

Trust Management of Ubiquitous Multi-Display Environments

Ekatarina Kurdyukova, Elisabeth André, and Karin Leichtenstern

Abstract While a lot of research has been devoted to improving the security and the reliability of ubiquitous display environments, work on the user experience factor of trust is still scarce. To ensure that ubiquitous environments find acceptance among users, the user experience factor of trust should, however, not be underestimated. In this paper, we present a decision-theoretic approach to trust management that we consider particularly appropriate when a system has to balance the benefits and risks of a decision carefully. In the paper, we define decision policies that help maintain trust in critical situations, such as the loss of sensor data or the presence of unknown people. The approach has been employed in three interactive applications that have been developed as part of a university-wide ubiquitous displays management system.

1 Introduction

Recent years have produced a large variety of interactive displays that are installed in many public places. Apart from simply providing information (e.g., news or weather) to people in public places, such as coffee bars or airports, public displays make it possible for passing individuals to view, edit and exchange with each other specific data. Mobile phones represent a popular interaction device for interacting with these displays since they have been widely adopted by people as an everyday companion and can be customized to individual interaction preferences.

A social setting with the possibility of viewing personalized information in the presence of other people inevitably raises privacy issues. For example, a user may

E. Kurdyukova (✉) • E. André • K. Leichtenstern
Human-Centered Multimedia, Institute of Computer Science, Universitätsstr. 6a, Augsburg 86159, Germany
e-mail: Kurdyukova@informatik.uni-augsburg.de <https://hcm-lab.de>; Andre@informatik.uni-augsburg.de <https://hcm-lab.de>; Leichtenstern@informatik.uni-augsburg.de <https://hcm-lab.de>

be viewing personal information on a public display as other people pass by. The basic question that arises is how the system should react in such a situation. Should it trust the user will take the necessary actions herself? Or should it adapt autonomously to the changing social context, for example, by masking personal information on the public display? In the first case, there is the risk that the user expects the system to protect her privacy and is upset if no appropriate actions are taken. The limitation of the second approach is that it leads to an interruption of the user's work flow and might give her the feeling that she no longer has the system under control. Furthermore, there is the danger that the user does not understand the rationale behind the system's behavior and perceives the system as being only a little transparent.

The example illustrates that a system needs to balance the benefits and drawbacks of its actions carefully in order not to entail the risk that a user loses trust in its workings and does not use it any more.

In this paper, we present a decision-theoretic approach to a trust management system for ubiquitous display environments that assesses the user's trust in a system, monitors it over time, and applies appropriate measures to maintain trust in critical situations [12]. Such situations arise when, among other things, other people enter the user's private space [9], the system has to generate presentations based on inaccurate user or context data [4] or the system's adaptation behavior mismatches the user's expectations [3].

As a test bed for our research, we employ three applications that have been developed as part of a university-wide displays management system. Two applications run on public displays located in public rooms at Augsburg University. They can be operated and assisted by mobile phones. The first application, Friend Finder, is an interactive campus map that shows the current location and status of the user's friends. Since many students have difficulty orienting themselves on the campus (especially in new buildings), Friend Finder also supports a routing function, showing a detailed path to a selected friend. The second application, Media Wall, fosters exchange between the students. It represents a gallery of media items (pictures or videos) uploaded by students or scientific staff. Users can rank the media items, upload new items, and view their favorite ones. The third application created for a personal projector, On-Campus Navigator, is designed to enable indoor navigation within university buildings. Users can switch between two views on the mobile projector: an overview map with self-updating user position, and an arrow view that is projected in the user's physical environment and that points to the direction in which the user should move.

All three applications require sophisticated mechanisms to adapt to various trust-critical events. Since Friend Finder may disclose private information about a user's social network, it should be able to adapt intelligently to the surrounding social context in order to avoid possible threats to privacy. Ranking the media on Media Wall again may threaten the user's privacy in case of observation. Several users may interact with Friend Finder simultaneously, rendering their networks on the same campus map. Therefore, the system should be able to accommodate the data and interaction coming from multiple users. Systems such as On-Campus Navigator

rely on sensor data to update the direction. Here, incomplete or possible incorrect sensor data may cause wrong or missing directions. The system should be able to cope with such deficits and adjust the presentation appropriately.

In the remaining chapter, we first present a scenario that served as inspiration for our applications. We then discuss work related to increasing the user's trust in ubiquitous display environments by appropriate interface design. After that, we present a model of a trust management system based on Bayesian Networks and influence diagrams and show how this model has been used within our applications. Finally, we present the first results of a user study.

2 Scenario

To illustrate our ideas, let us have a look at a scenario which inspired our applications.

