaving of the user-agent system as a whole. Examples of this kind of agent can be found in [8,17,9].

User-agent relationship is also conditioned by human computer interaction issues [14]. HCI people have criticized agent-based methodologies that seem to produce systems not easily accepted by users, mainly because of the loss of control caused by the autonomy of the agents. In this regard, some aspects to be considered to increase acceptability includes: giving users time for trusting their agents, allowing users to resume control, and leaving users the possibility of verifying what agents are doing.

[1] See "Microsoft's paper-clip assistant killed in Denver" at http://www.cnn.com/TECH/computing/9810/16/clipdeath.idg/; "Born again: Clippy pops up in Office XP" at http://zdnet.com.com/2100-11-267631.html and http://www.winsupersite.com/reviews/office10-whatsnew.asp

3 Empirical Study

In this study we analyzed a critical aspect of user-agent interaction such as the consequences of different manipulation metaphors in user perception of software. Particularly, the level of satisfaction of users regarding software and their preferences concerning manipulation metaphors were examined under both metaphors. Our work hypothesis is described in Section 3.1. As the main purpose of our study was to understand the effect of synthetic agents in user-interfaces, the agent employed to carry out the experiment is presented in Section 3.2. The interfaces users employed to access its functionality are detailed in Section 3.3 and Section 3.4. Finally, Section 3.5 reports the results we obtained.

3.1 Work Hypothesis

Interface agents can be considered as a supplementary class of software that acts on behalf of a user in an application domain. The hypothesis tested in this work is that agents acting in the interface influence user perception of software capabilities, which can in turn impact on software satisfaction. Moreover, we argue that such perception might influence the level of engagement of users and, thus, it can be a major differentiating factor in learning success.

In order to prove our hypothesis, we carried out an experiment with real users in which they were asked to interact with a same application accessible through two different interfaces in contrasting paradigms: a traditional direct-manipulation interface where users have complete control over interaction while the application access to agent-functionally; and a more flexible interaction strategy given by a mixed-initiative interface where the agent contributes to the task being performed by the user in the application.

Both interfaces were designed in the context of a Web search application where an agent, called $\mathcal{P}ersonal\mathcal{S}earcher$, is assisting users by tailoring search results according to their interests. In this case, the application functionality relates specifically to searching documents on the Web which includes posting queries to one or more search engines on the Web, fusing results and presenting a unified view to users. In this application, the agent functionality consists in learning user interests by observing user browsing behavior, building a user profile modeling these interests and assessing the relevance of Web pages according to the user interests in order to suggest the best pages to users.

Eighteen participants in this experiments were divided into two groups. The first group was provided with access to the application with a direct-manipulation interface, whereas the second group had access to the mixed-initiative interface of the same application in which users interact with a synthetic agent appearing in the screen. All participants used the different versions of $\mathcal{P}ersonal\mathcal{S}earcher$ during three weeks. At the end, participants filled a survey with a number of questions regarding their experience.

