Interacting with Computers, vol. 14, no. 5 (2002), pp. 575-599.
Preprint version

Trust in Information Sources: Seeking Information from People, Documents, and Virtual Agents

Morten Hertzum, Hans H. K. Andersen, Verner Andersen, and Camilla B. Hansen

Centre for Human-Machine Interaction, Risø National Laboratory, Denmark

Abstract. The notion of trust has been virtually absent from most work on how people assess and choose their information sources. Based on two empirical cases this study shows that software engineers and users of e-commerce websites devote a lot of attention to considerations about the trustworthiness of their sources, which include people, documents, and virtual agents. In the project-based software engineering environment trust tends to be a collaborative issue and the studied software engineers normally know their sources first hand or have them recommended by colleagues. Outside this network people are cautious and alert to even feeble cues about source trustworthiness. For example, users of e-commerce websites – generally perceived as single-user environments – react rather strongly to the visual appearance of virtual agents, though this is clearly a surface attribute. Across the two cases people need access to their sources in ways that enable them to assess source trustworthiness, access alone is not enough.

Keywords: Trust, information-seeking behaviour, information sources, personified virtual agents, usability criteria

1 Introduction

In seeking information, people choose to rely on some information sources while others are left unexploited. The factors that enter into people's assessment and choice of information sources include whether the source is oral or written, human or virtual, in-house or external, lay or authoritative, easily accessible or hard to get at, and whether the source contains information or pointers to information. Cutting across this multitude of information sources and an equally varied web of information needs, several studies of engineers' information-seeking behaviour find that the cost associated with the use of an information source is the most important determinant of its use (King, Casto, & Jones, 1994; Pinelli et al., 1993). Without implying that cost is unimportant, this study provides evidence of the importance of trustworthiness – which is closely linked to perceived quality – in people's assessment and choice of information sources.

People engaged in cooperative tasks spend a substantial amount of their time exchanging information. Several studies provide evidence that engineers spend 40%-66% of their time communicating in order to get input to their work and to output results from their work (King, Casto, & Jones, 1994). These studies also indicate that the individuals who spend more time communicating perform better, just as projects where the staff communicates more reach better outcomes. This emphasises the importance of information sharing and, by implication, the potential of basing information systems on a sound understanding of the ways in which people assess and choose information and information sources. To contribute to such an understanding this paper provides an analysis of data from two empirical studies. The first is a field study of the early stages of a large software engineering project. Based on observations of project meetings this study investigates how trust enters into the software engineers' discussions of potential people sources and document sources. The second empirical study is a focus-group study of people's attitudes to virtual agents on e-commerce websites. This study investigates what people require in order for them to seek information from virtual agents and how the appearance of virtual agents affects people's readiness to place trust in them. Collectively, the two studies provide an opportunity for comparing and contrasting how trust affects people's infomation-seeking behaviour in relation to people, documents, and virtual agents.

Section 2 outlines the concept of trust in relation to information sources and emphasises the close links between trust and perceived quality. Section 3 briefly reviews previous work on the factors that enter into people's assessment and choice of information sources. With regard to engineers' choice of people and document sources this work has been summarised in the least-effort principle, with regard to document sources in general this work concerns the concept of relevance, and with regard to virtual agents (and information systems) such work is performed under the heading of usability. Section 4 concerns the study of trust in software engineers' discussions of their information sources, and Section 5 concerns the study of trust in virtual e-commerce agents. Section 6 provides a discussion of how trust in different types of information sources is built and destroyed and how access to information sources must be accompanied by opportunities to assess source trustworthiness.

2 The concept of trust

The concept of trust entails an assumption of risks, which make it impossible to attain certainty. Depending on the nature of the risk being assumed, trustworthiness may mean discretion, reliability, competence, integrity, empathy, and so forth (Sheppard & Sherman, 1998). As evidenced by the wide range of fields in which trust is studied, some form of trust is inherent in all human relationships (see, e.g., Barber, 1983; Kramer & Tyler, 1996; Rosenbloom, 2000; Rousseau et al., 1998). Specifically, trust is at play whenever people exchange information. While an information source – the trusted party – may have a moral responsibility toward an information seeker – the trusting party – there is a risk that the source is not able or willing to live up to this responsibility. To the information seeker trust involves an assessment of whether the other person possesses the required knowledge and skills (does she know?) and is likely to give a truthful and unbiased account of what she knows (will she tell?). Further, the decision to place trust in something requires that the information seeker has confidence in his ability to correctly interpret the communicated message (do I understand?).

Siemieniuch and Sinclair (1999) identify trust – rather than technical quality, expertise, or management – as the real glue that binds a company together. Trust becomes a key issue because properties such as quality and expertise are perceived properties (Tseng & Fogg, 1999). Thus in looking for information of high quality, people are looking for information that is (1) accessible in a way that enables the person to form a perception of its quality and (2) perceived to be of high quality. The first step is necessary because the quality of the information does not reside in the information as a label that can be read but has to be established actively by the individual person. Establishing the perceived quality of a source or piece of information is essentially a matter of establishing to what extent one is willing to place trust in it. People place trust in each other to varying degrees, depending on numerous situational factors. They know their sources to be reliable in some domains but not in others, they collect multiple sources of evidence to safeguard themselves against actors with inadequate capabilities or deceiving intentions, and they engage in a lot of communication to build and maintain a network of people they can turn to for advice and inside information. Tseng and Fogg (1999) distinguish four types of trust by means of the evidence on which the trust is founded:

- First-hand experience (e.g., interacting with people over time, we assess their expertise and trustworthiness).
- Reputation; that is, what third parties have reported (e.g., asking someone for advice based on having her recommended by a colleague).
- Simple inspection of surface attributes (e.g., assessing people by the way they dress or the language they use).
- General assumptions and stereotypes (e.g., believing that your friends tell the truth, whereas car salespeople do not).

The four types of trust differ with respect to the amount of evidence involved. Thus, knowing an information source first-hand or knowing someone who knows it first-hand provides people with a more solid basis for assessing the trustworthiness of the source. This shows that time and previous encounters are key ingredients when firm trust is to be built. In contrast, trust may break down instantaneously (e.g., if the trusting party encounters that he is being cheated) and then be nearly impossible to rebuild. This asymmetry makes it particularly challenging to create virtual agents since one breach in the perceived trustworthiness of an agent may permanently damage users' trust in it.

3 Previous work on people's assessment of information sources

Previous work on the factors that enter into people's assessment of people, documents, and IT systems as sources of information has been conducted under the headings of the least-effort principle, relevance, and usability. Common to these three strands of research is a development from a belief in and focus on objective factors toward recognition of increasingly subjective and situation-dependent factors. Until recently the concept of trust has been virtually absent from this research.

3.1 The least-effort principle

Studies of engineers' information-seeking behaviour have investigated the factors that affect engineers' choice of information sources – primarily people but also documents. These studies find that engineers tend to rely on internal information sources such as their colleagues and personal files, as opposed to libraries and sources external to their organisation (King, Casto, & Jones, 1994; Pinelli et al., 1993). In explaining this behaviour the studies make use of a distinction between two aggregate factors that affect engineers' choice of information sources:

- *Cost*, which concerns the amount of time, effort, and other resources that must be expended to access and make use of a source.
- *Quality*, which concerns the technical quality of the information that can be obtained from a source.

Based on this distinction between cost and quality, the studies report that cost is the most important determinant of source use. For example, one major review of engineers' information-seeking behaviour states that "*clearly, most communication researchers have shown that ease of use or time required dictates information sources used by engineers*" (King, Casto, & Jones, 1994). That is, in selecting among information sources engineers seem to follow a principle of least effort and thus do not use information quality as their main criterion for source selection. Gerstberger and Allen (1968) studied how the frequency with which 19 electronics engineers used a range of information sources was affected by four factors: accessibility, ease of use, technical quality, and degree of experience. The results of their study show a strong correlation between accessibility and frequency of use, a somewhat weaker correlation between ease of use and frequency of use, and a weak correlation between technical quality and frequency of use. Additional evidence in support of the least-effort principle is, for example, reported by Chakrabarti, Feineman, and Fuentevilla (1983) and Rosenberg (1967).

