Available online at www.sciencedirect.com
ScienceDirect
Int. J. Human-Computer Studies 64 (2006) 1084–1098
International Journal of Human-Computer Studies
www.elsevier.com/locate/ijhcs

The use of interface agents for email notification in critical incidents

Alexander Serenko*

Faculty of Business Administration, Lakehead University, 955 Oliver Road, Thunder Bay, Ont., Canada P7B 5E1

Received 16 November 2005; accepted 15 June 2006
Communicated by C. Sierra
Available online 7 August 2006

Abstract

This study reports on several typical scenarios of the use of email notification interface agents under the influence of critical incidents. An interface agent is a reactive, collaborative, and autonomous visual computational system, which communicates directly with a person offering assistance and advice in performing computer-related tasks. The critical incident technique was employed to survey the actual users of an interface agent-based email notification application. Respondents were asked to provide the last most significant either positive or negative incident of the usage of interface agents in their email application. They were also asked to offer recommendations for designers and marketers of this technology. Sixty critical incidents were obtained and analysed. With regards to positive-outcome situations, one representative scenario was constructed. With respect to the negative-outcome events, three distinct scenarios were identified. Based on the critical incident technique, it is concluded that users acknowledge the quality of an agent when it acts reliably, an agent's intrusive behaviour results in an immediate agent usage termination, operability issues sometimes force people to reject the technology, and users attempt to preserve the employment of an agent under the negative impacts of external factors. A number of other practical recommendations for manufacturers and marketers are also outlined.

Keywords: Interface agents; Human–agent interaction; Email notification; Critical incident technique

1. Introduction

The purpose of this paper is to investigate how individuals use interface agents for email notification in their electronic mail applications under the influence of critical incidents and to offer recommendations for designers and marketers of this technology. An interface agent for email notification is a reactive and continuous visual computational system that keeps track of the current state of an email application and informs users about any changes, for example, the arrival of a new, important message. As any other type of an interface agent, it is in charge of interacting with the user; it should directly communicate with the person through the input and output of the user interface (Laurel, 1997; Lieberman and Selker, 2003; Detlor, 2004; Serenko et al., 2006). Email notification agents add graphics or animation to the interface, use speech input and output, and communicate via other sensory devices. Interface agents for email notification differ from autonomous agents or multi-agent systems (Jennings et al., 2000; Wooldridge, 2002) in that their goal is to serve as an interface to an already existing software system rather than automate complex or repetitive tasks working in the background and communicating with other agents.

One of the most salient reasons for incorporating interface agents in email systems is their ability to transform the way people utilize this text-based computer telecommunications medium. As reflected by the extensive number of research projects in this area, there has been a strong interest in incorporating interface agents in email applications in the last years (Maes and Kozierok, 1993; Lashkari et al., 1994; Maes, 1994; Payne and Edwards, 1997; Brzezinski and Dain, 2001; Bergman et al., 2002). At the root of this interest is the vision that agents will become a long-term solution for providing user assistance in tackling the currently challenging task of

*Tel.: +1 807 343 8698; fax: +1 807 343 8443.
E-mail address: aserenko@lakeheadu.ca.

1071-5819/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.ijhcs.2006.06.004

email management. For example, agents may reduce information overload associated with electronic communication (Segal and Kephart, 1999, 2000), speed up information exchange by acting as an intelligent information acceleration tool (Karnouskos and Vasilakos, 2002), and serve as intelligent and personalizable interfaces between users and other seamlessly integrated service agents that work in the background performing communication-related tasks (Kautz et al., 1994). It is these potential benefits for both end-users and organizations that raise awareness and interest in interface agents for email.

Despite the extensive work underway in the incorporation of interface agents in email applications, most previous studies and projects have been realized in forms of conceptual discussions, preliminary empirical investigations, and pilot systems (Lashkari et al., 1994; Maes, 1994; Florea and Moldovanu, 1996; Gruen et al., 1999; Bergman et al., 2002; Griss et al., 2002; Voss, 2004; Dabbish et al., 2005), and very few end-user products exist on the software market. Interface agents for email notification are one of the first commercial applications that utilize interface agent technologies in the electronic mail environment. In general, email notification systems have a long-standing tradition in computing; their goal is to inform users about changes in the state of their email clients (Libes, 1997). Email notification facilities have been incorporated in the early versions of BSD[1] UNIX OS by utilizing classic biff-like programs.[2] Biff is a small programme that enables OS to print notices directly on the terminal of an email recipient (Nemeth et al., 2000). Over time, the technology has been improved, and several independent software developers have already developed the commercial versions of interface agent-based email notification applications. To identify emerging issues associated with email notification agents and to ensure that agent manufacturers will be able to deliver products that meet user requirements, it is important of study user interaction with email interface agents at an early stage of technology development. Particularly, it is believed that surveys of actual users may generate useful findings important for both academics and practitioners.

Currently, there seems to be a gap between the expected user adoption of interface agents and their actual acceptance in both electronic messaging systems and other applications. Many projects have been technology-centered rather than problem-focused. These studies tend to look at the technical characteristics and capabilities of interface agents and value technical realizations of agent-based systems over that of user testing. The literature fails to provide clear evidence of the benefits of utilizing interface agents, and the results of past empirical studies on the usefulness and user adoption of interface agent technologies appear to be mixed and inconsistent (Dehn and van Mulken, 2000).

One of the factors that lie at the root of this issue is that most prior investigations were conducted in laboratory settings. On the one hand, the applicability of laboratory experiments has been successfully addressed in several human–computer interaction (HCI) areas, for example, in usability studies (Rubin, 1994). For instance, the laboratory experiment may produce statistically valid and generalizable results with respect to a new computer interface because the perceptions of its usefulness and ease of use may be established during a brief tutorial or an experiment with this technology (Davis, 1989). On the other hand, laboratory investigations may not produce results generalizable to the entire population of interface agent users. As hypothesized by Dehn and van Mulken (2000), in contrast to other information technologies, adequate perceptions of interface agents may take some time to establish. For example, the usage behaviour towards an interface agent of a person who explored it during a few-hour experiment may differ from that of an individual who utilized this agent for several weeks or months. More research is needed to bridge the gap between unconvincing and conflicting conclusions of past experiments.

With respect to email interface agents, the majority of recent projects were published in the form of technical reports rather than journal articles. Those technical reports explore technological aspects of a system leaving out user experience, perception, satisfaction, and adoption of the application (Helfman and Isbell 1995; Bergman et al., 2002; Rohall and Gruen, 2002). More importantly, the extant literature does not offer insights on how actual users employ interface agents in various software environments including electronic mail systems. Little is known about the actual end-use experience.

As one of the first attempts to fill that void, this study reports on how real-life users employ notification interface agents in their email applications. Particularly, given that it is relatively difficult to analyse all possible usage forms and circumstances in a single project, this investigation offers a detailed description of email notification interface agent usage in cases of positive and negative critical incidents. A critical incident is an event that a user perceives highly significant because it led to either success or failure. For example, a positive incident might be when an interface agent helped a person to complete an email-related task effectively, efficiently, or enjoyably. A negative incident might be when an interface agent hindered the completion of an email task. It is believed that this information may potentially shed some light on actual user behaviour that may lead to the creation of really useful email interface agents accepted by end-users. The following research question is suggested:

> What are the typical scenarios of behaviours of users who employ interface agents for email notification

[1]BDS—Short for Berkeley Software Design, Inc., a commercial supplier or BDS UNIX OS.

[2]The original biff program was written by Bill Joy. Biff was named after the Heidi Stettner's dog who barked at the postman. Heidi suggested that since a dog barks at the postman, so can software, when a new message arrives (i.e., bark at incoming email).

under the influence of positive and negative critical incidents?

In order to answer this research question, the critical incident technique is utilized. The following section of this paper describes this approach in detail.

2. The critical incident technique

The critical incident technique (CIT) (Flanagan 1954)[3] is a flexible set of principles for gathering certain important facts concerning behaviour in defined situations to facilitate the potential usefulness of obtained information in solving practical problems and developing broad psychological principles. An incident is "any observable human activity that is sufficiently complete in itself to permit inferences and predictions to be made about the person performing the act" (Flanagan, 1954, p. 327). Each incident should have a high degree of significance for an individual's success or failure in a task (Andersson and Nilsson, 1964). The technique is based on the assumption that people can report critical incidents on their own (Koenemann-Belliveau et al., 1994). As such, it involves the collection and analysis of brief, written, factual reports of actions in response to explicit situations or problems in a certain field. Incident reports may be written by individuals who actually took an action or by qualified observers. A critical incident is considered effective (or positive) if it helps to solve a problem, or ineffective (or negative) if it fails to solve a problem, creates new problems or facilitates the need for further actions.