Oscar is a student in the Informatics faculty. Currently he is having a lunch break and looking for his friends to join him for lunch at the University canteen. On a University floor, he finds a large public display. He has already used it once and has a mobile client installed on his PDA to operate the display. By means of one application, Oscar can find his friends on the campus map. Since Oscar is currently alone on the floor, he loads his personal social network on the public display. He notices that his best friends Barbara and Marina are both at the law library. Since he seldom visits the law faculty and does not know that campus area very well, he decides to download the route onto his mobile phone.

Suddenly a group of other students approach the display. Oscar is a bit irritated: he is not willing to expose his personal friends to the strangers, neither does he want to disclose his intention to meet specific persons. Apart from that, he considers the locations of his friends as private information not supposed to be shown to everybody. Surprisingly, he notices that the display masks the pictures of his friends, representing them by icons. The portraits and details have migrated now to his mobile device.

One of the approaching students wants to use the display as well. Thus, he stands next to Oscar and loads his personal social network on the campus map as well. The friends of the stranger are also rendered on the map; however, the icons are different. After Oscar selects Barbara on his mobile display, her icon is highlighted on the public screen. Oscar chooses the option to show the route, and the route is drawn on the large screen. Oscar then downloads the route to his mobile phone, and leaves the public display. His personal network disappears from the screen.

On his way, Oscar meets a friend, who is also willing to join him, Barbara and Marina for lunch. They both proceed to the law library, being navigated by the mobile application. For convenience, Oscar uses a personal projector attached to his mobile phone which facilitates on-campus navigation. The application displays an arrow pointing in the direction in which the user should move. Suddenly the arrow is disabled. The friends are irritated for a second. However, the application gives

a short explanation that the signal is currently too low, and the update will take a couple of seconds. Indeed, in a moment, the arrow is revived, and the friends proceed to their destination.

3 Related Work

Most work that investigates trust issues in the context of ubiquitous displays environments focuses on the distribution of private and public data over various displays. Often mobile phones are used as private devices that protect the personal component of interaction from public observation.

Röcker et al. [9] conducted a user study to identify the privacy requirements of public display users. Based on the study, they developed a prototype system that automatically detects people entering the private space around a public display using infrared and RFID technology, and adapts the information that is visible based on the privacy preferences of the users. An evaluation of the system revealed that users are willing to use public displays when there is a mechanism for privacy protection.

Based on the evaluation of two mobile guides, Graham and Cheverst [4] analyzed several types of mismatch between the users' physical environment and information given on the screen and their influence on the formation of user trust. Examples of mismatches include situations where the system is not able to detect the user's current location correctly or situations where the system conveys a wrong impression about the accuracy of its descriptions. To help users form trust, Graham and Cheverst suggest employing different kinds of guide, such as a chaperone, a buddy or a captain, depending on characteristics of the situations, such as accuracy and transparency. For example, the metaphor of a buddy is supposed to be more effective in unstable situations than the chaperone or the captain.

Cao et al. [1] introduce the notion of crossmodal displays that enable users to access personalized information in public places while ensuring their anonymity. The basic idea is to display the main information publicly, but to add cues for individual users to direct them to information that is relevant to them.

All in all, there is a vivid research interest in the design of novel user interfaces for heterogeneous display environments. However, the few approaches that address the user experience factor of trust in such environments do not attempt to model the user experience of trust explicitly as a prerequisite for a trust management system. A number of approaches to modeling trust in computational systems have been presented. Especially in the area of multi-agent systems (MAS), trust models have been researched thoroughly (see, e.g., Castelfranci's and Falcone's introduction [2] to a formal modeling of trust theory and its applications in agent-based systems). However, these approaches either focus on trust in software components or aim at modeling trust in human behavior.

4 Dimensions of Trust

Much of the original research on trust comes from the humanities. Psychologists and sociologists have tried for a very long time to get a grasp on the inner workings of trust in interpersonal and interorganisational relationships. Other fields, such as economics and computer science, relied on their findings, but adapted them to the special requirements of their respective fields and the new context to which they are applied. There is consensus that trust depends on a variety of trust dimensions. However, there is no fixed set of such dimensions.

Trust dimensions that have been researched in the context of internet applications and e-commerce include reliability, dependability, honesty, truthfulness, security, competence, and timeliness (see, for example, the work by Grandison and Sloman [5] or Kini and Choobineh [6]). The more sociologically inclined authors [11] introduce willing vulnerability, benevolence, reliability, competence, honesty, and openness as the constituting facets of trust. Researchers working on adaptive user interfaces consider transparency as a major facet of trust (see, for example, the work by Glass et al. [3]).

Our set of trust dimensions is based on interviews with 20 computer science students who were asked to indicate the trust factors of user interfaces that they felt contributed to their assessment of trustworthiness. The most frequent mentions fell into the following categories: comfort of use (“should be easy to handle”), transparency (“I need to understand what is going on”), controllability (“want to use a program without automated updates”), privacy (“should not ask for private information”), reliability (“should run in a stable manner”), security (“should safely transfer data”), credibility (“recommendation of friends”) and seriousness (“professional appearance”).