The extent to which an engineer is willing to place trust in an information source can, however, only be determined if the source is accessible in a way that allows the engineer to form a perception of its quality. Survey-based studies (such as Chakrabarti, Feineman, & Fuentevilla, 1983; Gerstberger & Allen, 1968; Rosenberg, 1967) tend to under-recognise the extent and richness of the process that goes into making sense of objects and events. Instead, these studies adopt a more aggregated, context-free notion of technical quality. Several studies of trust analyse issues of key importance to the choice of information sources. These studies report that people consider others to be more trustworthy when they are from the same organisation (Zucker et al., 1996), the same function within an organisation (Moenaert et al., 1992), and physically co-present (Valley, Moag, & Bazerman, 1998). Thus, choosing the information sources that people perceive to be the more trustworthy seems to lead to a preference for the same sources as a choice based on the least-effort principle. This challenges the least-effort principle but, so far, studies of trust have not had much impact on the literature on engineers' information-seeking behaviour.

3.2 Relevance

Within information science people's assessment and choice of information, primarily in textual form, have been discussed under the heading of relevance. When people look for information they are not looking for any kind of information but for information relevant to their problem or information need. Thus, the effectiveness of information retrieval processes is conventionally measured by the extent to which they succeed in retrieving all relevant documents and only relevant documents. Though we intuitively and readily understand what 'relevant' means the various attempts to pin this understanding down in definitions have spurred lots of discussion (see, e.g., Mizzaro, 1997; Saracevic, 1975; Schamber, Eisenberg, & Nilan, 1990). In a recent review, Saracevic (1996) delineates five kinds of relevance:

- *Algorithmic relevance*, which concerns the relation between a corpus of documents and a query as determined by the retrieval algorithm of a system. This kind of relevance determines whether a document is returned as a hit from the query and, possibly, the rank-order of the hits.
- *Topical relevance*, which concerns the relation between the subject/topic covered by a corpus of documents and the subject/topic expressed in a query. The criterion by which this kind of relevance is inferred is aboutness.
- *Cognitive relevance (pertinence)*, which concerns the relation between a corpus of documents and a user's information need or state of knowledge. The criteria by which this kind of relevance is inferred are informativeness, novelty, information quality, and the like.
- *Situational relevance (utility)*, which concerns the relation between a corpus of documents and a situation, task, or problem. The criteria by which this kind of relevance is inferred are usefulness in decision making, appropriateness in problem resolution, reduction of uncertainty, and the like.
- *Motivational relevance*, which concerns the relation between a corpus of documents and a user's intentions, goals, and motivations. The criteria by which this kind of relevance is inferred are satisfaction, success, accomplishment, and the like.

Whereas algorithmic relevance is objective in that the relevance assessments follow deterministically from the retrieval algorithm, the other kinds of relevance are subjective in that the individual user has the final say in the assessment of these four kinds of relevance. Numerous studies report modest correspondence between the hits returned by information retrieval systems – algorithmic relevance – and users' assessments of which documents are relevant to their needs (Blair & Maron, 1985; Swanson, 1988). Further, Tang and Solomon (1998) demonstrate how the same person may assess the relevance of a document differently at different points in time and, thus, illustrate

that relevance is not only multidimensional and subjective but also dynamic. Finally, Harter (1992) sees a contradiction between topical relevance, which is the most common definition of relevance, and pertinence. He argues that the most important documents are the ones that cause cognitive change by allowing new intellectual connections to be made. These documents will often not be on the topic expressed in the user's query but instead introduce novel aspects relevant to the user's information need but not originally considered by the user.

Barry and Schamber (1998) conducted two independent, empirical studies of the relevance criteria users apply in assessing documents and compared the results. Their studies confirm that pertinence, utility, and motivational relevance are important to users, whereas topical relevance did not come out as an explicit criterion. In addition, the users applied criteria relating to the cost (e.g., accessibility and availability) and quality (e.g., accuracy, depth, and tangibility) of the documents. Such criteria are well known from the studies of engineers' information-seeking behaviour (see Section 3.1) but since they are not directly concerned with the user's information need they do not fit into the definition of relevance as a relation between a corpus of documents and different aspects of an information need. This suggests that in assessing documents users are concerned with other factors apart from relevance and, referring to the argument in Section 3.1, trust may be one of these factors. Studies of how users' assessments of document relevance are affected by the amount of information available – bibliographic information, keywords, title, or abstract – show that abstracts are by far the most important and have the greatest impact (e.g., Janes, 1991). This provides some evidence that users attach more importance to the increased trustworthiness with which they can assess documents when abstracts are available than to the increased cost of taking in the extra information.

3.3 Usability

Within human-computer interaction (HCI) people's assessment and choice of information sources such as information systems and virtual agents have been discussed under the heading of usability. Although the importance of usability is gaining widespread recognition, considerable confusion exists over the actual meaning of the term. Sometimes usability is defined quite narrowly and distinguished from, for example, utility (Nielsen, 1993). On other occasions usability is defined as a broad concept synonymous to quality in use (Bevan, 1995). Notably, a considerable portion of the work on usability has been concerned with (the detection of) usability problems rather than with the concept of usability as such (e.g., Lavery, Cockton, & Atkinson, 1997; Lewis, 2001; Olson & Moran, 1998). According to the ISO definition usability is a broad concept consisting of three distinct aspects (ISO, 1998):

- *Effectiveness*, which is the accuracy and completeness with which users achieve certain goals. Indicators of effectiveness include quality of solution and error rates.
- *Efficiency*, which is the relation between (1) the accuracy and completeness with which users achieve certain goals and (2) the resources expended in achieving them. Indicators of efficiency include task completion time and learning time.
- *Satisfaction*, which is the users' comfort with and positive attitudes toward the use of a system. Users' satisfaction can be measured by attitude rating scales such as the software usability measurement inventory, SUMI (Kirakowski & Corbett, 1993).

While it is tempting to assume simple, general relations between effectiveness, efficiency, and satisfaction, this does not seem to be the case. In a study of information retrieval, Frøkjær, Hertzum, and Hornbæk (2000) found only a weak correlation between efficiency, measured by task completion time, and effectiveness, measured by quality of solution. Only two percent of the variation in quality of solution could be predicted from the 83 subjects' task completion time. This suggests that, at least for complex tasks, efficiency is useless as an indicator of effectiveness. Just as some studies assume that faster is better (i.e., that efficiency can be used as a measure of overall usability) several of the instruments for measuring satisfaction seem to conflate satisfaction with effectiveness and efficiency. For example, SUMI (Kirakowski & Corbett, 1993) contains five sub scales: efficiency, helpfulness, control, learnability, and positive affect. Except for positive affect, all of these sub scales concern the effectiveness and efficiency aspects of usability. Nielsen and Levy (1994) find that in 25% of the cases they studied the users did not prefer the system they were more efficient in using. This invalidates the use of efficiency as a general indicator of satisfaction and questions the possibility of finding any simple relationship between satisfaction and performance indicators such as effectiveness and efficiency. Hassenzahl (2001) finds that hedonic quality ('fun factors' such as originality and innovativeness) is a perception that is independent of the subjects' perceptions of effectiveness and efficiency, and that hedonic quality contributes to the subjects' judgement of overall system appealingness. Hedonic quality seems to include additional factors that have not hitherto been captured in measurements of usability.

