

Experiences from Real-World Deployment of Context-Aware Technologies in a Hospital Environment

Jakob E. Bardram, Thomas R. Hansen, Martin Mogensen, and Mads Soegaard

Centre for Pervasive Healthcare
Department of Computer Science, University of Aarhus
Aabogade 34, 8200 Aarhus N, Denmark
{bardram, thomasr, spider, madss}@daimi.au.dk

Abstract. Context-aware computing is a central concept in ubiquitous computing and many suggestions for context-aware technologies and applications have been proposed. There is, however, little evidence on how these concepts and technologies play out in a real-world setting. In this paper we describe and discuss our experiences from an ongoing deployment of a suite of context-aware technologies and applications in a hospital environment, including a context-awareness infrastructure, a location tracking system, and two context-aware applications running on interactive wall displays and mobile phones. Based on an analysis of the use of these systems, we observe that many of the ideas behind context-aware computing are valid, and that the context-aware applications are useful for clinicians in their work. By reflecting on the nature of the designed context-aware technologies, we present a model which states that the triggering of context-awareness actions depend upon the accuracy of the sensed context information, the degree to which you know which action to perform in a given situation, and the consequence of performing the action.

1 Introduction

The idea of *context-aware computing* was one of the early concepts introduced in some of the pioneering work on ubiquitous computing research [23,16] and has been subject to extensive research since. ‘Context’ refers to the physical and social situation in which computational devices are embedded. The goal of context-aware computing is to acquire and utilize information about this context, to display relevant information, or to provide services that are appropriate to this particular setting. For example, a cell phone will always vibrate and never ring in a concert hall, if it somehow has knowledge about its current location and the activity going on (i.e. the concert) [19]. The core premise of context-aware computing is to create more usable technology which adapt to the user’s context. Over the last decade, numerous context-aware applications have been reported for use in e.g. hospitals [3,20], shops [1], museums [14,21], tourism [2,9], zoological fieldwork [22], on universities [8,15], conferences [11,10], and offices [24]. Furthermore, numerous infrastructures and application frameworks have been proposed to enable the easy development of context-aware applications, each having different technological approaches and application focuses.

To our knowledge, however, only few of the proposed context-aware technologies and applications have been deployed in a real use-setting outside the computer science

lab for a longer period of time. Thus, despite the massive amount of research done within context-aware computing, little evidence actually exists of its utility. On a practical/empirical level, we therefore have little knowledge about how users perceive context-aware technologies, how they use them, what kind of automatic adaptation they find beneficial, and to what degree they find sensing technology privacy invading – just to name a few questions. Moreover, on a conceptual/theoretical level, we have a limited number of ‘real-world’ arguments for judging the scientific validity inherent to the notion of ‘context-aware computing’. Recently, arguments have been aired that the current notion of ‘context’ and hence ‘context-aware computing’ builds on a positivist philosophical stance, where ‘context’ is stable, delineable, and sense able information separated from human activity. The argument is that the notion of ‘context’ – as referring to the ‘usage context’ for a specific person using some technology – cannot be separated from the human activity. ‘Context’ then becomes firmly tied to ‘meaning’ – i.e. that context cannot be seen (and much less sensed) as an objective entity in the world, but only exists in connection with subjective meaning in an activity [13]. But what does this argument implies when designing for real-world use of context-aware computing (or its like)?

In this paper we report from an ongoing deployment of a suite of context-awareness technologies and applications in an operating ward at a medium-size European hospital. The deployed systems consist of an indoor location tracking system; a context-awareness infrastructure for acquiring, managing, and distributing context information; a context-aware operation scheduling application; and a context-aware mobile phone application helping users to keep a social awareness on colleagues and the activities inside operating rooms. At the time of writing, the system has been deployed and used extensively for three month; is used by approximate 30 users each day; is handling thousands of context events for each user or place inside the hospital; and runs with very limited support from the developers. The system has been completely adopted by the users, and after one month of usage the context-aware scheduling application replaced the former paper-based operating schedules. Section 3 describe the deployed systems and section 4 describes the deployment at the operating ward.