A follow-up study revealed that there are statistically significant positive correlations between trust and the identified factors. The better the ratings for the trust dimensions, the better were also the ratings for trust. In addition, we observed a statistically significant positive correlation between the users’ ratings of their general trust in software and their reported trust in the presented system. Furthermore, we found that a missing feeling of trust was accompanied by negative emotions, such as irritation, uneasiness, and insecurity. More information regarding this experiment and the exact findings can be found in Ref. [7].

5 Using a Decision-Theoretic Approach to Trust Management

In the following, we describe a model to assess the user’s trust in a computer system. Our model of trust should account for the following characteristics of trust:

Trust as a subjective concept: There is a consensus that trust is highly subjective. A person who is generally confiding is also more likely to trust a software program. Furthermore, users respond individually to one and the same event. While some

users might find it critical if a software asks for personal information, others might not care. We aim at a computational model that is able to represent the subjective nature of trust.

Trust as an uncertain concept: The connection between events and trust is inherently uncertain. For example, we cannot always be absolutely sure that the user notices a critical event at all. Furthermore, it may also happen that a user considers a critical event as rather harmless. As a consequence, it is not possible to predict with 100% certainty which level of trust a user has in a particular situation. A computational model of trust should be able to cope with trust as an uncertain concept.

Trust as a multifaceted concept: As shown in Sect. 4, trust is a multi-faceted concept. We therefore aim at a computational model that is able to represent explicitly the relative contribution of the trust dimensions to the assessment of trust. In particular, the model should help us predict the user's level of trust based on dimensions, such as the perceived transparency and controllability of a user interface. Furthermore, the model should allow us to add trust dimensions easily based on new experimental findings.

Trust as a dynamic concept: Trust depends on experience and is subject to change over time. Lumsden [8] distinguishes between immediate trust dimensions and interaction-based trust dimensions. Immediate trust dimensions, such as seriousness, come into effect as soon as a user gets in touch with a software system, while interaction-based trust dimensions, such as transparency of system behavior, influence the user's experience of trust during an interaction. To model trust as a dynamic concept, we need to be able to represent how the user's level of trust depends on earlier levels of trust.

Based on the considerations above, we have chosen to model users' feelings of trust by means of Dynamic Bayesian Networks. The structure of a Bayesian Network is a directed, acyclic graph (DAG) in which the nodes represent random variables while the links or arrows connecting nodes describe the direct influence in terms of conditional probabilities (see Ref. [10]).

5.1 Modeling the Determinants of Trust

Dynamic Bayesian Networks meet the requirements listed above very well. First of all, they allow us to cope with trust as a subjective concept. For example, we may represent the system's uncertain belief about the user's trust by a probability distribution over different levels of trust. Secondly, they enable us to model the non-deterministic nature of trust. In particular, we are able to make predictions based on conditional probabilities that model how likely it is that the child variable is given the value of the parent variables. For example, we may model how likely it is that the user has a moderate level of trust if the system's behavior is moderately transparent. Furthermore, Bayesian Networks enable us to model the relationship between trust and its dimensions in a rather intuitive manner. For example, it is

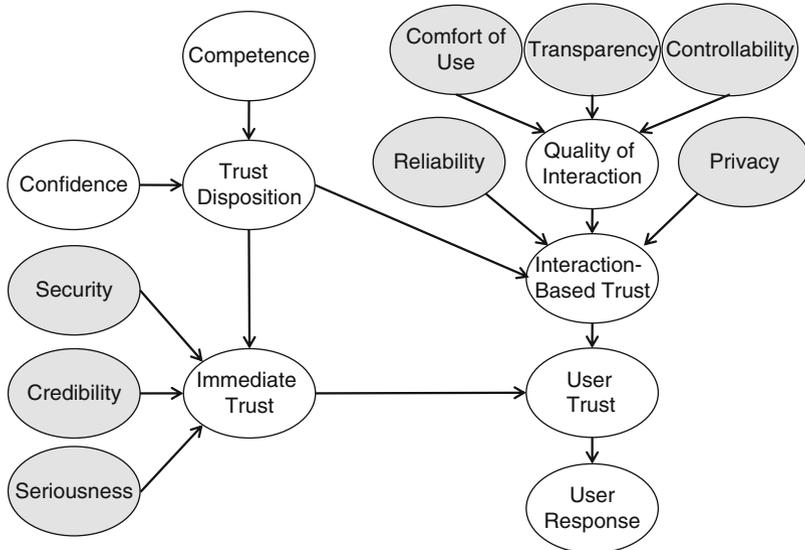


Fig. 1 Modeling trust by means of a Bayesian Network

rather straightforward to model that reduced transparency leads to a decrease in user trust. The exact probabilities are usually difficult to determine. However, the conditional probabilities can also be (partially) derived from the user data we collected in the experiment described in Ref. [7].