Although performance-biased perspectives on usability are giving way to more user-centred, satisfaction-aware perspectives the notion of trust has been notably absent until the concept of usability started to appear in discussions of websites and e-commerce (e.g., Light & Wakeman, 2000; Rosenbloom, 2000). As an example, Light and Wakeman (2000) demonstrate that personalisation of a website by putting statements in the mouth of individuals is a subtle issue. On the one hand users find that personalisation improves the accountability of the site by increased

relationship building. On the other hand users find it compromising to hand over personal or financial information to a specified person at the other end. Thus especially when users are requested to enter information, websites are safest using a neutral corporate voice. Personalisation appears to lessen perceptions of professionalism and this divides users into those who mistrust the service and those who welcome the informality. The study provides evidence that users' willingness to engage with a website is strongly dependent on the extent to which the site succeeds in conveying accountability and trustworthiness.

4 The CSA project: Trust in human actors

The first case, which is based on Hertzum (2002)[1], concerns the assessment and choice of people sources in a software engineering project. The software engineers involved in the project devoted a lot of attention to determining the trustworthiness of their people sources and comparatively less attention to the cost of using them. This use pattern is contrasted with the software engineers' use of document sources.

4.1 Introduction

The company where the field study took place is a large software house, which has developed and marketed a range of systems for use in municipal institutions. The studied project concerns a system to support municipal authorities in the handling of cases concerning child support and alimony (CSA). The CSA project is to completely redevelop the company's existing CSA system, which has been in operation for almost two decades. According to the project plan, the CSA project will last three years. The first eight months of the project, the period analysed in this study, concerned the requirements specification and the business modelling. During this period the project was staffed with a project manager, eleven designers/developers, two service consultants, a methods & tools consultant, a usability specialist, and a secretary. The project manager and six of the designers/developers worked full time on the CSA project, the remaining ten persons were assigned to the CSA project on a part-time basis. This multidisciplinary group – in the following termed the CSA engineers – is comprised of people with an average of more than ten years of professional experience.

To accomplish their task, the CSA engineers have to interact with external stakeholders such as user representatives and the governmental bodies responsible for the legislation regarding child support and alimony. Naturally, they also have to interact with management, marketing, technical services, the quality function, and other stakeholders internal to the company. Whereas the existing CSA system contains substantial amounts of code that duplicate functionality from other systems made by the company, the new CSA system will distribute this functionality onto components that are to be developed by other project groups in the company. This adoption of component-based design means that the CSA engineers have to cooperate closely with a number of people outside the project to negotiate, settle, and follow up on component definitions and how the development of components progresses. In addition to component-based design, the project involves a move to a technological platform with which the CSA engineers are less familiar – the Web.

4.2 Method

The study is based on observation and analysis of the 16 fortnightly project meetings that took place during the formative eight-month period from the initiation of the CSA project, through the requirements specification, to the completion of the business modelling. The main purposes of the meetings were to provide a forum for sharing information about the status of the project, maintaining awareness of the entire project, coordinating activities, discussing problems and progress, making decisions, and reviewing major project documents. During the meetings a researcher was seated at the meeting table with the other people present. From their point of view the researcher was invisible in that he was not to be spoken to and has himself remained silent. The meetings were recorded on tape and transcribed. Supplementary information was collected by participating in the two-day, project start-up seminar, conducting interviews with eleven of the core project participants, and inspecting various project documents.

The data analysis involved two passes. First, nine transcripts were examined sentence by sentence and all references to information sources were marked up and annotated. This bottom-up analysis, combined with findings from the literature (e.g., Allen, 1977; King, Casto, & Jones, 1994), provided the input for creating a coding scheme. Second, all 16 transcripts were examined to identify the incidents involving information sources and categorise them according to the coding scheme. The information-seeking incidents comprise the situations where the project participants make each other aware of available sources, discuss the quality of sources, decide on the sources to

[1] The first case is reprinted, in revised form, from *Information and Organization*, Vol. 12, M. Hertzum, The importance of trust in software engineers' assessment and choice of information sources, pp. 1-18. Copyright (2002), with permission from Elsevier Science.

approach in a variety of situations, and follow up on whether appropriate sources have been involved in settling complicated issues.

The incidents were coded with respect to the reasons for discussing, selecting, and referring to the information source. This coding employs a primary distinction between *quality-related factors* and *cost-related factors*. The quality-related factors concern technical quality, the appropriateness of the source's formal or practical background, and the fit between this background and the current situation (see Table 1). The cost-related factors concern the resources and/or difficulties involved in accessing and using a source: accessibility, ease of use, and cost to use. In addition, the incidents were coded with respect to whether they involved: (1) *People* or *documents*. People sources include individuals, project groups, and organisations. Document sources include electronic and paper documents as well as information systems. (2) *Information* or *commitment*. In the case of information, the source delivers insights, opinions, and other pieces of information. In the case of commitment, the source takes on a task, that is, agrees to perform a specified piece of work.

To give a feel for the data, two sample incidents are reproduced below. The first involves a *people* source, which has *committed* to do something but is not the *appropriate person for this task*. The second involves a *document* source, which provides *information* and is brought up because it is easily *accessible*.

1. *You have had the wrong consultant – at the wrong point in time. You should have had a business modeller in the beginning, to help you make the business model. And you haven't had that.*
2. *It is those CICS statistics I'm thinking of. Can we use them? We made those statistics for [a government institution] so they are readily available.*

Table 1. The seven quality-related factors used in coding why information sources were discussed, selected, and referred to.

Quality-related factor	Description
Appropriate organisational unit	The source comes from the organisational unit formally vested with the right authority and competence
Appropriate project experience	The source has hands-on experience with the issue from past or ongoing projects
Appropriate external body	The source comes from the external organisation formally vested with the right authority and competence
Appropriateness to task	The goodness/badness of the fit between the source and the task; i.e., focus is on the relation between source and task, not on the source as such
Technical quality	The high/low technical quality of the information that can be obtained from the source, irrespective of the genesis and appropriateness of the source
Up-to-dateness	The up-to-dateness/outdatedness of the information provided by the source
Representativeness	The extent to which the source is representative of the group it belongs to

4.3 Results and discussion

The 16 meetings included 362 information-seeking incidents. This amounts to an average of one incident every 5.7 minutes. Often, several CSA engineers made different contributions to the description and assessment of a source and consequently the total number of quality-related plus cost-related factors is larger than the number of information-seeking incidents (see Table 2). A total of 62% of the factors brought up during the incidents were concerned with quality, whereas the remaining factors were divided about evenly between cost and other factors. This tendency to give more consideration to quality-related factors was significant for the incidents concerning people sources (Binomial test (281, 436, 0.5), $p<0.0001$) but not for the incidents concerning document sources (Binomial test (78, 144, 0.5), n.s.). Thus, the assessment and choice of people sources was dominated by quality-related factors whereas the document sources displayed a more even balance between quality-related factors and factors related to costs and other things. This could be an indication that the project primarily relied on its people sources and, thus, gave much consideration to their trustworthiness. Alternatively, it could be an indication that people sources were brought to the group's attention in large numbers and thus some pruning was needed, whereas documents were brought up more sparingly and only when their appropriateness was more or less evident. While we cannot make strong claims about why the quality-related factors played a more prominent role in relation to the people sources this difference does reflect that people and documents were experienced as different types of sources, which had to be treated differently.

Table 2. Frequencies of quality-related and cost-related factors mentioned in relation to the information sources.

	People sources	Document sources	Total
Quality-related factors	281 (64%)	78 (54%)	359 (62%)
Cost-related factors	61 (14%)	43 (30%)	104 (18%)
Other factors	94 (22%)	23 (16%)	117 (20%)
Total	436 (100%)	144 (100%)	580 (100%)

Note: Often, several CSA engineers made different contributions to the description and assessment of a source. Thus, the total number of instances of quality-related, cost-related, and other factors exceeds the 362 information-seeking incidents.