The paper aims at two main contributions. First, we present ‘real-world’ evidence on the use of context-aware technologies in a non-trivial, hectic, life-critical work environment. Our findings indicate that the notion of context-aware computing is indeed useful for users and can be achieved on a large scale. Our study indicates that displaying (and distributing) context information was the most useful part of the gathered context information. Furthermore, the study revealed that privacy concerns regarding the display of personal context information (e.g. location, status, and activity) were fully acceptable and that the benefit outweighed privacy concerns. This contribution is based not only on the technology we built and deployed, but it is also based on several discussions with the users about a number of context-aware ‘features’ which we designed, but did not implement. It is interesting to note that most of the context information in the applications was used passively; i.e. primarily for displaying context information or displaying relevant clinical information in a specific usage context. Despite numerous attempts to build more active context-aware features into the systems, the clinicians rejected them all.

This leads to our second, more conceptual, contribution which is a thorough discussion of how to use context information. When is it beneficial to display context information? When would we like to find contextual ‘relevant’ material? When could we automatically execute an action or a service? We propose a simple general model for answering these questions, which takes into consideration issues like the accuracy of the sensed context information, the degree to which the system knows what the users wants in a specific ‘context’, and the consequences of displaying information or execution an action.

2 Related Work

Suggestions for context-aware infrastructures and applications are numerous in Pervasive and Ubiquitous Computing research, including suggestions for the use of context-aware technologies in a hospital setting. For example, the ‘Intelligent Hospital’ application [18] uses location awareness to enable video communication to follow clinicians around inside a hospital, and others have been adding context-awareness to the use of mobile technologies in hospitals [17,20]. Far less research has, however, been done in the study of deployed context-aware infrastructures and applications – to date, most ‘ubiquitous’ computing experiments have been restricted to specific laboratories and buildings, or to specific research groups [6]. Consequently, little evidence exists on real-world usage of context-aware technologies. There are, however, a few notable exceptions.

Pascoe et al. [22] describe the design and deployment of a context-aware PDA application made to support ecologists in field studies in Africa. In the prototype, context awareness enables the mobile device to provide assistance based on knowledge of its environment in terms of time and position (using GPS). The paper gives few details on the lessons learned from the deployment, but in relation to this paper it is interesting to note that context information is only used passively, since the whole purpose of the application was to minimize user attention.

Active Campus, on the other hand is an example of a much larger deployment of location-based services [15]. The Active Campus Explorer application supports context-aware activities, such as instant messaging and location-aware maps annotated with dynamic hyperlinked information. Active Campus Explorer has been deployed for a year (2002–3) at the UC San Diego campus, encompassing hundreds of users. The study showed that students are willing to share location with buddies and even non-buddies for location-aware social computing; they are more likely to message each other when they are in close proximity to one another; and that just one percent of the users changed their default privacy settings to hide location from buddies.

The GUIDE system provides city visitors with a hand-held context-aware tourist guide. The system has been successfully deployed in a tourist setting and used over a period of four weeks by 60 people [9]. The major conclusions from the study were that visitors found the location-aware navigation and information retrieval mechanisms both useful and reassuring; and that the majority of visitors appreciated that the system were aware of their location to within a certain area.

In general, it is interesting to note that all of these deployed context-aware applications only use context information in a ‘passive’ way, i.e. for tagging [22], for displaying

context information, like location [15], or for displaying relevant information in a specific context [9]. There are no examples reported on the use of ‘active’ use of context information for automatically triggering actions or services. A second interesting observation is, that all of these systems use *location* as the only context information – hence, there seem to be no evidence for the usefulness of other kinds of context information in real world deployment.

