Trust-based decision-making for smart and adaptive environments

Stephan Hammer, Michael Wißner, Elisabeth André

Angaben zur Veröffentlichung / Publication details:

Hammer, Stephan, Michael Wißner, and Elisabeth André. 2015. "Trust-based decision-making for smart and adaptive environments." User Modeling and User-Adapted Interaction 25 (3): 267–93.
https://doi.org/10.1007/s11257-015-9160-8.

User Model User-Adap Inter (2015) 25:267–293
DOI 10.1007/s11257-015-9160-8

Trust-based decision-making for smart and adaptive environments

Stephan Hammer[1] · **Michael Wißner**[1] · **Elisabeth André**[1]

Received: 2 August 2014 / Accepted in revised form: 10 February 2015 /
Published online: 16 April 2015

Abstract Smart environments are able to support users during their daily life. For example, smart energy systems can be used to support energy saving by controlling devices, such as lights or displays, depending on context information, such as the brightness in a room or the presence of users. However, proactive decisions should also match the users' preferences to maintain the users' trust in the system. Wrong decisions could negatively influence the users' acceptance of a system and at worst could make them abandon the system. In this paper, a trust-based model, called User Trust Model (UTM), for automatic decision-making is proposed, which is based on Bayesian networks. The UTM's construction, the initialization with empirical data gathered in an online survey, and its integration in an office setting are described. Furthermore, the results of a live study and a live survey analyzing the users' experience and acceptance are presented.

Keywords Computational trust · Context awareness · Proactive systems · Energy saving

✉ Stephan Hammer
hammer@hcm-lab.de

Michael Wißner
wissner@hcm-lab.de

Elisabeth André
andre@hcm-lab.de

[1] Human Centered Multimedia, Augsburg University, Universitätsstr. 6a, 86159 Augsburg, Germany

1 Introduction

Recent advances in sensor technologies enable us to capture the users' physical context continuously and to personalize information and services to them in real-time. One area that might greatly benefit from new forms of context-aware system behaviors are smart energy systems that exploit information on the users' environmental context to support them in saving energy—either by controlling energy-consuming devices proactively or by giving the users personalized advice.

Reducing energy consumption has been a major concern for more than four decades, and many approaches aimed at supporting sustainability were developed during this time (Hazas et al. 2011; DiSalvo et al. 2010). Some tried to improve people's environmental awareness by providing detailed feedback on their energy usage (Gamberini et al. 2012). Others tried to persuade people to reduce their energy demand by exploiting social factors and utilizing, for example, cooperative pervasive games (Simon et al. 2012).

A number of energy management systems allow users to control devices, such as displays or lights, remotely or by setting up time tables. Furthermore, attempts have been made to adjust the energy consumption implicitly based on various context information that describes the users' and the system's surroundings (Cheverst et al. 2005). For example, displays or lights can be switched off if they are not needed. On the one hand, a system that autonomously performs energy saving actions contributes to the users' convenience. On the other hand, proactive system actions are not always understood by users and limit their control over the system. As a consequence, users might lose trust in such a system and give up using it.

For illustration, let us assume a lamp is burning in the user's office even though daylight suffices for performing the work. How should an energy management system react in such a situation? Should it assume the users are aware of their energy consumption and will take necessary actions themselves? Should it switch off the light autonomously? Or should it ask the user for permission via messages presented on the user's display or mobile phone?

In the first case, the system would leave the responsibility for energy reduction with the users, and there would be the risk that users do not see any benefit in the energy management system. The second approach bears the danger that the users do not understand the rationale behind the system's behavior and perceive it as not sufficiently self-explanatory or even as acting in a random manner. In the last case, the system's behavior might appear transparent. However, users might nevertheless be upset because permanent and obtrusive messages interrupt their workflow. The example illustrates that a system needs to carefully balance the benefits and drawbacks of possible actions so as not to risk the users losing trust in its workings.

In this paper, a decision-theoretic approach to a trust management system for smart and proactive environments based on Bayesian Networks, the User Trust Model (UTM), is presented. It assesses the users' trust in a system, monitors it over time, and applies appropriate system reactions to maintain users' trust in critical situations (Yan and Holtmanns 2008). Section 2 discusses prior work in modeling trust, considering work done in the area of agent-based modeling, social media and adaptive and personalized systems. After that, the UTM's construction (Sect. 3) and its integration

into an office setting and the initialization with empirical data (Sect. 4) are described. Sections 5 and 6 present a live study and a live survey investigating the users' experience with and acceptance of the system. Finally Sect. 7 gives a conclusion and an outlook on future work.

2 Related work

In the area of user modeling, research on computational models of trust has become very popular due to the obvious overlap between trust and reputation modeling in recommender systems and social media. Nevertheless, approaches that model trust as a user experience and focus on the affective dimension of trust are rare. This is unsurprising because the psychological aspects of trust are hard to measure directly. In this section, we will first give an overview of computational models starting from approaches that have been presented for agent-based societies, social networks and recommender systems. After that, we discuss how the concept of trust has been treated in ubiquitous computing.

2.1 Computational models of trust

Much of the original research on trust comes from the social sciences. 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 to come up with dedicated models of trust that are adapted to the specific requirements of their domains and the context they are applied to. Since trust is a social phenomenon, it seems to be promising to exploit models that have been developed to characterize trust in human societies as a basis for computational models of trust.

Especially in the area of multi-agent systems, computational models for trust-based decision support have been researched thoroughly. Pioneering work in this area has been conducted by Marsh (1994) who modeled trust between distributed software agents as a basis for the agents' cooperation behavior. Computational mechanisms that have been proposed for trust management in agent-based societies include Bayesian Networks (Wang and Vassileva 2005), Dempster–Shafer Theory (Yu and Singh 2002), Hidden-Markov Models (Vogiatzis et al. 2010), Belief Models (Jøsang et al. 2006), Fuzzy models (Castelfranchi and Falcone 2010), game-theoretic approaches (Sankaranarayanan et al. 2007) or decision trees (Burnett et al. 2011). There is empirical evidence that the performance of agent-based societies may be improved by incorporating trust models.

In contrast to the approaches above, work in the area of social media aims to model trust between human users, see Sherchan et al. (2013) or Bhuiyan et al. (2010) for a survey investigating trust in social networks. Using algorithmic approaches or machine learning techniques, trust between users is derived from objective observations, such as behavior patterns in social networks. For example, Adali et al. (2010) assess trust between two users based on the amount of conversation and the propaga-

tion of messages within Twitter. Other approaches derive trust that is given to users from community-based reputation or social feedback (e.g. Ivanov et al. (2013)).

Computational models related to trust have also been explored in the area of recommender systems. Obviously, it does not suffice to generate recommendations solely based on the users' profile and preferences. In addition, the trustworthiness of people, organizations and services involved in the recommendation process have to be taken into account. There is empirical evidence that computational models of trust may help improve the recommendation accuracy of traditional collaborative filtering approaches, see, for example, O'Donovan and Smyth (2005).

Our research focuses on trust which users experience when interacting with a software system. A system may be robust and secure, but nevertheless be perceived as less than trustworthy, for example, because its behavior appears less than transparent or hard to control. Following the terminology by Castelfranchi and Falcone (2010), our work focuses on the affective forms of trust that are based on the user's appraisal mechanisms. That is why we aim at the development of computational trust models that capture how a system—in this paper a smart environment for energy saving—is perceived by a user who is confronted with it.

Computational models that assess trust felt by a user while interacting with a system are rare. There is a large amount of work that aims to identify factors that impact user trust. For example, Glass et al. (2008) research trust-enhancing factors for adaptive and personalized applications. However, they do not implement a model of the user's trust into an adaptive and personalized system based on these factors. Starting from the observation that people respond to technology socially, Lee and See (2004) discuss psychological factors of trust, such as the visual appearance of the interface, that influence to what extent people rely on technology. Yan and Holtmanns (2008) model captures the trust which users experience when interacting with mobile applications. In order to present users with recommendations that help increase their trust, they identified various behaviors that can be monitored by a mobile device in addition to external factors, such as brand impact. The benefits of this approach have been shown by means of simulations. However, the approach has not been embedded in an adaptive and personalized mobile application to control the selection of system actions during an interaction with the user.

2.2 Trust in ubiquitous computing

In the area of pervasive computing, the topic of trust has attracted a significant amount of interest. This comes as no surprise since the high dynamics and openness of pervasive computing environments come not only with great benefits, but also a number of security risks. Due to the large variety of smart objects and devices that can exchange information, the underlying infrastructure is heavily imperiled by manipulations. Typically users interact with such environments on a short-term basis without having the possibility to verify the security of the underlying infrastructure. Vice versa access control in open environments which people can enter and leave at any time is a challenging task. To solve these issues, a number of research projects have investigated how to apply trust mechanisms from the area of network security to pervasive

computing. A common approach is to explicitly model trust relationships between physical devices and exploit this information to choose appropriate devices for cooperatively solving a task, see, for example, Denko et al. (2011).