Looking in more detail at the quality-related factors we get a picture of the multiple dimensions involved in assessing source quality (see Table 3). The two quality-related factors most frequently involved in the assessment and choice of an information source differ for people sources and document sources. For people sources the two top factors are *appropriate organisational unit* and *appropriate project experience*. The CSA engineers drew on several organisational units that have been established to support development projects and formally vested with the authority to handle certain issues. This included technical services, the usability laboratory, the legal department, and the quality function. While it is mandatory to involve some of these sources, most of them were involved because the CSA engineers decided to use them. In both cases CSA engineers had to form an opinion about whether they considered these people to be in possession of the required knowledge and skills. A number of the information-seeking incidents were directed at establishing whether the sources formally appointed to handle certain issues were at the same time sufficiently competent and, thus, trustworthy. In this regard there was a certain tension between formal expertise and hands-on project experience in that CSA engineers at times preferred colleagues who had experience with an issue from actually working with it over appointed experts. Hands-on experience was, in general, a strong way of gaining trustworthiness.

For the document sources the two quality-related factors that are mentioned most frequently are *appropriateness to task* and *up-to-dateness*. Documents are written within a specific context and reusing them in another context, such as the CSA project, involves an assessment of the extent to which conclusions drawn in the original context are also valid in the new context. When people are in a hurry or simply absorbed in their day-to-day work they are likely to focus on their own sense-making process and not spend time expanding their writings into documents understandable to unknown future readers. Thus, people often make extensive use of condensed forms of writing, which leave most of the context unsaid because the document will be understood by the primary readers as belonging to a certain ongoing activity. In fact, the primary readers can see an elaboration of the condensed forms of writing as redundant and annoying (Brown & Duguid, 1996). When the reader is not familiar with the context in which a document was written it becomes very important to devote the necessary attention to inferring the ways in which the context has affected the document and, thereby, whether the document is appropriate to the current task. Specifically, being out of date is a serious threat to the trustworthiness of documents but hardly an issue in relation to the CSA engineers' assessments of people sources.

Table 3. The quality-related factors mentioned in discussing, selecting, and referring to the information sources.

Quality-related factor	People sources	Document sources	Total
Appropriate organisational unit	64 (23%)	9 (12%)	73 (20%)
Technical quality	55 (20%)	11 (14%)	66 (18%)
Appropriate project experience	59 (21%)	2 (3%)	61 (17%)
Appropriateness to task	40 (14%)	20 (26%)	60 (17%)
Appropriate external body	45 (16%)	10 (13%)	55 (15%)
Up-to-dateness	10 (4%)	17 (22%)	27 (8%)
Representativeness	8 (3%)	9 (12%)	17 (5%)
Total	281 (100%)	78 (100%)	359 (100%)

At the CSA meetings, 77% of the information-seeking incidents involved people sources (see Table 4). This is significantly more than the incidents involving document sources (Binomial test (279, 362, 0.5), $p<0.0001$). Zipperer (1993) found that in her study a major reason for the engineers' strong preference for oral communication was that they often sought feedback on their ideas or designs, either as trusted opinion or as impetus for creative discourse. Numerous incidents in the CSA project support this finding. In preparation for meetings with people external to the CSA project the CSA engineers often wrote and distributed documents to allow the participants to prepare themselves. That is, the documents were supplementary sources intended to improve the quality of the primary oral

communication. Also, there were several incidents where CSA engineers sought expert advice after having studied the authoritative written legislation. Although the handling of CSA cases is prescribed in detail in legislation, there is a large gap between the terse texts and the richness of real-world cases. To close this gap, people handling CSA cases have to interpret the legislation relative to the concrete cases with which they are confronted. Over time, this leads to a practice that is based on the legislation but not inherent in it. Thus, to acquire a trustworthy understanding of CSA work the CSA engineers had to talk to users and other experts in the handling of CSA cases. It was not possible for the CSA engineers to simply read the legislation and form a credible interpretation of it – the documents were merely supplementary sources.

The CSA engineers' preference for people sources was also a result of their dependence on a number of activities performed outside their project. To succeed with the strategy of having several project groups develop components for the CSA system, the CSA engineers were often looking just as much for commitment to future actions as they were looking for information. A total of 27% of the information-seeking incidents concern commitment (Table 4). Many of the commitments have to do with defining the functionality of software components and the schedule for their development. This is a delicate process since the people who are being tasked with developing the components are occupied with other tasks and thus have to reschedule part of their activities. Consequently, negotiations are needed to convince other project managers that the CSA project should be given priority and to provide the CSA engineers with an opportunity to determine the strength of the commitments made by these project managers. If a project manager makes what the CSA engineers perceive as a half-hearted commitment they will regard the project manager's promise as less trustworthy than if he or she had made a whole-hearted commitment, though even the strongest commitment is no guarantee of fulfilment. The CSA engineers are very aware of the situations where they ask for their colleagues' commitment to time-consuming activities because that usually involves negotiation. Apart from these situations, the distinction between information and commitment is easily blurred. A request for information may, for example, turn into a commitment if the source is unable to provide the information directly but takes on the task of finding out.

Table 4. The information-seeking incidents divided according to information or commitment.

	People sources	Document sources	Total
Information	182 (65%)	82 (99%)	264 (73%)
Commitment	97 (35%)	1 (1%)	98 (27%)
Total	279 (100%)	83 (100%)	362 (100%)

In the beginning of the CSA project many of the projects, departments, and other information sources that were mentioned at the meetings were unknown to some of the project participants. On a number of occasions speakers were asked questions such as who worked on that project or whether that was the department in which this person worked. Similarly, several speakers volunteered such contextual information. It appeared to be a well-established conversational practice to accompany the mentioning of sources that might be new to some project participants with information that put these sources in context. This served to inhabit new information sources, such as project groups, with known people who could lend the source an initial face and grounding. Thereby otherwise unknown sources inherited an initial level of trustworthiness from people of whom the CSA engineers already had an impression.

5 The COGITO experience: Trust in virtual agents

The second case, which is based on Andersen, Hansen, and Andersen (2001), concerns the elicitation of user needs for personified virtual agents in e-commerce. The users involved in the project attached particular importance to the trustworthiness of agents in relation to providing information that could solve the users' information and purchasing needs.

5.1 Introduction

The increasing amounts of data and functionality available on the Web mean that more and more information needs can be answered by searching the Web but also that this search task is becoming increasingly demanding. Virtual agents are means for assisting users in finding – or becoming aware of – needed information in a swift and situation-specific way. The aim of the user requirements study of the COGITO[2] project was to specify, in general, the appearance and functionality of personified, virtual agents supporting users in information seeking and decision

[2] COGITO "E-commerce with guiding agents based on personalised interaction tools" is a European project partly funded by the European Commission.

making on e-commerce websites. Within the COGITO project the user requirements will feed into the implementation and evaluation of a virtual agent assisting people in using a large Web-based bookstore. The purpose of the virtual agent is to help users of the site – potential customers – navigate the site in an effective and pleasing way and, thereby, increase the possibilities of satisfying the users' information needs as well as the vendor's need for selling products and services. Specifically, the agent will be able to assist users in finding products of interest, searching for specific products, browsing special offers, handling orders and payments, and getting help.