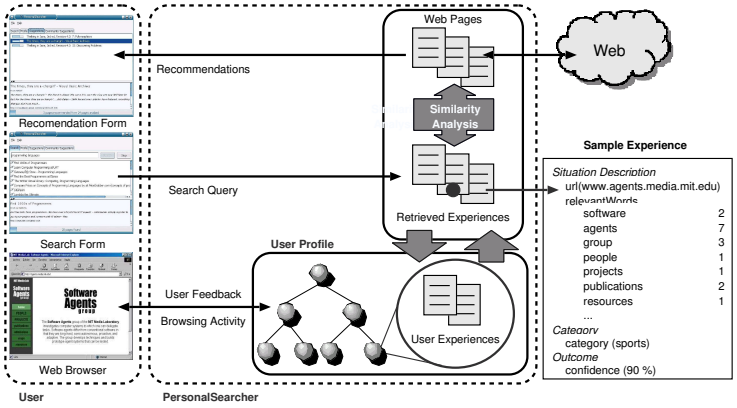

Figure 1. $\mathcal{P}ersonal\mathcal{S}earcher$ functionality

3.2 $\mathcal{P}ersonal\mathcal{S}earcher$

In order to provide personalized advice, personal agents rely on having some knowledge about users contained into user profiles i.e. models of user preferences and interests by which agents can assist user activities. $\mathcal{P}ersonal\mathcal{S}earcher$ [7] is an intelligent agent that learns about user interests by observing user behavior while users are carrying out regular activities on the Web. By a content-based analysis of the information extracted by observation, this agent is able to deduce the topics a user is interested in to create a user profile. Figure 1 depicts the complete agent functionality.

$\mathcal{P}ersonal\mathcal{S}earcher$ carries out a parallel search in the most popular search engines and filters the resultant list of pages according to profiles it builds based on the observation of user browsing on the Web. For each reading in the standard browser the agent observes a set of indicators in order to estimate the user interest in a given Web page. By means of this mechanism, each instance of this agent obtains pages relevant to a user without distracting him from his regular activities. Indicators include the time consumed in reading the page, the amount of scrolling and whether it was added to the list of bookmarks.

To build a user profile, pages considered interesting to the user are taken as input to a clustering algorithm, which output is a user interest hierarchy. This hierarchy models user interests in several domains (e.g. *sports*, *finances*, etc.) and at different level of abstraction (e.g. *tennis* and *football* within *sports*). The clustering algorithm allows agents to incrementally acquire profiles without an a priori knowledge of user interest categories, so that the learning process is completely unsupervised. In addition, it belongs to the conceptual clustering paradigm which includes not only clustering, but also characterization i.e. the formation of intentional concept descriptions for each extensionally defined cluster. In consequence, the clustering process offers comphensible profiles that can be understandable for both users and other agents.

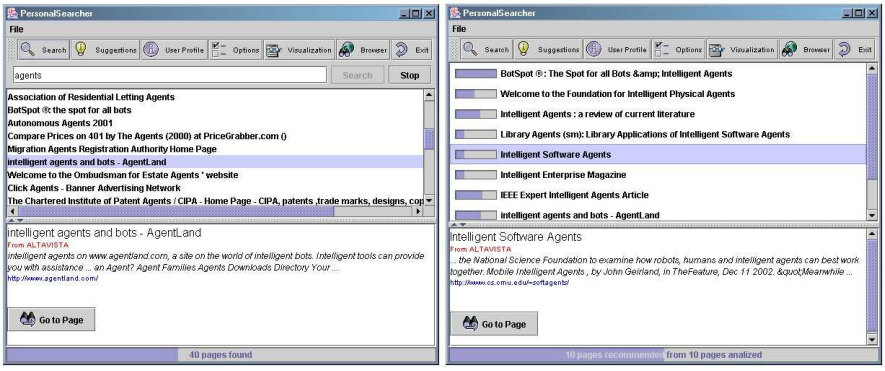

(a) Web search (b) Suggested Web pages

Figure 2. *PersonalSearcher* screenshots

Users interact with *PersonalSearcher* expressing their information needs by keywords. The agent posts these queries to the most popular search engines, receiving documents that cover a wide portion of the Web. *PersonalSearcher* determines the convenience of suggesting a Web page to the user by computing its relevance degree regarding to the user interest hierarchy. Those pages that surpass a user relevance threshold as regards to some category in the profile are sent back to the user as a result of his query. *PersonalSearcher* allows the user to customize the desired level of assistance at any moment, by adjusting the relevance threshold from the user interface. After the agent has presented suggestions, the user behavior is again observed to perform adaptations of the profile in terms of the user approval to the agent suggestions.