In Fig. 1, a Bayesian Network for modeling trust is shown. For each trust dimension, we introduced a specific node (highlighted in gray). Following Lumsden [8], we distinguish between immediate and interaction-based trust dimensions. Immediate trust dimensions include security (conveyed, for example, by the use of certificates), seriousness (reflected, for example, by the system's look-and-feel) and credibility (supported, for example, by company profile information). In this context, we would like to emphasize that trust dimensions may only affect the user's trust if the user is aware of them. For example, high security standards will only have an impact on user trust if the user knows that they exist. For the sake of simplicity, we assume that immediate trust dimensions do not change over time. That is we do not consider the fact that a user might notice references to security certificates only after working with a system over a longer period of time. To describe the determinants of interaction-based trust, we further distinguish between the quality of interaction, privacy and reliability. The quality of interaction is characterized by transparency, controllability and comfort of use. Both the establishment of immediate trust and interaction-based trust depend on the user's trust disposition which is characterized by his or her competence and their general confidence into technical systems.

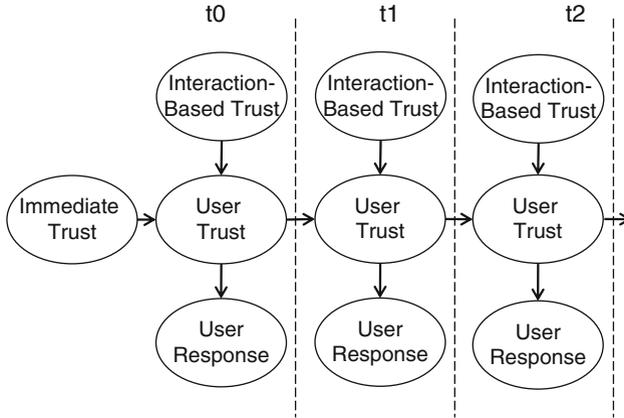


Fig. 2 Modeling the dynamics of trust by means of a Dynamic Bayesian Network

5.2 Monitoring Trust over Time

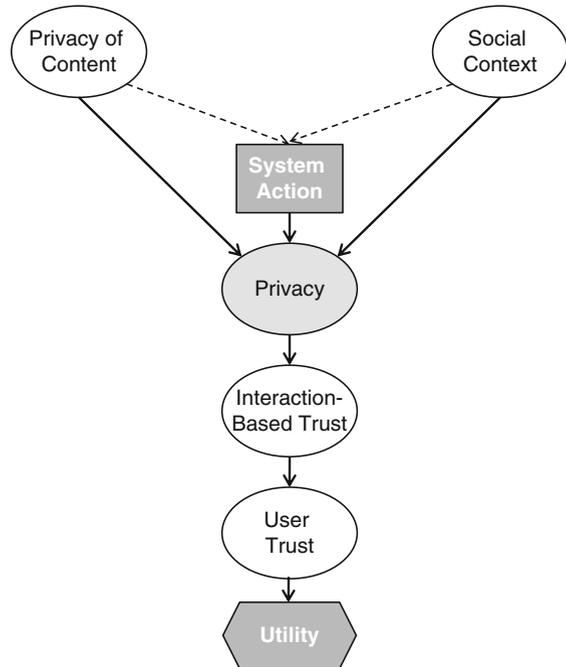
After smoothly interacting with a system over a longer period of time, the users' trust in a system is likely to increase. However, it may also happen that an unexpected system event, such as a sudden breakdown of the system, a substantial delay in the transfer of data or a serious leakage of data, causes a sudden loss of trust. All in all, the development of user trust must be continuously monitored at runtime in order to detect critical situations that require optimizations of the system to re-establish trust. As a consequence, we need not only a model that describes the relationship between user trust and its dimensions, but also one that explains the dynamics of trust.

Dynamic Bayesian Networks allow us to model the dependencies between the current states and the earlier states of variables. In particular, we are able to represent how the user's current level of trust is influenced by earlier levels of trust. In Fig. 2, a small portion of a Dynamic Bayesian Network is shown that illustrates how trust develops over time depending on the user's immediate level of trust and her interaction-based trust at time $t = 0$. Due to space limitations, we present only the time plates from ($t = 0$) to ($t = 2$). The arrow pointing from the node for user trust to the time plate for ($t = 2$) indicates that the user's trust at time ($t = 1$) influences the user's trust at time ($t = 2$). For simplicity, we consider only the user's level of trust at time t_{i-1} to determine the user's level of trust at time t_i .

5.3 Taking Decisions to Maximize Trust

So that we can use the Bayesian Network formalism for decision-making, it has to be extended to an influence diagram by adding a decision node and a utility node. The decision node represents all system actions that the system can perform, while the utility node encodes the utilities of all possible outcomes.