The COGITO agent is personified. Personified agents are virtual agents that mimic humans in terms of their visual appearance and the conversational way in which they interact with their users. The prospects of providing virtual agents with a human-like representation in the user interface are currently subject to much research, and a fixed terminology has yet to emerge. *Personified agents*, the preferred term in this paper, is used by several authors but often interchangeably with terms such as *anthropomorphised agents*, *humanoid agents*, and *conversational agents*. Further, authors interested in a specific modality in the representation of such agents talk about, for example, *facial displays*. In addition to being personified, the COGITO agent is also adaptive in that it attempts to tailor its advice, recommendations, and other responses to the individual user by making inferences from previous interactions with the user. Finally, the COGITO agent will aim at being proactive by suggesting products to users on the basis of user profiles the agent builds up over time.

5.2 Method

The focus of the user requirements study was on why agents are perceived and used the way they are, rather than on what people do when they interact with agents. The data collection consisted of "interviewing while doing" sessions, a focus group meeting, and a word association test, all of which involving the same eight participants. The participants were six females and two males, they were 24 to 38 years of age, all used the Web on a daily basis, and though six of them seldom shopped online all but one had done so at least once. Five of the participants had a background in library and information science; three had university degrees in business or economics.

The "interviewing while doing" sessions exposed the participants to agents on the Web. Each participant individually performed a series of tasks on pre-selected websites. The participants were asked to think out loud while solving the tasks and they were probed about their interaction behaviour, their reactions to the virtual agents, and their attitude toward general e-commerce issues. The focus group meeting came after the "interviewing while doing" sessions and was focused on general e-commerce problems, the visual appearance of agents, and in particular problems and requirements relating to the use of agents. Focus group meetings provide for organised discussion in a multi-person forum (for reviews of the advantages and disadvantages of focus groups see, e.g., Goss & Leinbach, 1996; Morgan, 1997; Nielsen, 1997). The word association test was conducted in association with the focus group meeting and concerned the participants' spontaneous reactions to the visual appearance of personified agents. Images of ten agents were displayed for 90 seconds each and the participants were instructed to write down their immediate associations or feelings toward the agents. The participants were encouraged to comment on any aspect of the agents and to write down their associations and feelings using the words that immediately came to their mind. After an agent had been displayed the participants' comments were collected and the next agent displayed. Discussion of the appearance of the agents did not occur until after all ten agents had been displayed. Seven of the displayed agents were personified virtual agents from different sites on the Web, two were photos of well-known actors rather than virtual agents, and the last was a virtual but not personified agent.

The data analysis consisted of eliciting user requirements and other user reactions toward virtual agents. This was done on the basis of transcripts of the data-collection sessions, which were initially captured on video. Inspired by cognitive systems engineering (Rasmussen, Pejtersen, & Goodstein, 1994), three interrelated levels of user requirements were used in the data analysis: strategic, procedural, and operational. The levels constitute a means-ends hierarchy where the strategic requirements define the ends that give rise to the procedural requirements and the operational requirements comprise the means that are, in turn, needed to achieve the procedural requirements. The strategic requirements reflect the goals and constraints governing the interaction between users and agents, including issues such as trustworthiness. The procedural requirements describe the general functions and capabilities that users expect of virtual agents. The operational requirements represent the actual activities performed by virtual agents, such as asking users for additional information about their information needs. The operational requirements also include the users' requirements toward equipment and physical configurations necessary to provide the services offered by agents, for example cookies and software plug-ins (Andersen, Andersen, & Hansen, 2001). Finally, the data from the word association test was analysed by grouping the participants' reactions to the visual appearance of the agents into three categories: positive, neutral, and negative. No distinctions were made between different degrees of positivity, neutrality, and negativity; rather, the categorisation was accepted as a rough indication of the participants' overall reaction to the agents.

5.3 Results and discussion

Based on the "interviewing while doing" sessions and the focus group meeting we identified six strategic requirements for how personified virtual agents can contribute to effective e-commerce. Trustworthiness came out as a recurrent and crucial strategic requirement. The other strategic requirements were seriousness or reliability, mediation quality, flexibility or tailorability, value-added services, and entertainment. In conjunction with their related procedural and operational requirements several of these other strategic requirements also touch upon the notion of trust. In relation to e-commerce trust is, among other things, a matter of trusting that information given by the customer (e.g., credit card information) will not be misused, that commitments made by the vendor will be fulfilled, and that the purchased products will be of the expected quality (see also Friedman, Kahn, & Howe, 2000; Jones et al., 2000; Shneiderman, 2000).

Virtual agents frequently require that users submit personal or financial information in order to get tailored recommendations, purchase products, and be alerted to new offers. The participants generally doubted that the potential benefits of submitting such information outweighed the risks, and they were thus reluctant to trust virtual agents with information about their personal preferences, email address, and so forth. Several participants were, for example, not prepared to risk that submitted information was sold to other organisations and resulted in loads of junk mail: *"the question is whether it is possible to refuse the subsequent marketing. If not, I would prefer not to [give information]*." In addition to this general cautiousness the participants felt an even greater uncertainty about submitting credit card information: *"I don't buy online if I have to give my credit card number on the net*." Some participants explicitly stated that their reluctance to submit information was a matter of a lack of trust: *"If one had faith [in the company] it could be a good idea to submit [personal] information*." Information about the individual user's preferences is a prerequisite for a virtual agent's possibilities of guiding the user. It is, therefore, essential that agents succeed in establishing a level of trust where the user is willing to submit at least information about personal preferences. It should also be noted that the participants did not maintain a strict distinction between agents as such and the Web at large. Rather, the participants' concerns regarding the Web at large carried over into their perception of virtual e-commerce agents.

If a virtual agent is to appear trustworthy the participants request that the agent is transparent with respect to how it uses submitted information, that it is only proactive when appropriate, and, if possible, that it is recommended by independent third parties:

- Transparency includes that it is made clear who owns the site and how this organisation can be contacted: *"One feels more reassured if a hotline […] is available for contacting a person*." This serves to anchor the virtual agent in the real world. Several participants also find that access to a code of ethics and to information about customer rights can provide a trust-increasing transparency about how the organisation behind the agent is running its business. Finally, it is important not to ask for information that is perceived by the customer as unrelated to the purpose. If customers do not provide all the requested information they should not be rejected but rather be assigned – based on the information available – the best possible general profile instead of a person-specific one.
- The participants were sensitive to the role assumed by agents but no consensus emerged on a preferred role. A passive, neutral mediator may be perceived as trustworthy because the agent, by waiting until called upon and not trying to exert influence on the user, lets the user remain in charge. This, however, leaves the agent little room for making proactive suggestions and assisting users who have less articulate ideas about what they want from the site. If the agent is allowed a more active role it also gets more opportunities for assisting as well as influencing the user. The participants agreed that they would appreciate if the agent – based on the information submitted by the user – could sense when there was a need for assistance. A more active agent may, however, appear less trustworthy to some users because its active involvement in the users' interaction with the site can be experienced as obtrusive and manipulating: *"It reminds me of the annoying salesperson who can't leave me alone."* This can instil the impression that the agent promotes a certain sales strategy on behalf of the vendor, rather than serves the user's best interest.
- Until users get some experience with virtual agents they have little knowledge about what they can expect and what they should look out for. Rather than trusting the various virtual agents at face value, the participants requested ways of getting agents recommended by independent third parties: "*It would help if a newspaper had a user test that gave an authoritative indication of the quality of the site*." A quality indication such as a guarantee from an authorised institution, like a bank, would also be most helpful in building trust in the vendor and encouraging users to submit credit card information. In the future users may also form Web-based communities for the exchange of personal experiences with the use of virtual agents.

