Managing User Trust for self-adaptive Ubiquitous Computing Systems

Karin Leichtenstern, Elisabeth André and Ekatarina Kurdyukova
Augsburg University, Universitätsstr. 6a, 86159 Augsburg
{leichtenstern, andre, kurdyukova}@informatik.uni-augsburg.de

ABSTRACT

Ubiquitous computing systems can cause serious problems for user trust. In particular if the system is self-adaptive and situations appear which are poorly self-explanatory. In this paper we aim at the trust management of adaptive systems. We present a user study that covers the correlation of trust dimensions and user feelings on user trust. As results of this study, a Bayesian Network is introduced that, at the design time and runtime of the system, provides knowledge about the interplay between a truster's disposition, system events and actions, trust dimensions, user trust and user response.

1. INTRODUCTION

Self-adaptive ubiquitous computing systems modify the content, dialogue, layout or modality of the user interface after critical variable changes, such system failures or different contextual situations of the user [12]. These adaptations are often not self-explaining since the users do not always recognise the reason for the adaptations. In these situations user trust can be impaired which can lead to disuse of the system in the worst case.

We aim at the problem of user trust in adaptive systems in the context of ubiquitous display environments. These environments make it possible for passing individuals to view, edit and exchange specific data between each other. Mobile phones represent a popular interaction device for interacting with these displays. They have become an everyday companion which maintains all kind of personal data, such as music, videos and photos. Transferring such data to large screens comes with a lot of benefits (e.g. usage of full screen mode) but also with a lot of risks, such as the loss of data due to unstable transmission technologies. Bluetooth is often used for the communication between mobile phones and ubiquitous display environments (e.g. [3]). Typical problems of Bluetooth emerge in the discovery process and the data transmission because they can unexpectedly require more time or even fail completely. Such a behaviour can seriously affect trust in a system since it is no longer considered as reliable and secure. The problem is aggravated by the fact that people usually interact with public displays on a short-term basis without having the possibility to verify the security of the underlying infrastructure.

In this paper we first describe a scenario that provides more insights to problems in terms of user trust when interacting with ubiquitous display environments. After that, we aim at related work and relevant trust triggers. Finally, we describe a user study that addresses the interplay between the trust triggers and user trust as well as a first version of a Bayesian Network to dynamically manage user trust at the design time and runtime of an adaptive system.

2. SCENARIO

On Friday afternoon, Emily and her friend Olivia are in a café in the old town. After they have found a vacant table and sat down, they realise that it is not an ordinary table and they wonder what it can be used for. When ordering, they ask the waiter and he explains to them that the table has a touch-sensitive display and that they can interact with it using their fingers. Furthermore, the waiter tells them that they also can transfer data, such as images or video clips, from their mobile devices to the table in order to view, edit and exchange them.

Since Emily has just returned from her vacation in Spain with her boyfriend Diego, she has a lot of pictures on her mobile phone which she wants to show to Olivia. Thus, Emily decides to use this new touch-sensitive table. But in the same moment, Emily is afraid that her private pictures - showing Emily and her boyfriend at the beach and drinking liquor - can be seen by other people than her friend Olivia. In addition, she does not fully trust the system since she cannot be sure that only her selected data will be transferred to the table because her mobile phone comes with a lot more intimate data, such as text messages she was exchanging with her boyfriend, which are not meant to be seen by anyone - not even her best friend. In addition, she is concerned that her data could get lost by misuse of the application. Emily is in the dilemma of initial trust. She wants to use the ubiquitous multi-display environment because it provides several benefits, but at the same time she needs to take a risk and rely on a system which she does not own and which she has little knowledge about.

Despite these concerns, Emily decides to send some pictures to the table in order to view them in full size since the table looks rather professional which helps Emily to form immediate trust. Now, Emily first selects the pictures on the mobile phone. Afterwards she places the mobile phone

MoMM 2010 November 10-12, 2010, Paris, France.
Copyright 2010 ACM 978-1-4503-0440-5

on the table. After establishing a connection between the mobile phone and the interactive table, she confirms the transfer of the selected images and the progress is visualised on the phone. At that moment, Emily is hoping that everything goes well and that her data does not get lost. Finally, her images become visible on the table and she confirms the successful transfer on the mobile phone. Now, she realises that some critical parts of the pictures became unrecognisable. The system uses a built-in privacy mechanism that recognises issues when other people are close to the table. Emily likes that support and thus she is even more confident with the system. She enlarges some of the pictures and tells her friend her holiday stories.