Fig. 3 Modeling decision making by means of influence diagrams



To make a decision, the system evaluates the utility of all possible options in terms of user trust and chooses the action with the highest utility. In Fig. 3, a small portion of the influence diagram is shown. We illustrate the basic idea by means of one trust dimension, namely privacy.

Privacy is handled as a hidden variable with three discrete values: low, medium and high. That is, its value cannot be directly observed, but has to be inferred from observable variables, such as *Privacy of Content* and *Social Context*. For example, the likelihood that the variable *Privacy* has the value *Low* would be high if *Privacy of Content* has the value *Private* and *Social Context* has the Value *People Approaching*. These dependencies are indicated by the arrows going from *Social Context* and *Privacy of Content* to *Privacy*. Associated with the decision node *System Action* is a table that describes the system's decision policy for each combination of the variables *Social Context* and *Privacy of Content*. The arrow going from *System Action* to *Privacy* represents the impact a particular system action, for example, the masking of private information, has on privacy.

6 Applying the Approach to Ubiquitous Displays

In the previous section, we described the general structure of an influence diagram as the basis for the implementation of a trust management system. In the following, we illustrate how to set the probabilities for a concrete application.

Table 1 Possible system reactions to trust-critical situations

Situation	Possible reaction
A – Changed social context	
User interacts with some private data on a public display. Another user approaches the display and starts to interact as well. Other people are around the display	A1: Do not take any action A2: Move all data from the public display to the mobile display A3: Mask private data on the public display and move the details to the mobile screen A4: Notify the user about potential privacy issues and offer options to protect data
B – Space conflicts	
Several users are approaching the public screen and want to interact with it. As a result, conflicts of space occur	B1: Allow just one user to interact at a time B2: Divide space among users B3: Move part of the data to mobile display, e.g., present final results on public displays and intermediate results on the mobile phone B4: One user interface integrates data of several users
C – Incomplete information	
The information the system needs to adapt a presentation to the user is not correct or incomplete	C1: Give textual explanations C2: Give visual feedback C3: Reduce level of detail C4: Give presentation based on available information

In an earlier experiment [7], we collected data for trust and its dimensions that we used for the creation of conditional tables employing the GeNIe (see <http://genie.sis.pitt.edu>) built-in algorithm for learning Bayesian Networks. To determine the parameters for the complete influence diagrams, we could, in principle, reply on empirical data as well. However, since the acquisition of these data turned out to be rather time consuming, we decided to set the probabilities that represent the dependencies between system actions and trust factors based on usability guidelines. For example, we know from the literature that there is a correlation between a system's transparency and user trust (see, for example, Ref. [3]).

In the following, we analyze the impact of various system reactions to typical trust-critical situations on relevant trust factors, providing illustrations from Friend Finder, Media Wall and On-Campus Navigator. Let us assume that the user is viewing private data on the public screen as other users pass by. Such a situation may occur in Friend Finder when users load a map of the university campus with friends on a public screen. The locations and pictures of friends are considered as private information, not supposed to be observed by just any one. Within the influence diagram shown in Fig. 3, this situation is described by the values of the variables *Social Context* and *Privacy of Content*.

In Table 1 (upper part), four possible responses to the described situation are listed. Basically, the system has to decide whether it should trust the user to take appropriate steps herself (Option A1), whether it should adapt the display of

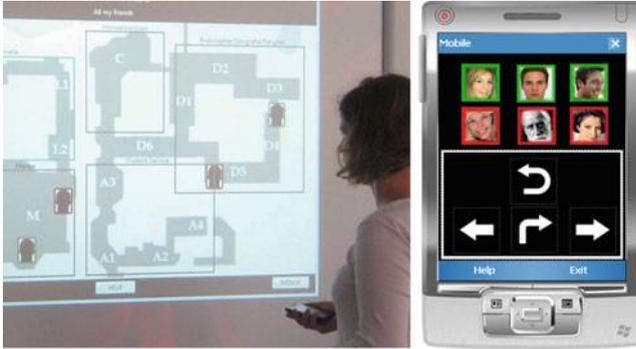


Fig. 4 Friend Finder: System masks private data on public display and displays the private details on the mobile screen

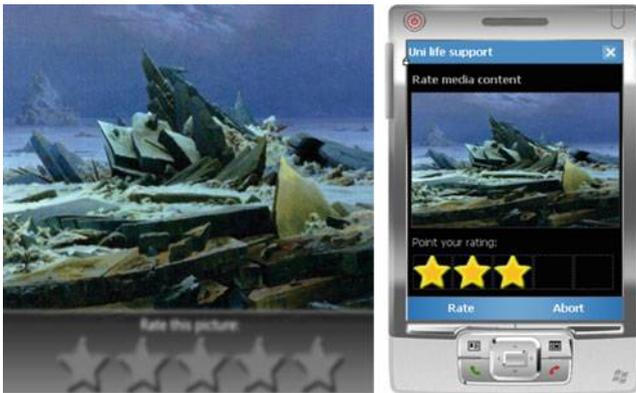


Fig. 5 Media Wall: The system masks user ranking and moves it to the mobile phone

information to the changed social context (Option A2 and A3) or whether it should offer the user options to protect private data (Option A4). Within the influence diagram shown in Fig. 3, the available options are represented by the decision node *System Action*.