At the same time, a significant amount of private data is collected silently using sensors worn on the user's body as well as external sensors smoothly integrated into the user's environment. On the one hand, the comprehensive collection of user data contributes to a better personalization of information and services. On the other hand, ubiquitous user modeling may be considered as a threat to privacy. To mitigate this threat, a variety of mechanisms has been presented to preserve the user's privacy and hide confidential information from others, such as preventing the tracking of tagged consumer items or displaying private information on the user's personal device. In our earlier work (Wißner et al. 2014), we presented a trust management system that assesses the users' trust in ubiquitous display environments and decides how and where to present personal information based on location-based context factors, such as other people in the user's immediate neighborhood. While our work did not distinguish between trusted and non-trusted people in the user's physical environment, the approach by Arimura et al. (2014) makes use of physical trust relationships between users. First, face-to-face communication between users is triggered by the authentication system in order to visually confirm users. Only in cases where visual authentication was not successful, the system would ask for additional authentication information, such as passwords. The interesting idea behind this work is the face-to-face communication between users that the authentication process encourages as a basis for the creation of trust between people.

Another factor that may affect the users' trust is the high heterogeneity, uncertainty and unpredictability of pervasive environments. Despite the large number of sensors that are employed to capture user and context information, the analysis and interpretation of the sensor data are error-prone. In our earlier work (Kurdyukova et al. 2012), we aimed to personalize recommendations based on the composition of social groups that were detected using video-based face detection software. However, in natural environments, the accuracy rates of the recognition process were often affected by noisy conditions. Adjusting recommendations to an audience based on incorrect context information may result in system behaviors that appear less than transparent to users. For example, interviews with users of an adaptive digital signage system that automatically adapted to the assumed interest of an audience revealed that some users did not understand the adaption mechanism, but rather had the impression that the system was presenting randomized information (see a study by Müller et al. (2009)). To counter comprehensibility issues caused by inaccurate context information, a number of researchers propose to display confidence values to users, see, for example, the work by Antifakos et al. (2005) or Yan et al. (2010).

Lim and Dey (2010) present a toolkit for generating eight different kinds of explanations automatically (such as what-if, why, how-to etc.), in order to increase the transparency of context-aware systems. Even though the connection to user trust is emphasized in their paper, they do not provide a mechanism to computationally model user trust. Cheverst et al. (2005) investigate techniques to increase the transparency of a system and to give users a higher level of control in a smart office environment. Their work is similar to ours since it investigates the tension between proactive system

behavior and user control and aims at improving the transparency of system behavior. Even though the topics they address have a tight relationship to user trust, they do not explicitly model user trust to decide on appropriate system behaviors.

3 The User Trust Model

The main idea underlying our approach to model the users' trust in a computer system is to derive the trust from a set of intermediate dimensions, the so-called trust dimensions. These trust dimensions describe relevant properties of the system in question. Their definition is based on an earlier survey (Steghöfer et al. 2010) where we elaborated on the determinants of trust in highly dynamic computing systems and interviews with users (Leichtenstern et al. 2010) in order to identify trust factors that are of relevance to user interfaces. In these interviews, users were asked to indicate factors of trust that they felt contributed to their assessment of the trustworthiness of a user interface. The most frequent mentions felt into the following categories that formed the basis of our User Trust Model (UTM) (Kurdyukova et al. 2012):

- Comfort of Use ("The system should be easy to handle")[1]
- Transparency ("I need to understand what the system is doing")
- Controllability ("I want to be in control of the system's actions")
- Privacy ("The system should neither ask for nor reveal private information")
- Reliability ("The system should run in a stable manner")
- Security ("The system should safely transfer data")
- Credibility ("The system should have been recommended by others")
- Seriousness ("The system should have a professional appearance")

We have chosen to model the users' feelings of trust by means of Bayesian networks. A Bayesian network (BN) is a directed, acyclic graph in which the nodes represent random variables while the links connecting nodes describe the direct influence in terms of conditional probabilities (Russell and Norvig 2009). BNs were chosen because they very well meet requirements that should be accounted for by models aimed at assessing users' trust towards computer systems:

Trust as a subjective concept: Throughout literature, there is a consensus that trust is highly subjective. Different users respond individually to one and the same event. While some might find it critical if a system acts autonomously, others might not care. Also, a generally trusting person is also more likely to trust a computer system. To represent this subjective nature of trust in the BN, the model's uncertain belief about the user's trust can be represented by a probability distribution over different levels of trust.

Trust as a non-deterministic concept: The connection between events and trust is inherently non-deterministic. For example, we cannot always be sure that the user notices a critical event at all. Users may also consider a critical event as rather harmless.

[1] Typical statements made by the participants of our earlier study (Leichtenstern et al. 2010) are indicated in brackets.

BNs allow us to make predictions based on conditional probabilities that model how likely the value of a 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. This allows for a much more flexible approach than, for example, rigid rules that exactly predict how a certain event or situation changes the user's trust.

Trust as a multifaceted concept: Computational models should be able to explicitly represent the relative contribution of different trust dimensions to the assessment of trust and should help predict the user's trust based on these dimensions. Furthermore, it should be easy to alter the model by adding or removing trust dimensions based on new experimental findings or if a certain dimension is not applicable in a given system. With BNs the modeling of relationships between trust and its dimensions is rather intuitive. For example, it is rather straightforward to model that reduced transparency leads to a decrease of trust. In the BN in Fig. 1 each trust dimension is represented by a specific node. Since exact probabilities are difficult to determine, the conditional probabilities were derived from empirical data collected in an online survey, see Sect. 4.1.

Trust as a dynamic concept: Trust depends on experience and changes over time. Following Lumsden (2009), we distinguish between *Initial Trust* and *Interaction-Based Trust*. Both contribute to the user's overall trust in the system. Initial trust dimensions, such as seriousness, come into effect as soon as a user gets in touch with the system while interaction-based trust dimensions, such as transparency of system behavior, influence the users' experience of trust during the interaction. Again, BNs

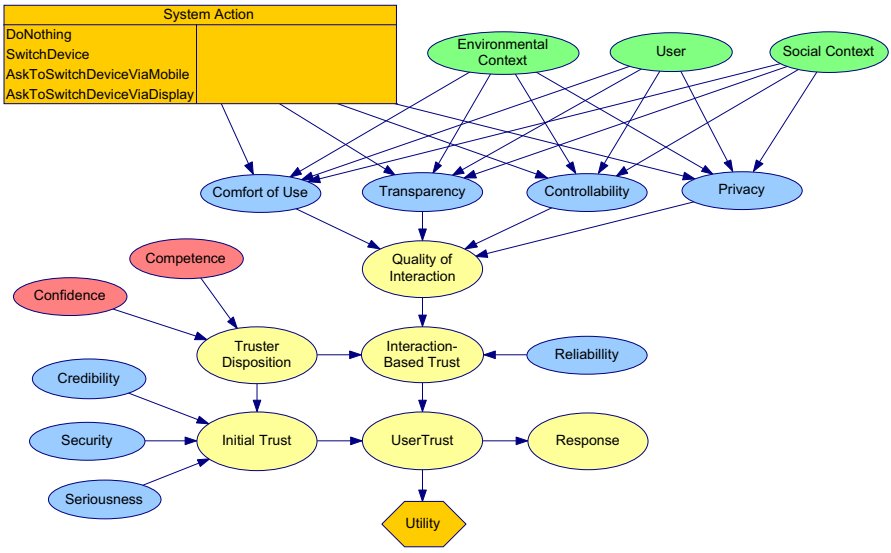

Fig. 1 Generic User Trust Model for a Smart Energy System. *Green* context information; *Red* user traits; *Blue* trust dimensions; *Orange* decision nodes (system actions and utility node). (Color figure online)

allow us to model this distinction: Which trust dimension affects which aspect of trust and how both aspects influence the overall trust.

In Fig. 1, a generic BN for modeling trust in the smart energy system is shown. As mentioned above, each trust dimension is represented by its own node (shown in blue). Each trust dimension either affects the *Initial Trust* or the *Interaction-Based Trust*. The determinants of *Initial Trust* are *Security*, *Seriousness* and *Credibility*. *Security*, for example, could be conveyed by the use of certificates. A system's *Seriousness* is reflected, for example, by its look-and-feel. *Credibility* could be supported by additional information, such as a company profile.

Strictly speaking we do not model the relationship between user trust and actual system features, but the relationship between user trust and the user's subjective reflection of them. Even the best security standards will not put a user at ease if he is unaware of their existence. Visual indicators of security, such as closed padlock icons, might fail their purpose because users might be missing the knowledge to interpret them correctly (Dhamija et al. 2006). A recent paper by Florencio et al. (2014) shows that generally accepted rules for password security have to be revisited. Thus not even experts might be able to correctly assess a system's security. As a starting point, we do, however, not further distinguish between actual system features and the user's subjective reflection of them.