3 Context-Awareness Technology in Hospitals

The context-aware technologies discussed in this paper address the inherent problem of coordination and collaboration between clinicians working distributed within a hospital. The systems use context awareness to facilitate information about the location, status, and current activity of clinicians and different core locations, including operating rooms. The research is rooted in a wide range of field studies (see e.g. [5]) and a user-centered design process involving a wide range of clinicians in the design, development, and evaluation of the systems. The deployed system consists of the following four main components, which are illustrated in figure 1:

- A *context-awareness infrastructure* responsible for the acquisition, distribution, and modeling of context-awareness information.
- A *location tracking system* responsible for tracking clinicians’ and patients’ location.
- The *AwareMedia* application which is a context-aware scheduling and awareness system running on large interactive displays at the operating ward and inside the operating rooms.
- The *AwarePhone* application running on mobile phones which facilitates social awareness amongst clinicians.

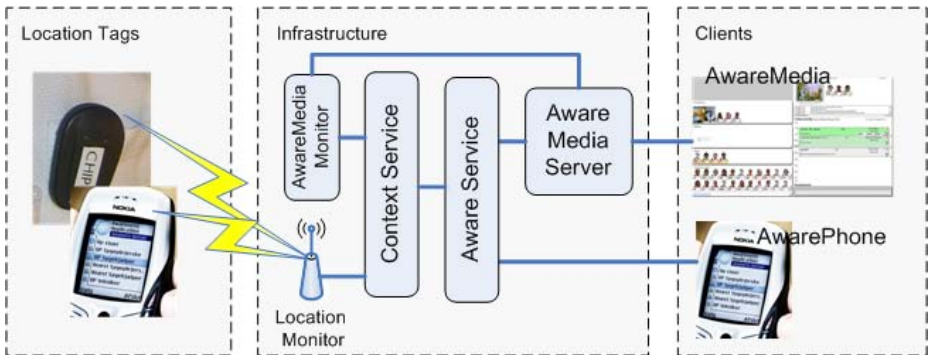


Fig. 1. The four main components deployed at the hospital: (i) location tracking (left); (ii) context-awareness infrastructure (middle); (iii) AwareMedia (top right); and (iv) AwarePhone (bottom right)

3.1 Context-Awareness Infrastructure

The context-awareness infrastructure is based on the Java Context-Awareness Framework (JCAF) [4]. In the deployment we are using JCAF for modeling entities and context information at the surgical ward, and for creating different types of context monitors. The context model includes entities like ‘Personnel’, ‘Patient’, and ‘Operating Room’. The main context information modeled is ‘location’, ‘status’, and ‘scheduled activity’; where status is self-reported status like ‘operating’ or ‘ward round’, and activity maps to the operation which the person is scheduled for.

The infrastructure is implemented following a three-tier architecture. First, a set of monitors are responsible for gathering location information from the location tracking system and activity information from the operating schedules. Second, the monitor layer forwards the information to a Context Service which is responsible for managing, storing and grouping context information. Third, an awareness layer (the Aware Service) is responsible for distributing relevant context information to clients subscribing to different kinds of context information. The awareness layer forwards information to the AwareMedia and AwarePhone applications used by the clinicians.

3.2 Location Tracking

In a large hospital, much time and effort is spent in order to locate staff, patients, and equipment. Hence, a hospital is a good candidate for indoor location-tracking technologies. In the hospital deployment we use a Bluetooth-based location-tracking system. Using Bluetooth has the benefits of being cheap and widely available in e.g. telephones, PDAs, and tablet PCs. A slightly modified USB stick deployed in stationary PCs is used to track mobile devices. The modification implies that its range is reduced to below 10 meters. The drawback of using Bluetooth is its low precision. But in most cases we were able to reveal which room the person was in, especially in operating rooms since such rooms have rather solid walls with X-ray shielding. Mobile phones are used to track surgeons, and clinicians. Persons who are not using a mobile phone are tracked via a small Bluetooth tag. Tags can be assigned and re-assigned to users via AwareMedia. Figure 2 (left) shows a technician with a Bluetooth tag.