The illustrated incident was self-explanatory and positively perceived by Emily. Consequently user trust was not impaired. But other adaptations could happen without being self-explanatory, such as whenever some of the pictures would suddenly disappear which could be seen as a system error. In this situation user trust could be harmed since the adaptation might be perceived as negatively. All in all, user trust is highly situation-dependent and uncertain. A trust management is required to understand the relationship between all facets of user trust and the user response.

3. TRUST IN UBIQUITOUS DISPLAY ENVIRONMENTS

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 and colleagues [11] conducted a user study to identify 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 in case there is a mechanism for privacy protection.

Based on the evaluation of two mobile guides, Graham and Cheverst [7] analyse 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 correctly detect the user's current location 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 and colleagues [1] introduce the notion of crossmodal displays that enable users to access personalised information in public places while ensuring their anonymity. The basic idea is to publicly display the main information, but to add cues for individual users to prompt them to information that is relevant to them.

As a conclusion, 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 explicitly model the user experience of trust as a prerequisite for a trust management system.

A number of approaches have been presented to model trust in computational systems. 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 modelling of trust theory and its applications in agent-based systems). However, these approaches either focus on trust in software components or aim at modelling trust in human behaviour.

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 of 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 they are applied to. 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 [8] or Kini and Choobineh [9]. The more sociologically inclined authors [14] 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 and colleagues [6].

Our set of trust dimensions is based on interviews with 20 students of computer science who were asked to indicate trust factors of user interfaces that they felt contributed to their assessment of trustworthiness. The most frequent mentions felt 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”), security (”should safely transfer data”), privacy (”should not ask for private information”), seriousness (”professional appearance”) and reliability (”should run in a stable manner”).

The interviews gave a first impression on which factors influence the user's trust in a user interface. However, they do not provide any concrete information regarding their relative importance. To acquire more quantitative data, we conducted an empirical study which is described in the subsequent section.

5. EMPIRICAL VALIDATION OF TRUST DIMENSIONS AND USER FEELINGS

In order to determine the relative importance of trust triggers above in an ubiquitous display environment, we prepared an experiment that was inspired by the scenario described in Section 2. In particular, we presented our users with a setting consisting of a mobile phone and an interactive table (Microsoft Surface). The table served as the central medium for showing and editing multimedia data whereas the mobile phone was used to send data to or receive data from the table. Thereby, the transmission and the point of time of the presentation of the data on the table are critical moments for the user to trust.

The first objective of our study was to investigate the relationship between trust and trust triggers by means of concrete user data. In particular, we hypothesised that there was a positive correlation between trust on the one hand and basic usability, controllability, transparency, privacy, security and seriousness on the other hand.

A second objective of our study was to find out whether a low level of trust is reflected by negative user feelings. Previous research investigates how the emotional state of a user influences the establishment of trust (e.g. [5]). There is empirical evidence that positive emotions foster the establishment of trust while negative emotions tend to decrease trust. Prior experiments focused in most cases on emotions that were not related to the subsequent trust judgement task, see [4]. We assume that emotional states can also be directly associated with trust-related stimuli. In particular, we hypothesise that uneasiness, uncertainty, irritation and surprise are negatively correlated to trust.

5.1 Experimental Setting

In order to get a sufficient variety of user ratings, we built a number of prototypes where we manipulated the following variables: self-explainability, transparency, controllability, privacy. That is we produced a prototype that was less self-explainable (interface included no help function and no descriptive labels), a second prototype that was less transparent (system gave no reasons for its behaviour), a third prototype that was less controllable (system did not ask for user confirmations before executing an action), a fourth prototype that followed as less stricter privacy policy (system displayed all kinds of data on user request on the table disregardless of whether they were private or not) and finally a system that did not show any of these problems. In our first study, we decided not to manipulate the reliability of the prototypes and to present users only with prototypes showing a proper behaviour.

Figure 1 illustrates the most important screens of the unproblematic prototype during the data transfer from the mobile phone to the interactive table. At the beginning, the images are selected by the user on the mobile phone (see Screen 1). Afterwards the mobile phone is placed on the table. After establishing a Bluetooth connection between the mobile phone and the interactive table, the user confirms the sending of the selected images (Screen 2 - Do you really want to send these pictures?) and the progress (Screen 3 - Sending image 1/3...) is visualised on the phone. Finally, the images become visible on the table and the user confirms the successful transfer (Screen 4 - transfer successful) on the mobile phone. For the reverse procedure (transferring data from the table to the mobile phone), the same screens (Screen 2 to 4) as in Figure 1 are used. Instead of confirming the transmission of the data, the user now confirms the reception.