For the sake of simplicity, we assume that the initial trust dimensions do not change over time, i.e. we do not consider that a user might only notice a security certificate after having worked with the system for a longer time. The determinants for *Interaction-Based Trust* are *Quality of Interaction* and *Reliability*. The *Quality of Interaction* is an aggregation of *Transparency*, *Controllability*, *Comfort of Use* and *Privacy*.

Both the establishment of *Initial Trust* and *Interaction-Based Trust* are influenced by the users' *Trust Disposition* which is characterized by their *Competence* and general *Confidence* towards technical systems (shown in red), thus allowing to model the subjectivity of trust mentioned above.

By *Confidence*, we mean the propensity of individuals or a group of individuals to trust technology in general. Highly confiding people are more likely to trust a particular system than wary people. In our case, *Competence* refers to the user's general technical knowledge. A lack of knowledge about the performance of technical systems may lead to a miscalibration of trust (Lee and See 2004). For example, non-expert users might be tricked by a shiny interface and overestimate a system's ability, i.e. build up an inappropriately high level of initial trust. However, their interaction-based trust is likely to be seriously affected by system failures—in particular if they are not able to explain them.

We treat the trust dimensions as hidden variables, i.e. they cannot be observed directly, but may be inferred from observable context variables that depend on the specific system (shown in green). For example, the Smart Energy System currently considers the *User* state, the *Social Context*, and the *Environmental Context*.

Finally we included a node called *System Action* (shown in orange in the upper left corner), representing the different actions the system could take to react to context changes, such as "Switch the light on automatically" if the "User is arriving". Knowing the contextual situation, the BN can estimate the impact of the different system

reactions on the trust dimensions and thus on the user's overall trust. As an example, *Transparency*, on the one hand, could be negatively affected if the system automatically turns off the light when it is dark outside, while, on the other hand, *Controllability* could be negatively affected if the system switches the light on autonomously when the user is arriving.

In order to use the BN for decision-making, it was extended to an influence diagram by modeling *System Action* as a decision node and adding a *Utility* node that computes the utility of all possible actions and their consequences and returns the action with the highest utility. Since the goal of our work is to maintain and maximize user trust, the *Utility* node is attached to a node representing the *User Trust* and measures the utility of each single decision in terms of the resulting user trust.

4 Building a smart office

In the following, we demonstrate how the UTM can be used to guide decision-making in an energy-aware device management system that controls the displays and the light in an office occupied by several people. For each type of device, a BN was constructed from the generic model described in the last section. Modeling the BN and integrating it into the system was done by using the GeNIe modeling environment and the SMILE reasoning engine respectively.[1] Note that the terms UTM and BN will be used interchangeably from now on.

In the BN for operating the light, whether and, if so, which action the system takes to control the light basically depends on the luminance outside, the user's presence and whether his coworker is present. If the system recognizes a situation in which the light might be adjusted, it may perform the corresponding action autonomously or ask the user for permission via the mobile phone or via the display of the user's PC. In order to not risk disturbing the user, the system might even decide to do nothing, even if there was an action that could have saved energy. The BN for the display has a similar structure. However, it relies on a more fine-grained representation of the user's current activity to distinguish, for example, whether the user is sitting in front of a PC and working with it or engaged in other activities, such as reading a book or leaving his desk while staying in the room. An overview of possible system actions and the utilized context information in both BNs is given in Table 1.

Figure 2 shows the overall architecture of the Smart Office System (numbers in circles refer to the example below). The system runs on a central server which also stores the two UTMs for the light and display. The data needed to recognize the context information for both UTMs is gathered by Arduino-Sensors[2] that are distributed in the office. We utilize light sensors to measure the outdoor luminance, ultrasonic sensors on the desks to detect the presence of persons and a flex sensor at the door to determine whether the office is empty (assuming that the door would be closed in this situation). The control of the devices (display and light) is conducted via a HomeMatic[3] system

[1] http://www.genie.sis.pitt.edu/.

[2] http://www.arduino.cc/.

[3] http://www.homematic.com/.

Table 1 Possible system reactions in different contextual combinations

Device	Situation			System reaction
	User	Social context	Luminance outside	
Display	Working at PC	–	–	Switch display automatically
	Idle at PC	–	–	Ask to switch via smartphone
	Away from desk	–	–	Do nothing
	Out of room	–	–	
Light	Arriving	Coworker present	Dark	Switch light automatically
	Present	Coworker away	Bright	Ask to switch via smartphone
	Leaving			Ask to switch via display
				Do nothing

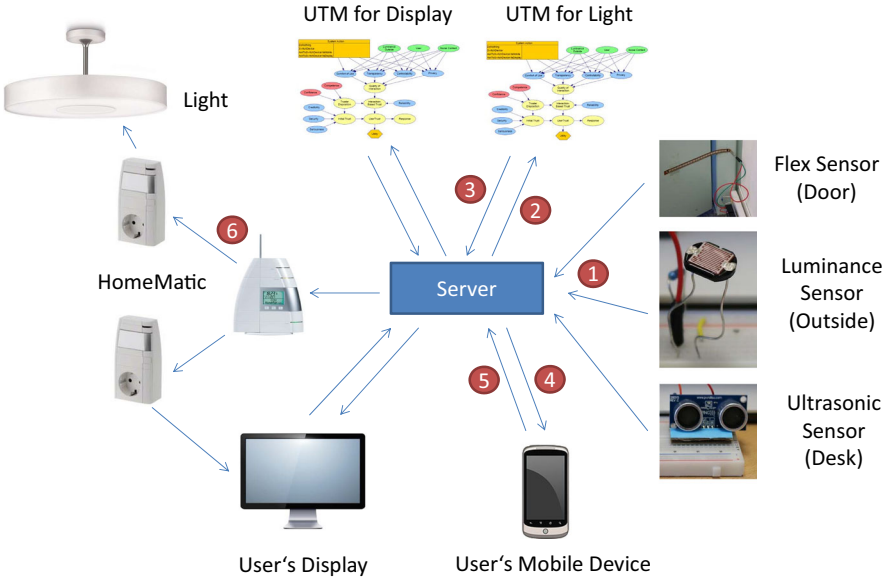

Fig. 2 Architecture of the Smart Office System

and remote controlled plugs. Finally, the server can send messages to both the user's display and their mobile device.

As an example of how the system would dynamically adapt the light, let us consider the following example (see circled numbers in Fig. 2). The user is alone in the office, working on the computer. It is still early in the morning, so it is dark outside and the light is on. The system knows this since it is aware of the states of all devices,

and it regularly polls the sensors for the most recent context data (1). As long as the situation does not change (i.e. the context data stays the same), nothing needs to be done. However, a bit later it gets bright outside, which the system registers as a new situation. This new situation is then entered into the appropriate BNs (2). In our case, only the one for the light is affected, outside luminance does not concern the display.

The UTM now has to act and decide on a system action and may thus consider four possibilities to cope with this changed situation: (a) Do nothing, (b) Switch off the light automatically, (c) Ask the user via a message shown on their display whether the light should be switched off, or (d) Ask the user via a message shown on their mobile device. Considering the example, option (b) offers no control to the user and may confuse (i.e. be less transparent) since the action happens automatically and without explanation, but offers a high comfort of use. Option (c) offers more control and transparency but might disturb the user, but at least they can respond immediately, while with option (d) they can ignore the message on their mobile device and thus continue working undisturbed, but if they want to respond they have to pick up the device first. Finally, option (a) certainly leaves the user undisturbed and in control but might not be the proper reaction one expects from a smart office system (and would thus negatively impact comfort of use and transparency). For each of the four possible actions, their impact on the different trust dimensions and the resulting user trust is now calculated. The user trust directly translates to utility (the higher the trust, the higher the utility) and the action with the highest utility is chosen and communicated back to the server (3). Let us assume that it is option (d). Hence, a message is sent to the user's mobile device, giving them the choice of switching off the light or not (4). They choose yes, the answer is sent back to the system (5) which then switches off the light via the HomeMatic (6).

4.1 Gathering empirical data (online survey)

In order to be able to generate decisions, the BNs had to be initialized with data. Both for the light and the display, we collected data in a web-based survey. In both surveys, participants were confronted with textual descriptions of typical situations during daily office routines. For each situation, possible system actions were proposed for the respective device that were supposed to improve the energy consumption of the users. Table 1 summarizes the situations represented by different settings of contextual variables and the possible system reactions.