Figure 4 shows Friend Finder's implementations of Option A3. The photos of the user's friends are masked with icons and the private data migrates to the mobile screen. Figure 5 shows the implementations of Option A3 in Media Wall. Here, the system masks the ranking and moves it to the mobile phone. Option A1 bears the risk that the user might expect the system to protect her privacy and will be upset if no appropriate actions are taken. Options A2 and A3 have the limitations that they cause an interruption of the user's work flow and might give her the feeling that she no longer has the system under control. The drawback of Option A4 is that there might not be enough time for the user to confirm the adaptations proposed by the system. Furthermore, it requires more effort from the user than the other options.

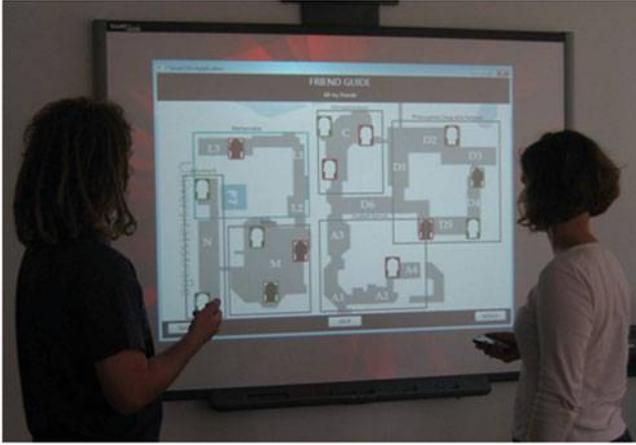


Fig. 6 Friend Finder: System masks the private data for both interacting users

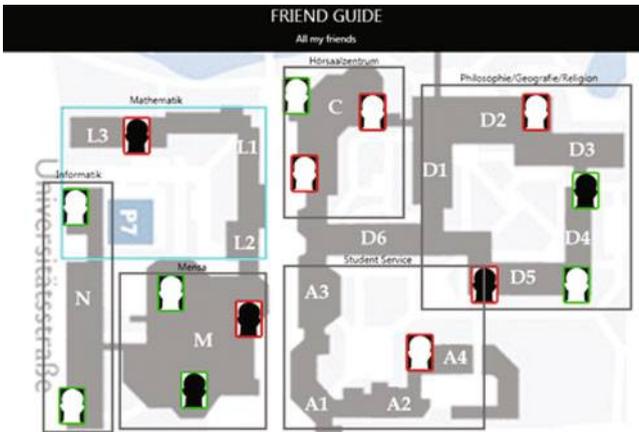


Fig. 7 Friend Finder: Sharing the public display by several users

Compared to Option A2, Option A3 has the advantage that the user still profits from the large real estate of the public screen, while preserving personal information. Furthermore, Option A3 allows several users to interact at the same time. In Fig. 6, two users browse their friends on a public display. Their network has been rendered on the same campus map. The display shows only the locations of the friends, and masks pictures with different icons for every user. The detailed information can be found on a mobile screen. The users distinguish their networks by the depiction of icons (Fig. 7).

When several people interact with public displays, conflicts of space are not uncommon. To solve such conflicts, the system may have just one user interact at the same time, allocate specific screen areas to the single users or move part of the

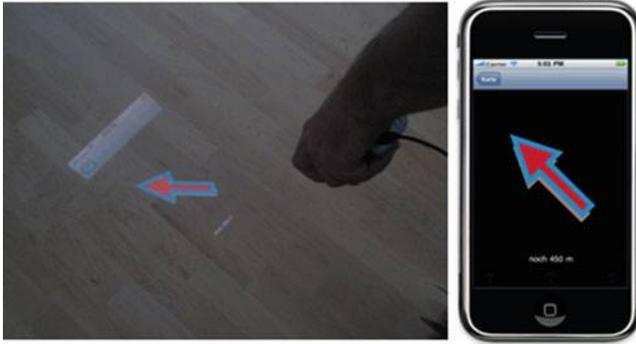


Fig. 8 On-Campus Navigator: Displaying a navigation arrow in the user's physical environment or on a mobile phone

data to mobile display (see the middle part of 1). Option B1 may affect the comfort of use negatively because the user may have to wait or postpone activities. On the other hand, the system's policy to handle conflicts of space may be easily understood. Option B2 may be advantageous when the display allows for a natural allocation of space that is immediately understood by users. In table top applications, a natural allocation of space is given by the user's position at the table. B3 requires the user to integrate public and personal information which may raise some usability problems. It can be advantageous when a common ground for user data can be found: for example, a map, or a table (see Fig. 6).