For our experiment a within subjects design was used. Thus, all subjects participated in all five conditions of the experiment. To prevent any ordering effects, we permuted the sequence of the different conditions with almost equal distribution for each prototype. After the successful completion of a condition with the prototype the subjects filled in an identical questionnaire.

In particular, the subjects had to rate the prototype according to the trust dimensions identified earlier (basic usability, controllability, transparency, privacy, security, seriousness and trustworthiness) as well as their emotions (uneasiness, insecurity, irritation and surprise) on a five point scale (from very low to very high). Afterwards, we used the results of the questionnaire to validate the relationship between trust and its dimensions as well as emotions.

5.2 Conducting the Experiment

We conducted the experiment with 20 people of which the majority (16 people) had a background in computer science. The average age of the subjects was 23.75 years (STD = 2.55) and except one person all subjects who participated in the test were male. The subjects rated their general trust into software systems with a mean value of 3.10 (STD = 0.79) and their knowledge about secure data transmission with a mean value of 3.5 (STD = 1.05). Before we started the experiment, each subject was introduced to the correct usage of the mobile phone and the interactive table which has a touch-sensitive display. Furthermore, we explained the subjects the purpose of our application running on the mobile phone.

During the experiment, each subject had to perform the following tasks with each of the five prototypes: (1) Select picture number one, three and five on the mobile phone and send them to the table. (2) Interact with the three pictures on the table and edit their size. (3) Send picture number three back to the mobile phone. Figure 1 shows a participant of the experiment while interacting with the mobile phone and the table.

5.3 Results and Discussion

To measure the degree of relationship between the ratings for trust and the ratings for the trust dimensions, we computed the Pearson product moment correlation coefficients. The test revealed a moderate to high positive correlation between the ratings for trust on the one hand and the ratings for seriousness (r = 0.724), controllability (r = 0.70), security (r = 0.62), privacy (r = 0.61) and transparency (r = 0.56) on the other hand. For all items, the correlation was very significant (p = 0.01). The better the ratings for controllability, transparency, privacy, security and seriousness, the better were also the ratings for trust. The strongest correlation was observed between the ratings for seriousness and the ratings for trust. Since the users were confronted with the system for the first time, they obviously had to rely on the first impression the system made on them when assessing the system's trustworthiness. As a consequence, there was a stronger correlation between the ratings for seriousness and the ratings for trust than between the ratings for the other items and the ratings for trust (which are too a larger extent based on experience). In our experiment, we did not observe any correlation between trust and basic usability ratings. As a potential reason, we indicate that no serious usability issues occurred when the users were interacting with the presented prototypes. Indeed our users rated the usability of the prototypes with a mean value of 4.01 (STD = 0.93) on a 5-ary scale. None of them thought the usability of any of the prototypes was very bad. There is a moderate positive correlation between the users' rating of usability on the hand and the users' rating of transparency

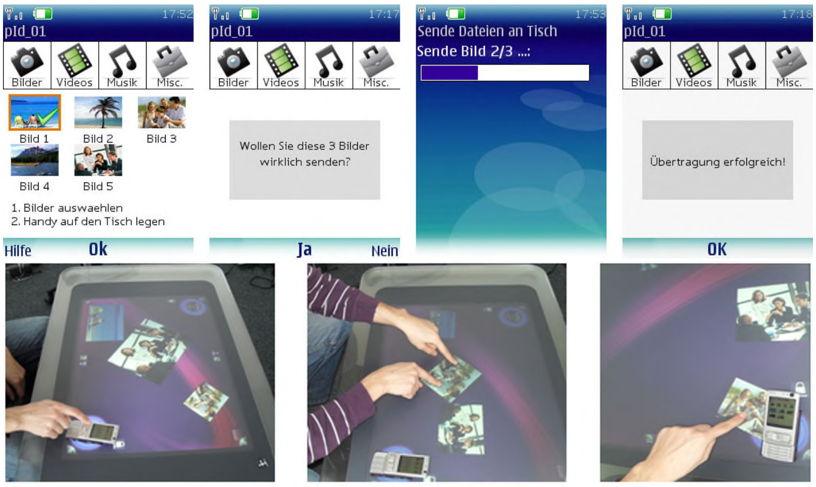

1. Bilder auswaehlen
2. Handy auf den Tisch legen

Wollen Sie diese 3 Bilder wirklich senden?

Sende Dateien an Tisch
Sende Bild 2/3 ...:

Übertragung erfolgreich!

Figure 1: Screen 1 to 4 of the Non-Problematic Prototype (first row) and the Interaction with the Non-Problematic Prototype on the Table (second row) .