In this study, the key goal of the employment of the CIT is to build several typical scenarios of email interface agent usage under the influence of positive or negative critical events. These scenarios may be presented graphically as a set of constructs and their relationships. The purpose of these scenarios is to form an understanding of user behaviour in cases of positive and negative critical incidents. It is expected to discover significant differences in user actions, feelings, and behaviour changes in cases of positive and negative incidents.

The CIT produces reliable, valid, and generalizable results (Ronan and Latham, 1974). Data may be analysed qualitatively by using a variety of methods. Mailed questionnaires produce the same results as those obtained by interview methods given that respondents are motivated to read the instructions and answer consciously (Andersson and Nilsson, 1964). Especially, self-administered surveys are acceptable for open-ended, self-reported items (Wang et al., 2000).

The CIT minimum sample size requirements depend on the nature of the phenomenon of interest. If an activity is relatively simple, it may be sufficient to collect only 50 incidents; most recent studies examined 50–100 incidents (Urquhart et al., 2003). CIT was applied in various business administration fields, for example, organizational behaviour (Cowie et al., 2002), marketing (Jones, 1999), and information technology (Tay and Ang, 1994; Muylle et al., 2004). The technique is particularly appropriate when the field is new and the goals of research include practical managerial problems and theory development (Keaveney, 1995).

As such, the CIT can be successfully utilized to achieve the purpose of this study. First, given that the technology of interest is relatively simple, it should be sufficient to collect around 50 valid responses. Secondly, data collection may be employed by means of self-administered questionnaires with results comparable to those obtained by interviews. Thirdly, any qualitative data analysis technique may be applied to perform data coding. Overall, it is believed that the use of the CIT may provide an adequate and realistic description of behaviours of email notification interface agent users, identify reasons why individuals continue or discontinue using this technology, find what individuals look for in an ideal email interface agent, and suggest strategies for agent designers and marketers.

3. Methodology

In order to answer the study's research question, an online self-administered survey of the actual users of email interface agents was conducted. The following subsections describe the methodology in detail.

3.1. Interface agent for email notification selection

An exhaustive online search for interface agents for email notification demonstrated that there are at least eleven products available on the market. All of them are referred to as 'email notification systems.' Their purpose is to inform users about incoming messages, to provide an intelligent interface between human users and other parts of a system, and to offer a variety of communicative, entertaining, and information management functions. Out of these applications, Email Announcer developed by Blind Bat Software (http://www.blindbat.com) was randomly chosen by the researcher. Fig. 1 offers the screenshots of an agent's interface and configuration environment.

Email Announcer adds two new features to an email application. First, it incorporates an interactive cartoon-like character, which announces incoming email messages by utilizing Microsoft Agent technologies. Depending on a user's preferences, it may also read the sender, subject line, or entire message. The agent also informs individuals if senders have requested a 'read receipt.' Email Announcer is capable of retrieving email from Hotmail email accounts. The major benefit of utilizing such an agent is that users do not have to be distracted from their current activities and switch from the application they are currently working

[3]Colonel John C. Flanagan is the founder of the American Institutes for Research, a not-for-profit organization dedicated to the study of human resources and their effective use. The critical incident technique evolved naturally from work in the Aviation Psychology Program of the United States Air Forces.

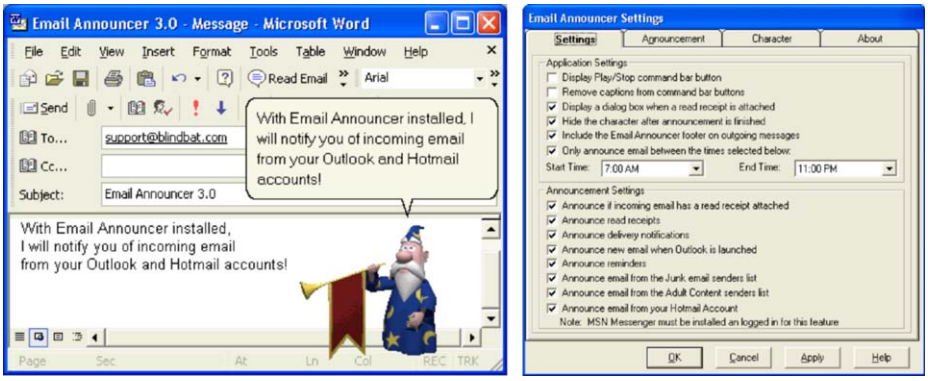

Fig. 1. Email Announcer developed by Blind Bat Software–agent interface and configuration environment.

with to their email system to see message details. Secondly, the Email Announcer is able to announce reminders, which pop up in the MS Outlook Calendar. This service is very convenient for users who are currently engaged in other activities, for example, web surfing, and cannot read messages. A simple graphical interface allows users to set a number of parameters of the programme, for example, to select a character, its voice and basic animation attributes, to personalize the MS Outlook toolbar, and to specify announcement details. The use of the application is very simple, and the employment of this system does not require coding. This is very important since most email users do not have any programming skills.

3.2. Questionnaire design and implementation

By drawing upon previous investigations that utilized the CIT (Johnson and Fauske, 2000; Wang et al., 2000), the study's participants were questioned on two distinct areas: (1) incidents that users deemed highly positive; and, (2) events that individuals considered extremely negative. Every respondent was asked to provide either one positive or one negative critical incident covering the following points: (a) incident category; (b) the nature of a specific memorable event in the process of human-interface agent interaction that captured or demanded a respondent's attention; (c) the outcome of this situation; (d) why an individual considered this specific event critical, important, and worth attention; (e) user feelings; (f) actions taken during the incident; (g) changes in future behaviours after the incident; (h) actions that an 'ideal' agent would take in addition or instead; and, (i) the regularity the similar situations occur. As such, one close-ended (i.e. item 'a') and eight open-ended (i.e. items 'b' to 'i') questions were constructed (see Table 1). It is proposed that by analysing these points a complete understanding of user behaviour can be obtained.

In addition to the items related to the CIT, two optional open-ended questions were offered. Empirical findings in innovation research suggest that users often play a leading role in the invention and improvement of new products and services (Biemans, 1991; Lüthje, 2004). Many commercial projects have succeeded because designers and manufactures involved users in the early stages of innovation development. A strong understanding of user needs is a key factor separating new product winners from losers (Cooper and Brentani, 1991). Therefore, the goal of these two questions was to allow end-users to provide recommendations for agent designers and marketers. Given that the solicitation of user suggestions was not a primary goal of this project, no research question was proposed. In addition, it was unknown whether users may offer valuable insights on the development and marketing of this technology. The questions were: (1) based on your experience with interface agents in your email applications, provide recommendations for interface agent designers; and (2) based on your experience with interface agents in your email applications, provide recommendations for interface agent marketers.

A list of potential respondents to the survey was randomly formed by the researcher from the company customer database. In order to reduce self-selection bias, Dillman's (1999) Tailored Design Method was adapted. All potential respondents were emailed an initial request to participate in the study, and three follow-up reminders. Demographic data were also obtained. The data were collected as part of a larger project (Serenko, 2005).

4. Results

The data collection phase was conducted during the period from April 14th to June 1st, 2004. An acceptable response rate was achieved that compares favourably to those of other similar studies.[4] Overall, 60 critical incidents were provided, 30 of them were positive and 30 negative. This split in the type of incidents was a desirable aspect of the data. If for example, most subjects described only positive incidents, it might be assumed that they liked email interface agents to such a great extent that they tended to ignore the negative outcomes of agent usage. In contrast, if most individuals offered only negative critical incidents, it

[4]Note that the actual response rate may not be revealed as per the non-disclosure agreement with Blind Bat Software.