3.3 Direct Manipulation Interface

In the direct-manipulation interface of *PersonalSearcher* both search results of a Web search and suggested pages are presented to the user after posting a query without explicit requirement. Figure 2(a) illustrates how search results are presented, whereas Figure 2(b) depicts a list of Web pages the agent found relevant to the user among those in the search results. Both options are directly accessible from a tool-bar in the upper part of the interface. In this interface, users are not aware of the actions the agent is performing in order to learn about their interests and personalize the search results.

3.4 Mixed-Initiative Interaction Interface

In the mixed-initiative interaction interface, a synthetic agent is displayed in the interface so that suggestions are only accessible through user-agent itera-

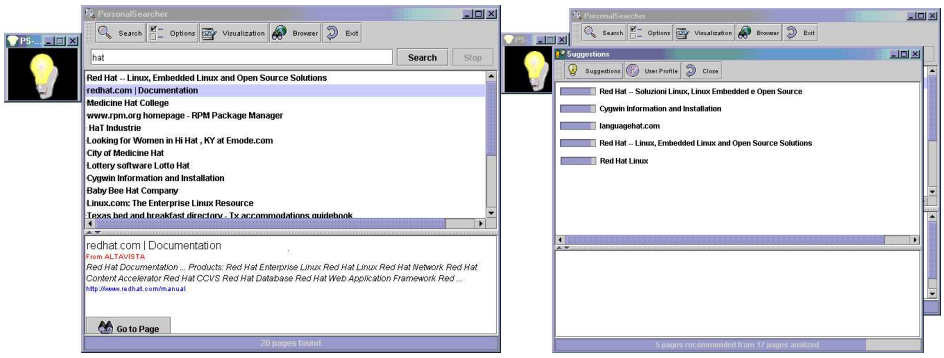

(a) Web search (b) Suggested Web pages

Figure 3. *PersonalSearcher* suggestions with agent screenshots

tion. Initially, users can only enter queries to obtain search functionally. A small window with a lamp image and soft colors which is shown in the upper corner of the application represents the agent. Users have to explicitly ask for agent functionality by clicking on the image of the agent to ask for suggestions.

In parallel to the search process, the agent is analyzing Web pages which are gathered from search engines and measuring their relevance according to the user model. As soon as the agent finds some interesting Web pages to recommend, it will try to attract user attention by lighting up the lamp. Figure 3(a) shows this situation. In contrast to the direct manipulation interface in which suggestions are given to users once the query is posted, in this interface users remain in control of the interaction. They can continue working with the searching application, ignoring the agent, or they can ask the agent to display its suggestions as it is shown in Figure 3(b).

We claim that the change in the manipulation metaphor can cause an increase in user expectation related to the interaction with the search assistant. High expectations are in most cases foster by anthropomorphic, personified and even animated agents which imitate human behaviors and, consequently, are expected to produce the same kind of results. An agent can also use suitable gestures or facial expressions to express its suggestions. However, in this experience the agent is limited to suggesting interesting pages to users so that it does not possess further ways of attracting user attention.

Additionally, human beings tend to pay attention to changes in the mobility, volume and contrast during interactions. Nevertheless, the techniques employed in this process have to be carefully designed so that we do not disturb the user when there is no interaction with the agent. This is the reason for choosing soft gray colors to represent the agent when it does not have suggestions to make.

3.5 Experiment Results

Participants in both groups responded a survey with questions about their experience with the application. Mainly, questions pursued the goal of assessing the level of satisfaction of users and also to analyze the preferences of users regarding to interaction metaphors.

Analyzing user answers, we found that most users using the mixed initiative iteration interface thought that the agent learned to make suggestions, at the most, quite well. However, part of the users using the direct-manipulation interface thought the application learned a lot and even too much. In both interfaces users can explore the user interest hierarchy which constitute their profiles in order to determine if the model matches their real interests. Figure 4(a) plots the survey results regarding this issue. Consistently, users of the mixed-initiative interaction interface also responded that suggestions they received were good but only to some extended, whereas an important proportion of user in the other group qualified suggestions as very good. Figure 4(b) shows the proportion of users according to their answers and the group they belong to.