In addition to these three main points about how to build trust in virtual agents a number of additional trust-related issues were brought up during the focus group meeting. The participants felt that agents should be able to introduce themselves as well as explain their communicative abilities and the neutral, proactive, or other role they would assume in assisting the user. This would support users in relating to the agent and in assessing its appropriateness relative to the users' information and purchasing needs. It was clearly stated that the participants did not expect one way or level of agent support to fit all situations and that they, in general, felt a need for opportunities to tailor an agent's behaviour to their personal preferences. Furthermore, the personification of the virtual agents meant that most of them required a conversational style of communication in which both agent and user communicated by means of full sentences. The "interviewing while doing" sessions showed that most of the participants made use of single keywords – as if interacting with a search engine – and considered it surprising and artificial that they were to use full sentences. Most of the agents to which the participants were exposed did not react reasonably or at all to input consisting of keywords only. This jeopardised the participants' trust in the agents and seems a patent example of inappropriate personification. Similarly, the participants disapproved of the chat abilities included in some of the agents and stressed that agents should be serious, reliable, and goal-directed in order to inspire confidence. Otherwise the participants would dismiss them as amusing toy features.

The word association test was carried out to unveil the participants' spontaneous reactions to the visual appearance of personified virtual agents. Figure 1 shows the ten agents presented to the participants. Examples of words the participants associated with these ten agents include *confidence inspiring* (associated with Actor2), *workman* (associated with Lego), and *silly* (associated with Tokiama and Cogito). Table 5 shows for each agent the result of breaking the words associated with the agent into positive, neutral, and negative.

Actor2, Eve, Nicole, and Lego were the only agents that had more positive than negative words associated with them. Three of these four agents are represented by a naturalistic, photo-like image of a person, and the COGITO participants perceived them as charming, confidence inspiring and friendly. The agent that got the highest percentage of positive words (Actor2) is actually not a virtual agent but a photo of a real-life actor, who has been heavily exposed in television series lately. This may influence the participants' reactions and could indicate that the use of a known character may be an effective means of shaping people's attitudes toward agents. The subtleties involved in doing this successfully was illustrated by the other photo of a real-life actor (Actor1). Both actors are 'good guys' in the same television series but they triggered quite different reactions from the participants. Actor1 was the only agent that was represented by a naturalistic, photo-like image of a person and did not score well in the word association test. Conversely, the Lego agent was the only agent that was not a naturalistic, photo-like image of a person and still had more positive than negative words associated with it. The Lego agent displays a close and readily perceived match between its visual appearance and the domain it acts within. The participants all agreed that this was a strong positive feature and they expressed a general preference for such appearance-domain matches, which were perceived to make the agent appear more relevant, knowledgeable, and trustworthy.

The agents that received the highest percentages of negative words were Tokiama, Virtual Friend, and Quest. They were perceived as weird, artificial, and silly. Tokiama and Virtual Friend are (as is Cogito) computer-modelled talking heads, which move their lips, eyebrows, and the like to mimic speaking and facial gestures. While the participants did not agree on the most suitable appearance of a personified virtual agent they tended to disapprove of these animated creatures. Quest, which is not personified but simply represented by a text field, was included because some participants claimed (during the "interviewing while doing" sessions) that no personification was needed; the user could just as well interact with a text field. However, when the participants were exposed to this situation during the word association test, they evaluated Quest as impersonal, dull, too much text, and trivial. This tentatively suggests that at least until users get acquainted with the behaviour and functionality of a virtual agent it can gain some initial goodwill by being personified.

Actor1	Actor2	Nicole	Virtual Friend	Lego	Quest	Eve	e-Cyas	Tokiama	Cogito
									

Figure 1. The agents used in the word association test. Note: Seven of the displayed agents are personified virtual agents from different sites on the Web (for copyright reasons the image of the Lego agent is a photo of a Lego figure identical to the Lego agent used in the word association test). The remaining three agents are not personified virtual agents but included to provide contrast: Actor1 and Actor2 are photos of real-world actors and Quest is a virtual but not a personified agent.

Table 5. Outcome of word association test (sorted by decreasing percentage of positive words).

Agent	Positive words	Neutral words	Negative words
Actor2	30 (73%)	6 (15%)	5 (12%)
Eve	22 (61%)	2 (6%)	12 (33%)
Nicole	18 (53%)	3 (9%)	13 (38%)
Lego	15 (52%)	8 (28%)	6 (21%)
e-Cyas	11 (24%)	6 (13%)	28 (62%)
Virtual Friend	7 (18%)	1 (3%)	31 (79%)
Actor1	7 (16%)	6 (14%)	31 (70%)
Cogito	6 (21%)	6 (21%)	16 (57%)
Quest	4 (14%)	3 (11%)	21 (75%)
Tokiama	3 (10%)	3 (10%)	24 (80%)

6 Concluding discussion

People and documents are established information sources. Compared to them virtual agents are just beginning to appear. Consequently, we must expect that people are uncertain about how to relate to virtual agents and alert to clues that can assist them in forming an opinion about virtual agents. While this may make trust issues more readily visible in relation to virtual agents, this study argues that trust is also fundamental to people's assessment and choice of people and document sources.

6.1 The emerging realisation of the importance of trust

Previous work on the factors that affect people's assessment and choice of information sources have embraced a gradually increasing number of factors – such as accessibility, effectiveness, and pertinence – but until recently trust has been virtually absent from this work. Work on, for example, hedonic quality (Hassenzahl, 2001) acknowledges that assessments of information sources are inherently personal. This is a radical extension of more objective factors such as task completion time and technical quality but it merely adds another factor to the list of factors that enter into people's assessment of information sources. Trust is different because it emphasises that even seemingly objective factors are in effect perceived factors. Thus, technical quality does not reside in sources as a label that can be read but has to be established actively by the individual person. This is essentially a matter of establishing to what extent the person is willing to place trust in the source. Thus, trust cuts across many factors and must be incorporated in our understanding of these factors rather than added as yet another factor.

The CSA and COGITO projects show that software engineers as well as potential users of e-commerce websites take on the task of forming an opinion about the trustworthiness of their sources and pay a lot of attention to this task. If a source is infallible there is no risk involved in treating the information it supplies as true, and trust is not an issue. In all other cases there is a risk that the source may be wrong and it then becomes an issue whether to trust it and act on the information it supplies. The CSA engineers and COGITO participants seem to accept that it is ultimately the information seeker's responsibility to assess source trustworthiness. It cannot generally be assumed that sources will inform about their trustworthiness or issue warnings when asked for information outside the areas they know about. Especially providers of e-commerce websites have an economic interest in appearing trustworthy irrespective of whether they actually are. Thus, to be useful as information sources people, documents, and virtual agents must be accessible in ways that enable information seekers to form an opinion about their trustworthiness. These ways vary across contexts, as illustrated by the two cases included in the present study.

6.2 Building trustworthiness

The engineers on the CSA project tend to either know their sources first hand or have them recommended by a colleague. In the beginning of the project the CSA engineers would state their views about sources and thereby contribute to a shared perception of sources important to the project. This seemed to be an element in establishing the CSA engineers as a group. Through this process CSA engineers who knew a source had an opportunity to compare and adjust their perceptions of the source and CSA engineers who did not know a source had an opportunity to get an impression of the source from colleagues who already knew it. This way the CSA engineers did not have to resort to assessing the trustworthiness of sources on the basis of either simple inspection of surface attributes or general assumptions and stereotypes (see Section 2). The COGITO project reflects an expectation that virtual agents will to a substantially larger extent be encountered by people who have no preconceptions of the encountered agents and will use or discard them on the basis of a few moments of interaction. If a virtual agent initially instils negative feelings there are few reasons to put up with it because it will usually be easy to shift to

another agent with similar facilities. This increases the importance of surface attributes such as the visual appearance of the agent. Moreover, several of the COGITO participants felt a need for up-front access to information such as a code of ethics, contact details for the organisation behind the website featuring the agent, and third parties' assessments of the quality of the agent. It is an open question whether the presence of such information would put users at ease but merely by requesting it the participants indicate a strong need for information that can support them in forming a perception of the trustworthiness of the agents.