Table 1
The critical incident technique questionnaire

Instructions	
Note: definitions of an email system and interface agent as well as screenshots of an interface agent developed by Blind Bat Software were provided. Participants were instructed to answer all questions based on their experience with the email notification programme developed by Blind Bat	
Please answer the questions below with respect to the last most significant *POSITIVE* or *NEGATIVE* incident of usage of interface agents in an email application (e.g. a positive incident might be when an interface agent helped you to complete a task in your email application effectively, efficiently, or enjoyably. A negative incident might be when an interface agent hindered the completion of a task in your email application)	
N	Question
(a)	Was this incident positive or negative? (positive/negative checkboxes)
(b)	Provide a complete and detailed description of this incident and indicate how long ago (e.g. days, weeks, months) it took place
(c)	What was the outcome of this incident?
(d)	Why do you consider this incident critical?
(e)	What were you feelings and perceptions of this situation?
(f)	What actions did you take during the incident?
(g)	Did you change the way you use interface agents after that? If yes, please specify
(h)	From your point of view, what are the most desirable actions that an 'ideal' interface agent would take in addition? (in the case of positive incidents) From your point of view, what are the most desirable actions that an 'ideal' interface agent would take instead? (in the case of negative incidents)
(i)	How often does a similar situation occur(ed) when you use(d) interface agents in your email applications (e.g. days, weeks, months, never again)?

Table 2
The Kripendorff's agreement coefficient

	Incident cause	User feelings	User actions	Behaviour change	Ideal agent
Coder agreement—positive incidents	0.93	0.81	0.81	0.80	0.90
Coder agreement—negative incidents	0.87	0.85	0.91	0.81	0.73

might be suggested that there were fundamental problems with the technology under investigation, for instance, it was highly unreliable.

Eighty and twenty percent of the surveyed users were male and female respectively. User age ranged from 20 to 65 years old. Over 65% of all users were between 31 and 50 years old. The 46–50 age category was the most frequent user group. Fifty-five percent of the respondents were occupied in the information technology sector, eight percent in the engineering industry, and 37% in other professions. The majority of users were well-educated; 81% of them had a college/university degree. Sixty-three percent of respondents resided in the USA, 12% in the European Union countries, followed by Canada, Australia, and New Zealand. They were very heavy email users, and were financially well-off. Overall, it was concluded that this user population corresponds to innovators (Rogers, 2003), who constitute 2.5% of all people that adopt a particular product.

Every incident was analysed along the following dimensions: (1) incident cause (why the incident took place); (2) user actions (what actions a user took during the incident); (3) user feelings (what a user felt about this situation); (4) behaviour change (whether and how a user changed the way he/she used email interface agents after the incident); and (5) ideal agent actions. Positive and negative incidents were analysed separately.

The codebook was developed within this study that is a valid technique in qualitative research (Miles and Huberman, 1994). First, based on the HCI and MIS literature, several a priori categories were formed. Second, up to ten rounds of codebook refinement, re-coding, new code development, code discarding, etc. were conducted. Third, the codebook was evaluated by an independent expert, and some codes were improved. After that, three independent coders analysed the open-ended items and achieved an acceptable level of agreement (Krippendorff, 1980) (see Table 2). The classical content analysis approach (Budd and Thorp, 1963; Budd et al., 1967; Kerlinger 1973; Riffe et al., 1998; Ryan and Bernard, 2000) was used, and the positivist paradigm was chosen (Benbasat et al., 1987; Myers, 1998). The same codebook was utilized by all coders. If coders failed to agree on an item's classification, the response was labelled as 'Other.' Several subjects left a few questions blank, therefore, totals sometimes do not add to 30.

4.1. Positive critical incidents

The results indicated that all users reported positive critical events that happened very recently, for example,

during the last incident of notification agent usage, today, or within a few weeks, and that similar events occurred very frequently, for instance, during every incident of use, daily, or weekly.

Regarding the *incident cause*, 24 incidents related to an event when an agent notified a user about the state of an email system. For instance, the agent presented an important incoming message or event in a timely manner:

> I use agents to monitor various mail folders in Outlook and either announce or read the mail based on set criteria. I have unattended processes that report failures or problems encountered that spawn emails to me. Three days ago my process emailed me that one of our databases was approaching its transaction log maximum size and this was announced by an agent. [In the result,] I was able to clear the log before it caused the database to stop processing due to a full drive error. Obviously it would have halted one of our critical business processes.

In another positive critical incident, an agent informed a user about a due event:

> A few weeks ago, [the agent] got me out of one meeting I was asked to join because it announced the meeting I was due to go to. The customer on the phone understood that I wasn't making the other meeting up to get out of the one I was attending. [I] left one conference call that was not scheduled and went to the one that was. [In the result, I] made it to my meeting in time.

Six incidents corresponded to the fact that the agent performed highly reliably over a certain period of time, for example:

> Yesterday 18/5—software performed as usual, stable, did not cause problems. [It] finished as usual, nothing happened.
> I have not had any problems with my [agent] e.g. it performs as usual informing me about important messages.

According to the original critical incident methodology, a critical incident is a discrete episode that contains sufficient detail to be visualized by a researcher (for example, see Wang et al., 2000). On the one hand, according to this viewpoint, responses pertaining to an agent's reliability may not be classified as critical incidents. On the other hand, users clearly emphasized the importance of the fact that the agent performed highly reliably over a certain period of time (e.g. yesterday). A partial explanation of this phenomenon lies in the imperfection of most contemporary software applications, including agent-based systems. Currently, computer users are so accustomed to bad design, poor usability, increasing complexity, and lack of important functionality of software that they tend not to complain about it (Lieberman et al., 2001). This is especially true regarding novel, agent-based technologies the usage of which is associated with high uncertainty and risk (Serenko and Detlor, 2004). At the same time, respondents to the survey were ready to acknowledge the quality of an agent when it reliably performed the required tasks, and they considered this event critical. Therefore, the responses pertaining to an agent's reliability were included in further analysis.

Fig. 2 offers *user feelings*. Most people had positive feelings towards the incident, such as satisfaction and enjoyment. Four of them indicated they felt that an agent was very helpful, and two felt indifferent.

With regards to *user actions* during the incident, 11 individuals said they immediately completed a task suggested by an agent, eight people continued doing a task they were doing before an agent's interference (i.e. they ignored the information presented by the agent), and three users made a better, more informative decision on a task they were working on (Fig. 3).

Regarding *behaviour change* after the critical incident, 22 users indicated that they did not change the way they used interface agents and three users said they did change. Out of those who reported behaviour change, one person indicted that he/she increased agent usage, and two individuals said they began to promote the agent among friends and colleagues:

> I demonstrated the product to many friends and coworkers who thought they would pursue adding agents to their email.

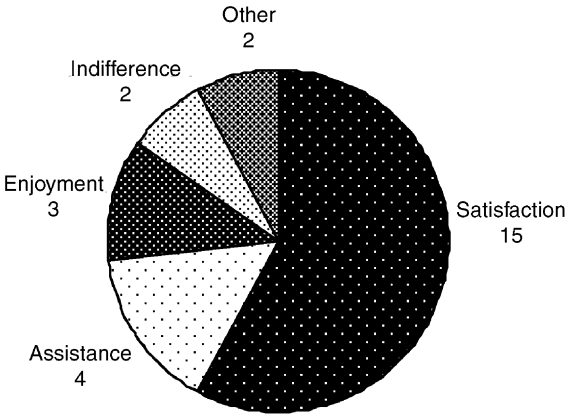

Fig. 2. Positive critical incidents—user feelings.

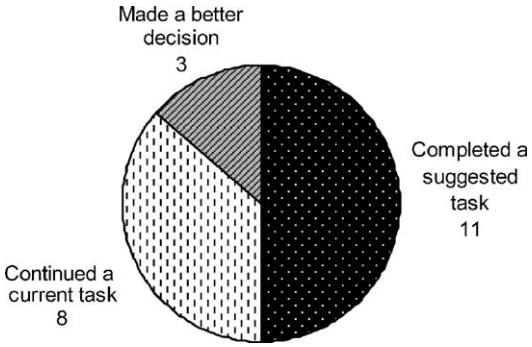

Fig. 3. Positive critical incidents—user actions.

Overall, this demonstrates that the occurrence of a positive-outcome situation either caused no changes in future user behaviour or caused positive changes.

In terms of the most desirable actions that an '*ideal*' *interface agent for email notification* would take in addition, 22 responses were obtained. Although the users were asked to provide an action, many of them described an extra feature that they wanted an agent to possess. Table 3 provides these features and actions. They demonstrate that many individuals, who experienced a positive-outcome critical incident, believed that their agent might adequately perform this task, and no additional features or actions were necessary. Most users wanted to increase the degree of an agent's intelligence, to improve the way it presents notification, and to be able to utilize more personalized features. Several additional actions and features were offered, for example, one person indicated that he/she would benefit if an agent was able to send text messages to cell phones (i.e. to be connected to the Short Messaging Service or SMS).