In respect to suggestions, users were also asked whether they perceive any change on suggestion accuracy over time. Both groups shared the same learning and suggesting mechanisms so that the application is expected to behave similarly under similar circumstances. Figure 4(c) depicts answers regarding to this issue. Surprisingly, users dealing with the direct manipulation interface did not perceive changes on suggestion accuracy over time. Although it was clear from previous answers that these users considered they have received good suggestions, they also claim to have received similar suggestions from the beginning. In contrast, users using the mixed-initiative interaction interface reported to have noticed some improvement in agent suggestions.

A more general question regarding user satisfaction with software reveled that regardless the interface, all users showed a good level of acceptance of the received functionality. Figure 4(d) depicts the levels of satisfaction in both groups of users. This level of satisfaction can be attributed to the fact that users using the direct-manipulation interface believed to have received the same level of suggestions from the beginning, whereas users in the mixed-initiative interface noted an improvement in suggestions.

In conjunction, answers to the previous questions reveled a better reception of users using the direct-manipulation interface to agent suggestions. In the other group, the presence of the lamp denoting an assistant collaborating in Web page search seems to have increased the expectations of users about the kind of help this assistant could have provided. In fact, these users realized the agent improvement which indistinctly occurs in both groups since $\mathcal{P}ersonal\mathcal{S}earcher$ improves with experience. As a result of high expectations, answers from the mixed-initiative interaction group demonstrate a more critical position concerning agent suggestions.

In order to establish the preferences of users concerning both manipulation metaphors, distinctive questions were addressed to users in each group. Users who used the direct manipulation interface were asked whether they would

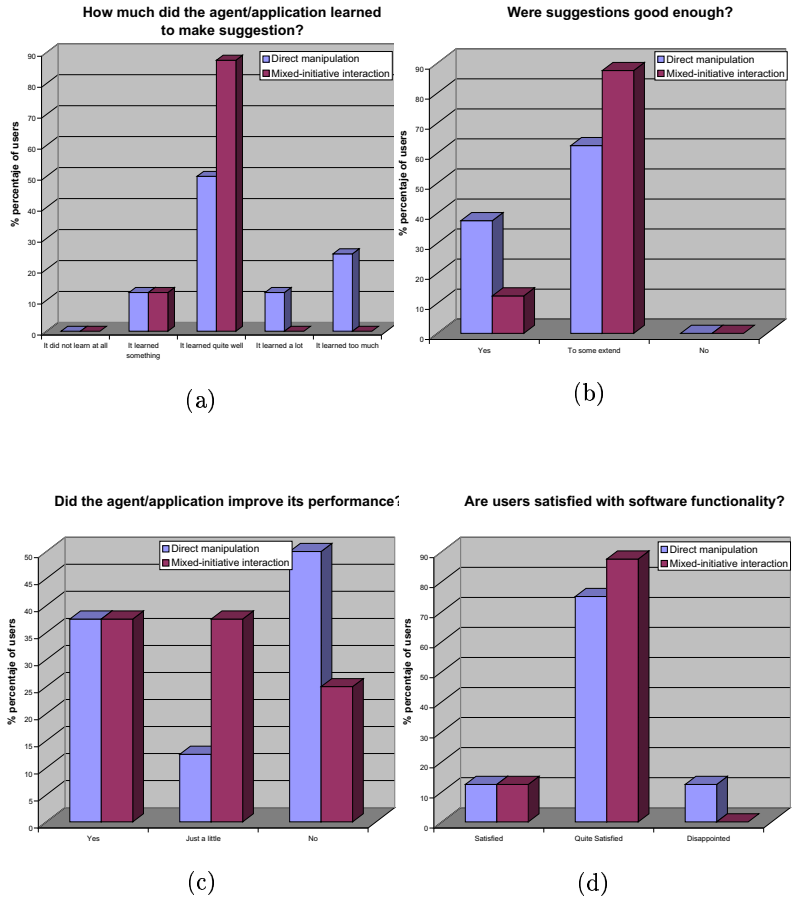

Figure 4. User perception of agent/application actions

have preferred not to receive the suggestions directly but ask for them. Figure 5(a) summarizes answers to this questions. Likewise, users who used the mixed-initiative interface were asked whether they would have preferred a direct access to suggestions without the agent intervention. Figure 5(b) summarizes answers given by the second group of users.