3.3 AwareMedia

AwareMedia is a context-aware application running on larger interactive displays. It is designed to support coordination, the scheduling of surgeries, and social awareness amongst clinicians involved in surgeries. In AwareMedia, all clinicians are represented by an icon revealing a picture of the user, the user’s location, self-reported status, and scheduled activity; all information coming from the context-awareness model in the infrastructure. Furthermore, AwareMedia has separate regions on the display representing different ‘important’ places in the hospital, like the different operating rooms, and the coordination central. The user’s icon is located inside this region based on the location reported by the location system.

Similar to clinicians, each operating room is treated as a context entity and hence possesses the context cues of ‘location’, ‘status’, and ‘activity’. The location of an




Fig. 2. The deployed systems: AwareMedia deployed in the coordination center (left), an operating technician wearing a location tag (right, top), and the AwarePhone (right, bottom) showing a list of users, their location, status, and scheduled operation

operating room is clearly never used, but the status of the operation inside the rooms and the activity, i.e. the ongoing operation are important context information which are displayed at the AwareMedia and AwarePhone. Examples of status for an operations are ‘patient arrived’, ‘patient anaesthetized’, ‘operation started’, and ‘patient left the operating room’. A message chat allows an unobtrusive and easy communication channel amongst clinicians involved in an operation. For example, the nurses at the patient ward can ask the surgeon to come to the ward to attend a patient, or the operating ward can broadcast a message about the last surgery being cancelled. A schedule on the interface shows planned and acute surgeries and changes made through an easy to use drag and drop interface are distributed to all clients. Figure 2 (left) shows AwareMedia deployed in the coordinating central.

3.4 AwarePhone

The AwarePhone system [5] runs on Symbian smart phones and provides an overview of the people at work, and the activity in each operating room. An interactive phone book displays each user’s location, current activity (i.e. scheduled operation), and self-reported status. This implies that the user can see where specific colleagues are located, what they are scheduled to be doing, and what their status is – see figure 2. Similarly, the location, status, and activity (i.e. operation) is shown for each operating room. The phone book can be used to place a call to one of the listed persons. Or if the person appears to be occupied, a prioritized message can be sent. Because users and operating rooms are treated equally, messages can be received from, and sent to operating rooms, and calls can be made to the operating rooms directly, if needed.

4 Deployment

The AwareMedia and AwarePhone applications presented above have been designed, refined, and evaluated through a user-centered design process involving a wide range of clinicians. The next step in this line of research was to deploy these context-aware technologies in a real-world setting, more specifically at a surgical ward and some of the associated wards. The first parts of system were taken to the hospital in November 2005 and by early January 2006 the whole system went into daily use. At the time of writing it is still in use.

4.1 The Deployment Setting

The operating ward employs around 130 clinicians in total and there are 30-50 clinicians present during a normal workday. Some of the associated departments are the patient wards where the patients are hospitalized before and after a larger surgery, and the recovery ward where a patient is closely monitored while waking up from anesthesia.

The operating ward has in total ten operating rooms. AwareMedia was deployed in three of these rooms and in the coordination central. The coordination central is the coordination hub for all operations unfolding during the day and from this central, the head nurse closely follows the unfolding of the daily schedule, makes adjustments to the operation schedules, and ensures a smooth flow of work by coordinating all the involved clinicians. Inside the operating room, AwareMedia is deployed on an interactive display and in the coordination central, AwareMedia is deployed on two large interactive displays.

Ten Bluetooth tags are available at the entrance to the operating ward, and fifteen Nokia phones have been handed out to some of key surgeons and anesthetists, the head nurse, the coordinating nurse, and to some of the operation technicians. Figure 3 shows where the different departments are situated at the hospital and an overview of the surgery ward. The dots indicate deployed AwareMedia clients and the shaded areas shows the coverage of location tracking system.