4.2. Negative critical incidents

With regards to negative *incident causes*, 11 distinct categories emerged. Fig. 4 offers these causes grouped on the first level of coding, and Table 4 on the second, lowest level of coding. The most frequently cited cause of a critical incident was an agent's incompatibility with other software systems, especially, with email environments.

A typical negative incident occurred when a user suddenly realized that the notification interface agent might not be used with the email application that he/she just installed or updated:

> [I am] having some trouble getting it to work with my newer version of Office software. [It] worked fine in 2000 version. In my case, it is preventing the use of an agent.

Other regular sources of negative incidents were interference, intrusiveness, and unreliability of an agent. Sometimes, an agent interfered with other systems or the entire computer that slowed down the CPU and consumed extra resources. During a particular task, several users found an agent disturbing and annoying that distracted them from other important activities. Three users complained about the event in which an agent behaved unreliably:

> [The agent] gives an error message when forwarding or replying to email and then stops announcing.

Several incidents were caused by other people, who complained about announcements of unnecessary messages, little control over an agent, company policies that did no allow the usage of unauthorized software, unclear agent voice, and the noise that an agent made:

> Other occupants in the house annoyed by the voice announcing new messages, read receipts, etc.
>
> I use my agent to tell me, verbally, what mail is about so I can decide whether or not to stop what I'm working on in another program to read that message. Since my agent looks at everything, it even reads SPAM and messages that I automatically file using rules, and that I don't need. [This happens] every day, all the time.

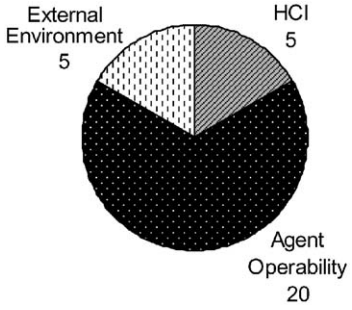

Fig. 4. Negative critical incidents—incident cause (Level 1).

Table 3
Positive critical incidents—an 'ideal' agent's actions

Number of responses	Code explanation
7	No additional actions are necessary, the agent does everything this user needs
5	An agent should have more intelligent features such as rule-based logic, machine learning capabilities (i.e. an agent follows user-specific rules), text analysis features, and autoreply. The agent's interface, actions, or behaviours should depend on user requirements and the type of incoming information (e.g. different agents for different senders, message contents, types of events, etc.)
2	An agent should notify a user about the state of an email system (e.g. present incoming messages or due events in a timely manner)
2	An agent should be more personalizable (e.g. it should have a large selection of characters, interfaces, voices, multi-language support, and verbal announcements)
1	An agent should be compatible with other software applications (e.g. email systems)
1	An agent should be easy to use and re-install
1	An agent should be able to send SMS (short messaging services) messages (i.e. to be connected to a cell-phone)
1	An agent should have good voice capabilities (i.e. clear, understandable, high quality voice)
1	An agent should incorporate voice recognition capabilities (i.e. a user may communicate with an agent in natural language)
1	The reason is not clear based on the respondent's comments
Total: 22	

Table 4
Negative critical incidents–incident cause (Level 2)

Number of responses	Code explanation
7	An agent's incompatibility with other software applications (e.g. email systems, Outlook Express)
5	Interference with the computer (e.g. slows it down, consumes extra resources) or applications (e.g. slows an email system down)
5	High intrusiveness or distraction caused by an agent (e.g. annoying or bad timing of notifications)
3	Poor reliability of an agent (e.g. crashes, bugs)
2	People, who know that a person utilizes an agent, attempt to abuse the use of this agent by sending irrelevant, obscene, or hard-to-read messages that upset or embarrass the individual (i.e. when the agent loudly pronounces those messages, the user is humiliated)
2	There is little functionality that gives a user more control over the agent, e.g. the user cannot stop the agent at any time or act before the agent completes a notification task
2	An agent bothers other people around because it communicates with a user in a natural voice and people around may hear it
1	Policies in the work environment do not allow the use of unauthorized software including email agents
1	The announcement of all incoming messages, including spam (if the filter failed to sort them out)
1	An agent is not intelligent enough to perform basic tasks that require some degree of reasoning capabilities
1	Voice capabilities need improvement (e.g. unclear speech)
Total: 30	

> I have to give the voice some attention and wait until I realize it's not relevant at all (as opposed to something I can delay till later). It sort of degrades the effectiveness of the tool.
>
> [I] couldn't get [an] agent to minimize and move out of my way. I shut the agent down [with] lots of clicking and swearing.
>
> [I] was told by Information Systems to remove the agent from my computer.

Two people reported cases in which other individuals, who were aware of the fact that someone utilized an email interface agent, abused the user by sending irrelevant, obscene, or hard-to-read messages and made fun of the situation when an agent announced either the subject or the entire message. These messages passed through email filters since senders were known to the recipient:

> Colleagues found out I was using the software at work and began sending me emails with vulgar subject lines. My co-workers every now & then send me messages that are read out & amusing or send me all capitalized letters that are spelt out individually & are annoying. [I get] laughter from all.

In one situation, a user realized that an agent was not as intelligent as he/she expected it to be and discontinued using it:

> [I] discovered that the agent is not intelligent enough for many tasks, needs more automation, [and I] stopped using [it].

Most people expressed negative *user feelings* towards the incident, such as frustration and annoyance (Fig. 5). Note that four people stated they felt indifferent towards the negative events when an agent behaved unreliably and incompatibly, or when it began to interfere with other systems. For example, in one case, an agent slowed down the CPU speed to an extent when a user had to temporarily turn it off. In another situation, when many messages arrived at a time, an agent began to behave unstably, and a user had to shut it down. However, these people made a conscious decision not to blame the agent for its inadequate performance:

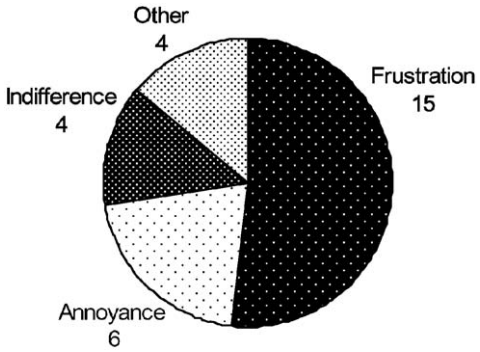

Fig. 5. Negative critical incidents—user feelings.

> [My feelings were] normal for using gadgets.

Figs. 6 and 7 outline user actions during negative-outcome incidents.

The analysis shows that those who chose to temporarily turn off the agent did so because of the agent's intrusiveness, unreliability, interference with other applications, unclear speech, and limited control over its actions. Out of four people who reinstalled the agent three did so because the agent became incompatible with a new email system, and they believed that a reinstallation or an upgrade might fix the problem. One person reinstalled it because of the agent's interference with other applications. Two individuals who received complaints from people nearby reduced the volume, and two users who experienced incompatibility and interference problems requested technical assistance. One person who was a victim of colleagues' abuse ignored

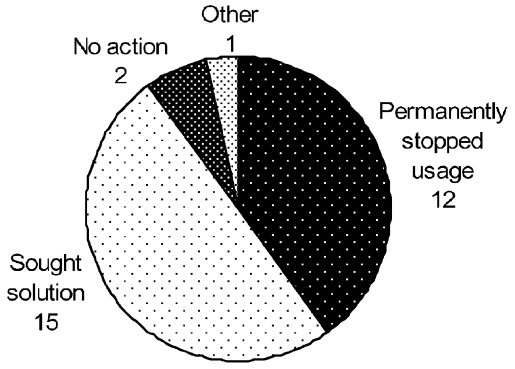

Fig. 6. Negative critical incidents—user actions (Level 1).