In both cases, more than a half users found no reason to change the interaction metaphors. However, the number of users preferring a change in the group of users using the the mixed-initiative interface was slightly higher than in the group of users using the direct manipulation interface. This difference can be explained by the level of satisfaction with suggestions of former group results lower than to the one of the second group.

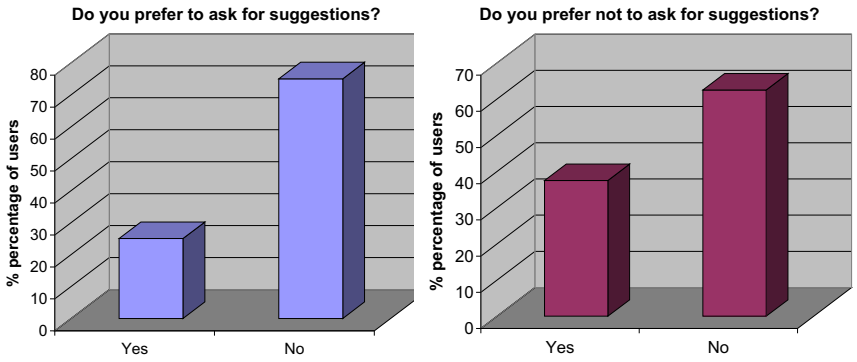

(a) (b)

Figure 5. Preferences about interaction metaphors

A further important factor in agent-user interaction and an essential component for agent learning is the relevance feedback user explicitly and/or implicitly provide about agent actions. A user gives explicit feedback by using one or more ordinal or qualitative scales, whereas implicit feedback is estimated by the agent according to observation of a group of interest indicators. Explicit feedback can be as simple as pressing a dislike/like button or more complex like scoring an item in a qualitative scale. It is generally assumed that implicit feedback, although less reliable and more difficult to obtain, does not burden the user with an additional cognitive load caused by the necessity of evaluating items.

In the survey we questioned users about the kind of feedback they preferred to give. In this case, results indicates that most of the users do not mind providing explicit feedback when they start their interaction with an agent, because they recognize the importance of training an agent in early stages. However, they decline to provide feedback when they have interacted with the agent for some

time. Other users do not mind giving explicit feedback provided that the feedback mechanisms are simple so that they do will not have to spend a lot of time and effort in this task. A considerable low number of users answered that they do not complain about giving feedback because they believe it is necessary for the agent to learn and improve its behavior. None of the users in this study refuses completely the possibility of giving feedback. Figure 6 depicts users reactions toward explicitly provide feedback.