In an engineering setting such as the CSA project trust is a collaborative issue. The whole project becomes dependent on a source when one CSA engineer makes a decision based on it, and the assessment and selection of sources is therefore an issue frequently raised by the project manager. One of the important roles of the project meetings is to provide a recurrent forum for reviewing the trustworthiness of sources that are being used or considered for use. Though the CSA engineers may not always reach group-wide consensus about the trustworthiness of a source, they will, at least, raise each other's awareness toward the issue in subsequent interactions with the source. In the COGITO project the use of virtual agents was perceived as an activity performed by individual users. As individual users normally know little about the other people using a website the agent may utilise its exclusive access to such information to provide users with information about how other users have assessed the items advertised on the website. As an example, recommender agents inform users that are buying item X that "other customers who bought item X have also bought item Y". The value of such an agent lies in mediating among users and thereby providing individual e-commerce users with some of the information the CSA engineers get from each other at the project meetings. Thus, recommender agents gain some trustworthiness by replicating a mechanism we use repeatedly in collaborative settings.

6.3 Threats to trustworthiness

The major threat to the trustworthiness of an information source is seemingly to be unknown to the user. This is not to say that users distrust unknown sources but that trust is absent – just as distrust is absent – because the user has not yet formed an opinion about the source. If it turns out to be difficult to get information about a source the absence of trust may become a serious concern or even turn into distrust. This is illustrated by an incident in the CSA project where a much needed upgrade of the CSA engineers' computers had not been performed. This was annoying in itself but the CSA engineers were seriously concerned because they were unable to get information about who were responsible for performing the upgrade. Without a name the CSA engineers started to distrust that anybody was responsible. In the virtual world of e-commerce users frequently encounter problems with respect to getting sufficient information about the realities behind a virtual agent. The virtual agents represent vendors but it is often unclear how this impacts on the agent's behaviour, and the possibilities for personal contact are normally scarce. Furthermore, the Web tends to favour the possibilities for shopping and seeking information from a vast number of dispersed vendors over the building of a reservoir of experience based on longstanding collaboration and communication with a few partners.

Other, more specific threats to source trustworthiness include lacking hands-on experience with the issue in question, being out-of-date, and – for personified virtual agents – having a visual appearance that is perceived as cold or out of place relative to the provided services. With respect to providing virtual agents with a personified visual appearance the COGITO experience yields mixed results. On the one hand, the participants claim that they do not need the personification and could just as well interact with 'a text field'. On the other hand, they find the better personifications more engaging, lively, and convincing than a virtual agent represented by a text field only. The finding that personified agents are perceived to be more engaging is in accordance with previous studies (e.g., McBreen & Jack, 2001; Walker, Sproull, & Subramani, 1994). However, previous studies also warn that personification may distract the users (McBreen & Jack, 2001), require more effort from the user (Walker, Sproull, & Subramani, 1994), or even be counterproductive (Shneiderman & Maes, 1997). Judging from the COGITO experience a text field is the safe but somewhat dull choice whereas personification generates more emotions but is also more risky because these emotions are difficult to predict and range from rather positive to distinctly negative.

6.4 Intertwining people, documents, and virtual agents

People, documents, and virtual agents are different types of information sources with different strengths and weaknesses. Consequently, situational factors affect the choice of information source and many situations are resolved through the use of a combination of several types of sources. For example, engineers often look for people in order to have them recommend which documents to read and they often look for internal reports in order to find names of colleagues to contact (Hertzum & Pejtersen, 2000). Analogously, virtual agents can mix information they provide themselves with referral services where the agent refers users to people, documents, or other virtual agents. This way virtual agents can be incorporated in a network of sources where people who consider using an agent get some indication of its trustworthiness from the source that refers them to the agent. If users trust the source referring them to an agent for the first time it is likely that this trust will be extended to the agent, at least until the users start

interacting with the agent and forming their own opinion about it. In the absence of references, users have no information about the trustworthiness of an agent upon first approaching it. This increases the impact of surface attributes such as the agent's visual appearance because any indication of the trustworthiness of the agent is perceived as an improvement over the initial absence of information. This is likely to make users receptive to feeble but readily available cues.

Potential customers take many aspects of the appearance and functionality of virtual agents into consideration when they judge the trustworthiness of this kind of information source. Consequently heavy demands are placed on the design of virtual agents. The CSA engineers' frequent use of people as information sources suggests a need for systems that can assist the process of people finding. A key aspect in designing such systems will be to devise effective means of communicating not just people's technical expertise but also their trustworthiness in relation to the users' specific information needs. In devising these means, experiences from communicating the trustworthiness of personified virtual agents may carry over to systems that provide access to real people. For example, if photos are provided – to make people less distant and easier to remember – small differences in people's look-and-feel may instil very different first impressions. Also, document databases, which are frequently established by removing documents from their work context and storing them in centralised company archives, will be of limited value unless information seekers are able to reinsert the documents in this work context. Information about the context is necessary in interpreting how it has affected the contents of the documents and, thereby, in assessing whether the documents are applicable to the information seeker's context and information need. Contextual information can, to some extent, be written into the individual documents or provided through links to project websites but people such as the CSA engineers typically acquire it by contacting the document author or another person who has been involved in the work reported in the document. The value of information sources is not simply to provide information but also to assist the information seeker in assessing the trustworthiness of other sources. This is done by pointing the information seeker to some sources, warning against others, and not least by allowing information seekers to extend their trust in a known source to an initial assessment of hitherto unknown sources. Outside this network people tend to be cautious and alert to even feeble cues about the trustworthiness of a potential information source.

7 Acknowledgements

This work has been supported by the Danish National Research Foundation through its funding of the Centre for Human-Machine Interaction and by the European Commission through its funding of the COGITO project on e-commerce with guiding agents. We thank the people who participated in the focus group interview. Special thanks are due to the members of the CSA project group who have put up with our presence in spite of their busy schedule.

8 References

Allen, T.J. (1977). Managing the flow of technology: Technology transfer and the dissemination of technological information within the R&D organization. Cambridge, MA: MIT Press.

Andersen, H.H.K., Andersen, V., & Hansen, C.B. (2001). Establishing criteria for evaluating intelligent agents in e-commerce. In S. Weibelzahl, D. Chin, & G. Weber (Eds.), *Proceedings of the UM2001 Workshop on Empirical Evaluation of Adaptive Systems* (pp. 51-60). Freiburg, Germany: Pedagogical University of Freiburg.

Andersen, V., Hansen, C.B., & Andersen, H.H.K. (2001). *Evaluation of agents and study of end-user needs and behaviour for e-commerce. COGITO focus group experiment*. Risø Report No. Risø-R-1264(EN). Roskilde, Denmark: Risø National Laboratory. Available at: http://www.risoe.dk/rispubl/SYS/ris-r-1264.htm (consulted March 8, 2002).

Barber, B. (1983). *The logic and limits of trust*. New Brunswick, NJ: Rutgers University Press.

Barry, C.L., & Schamber, L. (1998). Users' criteria for relevance evaluation: A cross-situational comparison. *Information Processing & Management*, 34(2&3), 219-236.

Bevan, N. (1995). Measuring usability as quality of use. *Software Quality Journal*, 4(2), 115-130.

Blair, D.C., & Maron, M.E. (1985). An evaluation of retrieval effectiveness for a full-text document-retrieval system. *Communications of the ACM*, 28(3), 289-299.

Brown, J.S., & Duguid, P. (1996). The social life of documents. *First Monday*, 1(1). Available at: http://www.firstmonday.dk/issues/issue1/documents/index.html (consulted March 8, 2002).

Chakrabarti, A.K., Feineman, S., & Fuentevilla, W. (1983). Characteristics of sources, channels, and contents for scientific and technical information systems in industrial R and D. *IEEE Transactions on Engineering Management*, 30(2), 83-88.