The purpose of the survey was to discover for each situation which of the system reactions succeeded in maintaining user trust and which did not. To this end, the participants had to rate the system reaction in terms of *Transparency*, *Controllability*, *Comfort of Use*, and *Trust* using a 5-point Likert scale:

- Q1: I understood why the system was reacting in this way.
- Q2: I had control over the system.
- Q3: I found the system comfortable to use.
- Q4: I found the system to be trustworthy.

In contrast to our earlier work Wißner et al. (2014) where we focused on the relationship between privacy and user trust, the current work investigates the tension between

controllability, transparency and comfort of use. The scenarios were designed in a way that privacy issues were less of a concern (even though they could not be completely excluded). Consequently, we did not include any questions related to privacy in the user study.

All in all, 16 participants (7 female, 9 male) evaluated the situations for the light and 22 participants (9 female, 12 male) rated the situations for the display. The participants were aged between 24 and 51 years (mean: 28).

4.2 Initializing the Bayesian network

The quantitative data obtained in the online survey enabled us to derive and model probability distributions for each trust dimension for all combinations of context and system reaction. The probability distributions for other node combinations that were not part of the data collected in the online survey (e.g. how *Confidence* and *Competence* influence *Trust Disposition*) were modeled after the results from a previous study (Bee et al. 2012). However, data for other user groups can be easily integrated into the BN by replacing the corresponding distributions in the BN. An interesting resource to explore is the work by Westin, who conducted a large number of studies to determine the percentage of people with certain levels of distrust or privacy concerns, see Kumaraguru and Cranor (2005) for a survey of these studies.

5 Live study in the lab

The BN has been trained with data from an online survey. Such data bear the advantage that data are relatively easy to obtain since the participants do not have to be presented with an interactive system. However, an online survey might not convey the experience of a real interaction and thus affect the ratings of the users.

Thus, the question arises of to what extent the BN is able to predict user trust and user preferences in a live setting based on training data that have been acquired by an online survey. To shed light on this question, we decided to conduct a live study in which the participants were presented with the smart office environment and actually able to interact with it.

The purpose of this study was to evaluate the decisions taken by the UTM focusing on two criteria: (1) Would the chosen system reactions affect the users' feelings of trust and the related trust dimensions in a positive way? (2) Would the system reactions match the actions favored by the users? Apart from evaluating the BN approach, we investigated the users' experience and acceptance of our smart office environment.

5.1 Experimental setting

During the study the participants had to run through different tasks and situations, all of which simulated the daily routine in an office occupied by several people. Changes in the participant's and the colleague's state (social context) were triggered by the participants themselves and by one of the experimenters who played the role of the participant's colleague. To ensure that all participants conducted the study under the same

conditions and in a realistic way, the room was darkened and changes in the outdoor luminance were simulated by a lamp and by covering and uncovering the light sensor.

5.2 Conducting the study

At first the participants had to provide general demographic information and information about their experience with home automation systems and their trust towards computer systems in general. Furthermore, the participants were asked whether they considered themselves to have a trusting nature.

After a short introduction to the setting and the scenario, the participants had to conduct the first task, and the system showed the reactions that were selected for both devices according to the UTM. After that, the participants had to fill in a short questionnaire for each of the reactions. Each questionnaire included questions Q1–Q4, which were also asked in the online survey. Furthermore, the users were asked to choose their preferred system action. For instance, the statement concerning the display and the first task was: "When I enter my office and sit at my desk, I prefer ...

- P1: ...no reaction from the display.
- P2: ...to switch the display on automatically.
- P3: ...to be asked via smartphone for permission to switch on the device.

After that, the procedure continued with the next task and the respective questionnaire. All tasks, the corresponding situations and the selected system reactions triggered by context changes are summarized in Table 2. To make the experiment more realistic, the tasks were embedded in a coherent story.

After rating the last task, the participants had to state what they liked and disliked about the system and to rate statements related to their experience during the usage and their attitude towards the system.

5.3 Results

Overall six women and 18 men aged between 23 and 33 (mean: 26) took part in the study. They studied and worked in all kind of professions related (88 %) and not related (12 %) to computer science. All statements in the questionnaires could be rated on a 5-point Likert scale. Ratings lower than 3 were interpreted as disagreement, ratings higher than 3 as agreement with a statement, and a rating of 3 as a neutral attitude. Only five persons reported a high or very high amount of experience with technology for controlling parts of their home environment, such as automatic timers or blind control systems. Eighteen people rated their experience with home automation technology as low or very low. One participant gave a neutral rating.

The participants also had to reflect on their confidence. They had to answer two general statements and one statement related to computer systems. Concerning the statement: *I act based on the saying "Trust, but verify"*, only one participant disagreed. 63 % of all participants agreed and 33 % had a neutral attitude. Concerning the statements *I am overly trusting* and *On most systems, you can be assured that they will do what they should*, one third each agreed, disagreed, or rated neutrally.

Table 2 Tasks, changed context variables and system reactions of the user study

Task	Situation			System reaction	
	User state	Social context	Outside luminance	Light	Display
Enter the office	Arriving	Coworker away	Dark	Phone	
Sit down at PC	Working at PC				Auto
It is getting light					
Check slides for mistakes			Bright	Display	
The participant's colleague enters the room and sits down at the desk					
Take book X off the shelf	Away from desk	Coworker present			Nothing
Come back and read chapter Y	Idle at PC				Auto
Add a slide about Z	Working at PC				Auto
It is getting dark					
			Dark	Phone	
The participant's colleague leaves the room					
Finish work and leave	Leaving	Coworker away		Phone	
Don't forget to close the door	Out of room				Auto

Nothing: do nothing, Auto: switch automatically, Phone: ask via smartphone, Display: ask via display

The participants consistently gave high ratings (between 4 and 5) for the criteria *Transparency*, *Controllability*, *Comfort of Use*, and *User Trust* when evaluating the reactions the system had chosen for the adjustment of the light (see Table 3).

However, some participants criticized that trust was impaired because of missing feedback when the light was switched off after they left the office. Despite these high ratings, in situations in which the system sent a message to the participants' phone, other system reactions were preferred by most of the participants (see Table 3—Preferred SR). These findings were in line with several statements of the participants. For example, several users mentioned that using a phone is inconvenient in many situations - either because it is not within reach or because they have to interrupt their work to read the message on the phone. Accordingly, some users preferred autonomous system actions instead of repeated messages on their phones because this would make the system less obtrusive.

In contrast, the automatically generated reactions for the display matched the preferences of most of the participants in all situations (see Table 3—Preferred SR). In most of the situations the participants clearly favored autonomous reactions for the display (as in the online condition), but at the expense of *Controllability* and *User Trust* (see Table 3). While the average trust ratings were still quite high (between 3.5 and 4)

Table 3 Results of the live study: user ratings (mean (M); standard deviation (SD)) for the trust dimensions and the perceived trust related to the selected system reaction (SR), and the preferred SR by most of the users

Situation	Selected SR	Transparency	Controllability	Comfort of use	User trust	Preferred SR (% of users)
Live study—device: Light						
Arriving, dark, C away	Phone	M = 5.00	M = 4.46	M = 4.17	M = 4.25	Auto
		SD = .00	SD = .87	SD = .99	SD = .72	(75 %)
Bright, C away	Display	M = 4.92	M = 4.58	M = 4.63	M = 4.42	Display
		SD = .28	SD = 1.04	SD = .70	SD = .64	(79 %)
Dark, C present	Phone	M = 4.67	M = 4.25	M = 4.13	M = 4.13	Display
		SD = .75	SD = 1.20	SD = .97	SD = .78	(58 %)
Leaving, dark, C away	Phone	M = 4.92	M = 4.13	M = 4.29	M = 3.92	Auto
		SD = .28	SD = 1.27	SD = 1.10	SD = .86	(67 %)
Live study—device: display						
Working at PC	Auto	M = 4.83	M = 2.83	M = 4.58	M = 3.75	Auto
		SD = .47	SD = 1.31	SD = .57	SD = 1.09	(54 %)
Away from desk	Nothing	M = 3.79	M = 2.79	M = 4.13	M = 3.75	Nothing
		SD = 1.35	SD = 1.55	SD = 1.20	SD = 1.13	(88 %)
Idle at PC	Auto	M = 4.58	M = 2.50	M = 4.00	M = 3.63	Auto
		SD = .91	SD = 1.29	SD = 1.08	SD = 0.95	(71 %)
Out of room	Auto	M = 5.00	M = 3.46	M = 4.46	M = 3.88	Auto
		SD = .00	SD = 1.44	SD = 1.15	SD = .88	(79 %)

C: coworker, Nothing: do nothing, Auto: switch automatically, Phone: ask via smartphone, Display: ask via display

the average ratings for *Controllability* were only mediocre (between 2.5 and 3.5). The ratings for the trustworthiness of autonomous reactions were affected, among other things, by a missing authentication mechanism after switching on the display and by a lack of feedback when leaving the room. The low ratings for "Controllability" could be explained by requests for functionality to set or disable the automatic control of the display.