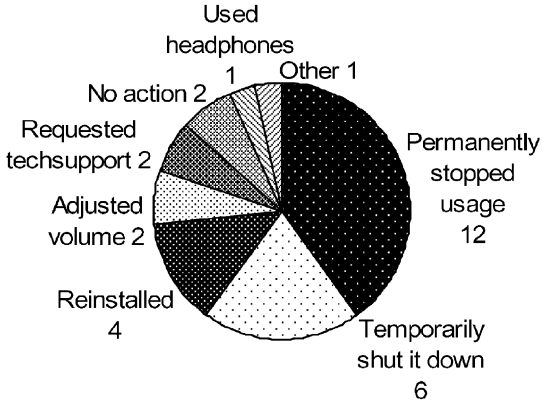

Fig. 7. Negative critical incidents—user actions (Level 2).

the incident, and one individual began to use headphones for privacy.

Fig. 8 presents *behaviour change* after the critical incident.

Recall that 20 critical incidents were caused by an agent's operability issues, five by the HCI process, and five by external factors. In cases that originated from an agent's operability problems, 11 users sought a solution, seven immediately abandoned the agent's employment, and one person took no action. Overall, out of them, 11 people discontinued the usage and eight individuals did continue.

Out of five people that experienced an incident in which an agent behaved highly intrusively, four immediately terminated the use, and only one person temporarily shut down the agent. In five critical incidents caused by an external environment, three individuals terminated future employment of this technology and two people continued using the agent. In response to these incidents, only one person removed an agent immediately because he/she was forced by the IS department of the company. Out of the other four people, three attempted to find a solution and one ignored the event. For example, to eliminate noise, they tried to utilize headphones or decrease the volume. The user, who was forced to terminate the usage because of

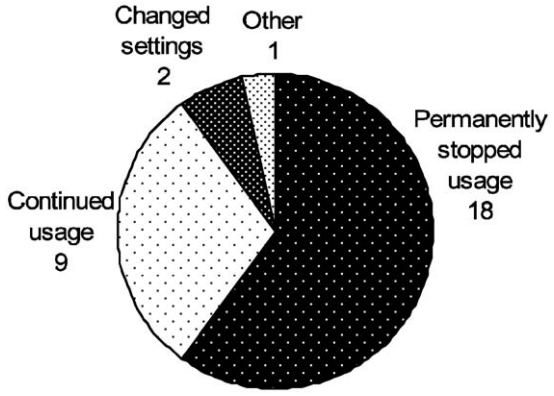

Fig. 8. Negative critical incidents—behaviour change.

noise constraints, indicated that he/she would attempt to use an agent again:

> [I] uninstalled the interface agent ... but ... planning on reinstalling it.

In addition, the respondents, who experienced a negative critical incident, provided features of an 'ideal' interface agent for email notification. Table 5 offers these characteristics. It should be noted that none of the users indicated that no additional features or actions were required. This contradicts the information presented by the users who experienced a positive critical incident. Recall seven of them were so satisfied with the agent's actions that they felt no need for extra assistance.

The analysis of incident frequency and timeline indicated that two groups of data emerged. The first category pertained to the incidents that occurred in the past, usually, over one year ago, and that had occurred only once or a few times. Typically, a user terminated the employment of an agent after that event. The second group related to the incidents that took place recently, and that appeared more often, for instance, monthly. Generally, users continued using an agent after that incident.

In order to ensure that the sample size employed in the present study was acceptable, a test for theoretical saturation was conducted. The purpose was to ensure that the increase in sample size will not allow discovering new data categories. For this, the data were sorted in the order they arrived, and the number of newly emerged categories was counted in every cluster of ten responses. In terms of positive situation causes, all causes were offered in the first cluster of ten. With respect to negative situation causes, 7, 3 and 1 new causes appeared in the first, second and third cluster respectively. As such, it is believed that an increase in sample size would allow identifying only a few, if any, new categories, and theoretical saturation was reached.

4.3. Recommendations for designers

Recall that respondents were asked two optional open-ended questions soliciting recommendations for agent designers and marketers. The goal was to obtain the user

Table 5
Negative critical incidents–an 'ideal' agent's actions

Number of responses	Code explanation
7	An agent should have more intelligent features such as rule-based logic, machine learning capabilities (i.e. an agent follows user-specific rules), text analysis features, and autoreply. The agent's interface, actions, or behaviours should depend on user requirements and the type of incoming information (e.g. different agents for different senders, message contents, types of events, etc.)
4	An agent should be more personalizable (e.g. it should have a large selection of characters, interfaces, voices, multi-language support, and verbal announcements)
3	An agent should be compatible with other software applications (e.g. email systems)
3	An agent should present the functionality that gives a user more control over the agent, e.g. the user should be able to stop the agent at any time or to act before the agent completes a notification task
2	An agent should not intrude in current user tasks, provide smooth information, and disappear when the task is complete
2	An agent should not interference with the computer (e.g. not to slow it down, or consume extra resources) or applications (e.g. not to slow an email system down)
1	An agent should be more attractive to a user (e.g. it should have a more attractive interface and/or voice)
1	An agent should be easy to use and re-install
1	An agent should be very reliable (e.g. it should not have bugs or crash)
1	An agent should incorporate voice recognition capabilities (i.e. a user may communicate with an agent in natural language)
3	The reason is not clear based on the respondent's comments
Total: 28	

views on this technology that were not identified in the previous questions. Only a minority of the respondents provided answers to these items. As such, 28 suggestions for agent designers were offered. No content analysis of these items was done since the responses were presented in form of ready-to-use recommendations that might be grouped together and summarized without the use of a codebook.

First, many respondents wished to have more agent characters available at their disposal. Currently, the Windows XP operating system supplies users with eight Microsoft Agent characters that are installed by default. In order to obtain new agent characters, people need to download them off different websites. These MS Agent compatible characters are created by Microsoft as well as by a variety of independent designers. Now, there are hundreds of MS Agent Characters available online, most of which were free.[5] However, the users were faced with two problems. The first issue was the lack of information about the availability of agent characters in the Internet. The second was the lack of facilities in the agent software that would allow automatically downloading and installing new agent interfaces. A quick review of other interface agents for email notification revealed that their users were faced with the same problems since none of them presented agent character upgrade facilities. Although the manual installation process is very simple, some users may be unaware of that opportunity or may not have time to search the Internet and install new characters.

Second, the respondents indicated that they would like to assign account-specific message processing rules to handle multiple email accounts consolidated in one email system. Consider, for example, a person who retrieved messages from several email accounts by using a single mail client such as MS Outlook. The interface agent application did not allow the design of account-specific rules. However, it might be beneficial for the user to utilize different agent characters, message introductions, and message processing rules for different email accounts.

Third, some users advised that the usage of animated, cartoon-like characters was inappropriate in the software that was often utilized as essential business tools. Indeed, not everybody might like these entertaining interfaces; some users might prefer a simple informative box that would possess some degree of intelligence and present users with new information.

Fourth, according to the users, the introduction of simple rules that specify the maximum number of message or event notifications over a certain period of time would be very useful. If, for example, ten emails arrived at once, people might be annoyed by the announcement of all of them. In this case, an agent should inform users about the delivery of ten messages and specify their location, such as the mailbox name. It would be more productive for individuals to open that mailbox and to determine message relevance, urgency and importance, and to ask the agent to read some of them in a certain order.

4.4. Recommendations for marketers

In addition to the recommendations for agent developers, users offered 12 suggestions for the marketers of interface agents for email notification. Again, the key points are presented as they were provided, and no data coding procedures was done.

First, the users advised that agent marketers should offer simple demos of their products online. All online sellers of

[5]Lists of available MS Agent characters may be found on the MS Agent Ring website at http://www.msagentring.org and the Agentry website at http://www.agentry.net.

interface agent-based notification applications offered free trials (or shareware) of their products for a limited period of time. The purpose of these trials was to allow potential customers to utilize agents risk-free. However, the usage of shareware required users to download and install agents on their machines. The respondents argued that those email users who were not familiar with agents were unlikely to do so. In order to motivate them, online agent distributors should offer demos of their agents in form of animated graphical images, Macromedia Flash movies, video clips, or sound files. This would increase the number of people who decide to try out this new technology and raise sales. In addition to shareware versions and online demos, potential users needed to be assured that an interface agent was compatible with their operating systems, and that it might be easily and safely removed from their computers at any time.

Second, the respondents suggested that agent marketers should emphasize the functionality of their products that was not available in conventional non-agent systems. As such, online sellers needed to highlight the usefulness, productivity, and entertainment potential of their agents. For example, the increase in email productivity might result from the reduction in unnecessary interruptions when new messages arrived, especially for the power users who favored multi-tasking or utilized two computers simultaneously. By stressing this information, online sellers might motivate people to try out their agent software.