Friedman, B., Kahn, P.H., & Howe, D.C. (2000). Trust online. *Communications of the ACM*, 43(12), 34-40.

Frøkjær, E., Hertzum, M., & Hornbæk, K. (2000). Measuring usability: Are effectiveness, efficiency, and satisfaction really correlated? In *CHI 2000: Proceedings of the ACM Conference on Human Factors in Computing Systems* (pp. 345-352). New York: ACM Press.

Gerstberger, P.G., & Allen, T.J. (1968). Criteria used by research and development engineers in the selection of an information source. *Journal of Applied Psychology*, 52(4), 272-279.

Goss, J.D., & Leinbach, T.R. (1996). Focus groups as alternative research practice. *Area*, 28(2), 115-123.

Harter, S.P. (1992). Psychological relevance and information science. *Journal of the American Society for Information Science*, 43(9), 602-615.

Hassenzahl, M. (2001). The effect of perceived hedonic quality on product appealingness. *International Journal of Human-Computer Interaction*, 13(4), 481-499.

Hertzum, M. (2002). The importance of trust in software engineers' assessment and choice of information sources. *Information and Organization*, 12(1), 1-18.

Hertzum, M., & Pejtersen, A.M. (2000). The information-seeking practices of engineers: Searching for documents as well as for people. *Information Processing & Management*, 36(5), 761-778.

ISO 9241-11 (1998). *Ergonomic requirements for office work with visual display terminals (VDTs) - Part 11: Guidance on usability*. Genève, Switzerland: International Organization for Standardization.

Janes, J.W. (1991). Relevance judgments and the incremental presentation of document representations. *Information Processing & Management*, 27(6), 629-646.

Jones, S., Wilikens, M., Morris, P., & Masera, M. (2000). Trust requirements in e-business. *Communications of the ACM*, 43(12), 81-87.

King, D.W., Casto, J., & Jones, H. (1994). *Communication by engineers: A literature review of engineers' information needs, seeking processes, and use*. Washington, DC: Council on Library Resources.

Kirakowski, J., & Corbett, M. (1993). SUMI: The software usability measurement inventory. *British Journal of Educational Technology*, 24(3), 210-212.

Kramer, R.M., & Tyler, T.R. (Eds.). (1996). *Trust in organizations: Frontiers of theory and research*. Thousand Oaks, CA: Sage.

Lavery, D., Cockton, G., & Atkinson, M.P. (1997). Comparison of evaluation methods using structured usability problem reports. *Behaviour & Information Technology*, 16(4&5), 246-266.

Lewis, J.R. (Ed.). (2001). Current issues in usability evaluation [Special issue]. *International Journal of Human-Computer Interaction*, 13(4).

Light, A., & Wakeman, I. (2000). "My name's Chris…": The effect of producer visibility on the use of interactive websites. In C. Paris, N. Ozkan, S. Howard, & S. Lu (Eds.), *OZCHI 2000 Conference Proceedings: Interfacing Reality in the New Millennium* (pp. 50-57). North Ryde, NSW, Australia: CSIRO.

McBreen, H.M., & Jack, M.A. (2001). Evaluating humanoid synthetic agents in e-retail applications. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, 31(5), 394-405.

Mizzaro, S. (1997). Relevance: The whole history. *Journal of the American Society for Information Science*, 48(9), 810-832.

Moenaert, R.K., Deschoolmeester, D., De Meyer, A., & Souder, W.E. (1992). Information styles of marketing and R&D personnel during technological product innovation projects. *R&D Management*, 22(1), 21-39.

Morgan, D.L. (1997). *Focus groups as qualitative research, Second edition*. London: Sage.

Nielsen, J. (1993). *Usability engineering*. Boston: Academic Press.

Nielsen, J. (1997). The use and misuse of focus groups. *IEEE Software*, 14(1), 94-95.

Nielsen, J., & Levy, J. (1994). Measuring usability: Preference vs. performance. *Communications of the ACM*, 37(4), 66-75.

Olson, G.M., & Moran, T.P. (Eds.). (1998). Experimental comparisons of usability evaluation methods [Special issue]. *Human-Computer Interaction*, 13(3).

Pinelli, T.E., Bishop, A.P., Barclay, R.O., & Kennedy, J.M. (1993). The information-seeking behavior of engineers. In A. Kent & C.M. Hall (Eds.), *Encyclopedia of Library and Information Science, Vol. 52* (pp. 167-201). New York: Marcel Dekker.

Rasmussen, J., Pejtersen, A.M., & Goodstein, L.P. (1994). *Cognitive systems engineering*. New York: Wiley.

Rosenberg, V. (1967). Factors affecting the preferences of industrial personnel for information gathering methods. *Information Storage and Retrieval*, 3, 119-127.

Rosenbloom, A. (Ed.). (2000). Trusting technology [Special issue]. *Communications of the ACM*, 43(12).

Rousseau, D.M., Sitkin, S.B., Burt, R.S., & Camerer, C. (Eds.). (1998). Not so different after all: A cross-discipline view of trust [Special issue]. *Academy of Management Review*, 23(3).

Saracevic, T. (1975). Relevance: A review of and a framework for the thinking on the notion in information science. *Journal of the American Society for Information Science*, 26(6), 321-343.

Saracevic, T. (1996). Relevance reconsidered '96. In P. Ingwersen & N.O. Pors (Eds.), *CoLIS2: Proceedings of the Second International Conference on Conceptions of Library and Information Science: Integration in Perspective* (pp. 201-218). Copenhagen, Denmark: Royal School of Librarianship.

Schamber, L., Eisenberg, M.B., & Nilan, M.S. (1990). A re-examination of relevance: Toward a dynamic, situational definition. *Information Processing & Management*, 26(6), 755-776.

Sheppard, B.H., & Sherman, D.M. (1998). The grammars of trust: A model and general implications. *Academy of Management Review*, 23(3), 422-437.

Shneiderman, B. (2000). Designing trust into online experiences. *Communications of the ACM*, 43(12), 57-59.

Shneiderman, B., & Maes, P. (1997). Direct manipulation vs interface agents. *ACM Interactions*, 4(6), 42-61.

Siemieniuch, C.E., & Sinclair, M.A. (1999). Organizational aspects of knowledge lifecycle management in manufacturing. *International Journal of Human-Computer Studies*, 51(3), 517-547.

Swanson, D.R. (1988). Historical note: Information retrieval and the future of an illusion. *Journal of the American Society for Information Science*, 39(2), 92-98.

Tang, R., & Solomon, P. (1998). Toward an understanding of the dynamics of relevance judgment: An analysis of one person's search behavior. *Information Processing & Management*, 34(2&3), 237-256.

Tseng, S., & Fogg, B.J. (1999). Credibility and computing technology. *Communications of the ACM*, 42(5), 39-44.

Valley, K.L., Moag, J., & Bazerman, M.H. (1998). 'A matter of trust': Effects of communication on the efficiency and distribution of outcomes. *Journal of Economic Behavior & Organization*, 34(2), 211-238.

Walker, J.H., Sproull, L., & Subramani, R. (1994). Using a human face in an interface. In *CHI '94: Proceedings of the ACM Conference on Human Factors in Computing Systems* (pp. 85-91). New York: ACM Press.

Zipperer, L. (1993). The creative professional and knowledge. *Special Libraries*, 84(2), 69-78.

Zucker, L.G., Darby, M.R., Brewer, M.B., & Peng, Y. (1996). Collaboration structure and information dilemmas in biotechnology: Organizational boundaries as trust production. In R.M. Kramer & T.R. Tyler (Eds.), *Trust in Organizations: Frontiers of Theory and Research* (pp. 90-113). Thousand Oaks, CA: Sage.