The concluding questions also showed promising results. Most participants were satisfied (83 %; M:3.96; SD: .68) and agreed that the system assisted them to improve their energy consumption (96 %; M: 4.71; SD: .54), that it behaved adequately (88 %; M: 4.38; SD: .70), and that it was transparent (100 %; M: 4.96; SD: .20). The lower, but still acceptable results for unobtrusiveness (58 %; M: 3.71; SD: 1.10) could be mainly explained by the fact that the users had to operate the mobile phone. Further results showed that most of the participants did not feel distracted (75 %; M: 2.00; SD: 1.00), restricted (88 %; M: 1.83; SD: 1.07), or observed (63 %, M: 2.33; SD: 1.18).

6 Further investigations

The results of the live study gave promising results for the control of the display. Most participants preferred the actions that were selected by the system. However, the system's decisions for controlling the light frequently did not match the participants' choices. In the live study, scores for the trust dimensions and user trust were only obtained for the selected system action. Thus, users might have been biased because they were only presented with the system's choice. Furthermore, user scores for the alternatives could help shed light on the question of why participants preferred a particular system reaction.

For these reasons, we decided to complement the live study by a live survey that was conducted under similar conditions as the live study. In the following, the design and the results of the live survey will be described.

6.1 Live survey: experimental setting and execution

The aim of the live survey was to acquire user ratings for all combinations of situations and possible system reactions under natural conditions. The experimental setting for the live survey was adopted from the live study (see Sect. 5.1) in order to obtain comparable results. However, instead of running the real system and confronting the participants only with the system actions selected by the UTM, all possible system actions were shown to the participants in each situation.

At the beginning of the survey the participants had again to provide demographic data. After a short introduction to the setting and the scenario, all possible system actions were presented to the participants. They then had to enter the unoccupied and dark office. After entering the room, they were immediately confronted with the first set of possible system actions designated for the light in the office and they had to rate statements related to perceived *Transparency*, *Controllability*, *Comfort of Use* and *Trust* for each of these actions (see Q1–Q4 in Sect. 4.1).

After rating all system actions, the users were also asked to indicate which system action they preferred. However, in this survey the participants were not to choose one of these actions. Instead, they had to rate the statement "I would prefer the system action..." for each action on a 5-point Likert scale in order to enable an easier comparison of preferences. Then the procedure continued with the next situation and the corresponding questionnaire. The entire sequence of tasks and situations is summarized in Table 2.

In total, eight men and two women aged between 23 and 35 (Mean: 28) took part in the live survey.

6.2 Performance of the User Trust Model

We further investigated how well the UTM performed when compared to the data gathered in the live study and live survey, both in terms of preference and user trust. Similar to the evaluation of the live study described above, we compared the decision generated by the UTM for each situation with the one the users found the most preferable or the most trustworthy.

For the display, most participants (73 % for the live study and 90 % for the live survey) preferred the actions chosen by the UTM. The selected system actions also received the highest trust ratings from 80 % of the participants. This means that the trust and preference ratings showed similar tendencies. As mentioned above, in the live study we only asked the participants whether the action taken by the UTM was trustworthy, but did not ask about the trustworthiness of the alternatives. Thus, we only have trust ratings for the actions taken, which were quite high with an average trust of 3.75 (display) and 4.18 (light).

The actions generated by the UTM for the light were much less in line with the preferences of the users. Only 34 % of the people participating in the live study would choose the same actions as the system. In the live survey, only 18 % of the participants expressed the highest preference for the selected system actions. Nevertheless, 80 % of the participants gave the selected system actions the highest trust ratings. The results show that the UTM was able to create trustworthy decisions, but also that trust was not the only factor that determines which action a user preferred.

6.3 A further analysis of the trust dimensions

Since the comparison of the UTM's assessments and the collected data revealed differences between the users' trust and their preferences, the collected data were analyzed in more detail. The aim was to identify major factors that affected the users' preferences.

At first, the trust ratings provided by the participants in the online survey as well as in the live survey were investigated. In both surveys, the system actions achieved a similar level of trust for most of the situations. In the live survey, the participants considered actions performed by the system as more trustworthy (4.18 on average on the 5-point Likert scale) than doing nothing (3.28 on average) albeit not significant. They also expressed a high preference for proactive system reactions (4.46 on average)

compared to doing nothing (2.47 on average), asking for confirmation on the mobile phone (2.59 on average) or asking for confirmation on the display (3.18 on average). We furthermore found that trust tended to get higher ratings in the live survey (3.91 on average) than in the online survey (3.16 on average), a trend we had already observed in Wißner et al. (2014). Apparently the users have more trust in a system that they can actually experience than in a system that is just verbally described to them. With the exception of doing nothing, most system actions got quite high mean values for trust (above 4.0). In our earlier work (Wißner et al. 2014), trust ratings for system actions showed greater variations. One reason is that our earlier work also investigated situations that introduced serious risks to privacy, such as viewing personal information or photos on a public display, depending on the system action chosen. Consequently, users gave a number of system actions rather low trust ratings.

The first trust dimension that was investigated in more detail was perceived *Transparency*. Both in the online and the live survey, most possible system actions achieved high mean ratings for *Transparency* (around 4.0 on the 5-point Likert scale). Obviously, most system actions seemed plausible to the participants. Remarkably, the system action "Do nothing" only achieved moderate ratings in 75 % (online survey) and 88 % (live survey) of the situations respectively. In a number of cases, it was even rated significantly less transparent than an automated system reaction (see Table 4). At first, this result may appear surprising. However, we believe that a proactive behavior is noticed by a user more easily than inactivity, which might be one reason for the less positive ratings the users gave to "Do nothing". Interestingly, the result is in line with the experience reported by a company described by Picard and Klein (2001). Customers who bought a product that did not work properly and got excellent support were more likely to keep buying their brand than customers who did not figure out any problems at all with the product.

Table 4 Investigated trust dimension: transparency—significant results of a repeated-measures ANOVA and a Bonferroni-Post-Hoc-Test

Device	Situation	Significances (A < B)	Ratings			
			M(A)	SD(A)	M(B)	SD(B)
Online survey—trust dimension: transparency						
Light	Arriving, dark, coworker away	Nothing < Auto*	2.75	1.71	4.31	1.16
	Leaving, dark, coworker away	Nothing < Auto*	2.44	1.46	4.00	1.54
Display	Working at PC	Nothing < Auto**	2.82	1.53	4.18	1.37
		Phone < Auto*	3.23	1.62		
	Out of room	Nothing < Auto**	3.18	1.59	4.36	1.15
		Phone < Auto*	3.64	1.52		

Nothing: do nothing, Auto: switch automatically, Phone: ask via smartphone, Display: ask via display
* $p < .05$; ** $p < .01$; *** $p < .001$

Next, the participants' perceived *Controllability* in the surveys was analyzed. In the live study, autonomous decisions by the UTM resulted in lower scores for this trust dimension. These findings could be confirmed for most situations of the online and live survey. In all situations, automatic system actions only achieved mean ratings lower than 3.0. They were perceived as less controllable than performing no system action at all or system actions that involved asking the users for confirmation before performing an action. In many cases, the differences were statistically significant. Consistently high ratings for *Controllability* were achieved by the system action "Ask the user for confirmation via her or his smartphone". However, a comparison of Tables 5 and 6 reveals that the low *Controllability* scores for automatic system actions did not negatively affect the participants' preferences for those actions.

Although automatic system actions resulted in a decreased perceived control over the system, they were rated as the most preferred actions in most of the situations in which a system reaction was expected. In contrast, "Ask the user for confirmation via his or her smartphone" only achieved mean ratings lower than 3.0 or less in most of the situations (see Table 6). Although this system action was mostly perceived as very controllable, in many situations, participants preferred to have the system respond automatically or to give confirmations via their display.

These findings raised the question of which trust dimension would affect the participants' preferences for a particular system action the most. The results of the live study and the statements of the participants indicated that perceived *Comfort of Use* could be a decisive factor. To confirm these findings, the ratings for perceived *Comfort of Use* obtained in the surveys were analyzed in detail. In both surveys, automatic system actions applied to the light and the display scored best. In the live survey, the ratings were in general above 4.0 on the 5-point Likert scale and in some cases even higher than 4.5. In comparison, "Do Nothing" and especially the action "Ask the user for confirmation via her or his smartphone" scored significantly worse (see Table 7). In many situations, participants gave the lowest score for the latter action. In most situations, the mean ratings were lower than 3.0, in individual cases even lower than 2.0.