4.2 Research Methods

The deployment of the system is motivated by the desire to investigate long term use and adaptation of advanced pervasive technology in a real world setting. Does ubiquitous computing technology provide a valuable tool or is the problem of managing complex technologies far greater than the benefits? How do you get an organization to incorporate advanced technology into its daily work practices? Are you able to improve the way work is done with context awareness technologies? To answer some of these questions we have used a wide variety of research approaches. Before, during, and after the deployment of the system we spent days performing field studies with the focus of analyzing work, identifying problems, and registering changes in work practices. Questionnaires have been used both before and after the deployment and a number of interviews have been performed with key personnel after the deployment, in order to identify the impact of the system on their daily work practices. The interviews were conducted in the clinicians' native language, but quotes from some of the interviews have been translated into English for this publication. In addition, the infrastructure




Fig. 3. The figure shows the parts of the hospital in which we installed the system and provide a detailed view of the operating ward

logs all interaction from AwareMedia clients and AwarePhones. The log files records 50.000+ entries a month; we have done 10 man-days of participant observations; and additional 12 semi-structured interviews with different types of clinicians working at the operating wards.

5 Context-Awareness in Use

The study has generated a large amount of data and we have hence chosen to discuss how context-awareness cues were used at the department in general, how context-awareness was used from the AwarePhones, and how privacy issues were conceived in the hospital since these issues illustrates central parts of the system usage. The discussion will be based on both logged data about the use of the system as well as data from the interviews.

5.1 Using Context-Awareness Cues

The top part of Figure 4 shows the number of discrete context-awareness events flowing through the system as a function of time and the bottom part of Figure 4 presents the number of unique persons, which has triggered events in the system as a function of time. Separate graphs for status, location, and scheduled activity events are shown in both figures.

There are several things that can be seen from the graphs:

- Usage is stable throughout the period. Overall we see the same usage pattern for all weeks in figure 4 except for the week following Feb. 11th, which is the general ‘winter vacation’ week. Since the number of ‘activity’ events correlates to the number of operations scheduled using AwareMedia, the patterns in Figure 4 imply first of all that the system has been used for all operations since it was deployed




Fig. 4. Top: Number of unique context events pr. day. Bottom: Number of unique persons triggering a context event.

in January. It can be noted that interest is not fading as the technology moves from being ‘new and exciting’ to being adopted into the work at the department.

- On an average day we see 2,500 location events, 80 activity events, and 70 status events. Even though there is 30 times more location events, location events only relates to 12-13 unique ‘entities’ (i.e. persons), whereas about 25 unique ‘entities’ (i.e. persons and ORs) are using the status events. The fact that the ‘status’ context cue is the most frequently used, in terms of involved number of persons, is due to the system automatically setting the status in some situations, e.g. a person is put to be ‘at work’ when first located, but also manual reported status contributed to the overall number. Another reason for the relatively low number of unique persons located is the fact that we did not have enough Bluetooth chips and mobile devices to equip everyone with a chip; hence only selected persons are being tracked by the system in the pilot study.

These findings were consistent with our qualitative observations and interviews. Most significantly, during the second week of deployment, the head nurse told us that

they had abandoned the use of their previous paper-based scheduling and awareness system and was entirely using AwareMedia, AwarePhone, and their underlying infrastructure. As she put it: “[now] when we don’t have the paper schedule anymore then we are feeling even more dependent on the system.” In our qualitative analysis, we were trying to understand how clinicians were using the context-awareness technology in general, and more specifically how the three context information about location, status, and activity was used by different types of clinicians in their work.

During our deployment we have found substantial evidence that providing location information to clinicians is indeed quite useful in a hospital setting. As stated by the head nurse:

It is something I can clearly use as the one responsible for the coordination. Especially the knowledge about where the replacements are located, because if they are carrying a chip or a phone it is really easy to get an overview about which rooms they are located in.