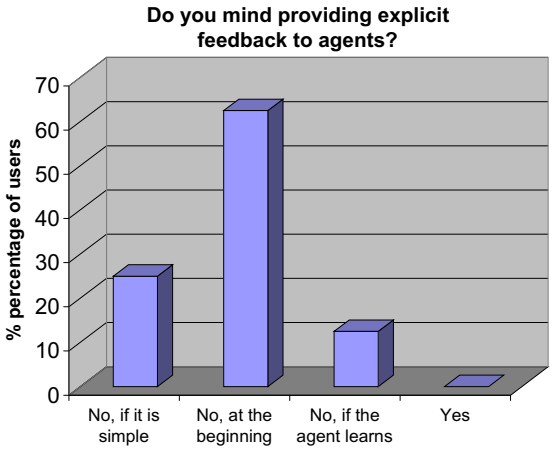

Figure 6. Opinions of users regarding feedback

4 Conclusions

In this paper we have discussed the impact the manipulation metaphors and, particularly, the interaction with synthetic agents in agent-based interfaces has on the user perception of software. The purpose of this study was to contribute to the better understanding of these issues in order to help designers in building agent-based interfaces. In the experiment we carried out with real users, two groups of users interacted with an application through interfaces presenting different interaction metaphors: direct manipulation and mixed-initiative interaction. Results revealed that the interaction with an agent generates a greater perception of the computer as a collaborative decision maker and increases the user demand for high quality results. Although the representation of agents in the user interface increases the level of engagement of users with the application, such representation can lead to misinterpretations about agent actions and expected results. This problem can be mitigated by the design of appropriate personification metaphors allowing users to assess the capabilities and limitations of agent-based systems.

References

1. R. Armstrong, D. Freitag, T. Joachims, and T. Mitchell. WebWatcher: A learning apprentice for the World Wide Web. In *Workshop on Information Gathering for Heterogeneous Distributed Environments*, pages 6–12. AAAI Spring Symposium Series, 1995.
2. P. Brusilovsky, A. Kobsa, and J. Vassileva. *Adaptive hypertext and hypermedia.* Kluwer Academic Publishers, 1998.
3. P. Brusilovsky and M. T. Maybury. From adaptive hypermedia to the adaptive Web. *Communications of the ACM*, 45(5):30–33, 2002.
4. D. Cordero, P. Roldan, S. Schiaffino, and A. Amandi. Intelligent agents generating personal newspapers. In *International Conference on Enterprise Information Systems, ICEIS'99*, pages 195–202, Setubal, Portugal, 1999.
5. A. Cypher. *Watch what I do – Programming by demonstration.* MIT Press, 1993.
6. D. D'Aloisi, A. Cesta, and R. Brancaleoni. Mixed-initiative aspects in an agent-based system. Technical report, Computational Models for Mixed Initiative Interaction, AAAI 1997 Spring Symposia Series, 1997.
7. D. Godoy and A. Amandi. PersonalSearcher: An intelligent agent for searching Web pages. In Springer-Verlag, editor, *Advances in Artificial Intelligence*, volume 1952 of *Lecture Notes in Artificial Intelligence*, pages 43–52, 2000.
8. D. Gruen, C. Sidner, C. Boettner, and C. Rich. A collaborative assistant for E-Mail. In *Proceedings of Human Factors in Computing Systems*, pages 196–197, 1999.
9. T. Koda and P. Maes. Agents with faces: The effects of personification of agents. In *Proceedings of Human-Computer Interaction*, pages 239–245, London, UK, 1996.
10. H. Lieberman. Letizia: An agent that assists Web browsing. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 924–929, 1995.
11. H. Lieberman. Autonomous interface agents. In *Proceedings of the ACM Conference on Computers and Human Interface*, pages 67–74, 1997.
12. P. Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7):31–40, 1994.
13. P. Maes and R. Kozierok. Learning interface agents. In MIT-Press/AAAI-Press, editor, *Proceedings of the 11th National Conference on Artificial Intelligence*, pages 459–465, Washington, DC, USA, 1993.
14. D. Norman. How might people interact with agents. *Communication of the ACM*, 37(7):68–71, 1994.
15. C. Okonkwa. Affective pedagogical agents and user persuasion. In *Proceedings of the 9th International Conference on Human-Computer Interaction*, pages 397–401, 2001.
16. M. Pazzani, J. Muramatsu, and D. Billsus. Syskill & Webert: Identifying interesting Web sites. In *Proceedings of the 13th National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference*, pages 54–61, Portland, Oregon, 1996. AAAI Press/MIT Press.
17. C. Rich and C. Sidner. COLLAGEN: A collaboration manager for software interface agents. *User Modeling and User-Adapted Inter-action*, 8(3/4), 1998.
18. S. Schiaffino and A. Amandi. User - interface agent interaction: personalization issues. *International Journal of Human-Computer Studies*, 60(1):129–148, 2004.