Overall, an analysis of the participants' preferences revealed that they preferred more comfortable system actions over more controllable actions in both surveys. This finding was confirmed by their statements during the live study as well as during the live survey. The participants liked the idea of being asked in some situations. However, they considered a message on their smartphone only reasonable when they entered or left the office. When they were seated at their desk, they preferred autonomous decisions by the system or messages that were shown on their displays. Using the phone was considered to be inconvenient, e.g. because it was often not within reach and it was also considered obtrusive because the participants would have to interrupt their work every time the phone received a message.

In our earlier work (Wißner et al. 2014), the mismatch between online and live data was less pronounced than in the current paper. We assume one reason to be the fact that the scenarios investigated in the earlier paper included situations in which the user's privacy was at stake. Consequently, the users' choice between different system actions was mainly based on privacy concerns both in the online and the live setting. According to the live survey presented in this paper, users seem to weight

Table 5 Investigated trust dimension: controllability—significant results of a repeated-measures ANOVA and a Bonferroni-Post-Hoc-test

Device	Situation	Significances (A < B)	Ratings			
			M(A)	SD(A)	M(B)	SD(B)
Online survey—trust dimension: controllability						
Light	Arriving, dark, coworker away	Auto < Display**	2.44	1.37	3.50	1.37
		Auto < Phone**			4.00	1.37
	Bright, coworker away	Auto < Nothing*	1.88	1.17	3.44	1.58
		Auto < Phone**			3.88	1.36
		Auto < Display**			4.00	1.22
	dark, coworker present	Auto < Display**	1.88	1.05	3.88	1.41
		Auto < Phone**			4.00	1.27
	Leaving, dark, coworker away	Nothing < Phone*	2.50	1.54	4.19	1.13
		Auto < Phone*	2.63	1.49		
Display	Away from desk	Auto < Phone**	2.27	1.17	3.73	1.39
		Auto < Nothing***			3.59	1.40
	Idle at	Auto < Nothing**	2.23	1.04	3.45	1.53
	PC	Auto < Phone**			3.64	1.30
Live survey—trust dimension: controllability						
Light	Arriving, dark, coworker away	Auto < Display**	2.40	1.36	4.50	.67
		Auto < Phone**			4.70	.46
	Bright, coworker away	Auto < Phone**	2.30	1.19	4.60	.49
		Auto < Display**			4.80	.40
	Dark, coworker present	Auto < Phone**	2.40	1.20	4.80	.40
		Auto < Display**			4.90	.30
	Leaving, dark, coworker away	Auto < Phone**	2.30	1.19	4.80	.40
Display	Working at PC	Auto < Phone**	2.80	1.17	4.50	.50
	Away from desk	Auto < Phone***	2.30	1.19	4.70	.46
	Idle at PC	Nothing < Phone*	3.30	1.55	4.70	.46
		Auto < Phone***	2.30	.78		
	Out of room	Auto < Phone**	2.60	1.28	4.80	.40

Nothing: do nothing, Auto: switch automatically, Phone: ask via smartphone, Display: ask via display
* $p < .05$; ** $p < .01$; *** $p < .001$

trust dimensions differently depending on whether they are confronted with a real system or just a verbal description of it. For example, participants expressed a stronger preference for proactive system behaviors in the live survey than was suggested by

Table 6 Investigated: preference—significant results of a repeated-measures ANOVA and a Bonferroni-Post-Hoc-test

Device	Situation	Significances (A < B)	Ratings			
			M(A)	SD(A)	M(B)	SD(B)
Online survey—preference						
Light	Arriving, dark, coworker away	Phone < Auto**	2.31	1.06	4.00	1.22
		Display < Auto**	2.00	1.06		
Display	Working at PC	Phone < Auto***	1.64	.93	3.86	1.14
		Phone < Nothing***			3.14	1.46
	Away from desk	Phone < Auto*	1.59	.94	2.68	1.26
	Idle at PC	Phone < Nothing***	2.05	1.36	4.09	.95
		Auto < Nothing*	2.91	1.16		
	Out of room	Phone < Auto**	2.55	1.50	4.09	1.20
Live survey—preference						
Light	Arriving, dark, coworker away	Nothing < Auto*	2.40	1.28	4.60	.92
		Display < Auto*	2.60	.92		
		Phone < Auto**	2.30	1.10		
	Bright, coworker away	Nothing < Auto**	1.70	.78	4.10	.94
		Nothing < Display***			4.50	.67
		Phone < Display**	2.70	1.19		
	Dark, coworker present	Nothing < Display*	2.30	1.35	4.50	.67
		Phone < Display**	2.30	1.10		
Display	Working at PC	Nothing < Auto**	2.80	1.08	4.80	.40
		Phone < Auto***	1.80	.87		
	Away from desk	Phone < Auto*	2.00	1.26	3.90	.94
		Phone < Nothing*			4.10	.94
	Idle at PC	Nothing < Auto**	2.70	1.27	4.80	.40
		Phone < Auto**	2.50	1.20		
	Out of room	Phone < Auto**	3.30	1.00	4.90	.30
		Nothing < Auto***	1.70	1.00		

Nothing: do nothing, Auto: switch automatically, Phone: ask via smartphone, Display: ask via display
* $p < .05$; ** $p < .01$; *** $p < .001$

the online data from which the Bayesian Network was trained. While our earlier work seemed to indicate that it is possible to train Bayesian networks from online data and employ them in live scenarios, the current paper shows that the reliance on online data is only possible to a limited extent.

Table 7 Investigated trust dimension: comfort of use—significant results of a repeated-measures ANOVA and a Bonferroni-Post-Hoc-Test

Device	Situation	Significances (A < B)	Ratings			
			M(A)	SD(A)	M(B)	SD(B)
Online survey—trust dimension: comfort of use						
Light	Arriving, dark, coworker away	Phone < Auto*	2.50	1.37	3.88	1.22
		Display < Auto**	2.25	1.09		
		Nothing < Auto**	2.06	.90		
Display	Working at PC	Phone < Nothing*	1.73	.86	2.64	1.30
		Phone < Auto***			3.86	1.36
		Nothing < Auto**	2.64	1.30	3.86	1.36
	Away from desk	Phone < Auto*	1.86	.97	2.64	1.33
	Idle at PC	Phone < Auto*	2.05	1.11	2.91	1.16
		Phone < Nothing*			3.36	1.26
	Out of room	Phone < Auto*	2.64	1.40	3.91	1.41
Live survey—trust dimension: comfort of use						
Light	Arriving, dark, coworker away	Phone < Auto*	2.50	1.12	4.50	.67
	Bright, coworker away	Nothing < Auto**	2.40	1.20	4.60	.66
		Phone < Auto*	2.70	.90		
		Nothing < Display*	2.40	1.20	4.20	.40
		Phone < Display**	2.70	.90		
	Dark, cworker present	Nothing < Auto**	2.50	1.12	4.60	.66
		Phone < Auto*	2.60	1.20		
		Nothing < Display**	2.50	1.12	4.30	.46
		Phone < Display**	2.60	1.20		
	Leaving, dark, coworker away	Phone < Auto*	3.50	1.12	4.80	.40
		Nothing < Auto**	2.50	1.12		
Display	Working at PC	Nothing < Auto**	2.50	1.20	4.70	.46
		Phone < Auto***	2.00	.63		
	Away from desk	Phone < Auto*	1.90	1.22	3.90	1.14
	Idle at PC	Phone < Auto*	2.50	1.20	4.30	1.00

Nothing: do nothing, Auto: switch automatically, Phone: ask via smartphone, Display: ask via display
* $p < .05$; ** $p < .01$; *** $p < .001$

7 Conclusion

We presented an approach for trust-based decision-making for smart and proactive environments based on Bayesian networks, the User Trust Model. It assesses the users' trust experienced while interacting with a system and performs appropriate (i.e.

trustworthy) system reactions to properly adapt to new situations. We described the construction of the UTM, its integration in an office setting, and its initialization with empirical data. The results of a live user study revealed that the system generally succeeded in performing appropriate actions in the investigated situations. Earlier work on computational models of trust either focuses on trust between users or trust between software or hardware components while our work attempts to explicitly model affective user trust towards a pervasive environment. We would like to note that the prediction quality for user trust has only been indirectly measured in this paper. Our approach selects the system action that is supposed to create the highest amount of user trust. Our experiments showed that the users give the selected system actions indeed high trust ratings. However, we did not explicitly address the question of how accurately the UTM predicts the user's current level of trust.

Even though the UTM approach has been developed and evaluated for an energy management system, the basic mechanism is applicable to other kinds of ubiquitous systems as well. Following a component-based software development approach for user modeling as suggested by Dim et al. (2015), the basic structure of the BN representing the dependencies between trust and its dimensions could be reused by other developers for their applications. Only the nodes representing the context and possible system actions would have to be adapted to the corresponding applications. In this way, the development of applications that make decisions based on user trust could be significantly facilitated.