Finally, let us have a look at a situation where the system is forced to present data based on incomplete or incorrect information. The problems with incomplete or incorrect data often arise from sensor issues or unstable service functionality. In such a situation, the system has to decide whether it should inform the user about the problem or whether it should try to adapt its behavior to available data, for example, by reducing the level of detail (see lower part of Fig. 1). The drawback of the first option is that the user might lose trust in the system. The risk of the second approach is that the system might no longer be able to hide the problem at a later point in time resulting in an even greater loss of trust.

Figure 9 shows one of the strategies to cope with inaccurate data in On-Campus Navigator. As described earlier, On-Campus Navigator is able to give the user directions either by placing arrows in the user's physical environment using a miniaturized projector or by placing arrows on a mobile phone (Fig. 8). The On-Campus Navigator may temporarily lose sensor signal, or be unable to calculate an immediate direction. Once the system is not able to compute the direction, the arrow is temporarily disabled (Option C2) to indicate that the system does not have accurate location information. Another option to cope with the problem of insufficient signal is to provide an alternative view based on available information (Option C4). Thus, if the arrow direction is not possible to calculate, On-Campus Navigator may switch to the map view with the last available location update (Fig. 9).



Fig. 9 On-Campus Navigator: Switching to the map view due to insufficient sensor data.

Table 2 Impact of possible system reaction on transparency (Tran), controllability (Contr), comfort of use (Comf), privacy (Priv) and reliability (Rel)

Action	Tran	Contr	Comf	Priv	Rel
Do not take any action (A1)	0	0	0	--	--
Complete migration to mobile phone (A2)	-	-	--	++	0
Partial migration to mobile phone (A3)	-	-	-	+	0
Offer options to user (A4)	+	++	--	++	0
Sequentialize interaction (B1)	0	0	-	0	0
Allocate space (B2)	0	0	-	0	0
Partial migration with use of common space (B3)	0	0	-	0	0
Textual or visual indication of error (B1, B2)	++	-	-	0	--
Adapt presentation to error (C3, C4)	-	-	-	0	0

-- very likely to decrease, - likely to decrease, 0 not likely to change, + likely to increase, ++ very likely to increase

The decisions to be made by the system have an influence on the trust dimensions. The corresponding conditional tables were set up based on general guidelines in user interface design. For example, asking a user for confirmation before transferring private data to a public display has a positive impact on controllability, transparency and privacy, but a negative impact on ease of use. Further examples are shown in Table 2.

7 First Evaluation of the Approach

In the following, we present a first evaluation of the decision-theoretic approach.

First of all, we wanted to know to what extent the system is able to predict the user's ratings of trust based on her ratings of transparency, controllability, ease of use, seriousness, credibility, and security. To this end, we ask 20 users to rate five variants of a prototype that combined a public display (in this case a table top

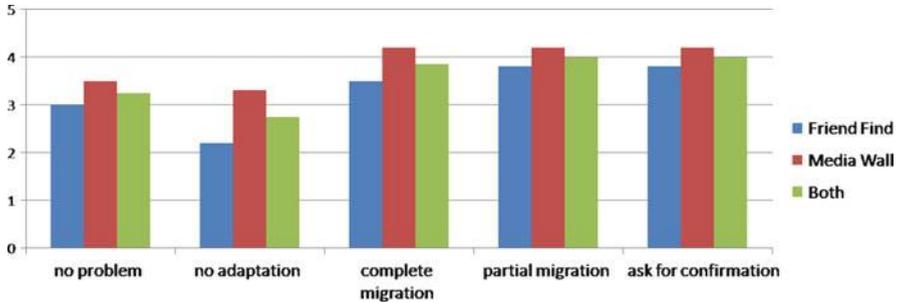


Fig. 10 Trust ratings for different responses

application) with a mobile phone. In particular, we produced a prototype that was less self-explainable (the interface included no help function and no descriptive labels), a second prototype that was less transparent (the system gave no reasons for its behavior), a third prototype that was less controllable (the system did not ask for user confirmations before executing an action), a fourth prototype that followed a less stricter privacy policy (the system displayed all kinds of data on user request on the table regardless of whether they were private or not) and finally a system that did not show any of these problems. The idea behind the variants was to obtain a sufficient variety of user ratings. To assess to what extent the user's level of trust can be predicted based on her ratings of transparency, controllability, ease of use, seriousness, credibility and security, we created a model using the Genie built-in algorithm for learning Bayesian Networks. When evaluating this model in tenfold cross validation, we achieved an accuracy rate of 73% for five classes (very low trust, low trust, medium trust, high trust, very high trust).