As a manager responsible for the daily coordination of work at the operation ward, this nurse has started to use location information quite extensively. As she says, she has a much better overview of where people are especially the replacements and operating technicians who roam around and help out in different operating rooms when needed. Hence, if one room is calling for help, the head nurse would need to locate a replacement immediately, a task which clearly is significantly simplified by using location tracking technology. Looking at it from the operating technicians’ point of view, they did not mind being tracked:

It is a way of using each other instead of saying “Where are you? What are you doing?” you can look at the display, [...] and then I don’t have to let go of something to answer a phone and tell them where I am. You can really quickly get this information from the display and get an overview.

Here the technician explains that because people calling him can see where he is, he would get less disruptive calls. There are actually two steps in his line of reasoning: first people can see where he is and then they can deduce the typical kind of activity he might be engaged in. For example, if the technician is inside an operating room, he is probably busy helping them. In this sense, location cues are not merely informative with respect to locating people, but are also an important source of information with respect to their current activity.

Revealing status information as a type of context information was also used in the coordination of work. When asking an operating nurse about the use of the status bar in the operating theatre and the relation to people at the recovery ward she says:

Interviewer: What about the collaboration or coordination with P5 [the patient ward] and the recovering department? Do you contact them from the ward or what happens? Operating nurse: No, it has more been the other way around. They have used us and asked questions such as how far we are with the surgery if we haven’t moved the arrow [updated status], so they are using it.

Hence, the operating staff inside the operating theatre – more or less to their own surprise – discovered that people outside the operating theatre and outside the department were actually relying on an updated status information in their own coordination of work. The realization of the benefit for others and the reduced number of misplaced and disruptive phone calls were some of the reasons the operating staff gave to explain why they kept the status bar up to date.

There were, however, also some challenging aspects to the use of context information. First of all, because not all persons had a personal location token, location tracking was not good enough. There was simply a too large overhead associated with picking up a tag and then going to the computer and associate this tag with a person. Furthermore, to our surprise, the doctors did not pick up on the use of the system. And finally, several complained that you had to set the status bar manually – they would prefer that it revealed some sensed real-world status, like when the patient enters the operating room, or when the cleaning is done.

5.2 AwarePhone Usage

Figure 5 shows the use of the fifteen AwarePhones and illustrates the number of distinct connection events of a phone and the number of unique phones connected to the Awareness Service during a day. Again we see a fairly stable usage pattern, except that the phones were not deployed during the first week, they were tested during the second week, and there is lower activity during the aforementioned winter vacation.




Fig. 5. The number of distinct connections from AwarePhones per day

In the original plan for deployment, the AwarePhone was targeted for the surgeons who typically roam around inside the hospital, and hence can be situated in many different places during a normal workday: the conference room, the bed ward, the operating ward, the operating room, in the recovery areas, and in their offices. Hence, mobile technology is particular well-suited for their type of work patterns and the AwarePhone

could help them keep an awareness on the unfolding of work at the operating ward while roaming around. During the deployment study, however, we observed that the phones were gradually being taken over by the head nurse and by the operation technicians, as illustrated in the quotes above. In the deployment plan we had envisioned that technicians would carry passive location tags, which they also did in the beginning of the deployment period (figure 2). But it turned out that it was important for technicians to be able to monitor the status of the work as well, in order for them to align their own work accordingly. As one of the technicians explained:

When I have the phone, I am updated on the program all the time. I don't need to go to the wall [referring to the screens in the coordinating central]. I can see how far the operating rooms are - if they are on the way out [if the patient is being moved out] or halfway. That, I think is really good for me when I am walking around out here - and the other way around ... when I am using the phone they also know where I am.

Here the technician explained how he was using the phone to maintain an awareness of the progress of the whole operating schedule while attending other duties inside the operating department. Moreover, what is interesting to note is that he is well aware that others can locate him, when he is using the AwarePhone – a fact that he (implicitly) see as a benefit, as illustrated in the quote from the technicians in the previous section.