In addition to the live study, we performed a live survey, which gave us not only ratings for the performed system actions, but also for alternatives. The live survey revealed that users had a high amount of trust in the chosen system action, but generally preferred a higher degree of proactive system behavior in order to increase the comfort of use. Obviously, the users' weighting of the trust dimensions in the online survey differed from that in the live survey. Since our results show a discrepancy between the user's trust and preferences, future work should investigate whether giving more weight to *Comfort of Use* when selecting a system action could rectify this. Another option to explore would be to collect a sufficient amount of live data by recruiting a larger number of users as a basis for the training of the Bayesian Network which was initialized with data obtained from an online survey.

To confront the users with realistic scenarios, we embedded the single tasks into a coherent story representing a working day in the life of the user. As a consequence, the sequence of tasks was not randomized, but determined by the story line. While this approach helped us create a plausible scenario for users, it might have led to an overfit of the Bayesian Network. Future data collection efforts should concentrate on longer scenarios with a larger number of tasks that can be presented in randomized order.

The Bayesian Network used in the paper was initiated by data from 38 individuals. Consequently, the Bayesian Network rather reflected the attitude of a variety of users as opposed to an individual user. In the future, we will investigate how to improve the accuracy of the UTM by incorporating knowledge about user-specific attitudes. Depending on their trust disposition, users might favor different system reactions. For example, users that tend to distrust technical systems might give more importance to a high level of control than to a high level of comfort. A promising approach might

be to distinguish between different categories of users based on multiple dimensions (Knijnenburg et al. 2013). In comparison to the live study, the questionnaire for the live survey contained more questions about participants' opinions and habits concerning sustainability and their trust towards other people and technical systems in general, so this data could be included in the model in the future. In order to achieve an even higher degree of personalization, the UTM could also be trained with data from individual users. In the ideal case, the UTM should not require extensive training before it can be used, but dynamically adapt to people's preferences by learning from their behavior during the interaction with the smart environment.

In our current work, we focused on trust as an attitude as opposed to reliance as a behavior (Lee and See 2004). To evaluate the system's ability to select actions that maximize user trust, we asked participants' to rate the trustworthiness of individual actions. In addition to investigating subjective user impressions, more objective measurements should be addressed, such as the user's reliance on a system as a result of user trust. That is we should investigate to what extent users are willing to relinquish control to the system.

Another important aspect is the decision making for more than one user. For example, some participants wondered whether they were the only person in control of the light. Therefore, the UTM should be extended to be able to consider the trust of all affected users.

So far, we focused on the question of how particular system actions influence the user's level of trust. However, the users' level of trust does not only depend on the momentary system behavior, but also on their experience with the system in the past. For example, users might forgive a minor bug if it occurs for the first time. However, if the system fails repeatedly, the users' trust will be affected significantly. In order to consider how user trust felt at a particular point in time influences user trust experienced at a later point in time, we intend to extend the Bayesian Network to a Dynamic Bayesian Network that allow us to model the dependencies between the current states of variables and earlier states of variables.

Acknowledgments This research is co-funded by OC-Trust (FOR 1085), funded by the German Research Foundation (DFG) and by IT4SE, funded by the German Federal Ministry of Education and Research (Grant number NZL 10/803 IT4SE) under the APRA initiative. The core of our implementation is based on the SMILE reasoning engine and the network shown in this paper was created using the GeNIe modeling environment. Both SMILE and GeNIe are developed and contributed to the community by the Decision Systems Laboratory, University of Pittsburgh and available at http://www.genie.sis.pitt.edu/.

References

Adali, S., Escriva, R., Goldberg, M.K., Hayvanovych, M., Magdon-Ismail, M., Szymanski, B.K., Wallace, W.A., Williams, G.T.: Measuring behavioral trust in social networks. In: Yang, C.C., Zeng, D., Wang, K., Sanfilippo, A., Tsang, H.H., Day, M.Y., Glässer, U., Brantingham, P.L., Chen, H. (eds.) Intelligence and Security Informatics (ISI), 2010 IEEE International Conference on, pp. 150–152. IEEE, Vancouver, BC, Canada (2010). doi:10.1109/ISI.2010.5484757

Antifakos, S., Kern, N., Schiele, B., Schwaninger, A.: Towards improving trust in context-aware systems by displaying system confidence. In: Proceedings of the 7th International Conference on Human Computer Interaction with Mobile Devices and Services, MobileHCI '05, Salzburg, Austria, pp. 9–14. ACM (2005). doi:10.1145/1085777.1085780

Arimura, S., Fujita, M., Kobayashi, S., Kani, J., Nishigaki, M., Shiba, A.: i/k-contact: A context-aware user authentication using physical social trust. In: Privacy, Security and Trust (PST), 2014 Twelfth Annual International Conference on, pp. 407–413 (2014). doi:10.1109/PST.2014.6890968

Bee, K., Hammer, S., Pratsch, C., André, E.: The automatic trust management of self-adaptive multi-display environments. In: Khalil, I., Mantoro, T. (eds.) Trustworthy Ubiquitous Computing, Atlantis Ambient and Pervasive Intelligence, pp. 3–20. Atlantis Press, Amsterdam (2012). doi:10.2991/978-94-91216-71-8-1

Bhuiyan, T., Xu, Y., Jøsang, A.: A review of trust in online social networks to explore new research agenda. In: Arabnia, H.R., Clincy, V.A., Lu, J., Marsh, A., Solo, A.M.G. (eds.) Proceedings of the 2010 International Conference on Internet Computing (ICOMP 2010), pp. 123–128. CSREA Press, Las Vegas, NV, USA (2010). http://www.eprints.qut.edu.au/41447/

Burnett, C., Norman, T.J., Sycara, K.: Trust decision-making in multi-agent systems. In: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence—Volume Volume One, IJCAI'11, Barcelona, Catalonia, Spain, pp. 115–120. AAAI Press, Menlo Park (2011). doi:10.5591/978-1-57735-516-8/IJCAI11-031

Castelfranchi, C., Falcone, R.: Trust Theory: A Socio-Cognitive and Computational Model, 1st edn. Wiley Publishing, Chichester (2010)

Cheverst, K., Byun, H., Fitton, D., Sas, C., Kray, C., Villar, N.: Exploring issues of user model transparency and proactive behaviour in an office environment control system. User Model. User Adapt. Interact. **15**(3–4), 235–273 (2005). doi:10.1007/s11257-005-1269-8

Denko, M.K., Sun, T., Woungang, I.: Trust management in ubiquitous computing: a bayesian approach. Comput. Commun. **34**(3), 398–406 (2011). doi:10.1016/j.comcom.2010.01.023

Dhamija, R., Tygar, J.D., Hearst, M.: Why phishing works. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06, Montréal, Québec, Canada, pp. 581–590. ACM (2006). doi:10.1145/1124772.1124861

Dim, E., Kuflik, T., Reinhartz-Berger, I.: When user modeling intersects software engineering: the info-bead user modeling approach. User Model. User Adap. Interact. (2015). doi:10.1007/s11257-015-9159-1

DiSalvo, C., Sengers, P., Brynjarsdóttir, H.: Mapping the landscape of sustainable HCI. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '10, Atlanta, Georgia, USA, pp. 1975–1984. ACM (2010). doi:10.1145/1753326.1753625

Florencio, D., Herley, C., Van Oorschot, P.C.: Password portfolios and the finite-effort user: sustainably managing large numbers of accounts. In: Proceedings of the 23rd USENIX Conference on Security Symposium, SEC'14, San Diego, CA, pp. 575–590. USENIX Association, Berkeley, CA, USA (2014). http://www.dl.acm.org/citation.cfm?id=2671225.2671262

Gamberini, L., Spagnolli, A., Corradi, N., Jacucci, G., Tusa, G., Mikkola, T., Zamboni, L., Hoggan, E.: Tailoring feedback to users actions in a persuasive game for household electricity conservation. In: Bang, M., Ragnemalm, E. (eds.) Persuasive Technology. Design for Health and Safety. Lecture Notes in Computer Science, pp. 100–111. Springer, Berlin Heidelberg (2012). doi:10.1007/978-3-642-31037-9-9

Glass, A., McGuinness, D.L., Wolverton, M.: Toward establishing trust in adaptive agents. In: Proceedings of the 13th International Conference on Intelligent User Interfaces, IUI '08, Gran Canaria, Spain, pp. 227–236. ACM (2008). doi:10.1145/1378773.1378804

Hazas, M., Friday, A., Scott, J.: Look back before leaping forward: four decades of domestic energy inquiry. IEEE Pervasive Comput. **10**(1), 13–19 (2011). doi:10.1109/MPRV.2010.89

Ivanov, I., Vajda, P., Korshunov, P., Ebrahimi, T.: Comparative study of trust modeling for automatic landmark tagging. Trans. Info. For. Sec. **8**(6), 911–923 (2013). doi:10.1109/TIFS.2013.2242889