(r = 0.22) and controllability (r = 0.26) on the other hand at the significance level of p = 0.05. Obviously, the subjects' ratings of transparency and controllability influenced their ratings of basic usability.

Finally, our results revealed a moderate negative correlation between trust on the one hand and uneasiness (r = -0.629), insecurity (r = -0.533), irritation (r = -0.484) on the other hand. For all items the correlation was very significant (p = 0.01). We conclude that poor transparency, poor controllability, poor security, poor privacy and poor seriousness result into a loss of trust which in turn leads to a feeling of uneasiness. Contrary to our expectations, we did not find any correlation between the users' ratings of surprise and the user's rating of trust.

6. TOWARDS AN AUTOMATIC TRUST MANAGEMENT SYSTEM

In the following, we describe first ideas regarding an automated trust management system that assesses the user's immediate trust in a system, monitors it over time and applies appropriate measurements to maintain trust (see [15]). The trust management system is based on findings from the literature (e.g. [8, 9, 14]) as well as our empirical study that investigated the relationship between trust and its dimensions. 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. However, it is hard to formulate rules that predict in a deterministic manner how a person will respond to a critical event. We therefore aim at a model that is able to represent uncertainties.

- *Trust as a multifaceted concept*
  As shown in Section 4, trust is a multi-faceted concept. We therefore aim at a computational model that is able to explicitly represent the relative contribution of the trust dimensions to the assessment of trust. In addition, the model should allow us to easily add trust dimensions based on new experimental findings.

- *Trust as a dynamic concept*
  Trust depends on experience and is subject to change over time. Lumsden [10] 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 users' experience of trust during an interaction.

6.1 Using Bayesian Networks to Model Trust

Based on the considerations above, we have chosen to model the users' feelings of trust by means of 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 [13]).

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. Furthermore, the connection between critical events and trust is inherently non-deterministic. For example, we cannot always be absolutely sure that the user notices a critical event at all. It may also happen that a user considers a critical event as rather harmless. Bayesian Networks allow us to make predictions based on conditional probabilities that model how likely the value of 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 to model the relationship between trust and its dimension in a rather intuitive manner. For example, it is rather straightforward to model that reduced transparency leads to a decrease of 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 Section 5.

In Figure 2, a Bayesian Network for modeling trust is shown. The left part of the Bayesian Network represents the factors that influence the establishment of immediate trust while the right part of the Bayesian network models the development of interaction-based trust. 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. 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 development of immediate trust and interaction-based trust depends on the user's trust disposition which is characterized by his or her competence and their general confidence into technical systems.

6.2 Monitoring Trust over Time

After smoothly interacting with a system over a longer period of time, the users' trust into 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 optimisations of the system to re-establish trust. As a consequence, we do not only need a model that describes the relationship between user trust and its dimensions, but also a model that explains the dynamics of trust. Dynamic Bayesian Networks allow us to model the dependencies between the current states of variables and earlier states of variables.

In the middle part of Figure 2, a fraction of the Bayesian Network is shown illustrating how trust develops over time depending on the user's immediate level of trust and events occurring at time $t = 1$. Due to space limitations, we only present one time plate ($t = 1$). For simplicity, we only consider the user's level of trust at time t_{i-1} to determine the user's level of trust at time t_i. We introduce a variable called *System Event* to represent for each point in time what (if any) kind of event occurred. The values of the variables *System Event* influence the values of the dimension variables *Comfort of Use*, *transparency*, *controllability*, *privacy* and *reliability*.

6.3 Maintaining User Trust

The Bayesian Network presented above supports us in making decisions on how to maintain trust in critical situations. Such situations arise, among other things, when other people enter the user's private space [11], when the system has to generate presentations based on inaccurate user or context data [7] or when the system's adaptation behavior mismatches the user's expectations [6]. Within the Bayesian Networks such situations can be handled by adding decision and utility nodes. A decision node represents all choices that can be made by the system while a utility node indicates the utilities of all possible outcomes. As an example, let us assume a user wishes to display data on a public display. To cope with such a request, the system may consider four options: (1) transferring all data to the public display, (2) filtering out data that the system considers as private or (3) asking the user for confirmation. In the Bayesian Network shown in Figure 2 we introduced a decision node called *System Action* to represent all actions the system may decide to execute. In the example, Option (1) may raise serious privacy concerns, option, (2) may confuse users and option and (3) is rather cumbersome. In addition, option (1) and (2) might give users the feeling that they have no longer the system under control. The arc between the decision node and the nodes for the dimensions of trust represents such influences. All decisions are evaluated based on the usefulness of their consequences. The utility node attached to the node called *user trust* indicates the utility of the single decisions based on user trust.