5. Discussion and conclusion

Recall the major purpose of the study is to develop the typical scenarios of email interface agent user behaviours under the influence of positive and negative critical incidents. Based on the findings from the CIT as well as additional user suggestions, recommendations for developers and marketers should be offered. For this, a survey of 60 actual users of email interface agents was conducted.

5.1. User behaviour scenarios

With respect to *positive* critical incidents, the following typical scenario of user behaviour is constructed (Fig. 9).

According to this scenario, a positive incident occurred when an agent presented a user with a notification in a proactive and reliable manner. For example, it read an important message from a colleague, and the user did not have to switch from a currently open application to an email system; this saved time and improved email efficiency. The person enjoyed his/her experience, utilized the information for decision making and continued employing the agent. In some cases, he/she began to promote the agent by demonstrating it to peers.

After a positive event, approximately one-third of the users believed no additional actions were required. Some people wanted an agent to be more intelligent, present better notification facilities, or offer more personalization features. In other words, when an agent works reliably, users acknowledge its quality and continue using it.

In terms of *negative* critical incidents, three distinct scenarios were identified because each situation was caused by a unique type of a critical incident. Fig. 10 offers the scenario of user behaviour that occurred because of agent operability problems.

Most such incidents happened because an agent was incompatible with a user's email client, interfered with other applications, or behaved unreliably. Users felt very frustrated; some of them made an immediate decision to terminate agent usage, whereas others tried to find a solution. Those, who tried to solve the problem, either terminated agent usage at a later date or ignored the incident and continued the employment of the agent.

Fig. 11 outlines the second scenario of a negative-outcome event that resulted from high perceived intrusiveness of an agent.

According to this scenario, an incident happened because an agent behaved highly intrusively. As a result, a user felt very frustrated and annoyed by the actions of the agent. Immediately, he/she decided to permanently terminate the usage of the agent.

Fig. 12 offers the third, last scenario of a negative-outcome incident that took place under the influence of external factors which a user could not control.

In terms of this scenario, other people or employers interfered. For example, peers abused the user by sending obscene messages, colleagues complained about noise, or company staff requested that the agent be removed from a computer. An individual was very frustrated and annoyed by the actions of the peers or company personnel. In response to their actions, the user attempted to solve the problem to preserve the usage of the agent, but, in most cases, he/she had to terminate the usage. Rarely, the person found a solution or ignored the incident.

All individuals who experienced a negative-outcome event felt their agent should possess some additional actions or features, for example, more intelligence, personalization, compatibility, and control functions.

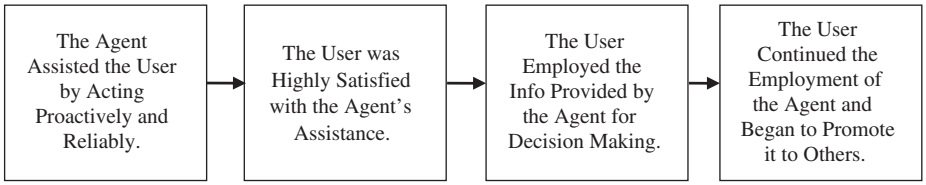

Fig. 9. A typical scenario of user behaviour—positive critical incidents.

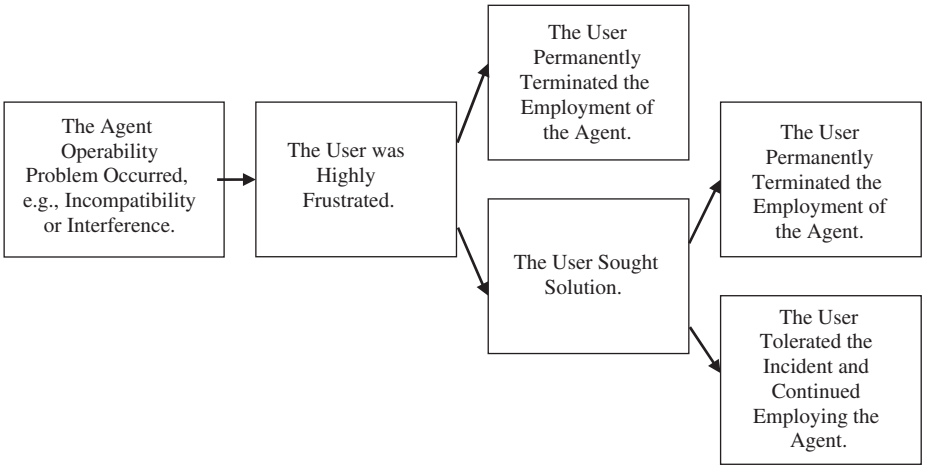

Fig. 10. A typical scenario of user behaviour—negative critical incidents—agent operability.

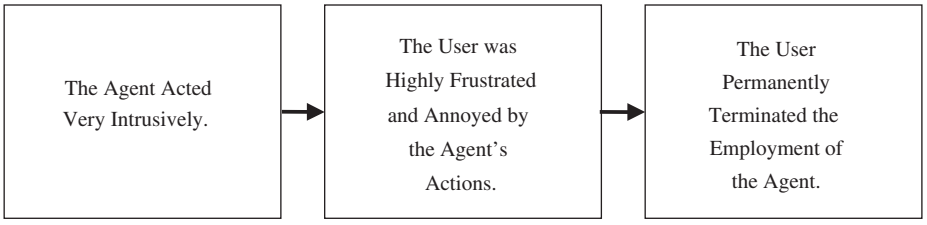

Fig. 11. A typical scenario of user behaviour—negative critical incidents—perceived intrusiveness.

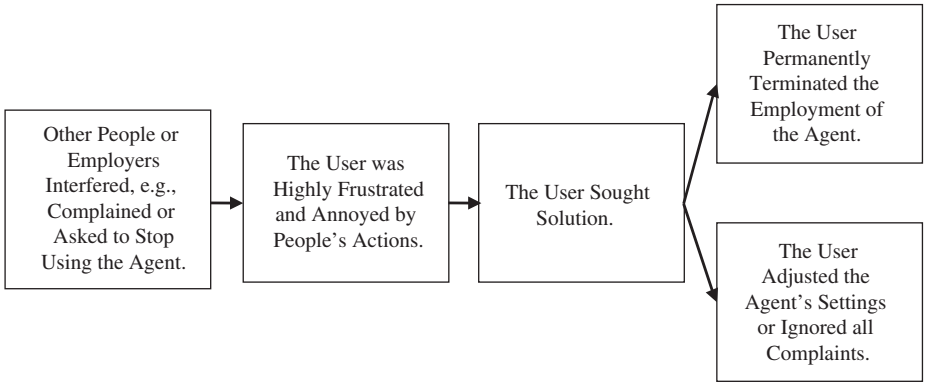

Fig. 12. A typical scenario of user behaviour—negative critical incidents—external factors.

Overall, it is believed that these scenarios show a realistic picture of user behaviours under the influence of positive and negative critical incidents.

5.2. Implications for theory and practice

Based on the findings, several theoretical and practical contributions are suggested. First, it is noted that agent users acknowledge the quality of an agent when it reliably performs the required tasks, and they consider this event critical. In addition, even when a negative critical incident occurs, some people feel indifferent because they take such situations for granted that still might negatively affect their level of satisfaction with an agent. Indeed, most modern software systems, including email interface agents, are not perfect, and people are used to encountering problems. At the same time, individuals highly appreciate reliable performance of interface agents. This demonstrates the importance of creating high-quality agent-based systems.

Second, an agent's intrusive behaviour results in an immediate agent usage termination, agent operability issues (e.g. bugs or crashes) may or may not force people to reject the technology, and users attempt to preserve the employment of an agent under the negative impacts of external factors. This shows that many real-life interface agent users feel optimistic about the employment of this technology, and they want to continue using agents for email notification, except for the cases when an agent

behaves extremely intrusively. Therefore, efforts should be directed to reduce the degree of perceived intrusiveness of email interface agents. As a short-term solution, more personalization features need to be introduced. For example, advanced options or visual programming environments for message or event processing rules would allow individuals to precisely specify an agent's actions depending on each particular situation. They may instruct the agent to ignore messages that are automatically filed, arrive from certain people, or contain special keywords.