Secondly, we aimed at evaluating to what extent the decisions of the system contributed to a higher level of trust. As a first test, we presented six users in Friend Finder and Media Wall with a situation in which the user is viewing private information as other users pass by (Situation A). We furthermore showed them how the system would respond in such a situation (Options A1–A4) and asked them to rate their trust in the system on a scale from 1 (very untrustworthy) to 5 (very trustworthy). The results of this test are shown in Fig. 10. There was no clear preference for a particular system action. Despite the small number of users, the test shows, however, that the users clearly preferred the system to take the initiative and either adapt automatically to the social context or ask the user to confirm the adaptation. We take this as evidence that our adaptation policies were generally appropriate even though additional research is needed to prioritize decisions.

Another interesting finding is the fact that user trust was higher in situations where the system automatically adapted to a trust-critical situation than in situations where no problem occurred. Furthermore, the differences in the case of the Media Wall were less extreme than in that of the Friend Finder. We hypothesize that the information displayed with Media Wall (ratings of photos) was considered as less sensitive than the information displayed with Friend Finder (location of people).

8 Conclusions

Ubiquitous displays environments require a high degree of flexibility due to the changing social context and the probably incomplete or inaccurate information on which a system has to base its presentations. In order to maintain user trust in such environments, a system needs to be able to evaluate carefully the consequences of its actions and the trade-offs between them. In this paper, we presented a decision-theoretic approach to trust management. The approach has been informed by guidelines on user interface design in order to assess the impact of system actions on trust dimensions, such as comfort or use, transparency, controllability and privacy. A first evaluation of the approach within two applications that have been developed as part of a university-wide public displays environment revealed that users preferred the adaptive to the non-adaptive system.

Our future work will concentrate on conducting further experiments considering a larger number of users and a greater variety of trust-critical situations. The experiments conducted so far seem to indicate that the decisions taken by the system to adapt to a trust-critical situation find general user acceptance. However, a greater amount of user data is necessary to evaluate the appropriateness of particular design decisions. So far, the presented decision-theoretic approach and the corresponding evaluation concentrated on short-term interactions with public displays. We need, however, to take into account that users interacting with a system for a longer period of time might become annoyed about repeated system explanations and requests to confirm system actions. Thus, a too high amount of controllability and transparency may even negatively affect user trust. In order to avoid such problems, our future work will also consider the history of interactions during trust management.

Acknowledgement This research is partly sponsored by *OC-Trust* (FOR 1085) of the German research foundation (DFG).

References

1. H. Cao, P. Olivier, and D. Jackson. Enhancing privacy in public spaces through crossmodal displays. *Soc. Sci. Comput. Rev.*, 26(1):87–102, 2008.
2. C. Castelfranchi and R. Falcone. *Trust Theory: A Socio-Cognitive and Computational Model*. Wiley, 2010.
3. A. Glass, D. L. McGuinness, and M. Wolverton. Toward establishing trust in adaptive agents. In *IUI '08: Proceedings of the 13th international conference on Intelligent user interfaces*, pages 227–236, New York, NY, USA, 2008. ACM.
4. C. Graham and K. Cheverst. Guides, locals, chaperones, buddies and captains: managing trust through interaction paradigms. In *3rd Workshop 'HCI on Mobile Guides' at the Sixth International Symposium on Human Computer Interaction with Mobile Devices and Services*, pages 227–236, New York, NY, USA, 2004. ACM.
5. T. Grandison and M. Sloman. A survey of trust in internet applications. *IEEE Communications Surveys and Tutorials*, 3(4):2–16, 2000.

6. A. Kini and J. Choobineh. Trust in electronic commerce: definition and theoretical considerations. In *Proc. of the Hawaii International Conference on System Sciences*, volume 31, pages 51–61, 1998.
7. K. Leichtenstern, E. André, and E. Kurdyukova. Managing user trust for self-adaptive ubiquitous computing systems. In *Proceedings of the 8th International Conference on Advances in Mobile Computing and Multimedia, MoMM '10*, pages 409–414, New York, NY, USA, 2010. ACM.
8. J. Lumsden. Triggering trust: to what extent does the question influence the answer when evaluating the perceived importance of trust triggers? In *BCS HCI '09: Proceedings of the 2009 British Computer Society Conference on Human-Computer Interaction*, pages 214–223, Swinton, UK, UK, 2009. British Computer Society.
9. C. Röcker, S. Hinske, and C. Magerkurth. Intelligent privacy support for large public displays. In *Proceedings of Human-Computer Interaction International 2007 (HCI'07)*, Beijing, China, 2007.
10. S. J. Russell and P. Norvig. *Artificial Intelligence a modern approach*. Prentice Hall, Upper Saddle River, N.J., 2nd international edition, 2003.
11. M. Tschannen-Moran and W. Hoy. A multidisciplinary analysis of the nature, meaning, and measurement of trust. *Review of Educational Research*, 70(4):547, 2000.
12. Z. Yan and S. Holtmanns. Trust modeling and management: from social trust to digital trust. *Book chapter of Computer Security, Privacy and Politics: Current Issues, Challenges and Solutions*, 2008.