7. CONCLUSION

In this paper we aimed at the management of user trust in ubiquitous computing systems. Our introduced Bayesian Network provides knowledge about the interplay between trust dimensions and user trust at the design time of a system. At the runtime of the system, the Network also assists the dynamic monitoring of user trust for estimating influences of incidents on user trust. In future work we will evaluate different parts of the Bayesian Network. In particular, we would like to validate aspects of immediate and interaction-based trust.

8. ACKNOWLEDGEMENT

This research is partly sponsored by *OC-Trust* (FOR 1085) of the German research foundation (DFG). A special thank to Georg Döhring, Jessica Eichberg, Dominik Hecht, Michael Kutsch, André Lohrenz and Thuy Linh Nguyen for their work on the surface application.

9. 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] K. Cheverst, A. Dix, D. Fitton, C. Kray, M. Rouncefield, C. Sas, G. Saslis-Lagoudakis, and J. G. Sheridan. Exploring bluetooth based mobile phone interaction with the hermes photo display. In *MobileHCI '05: Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*, pages 47–54. ACM, 2005.

[4] J. Dunn and M. Schweitzer. Feeling and believing: The influence of emotion on trust. *Journal of Personality and Social Psychology*, 88:736–748, 2005.

[5] J. R. Dunn and M. E. Schweitzer. Feeling and believing: The influence of emotion on trust. *Journal*

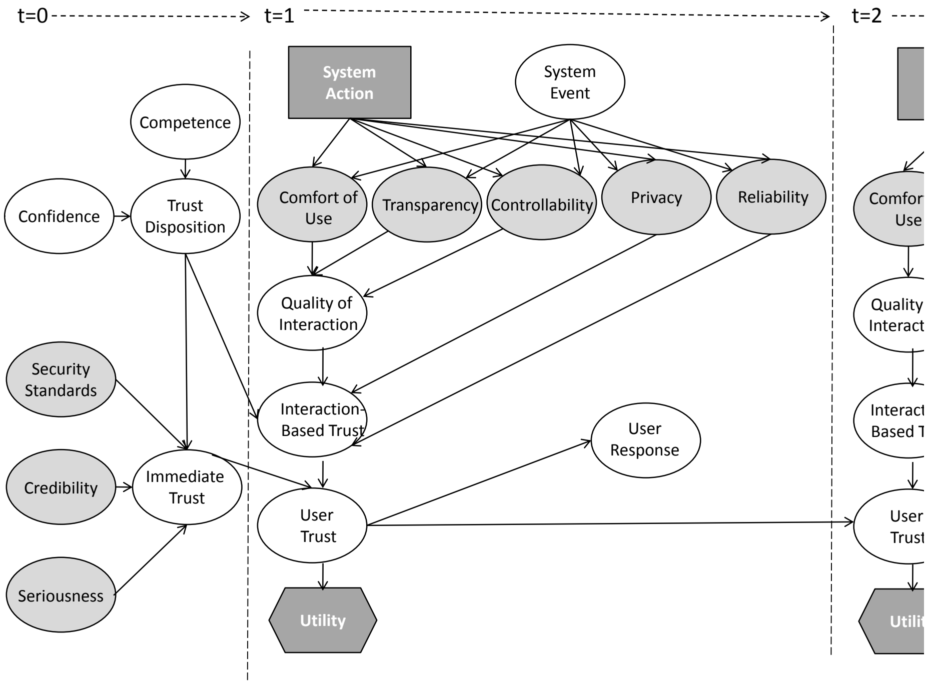

Figure 2: Modeling Trust by Means of a Bayesian Network

of Personality and Social Psychology, pages 736–748, 2005.

[6] 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. ACM, 2008.

[7] 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.

[8] T. Grandison and M. Sloman. A survey of trust in internet applications. *IEEE Communications Surveys and Tutorials*, 3(4):2–16, 2000.

[9] 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.

[10] 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. British Computer Society, 2009.

[11] C. Röcker, S. Hinske, and C. Magerkurth. Intelligent privacy support for large public displays. In *Proceedings of Human-Computer Interaction International 2007 (HCII'07)*, 2007.

[12] L. Rothrock, R. Koubek, F. Fuchs, M. Haas, and G. Salvendyk. Review and reappraisal of adaptive interfaces: Toward biologically inspired paradigms. volume 3, pages 47–84, 2002.

[13] S. J. Russell and P. Norvig. *Artificial Intelligence a modern approach.* Prentice Hall, Upper Saddle River, N.J., 2nd international edition edition, 2003.

[14] 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.

[15] 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.