Third, to improve user acceptance of this technology in the short-term, agent designers need to make this technology compatible with other software and email systems, eliminate an agent's interference with other applications, and focus on reliability issues. In the long-term, agent developers need to utilize the existing body of knowledge in the field of artificial intelligence and human–computer interaction to incorporate more 'intelligent' features in email interface agents, such as machine learning capabilities, basic text analysis with automatic message response mechanisms, run-time adjustments of an agent's behaviour, appearance and voice, and rule-based logic.

Fourth, interface agent-based applications for email notification should incorporate the facilities for the quick downloading and installing of new MS Agent characters that were accessible online that is feasible to implement. The agent character is an executable (i.e. name.exe) file, and there are websites that host those files. An agent-based system should automatically locate, download, and install new characters when a user wants so. Users also need to have a choice between cartoon-like agent interfaces and more conventional message presentation interfaces. Each email address should be also tied to its own agent character and rules, and users should be able to select the maximum number of messages delivered at a time.

Fifth, email interface agent marketers need to realize that offering free trial versions of agent-based software is not sufficient to ensure the awareness and diffusion of this novel technology. Internet and email users who are not familiar with interface agents are unlikely to try out this technology when they come across a website that offers agents. Many of those who wish to try it out cannot run an interface agent system because there are several additional components that need to be installed, for example, text-to-speech engines or agent characters. To facilitate the distribution of trial versions, marketers need to promote their software by offering various online demos that do not utilize the MS Agent technology. For example, these demos may be created in the form of graphical images, Macromedia Flash movies, or videos that can be viewed on most contemporary computer systems and that do not require the installation of additional software components. Marketers should also emphasize various aspects and features of this novel technology that are not available in other non-agent applications.

5.3. Limitations and directions for future research

The results of this investigation are constrained by several limitations. Perhaps the most salient is that users of only one interface agent-based system for email notification were surveyed. To strengthen the validity of the findings, a survey of users of an application developed by another manufacturer should be conducted. Future researchers may also replicate this project by looking at other types of agents, for example, electronic shopping or personal assistance agents.

The second limitation is that this study was restricted to the reactions of users, and it does not encompass two-way interactions between people and agents. In other words, from the user perspective, the human–agent interaction experience was a one way process when all communication was initiated by an agent, and a user had relatively little control over the process. As such, the contemporary interface agents for email notification are relatively simple. At the same time, it is important to study the real-life user experience at the early stages of technology development to form the foundation for the creation of better systems. It is suggested that future researchers investigate the user experience with more sophisticated interface agents that facilitate two-way human–agent interaction processes. This can be done either through surveys or laboratory experiments.

The third constraint is that each respondent offered either a positive or a negative critical incident. Presumably, most respondents experienced both types of incidents, but they reported only one. This was done to keep the questionnaire short and achieve an acceptable response rate. At the same time, that would be interesting to analyse positive and negative experiences reported by the same person and look for possible interrelationships across the data.

The fourth limitation of this project is that the suggested scenarios describe how individuals behave, but they do not explain why each person acts this way. For example, when an agent operability problem occurred (Fig. 10), some users permanently disabled the agent whereas others tried to find a solution. As such, the present investigation cannot offer insights on the disparity in this behaviour; more research is required to answer this question.

In this study, agent personalization is warranted to decrease the user perceptions of an agent's intrusiveness. At the same time, little is known about the degree of an agent's personalization that is required to accomplish this task and about the extent to which people are ready to configure the agent. For example, some individuals may not want to spend much time learning complex personalization features or frequently adjusting an agent. At some point, users may feel that personalization features themselves lead to further intrusiveness and annoyance. More research is needed to understand this issue.

Future researchers may also study the user perceptions of the importance of different types of negative incidents associated with the employment of email notification agents. By knowing the relative weight of each negative incident cause reported by the users in this study, designers would be able to focus their efforts on the most critical ones. It is also important to know what features people really like and dislike. For this, another experiment is required.

The present investigation discovered that if people perceive an email interface agent to be highly intrusive, they immediately stop using it. The issue of perceived intrusiveness of information technologies has already been studied by the HCI research community. For example, Perry et al. (2001) and Love and Perry (2004) report on the issue of perceived intrusiveness of mobile phones. At Microsoft Research, several projects have been initiated with the goal to minimize unnecessary interactions caused by computers (Horvitz et al., 1998, 1999, 2003). At the same time, the extant HCI literature does not provide a clear definition of perceived intrusiveness, misses measurement instruments, and lacks recommendations on the manipulation of user perceptions of technology intrusiveness. It is recommended that agent developers start investing in research projects that investigate the influence of perceived intrusiveness of interface agents and the methods to manipulate user perceptions of an agent's intrusiveness.

5.4. Conclusions

The goal of this study was to construct the typical scenarios of email interface agent user behaviours in cases of positive and negative critical incidents as well as to present several recommendations for agent designers and marketers. The successful employment of a self-administered questionnaire to obtain critical incidents with agents confirms the fruitfulness of this data collection procedure as part of the critical incident technique. Consistent with prior research, the utilization of the classical content analysis method is believed to be adequate, and the obtained critical events were presented in form of models that serve as a useful visual representation of incidents and their outcomes.

It is suggested that the results of this investigation may be of interest to both academics and practitioners. Especially, it is hoped that agent designers and marketers will be able to utilize a number of practical recommendations based on the findings. For example, to improve user acceptance of email interface agents, developers need to reduce the degree of an agent's perceived intrusiveness. Agent marketers should advertise agents by employing non-agent technologies; otherwise many people will not be able to try out agents. Overall, it is believed that this study may potentially improve our understanding of the human–agent interaction field and facilitate the creation of useful interface agents for email notification.

Acknowledgements

The author is grateful to three anonymous reviewers who provided very valuable feedback on a previous version of this paper. A very special thanks goes to Dr. B. Detlor, Dr. N. Bontis, and Dr. M. Head from the DeGroote School of Business, McMaster University for their support during the author's Ph.D. studies.

References

Andersson, B.-E., Nilsson, S.-G., 1964. Studies in the reliability and validity of the critical incident technique. Journal of Applied Psychology 48, 398–403.

Benbasat, I., Goldstein, D.K., Mead, M., 1987. The case research strategy in studies of information systems. MIS Quarterly 11, 369–386.

Bergman, R., Griss, M., Staelin, C., 2002. A personal email assistant. Technical Report HPL-2002-236, Hewlett-Packard Company.

Biemans, W.G., 1991. User and third-party involvement in developing medical equipment innovations. Technovation 11, 163–182.

Brzezinski, J., Dain, M., 2001. MEME—a society of virtual characters in an adaptive information sharing system. In: Proceedings of the 10th IEEE International Conference on Fuzzy Systems, IEEE, Melbourne, Vic., Australia, pp. 1343–1346.

Budd, R.W., Thorp, R.K., 1963. An Introduction to Content Analysis, Including Annotated Bibliography. University of Iowa, School of Journalism, Iowa City, Iowa.

Budd, R.W., Thorp, R.K., Donohew, L., 1967. Content Analysis of Communications. The Macmillan Company, New York.

Cooper, R.G., Brentani, U., 1991. New industrial financial services: what distinguishes the winners. Journal of Product Innovation Management 8, 75–137.

Cowie, H., Naylor, P., Rivers, I., Smith, P.K., Pereira, B., 2002. Measuring workplace bullying. Aggression and Violent Behavior 7, 33–51.

Dabbish, L. A., Kraut, R. E., Fussell, S., Kiesler, S., 2005. Understanding email use: predicting action on a message. In: Proceedings of the SIGCHI Conference on Human Factors and Computing Systems, The ACM Press, Portland, Oregon, pp. 691–700.

Davis, F.D., 1989. Perceived usefulness, perceived ease of use and user acceptance of information technology. MIS Quarterly 13, 319–340.

Dehn, D.M., van Mulken, S., 2000. The impact of animated interface agents: a review of empirical research. International Journal of Human–Computer Studies 52, 1–22.

Detlor, B., 2004. Towards Knowledge Portals: From Human Issues to Intelligent Agents. Kluwer Academic Publishers, Dordrecht, The Netherlands.

Dillman, D.A., 1999. Mail and Internet Surveys: The Tailored Design Method. Wiley, New York.

Flanagan, J.C., 1954. The critical incident technique. Psychological Bulletin 5, 327–358.

Florea, A.M., Moldovanu, A., 1996. An intelligent email agent. In: Proceedings of the Fourth Romanian Conference on Open Systems, Bucharest.