Jøsang, A., Hayward, R., Pope, S.: Trust network analysis with subjective logic. In: Proceedings of the 29th Australasian Computer Science Conference—Volume 48, ACSC '06, Hobart, Australia, pp. 85–94. Australian Computer Society Inc, Darlinghurst, Australia (2006). http://www.dl.acm.org/citation.cfm?id=1151699.1151710

Knijnenburg, B.P., Kobsa, A., Jin, H.: Dimensionality of information disclosure behavior. Int. J. Hum.-Comput. Stud. **71**(12), 1144–1162 (2013). doi:10.1016/j.ijhcs.2013.06.003

Kumaraguru, P., Cranor, L.F.: Privacy indexes: a survey of westin's studies. Technical Report CMU-ISRI-5-138, Technical Report, Institute for Software Research Int. (ISRI), Carnegie Mellon University (2005)

Kurdyukova, E., André, E., Leichtenstern, K.: Trust management of ubiquitous multi-display environments. In: Krüger, A., Kuflik, T. (eds.) Ubiquitous Display Environments, Cognitive Technologies, pp. 177–193. Springer, Berlin Heidelberg (2012). doi:10.1007/978-3-642-27663-7-11

Kurdyukova, E., Hammer, S., André, E.: Personalization of content on public displays driven by the recognition of group context. In: Patern, F., de Ruyter, B., Markopoulos, P., Santoro, C., van Loenen, E., Luyten, K. (eds.) Ambient Intelligence. Lecture Notes in Computer Science, pp. 272–287. Springer, Berlin Heidelberg (2012). doi:10.1007/978-3-642-34898-3-18

Lee, J.D., See, K.A.: Trust in automation: designing for appropriate reliance. Hum. Factors **46**(1), 50–80 (2004). doi:10.1518/hfes.46.1.50-30392. http://www.hfs.sagepub.com/content/46/1/50.abstract

Leichtenstern, K., André, E., Kurdyukova, E.: 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, Paris, France, pp. 409–414. ACM (2010). doi:10.1145/1971519.1971589

Lim, B.Y., Dey, A.K.: Toolkit to support intelligibility in context-aware applications. In: Proceedings of the 12th ACM International Conference on Ubiquitous Computing, UbiComp '10, Copenhagen, Denmark, pp. 13–22. ACM (2010). doi:10.1145/1864349.1864353

Lumsden, J.: Triggering trust: to what extent does the question influence the answer when evaluating the perceived importance of trust triggers? In: Proceedings of the 23rd British HCI Group Annual Conference on People and Computers: Celebrating People and Technology, BCS-HCI '09, Cambridge, UK, pp. 214–223. British Computer Society, Swinton (2009). http://www.dl.acm.org/citation.cfm?id=1671011.1671037

Marsh, S.: Trust in distributed artificial intelligence. In: Castelfranchi, C., Werner, E. (eds.) Artificial Social Systems. Lecture Notes in Computer Science, pp. 94–112. Springer, Berlin Heidelberg (1994). doi:10.1007/3-540-58266-5-6

Müller, J., Exeler, J., Buzeck, M., Krüger, A.: Reflectivesigns: digital signs that adapt to audience attention. In: Tokuda, H., Beigl, M., Friday, A., Brush, A., Tobe, Y. (eds.) Pervasive Computing. Lecture Notes in Computer Science, pp. 17–24. Springer, Berlin Heidelberg (2009). doi:10.1007/978-3-642-01516-8-3

O'Donovan, J., Smyth, B.: Trust in recommender systems. In: Proceedings of the 10th International Conference on Intelligent User Interfaces, IUI '05, San Diego, California, USA, pp. 167–174. ACM (2005). doi:10.1145/1040830.1040870

Picard, R.W., Klein, J.: Computers that recognise and respond to user emotion: theoretical and practical implications. Interact. Comput. **14**(2), 141–169 (2002). doi:10.1016/S0953-5438(01)00055-8. http://www.iwc.oxfordjournals.org/content/14/2/141.abstract

Russell, S., Norvig, P.: Artificial Intelligence: A Modern Approach, 3rd edn. Prentice Hall Press, Upper Saddle River (2009)

Sankaranarayanan, V., Chandrasekaran, M., Upadhyaya, S.: Towards modeling trust based decisions: a game theoretic approach. In: Biskup, J., Lpez, J. (eds.) Computer Security ESORICS 2007. Lecture Notes in Computer Science, pp. 485–500. Springer, Berlin Heidelberg (2007). doi:10.1007/978-3-540-74835-9-32

Sherchan, W., Nepal, S., Paris, C.: A survey of trust in social networks. ACM Comput. Surv. **45**(4), 47:1–47:33 (2013). doi:10.1145/2501654.2501661

Simon, J., Jahn, M., Al-Akkad, A.: Saving energy at work: The design of a pervasive game for office spaces. In: Proceedings of the 11th International Conference on Mobile and Ubiquitous Multimedia, MUM '12, Ulm, Germany, pp. 9:1–9:4. ACM (2012). doi:10.1145/2406367.2406379

Steghöfer, J.P., Kiefhaber, R., Leichtenstern, K., Bernard, Y., Klejnowski, L., Reif, W., Ungerer, T., André, E., Hähner, J., Müller-Schloer, C.: Trustworthy organic computing systems: challenges and perspectives. In: Xie, B., Branke, J., Sadjadi, S., Zhang, D., Zhou, X. (eds.) Autonomic and Trusted Computing. Lecture Notes in Computer Science, pp. 62–76. Springer, Berlin Heidelberg (2010). doi:10.1007/978-3-642-16576-4-5

Vogiatzis, G., MacGillivray, I., Chli, M.: A probabilistic model for trust and reputation. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, vol. 1, AAMAS '10, Toronto, Canada, pp. 225–232. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2010). http://www.dl.acm.org/citation.cfm?id=1838206.1838238

Wang, Y., Vassileva, J.: Bayesian network trust model in peer-to-peer networks. In: Moro, G., Sartori, C., Singh, M.P. (eds.) Agents and Peer-to-Peer Computing. Lecture Notes in Computer Science, pp. 23–34. Springer, Berlin Heidelberg (2005). doi:10.1007/978-3-540-25840-7-3

Wißner, M., Hammer, S., Kurdyukova, E., André, E.: Trust-based decision-making for the adaptation of public displays in changing social contexts. J. Trust Manag. **1**(1), 6 (2014). doi:10.1186/2196-064X-1-6. http://www.journaloftrustmanagement.com/content/1/1/6

Yan, Z., Holtmanns, S.: Trust Modeling and Management: From Social Trust to Digital Trust. IGI Global, chap. 13, pp. 290–323 (2008)

Yan, Z., Liu, C., Niemi, V., Yu, G.: Effects of displaying trust information on mobile application usage. In: Xie, B., Branke, J., Sadjadi, S., Zhang, D., Zhou, X. (eds.) Autonomic and Trusted Computing. Lecture Notes in Computer Science, pp. 107–121. Springer, Berlin Heidelberg (2010). doi:10.1007/978-3-642-16576-4-8

Yu, B., Singh, M.P.: An evidential model of distributed reputation management. In: Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 1, AAMAS '02, Bologna, Italy, pp. 294–301. ACM (2002). doi:10.1145/544741.544809

Stephan Hammer is a Ph.D. candidate in Computer Science at Augsburg University. He received his B.Sc. and M.Sc. in Computer Science and Multimedia from the Augsburg University in 2006 and 2008, respectively. His primary interests lie in the areas of human computer interaction and context-aware recommender systems. The research for his thesis work focuses on the generation and presentation of recommendations to support behavior change related to people's environmental-awareness, as well as elderly persons' well-being.

Michael Wißner is a Ph.D. candidate in Computer Science at Augsburg University. He received his diploma in Applied Computer Science from Augsburg University in 2006. His primary research interests are Pedagogical Agents and User Modeling. The research for his thesis work focuses on Natural Language Dialog Generation for Virtual Characters in Interactive Learning Environments.

Elisabeth André is a Full Professor of Computer Science at Augsburg University, and Chair of the Research Unit Human-Centered Multimedia. She received her Diploma and Doctoral Degrees in Computer Science from Saarland University. Elisabeth André has a long track record in multimodal human-machine interaction, embodied conversational agents, affective computing and social signal processing. She is on the editorial board of international journals such as the Journal of Autonomous Agents and Multi-Agent Systems, the IEEE Transactions on Affective Computing, the ACM Transactions on Intelligent Interactive Systems, and the AI Communications. She became a Fellow of the Alcatel-Lucent Foundation for Communications Research in 2007, and of the European Coordinating Committee for Artificial Intelligence 2014. In 2010, she was elected to the German Academy of Sciences Leopoldina, the Academy of Europe and AcademiaNet.