Griss, M., Letsinger, R., Cowan, D., Sayers, C., VanHilst, M., Kessler, R., 2002. CoolAgent: Intelligent digital assistants for mobile professionals—phase 1 retrospective. HP Laboratories Report HPL-2002-55(R1), Hewlett-Packard Company.

Gruen, D., Sidner, C., Boettner, C., Rich, C., 1999. A collaborative assistant for Email. In: Proceedings of the Conference on Human Factors and Computing Systems. The ACM Press, Pittsburg, Pennsylvania, pp. 196–197.

Helfman, J.I., Isbell, C.L., 1995. Ishmail: immediate identification of important information. Technical Report, AT&T Bell Laboratories.

Horvitz, E., Breese, J., Heckerman, D., Hovel, D., Rommelse, K., 1998. The Lumiere project: Bayesian user modeling for inferring the goals

and needs of software users. In: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, WI, pp. 256–265.

Horvitz, E., Jacobs, A., Hovel, D., 1999. Attention-sensitive alerting. In: Proceedings of the Fifteenth Conference on Uncertainty and Artificial Intelligence, Morgan Kaufmann, Stockholm, Sweden, pp. 305–313.

Horvitz, E., Kadie, C., Paek, T., Hovel, D., 2003. Attentive user interfaces: models of attention in computing and communication: from principles to applications. Communications of the ACM 46, 52–59.

Jennings, N.R., Norman, T.J., Faratin, P., O'Brien, P., Odgers, B., 2000. Autonomous agents for business process management. International Journal of Applied Artificial Intelligence 14, 145–189.

Johnson Jr., B.L., Fauske, J.R., 2000. Principals and the political economy of environmental enactment. Educational Administration Quarterly 36, 159–185.

Jones, M.A., 1999. Entertaining shopping experiences: an exploratory investigation. Journal of Retailing and Consumer Services 6, 129–139.

Karnouskos, S., Vasilakos, A., 2002. Neuro-fuzzy applications: active electronic mail. In: Proceedings of the 2002 ACM Symposium on Applied Computing. The ACM Press, New York, pp. 801–806.

Kautz, H.A., Selman, B., Coen, M., 1994. Bottom-up design of software agents. Communications of the ACM 37, 143–146.

Keaveney, S.M., 1995. Customer switching behavior in service industries: an exploratory study. Journal of Marketing 59, 71–82.

Kerlinger, F.N., 1973. Foundations of Behavioral Research. Holt, Rinehart and Winston, New York.

Koenemann-Belliveau, J., Carroll, J.M., Rosson, M.B., Singley, M.K., 1994. Comparative usability evaluation: critical incidents and critical threads. In: Proceedings of the ACM Conference on Human Factors in Computing Systems. The ACM Press, Boston, MA, pp. 245–251.

Krippendorff, K., 1980. Content Analysis: An Introduction to its Methodology. Sage Publications, Beverly Hills, CA.

Lashkari, Y., Metral, M., Maes, P., 1994. Collaborative interface agents. In: Proceedings of the Twelfth National Conference on Artificial Intelligence. AAAI Press, Seattle, Washington, pp. 444–450.

Laurel, B., 1997. Interface agents: metaphors with character. In: Bradshaw, J.M. (Ed.), Software Agents. The AAAI Press/The MIT Press, Menlo Park, CA, pp. 67–78.

Libes, D., 1997. Tcl/Tk-based agents for mail and news notification. Software: Practice and Experience 27, 481–493.

Lieberman, H., Selker, T., 2003. Agents for the user interface. In: Bradshaw, J.M. (Ed.), Handbook of Agent Technology. The MIT Press, Cambridge, MA.

Lieberman, H., Rosenzweig, E., Singh, P., 2001. Aria: an agent for annotating and retrieving images. IEEE Computer 34, 57–62.

Love, S., Perry, M., 2004. Dealing with mobile conversations in public places: some implications for the design of socially intrusive technologies. In: Proceedings of the ACM Conference on Human Factors and Computing Systems. The ACM Press, Vienna, Austria, pp. 1195–1198.

Lüthje, C., 2004. Characteristics of innovating users in a consumer goods field: an empirical study of sport-related product consumers. Technovation 24, 683–695.

Maes, P., 1994. Agents that reduce work and information overload. Communications of the ACM 37, 31–40.

Maes, P., Kozierok, R., 1993. Learning interface agents. In: Proceedings of the Eleventh National Conference on Artificial Intelligence. The MIT Press, Washington, DC, pp. 459–465.

Miles, M.B., Huberman, A.M., 1994. Qualitative Data Analysis: An Expanded Sourcebook. Sage Publications, Thousand Oaks.

Muylle, S., Moenaert, R., Despontin, M., 2004. The conceptualization and empirical validation of web site user satisfaction. Information & Management 41, 543–560.

Myers, M., 1998. Qualitative research in information systems. MIS Quarterly Discovery, third ed. Available online at 〈http://www.misq.org/discovery/MISQD_isworld〉. Retrieved January 11, 2004.

Nemeth, E., Snyder, G., Seebass, S., Hein, T.R., 2000. UNIX System Administration Handbook. Prentice-Hall, Upper Saddle River, NJ.

Payne, T.R., Edwards, P., 1997. Interface agents that learn: an investigation of learning issues in a mail agent interface. Applied Artificial Intelligence 11, 1–32.

Perry, M., O'Hara, K., Sellen, A., Brown, B., Harper, R., 2001. Dealing with mobility: understanding access anytime, anywhere. ACM Transactions on Computer–Human Interaction 8, 323–347.

Riffe, D., Lacy, S., Fico, F.G., 1998. Analyzing Media Messages: Using Quantitative Content Analysis in Research. Lawrence Erlbaum Associates, Mahwah, NJ.

Rogers, E.M., 2003. Diffusion of Innovations. Free Press, New York.

Rohall, S.L., Gruen, D., 2002. ReMail: a reinvented email prototype. Technical Report TR2002-13, IBM Watson Research Center.

Ronan, W.W., Latham, G.P., 1974. The reliability and validity of the critical incident technique: a closer look. Studies in Personnel Psychology 6, 53–64.

Rubin, J., 1994. Handbook of Usability Testing: How to Plan, Design, and Conduct Effective Tests. Wiley, New York.

Ryan, G.W., Bernard, H.R., 2000. Data management and analysis methods. In: Denzin, N.K., Lincoln, Y.S. (Eds.), Handbook of Qualitative Research. Sage Publications, Thousand Oaks, CA, pp. 769–802.

Segal, R.B., Kephart, J.O., 1999. MailCat: an intelligent assistant for organizing e-mail. In: Proceedings of the Third International Conference on Autonomous Agents. The ACM Press, New York, pp. 276–282.

Segal, R.B., Kephart, J.O., 2000. Incremental learning in SwiftFile. In: Proceedings of the Seventh International Conference on Machine Learning. Morgan Kaufmann Publishers, Standord, CA, pp. 863–870.

Serenko, A., 2005. User adoption of interface agents for email. Unpublished Doctoral Dissertation, DeGroote School of Business, McMaster University, Hamilton, Canada.

Serenko, A., Bontis, N., Detlor, B., 2006. Understanding the user experience with interface agents in everyday work applications. Behaviour & Information Technology, in press.

Serenko, A., Detlor, B., 2004. Intelligent agents as innovations. Artificial Intelligence & Society 18, 364–381.

Tay, C., Ang, S., 1994. User competence in information technology: a critical incident analysis. In: Proceedings of the IEEE Region 10's Ninth Annual International Conference, Singapore, pp. 25–32.

Urquhart, C., Light, A., Thomas, R., Barker, A., Yeoman, A., Cooper, J., Armstrong, C., Fenton, R., Lonsdale, R., Spink, S., 2003. Critical incident technique and explicitation interviewing in studies of information behavior. Library & Information Science Research 25, 63–88.

Voss, I., 2004. The semantic of episodes in communication with the anthropomorphic interface agent MAX. In: Proceedings of the Ninth International Conference on Intelligent User Interface. The ACM Press, Funchal, Madeira, Portugal, pp. 343–345.

Wang, K.-C., Hsieh, A.-T., Huan, T.-C., 2000. Critical service features in group package tour: an exploratory research. Tourism Management 21, 177–189.

Wooldridge, M., 2002. Introduction to MultiAgent Systems. Wiley, New York.