5.3 Privacy in Context-Aware Technologies

In the design of the context-aware technologies, it was an open question whether clinicians would accept that information about their location, status, and scheduled activity was sensed and broadcasted around for others to see. Most users, and especially the management of the ward, felt a need to openly discuss these aspects of the system. As the head nurse explained:

It was important from the beginning that the people did not feel that the system was a surveillance system observing what they are doing. But it is a surveillance system because you track the surgeries and where people are located. It is just important to tell the people, that we watch the work processes and not the persons as such. We did a lot of considerations about the issue in the beginning, but I have not at any time heard about people feeling kept under surveillance.

In the design of the system we expected privacy issues to emerge, it was discussed during the design of the systems, and was also an issue in the beginning of the deployment, as expressed by the head nurse above. However, in our observations and interviews we have found no evidence that privacy is a major concern. The main argument is that the context information displayed (location, activity, and status) are no different from the information that all clinicians already try to obtain by other means. For example, by calling the coordinating nurse asking for status and activity reports, or by monitoring the work and whereabouts of others. Hence, in a professional setting like the hospital, displaying your current location, status, and activity did not cross any 'privacy

border’ in general. There are, however, some important deployment details to consider here. First, the clinicians did point out co-called ‘non-tracking’ zones, i.e. areas where they did not want to be tracked. Examples were the coffee room used during breaks, the staff canteen, and the toilets. Furthermore, the deployed JCAF context service did only display the current situation and it did not store context information. The ‘history-less’ nature of the context acquisition was understood by the clinicians and was hence seen as a guarantee that upper management would not use the data for analyzing the workflow inside the operating ward, and would not be able to see where people have been during the course of a day.

The patients overall were not involved with the use of the system since they were anaesthetized during most of their time at the operating ward. Some of the benefits for the patients were better information about when they were going to be operated and especially the relatives also appreciated to be able to get detailed information from the nurses about the progress of the surgery of their loved ones.

6 Discussion: A Model for Triggering Context-Aware Actions

At this point, we conclude that the users at the operating ward found the context-aware applications useful, and that context information was used quite extensively in the daily coordination of operations. It is, however, interesting to note that similar to other deployed context-aware applications (see section 2), we have designed applications which primarily employ a ‘passive’ approach to context-awareness. Despite the fact that we quite deliberately presented the clinicians with several suggestions for more ‘active’ context-aware features, they more or less rejected them all. Examples include automatically starting and stopping of video broadcasts from the operating rooms, based on the people present in it and their privacy preferences; and automatically changing the status of the operation based on people (including the patient) present in the room. The only example of a somewhat active use of context information is to change status for a clinician to ‘at work’, if they are located inside the hospital (i.e. their mobile phone can be seen by a location monitor). This feature can, however, lead to errors; it does happen that clinicians – especially surgeons – arrive at the hospital outside working shifts for doing administrative tasks or research. Hence, even though a surgeon is at the hospital, he may not actually be ‘at work’ with respect to the operating ward. The consequence of such a ‘false positive’ where a surgeon’s status is set incorrectly was, however, deemed to be insignificant compared to the general benefit of the feature.

Observations like the ones above have lead us to consider a more general model, which can help designers of context-aware applications to establish which kind of actions to take based on context information. Basically, our model says that you should consider the following three factors in the design of context-aware computing systems:

- Accuracy of context information
- Confidence in the action to take
- The consequence of the action

The general design rule is that *the triggering of a context-awareness action depends upon the accuracy of the sensed context information, the degree to which you know what action to take in a certain situation, and the consequence of performing this action.*

6.1 Accuracy of Context Information

The basic premise in any context-aware system is that it makes decisions based on context information input – either raw sensed data, inferred data, or input from users. The core problem is, however, that this data will always be imperfect. A core challenge facing the development of realistic and deployable context-aware services, is therefore the ability to handle imperfect, or ambiguous, context – and yet this issue has often been ignored or glossed over [12]. It is hence important that context-awareness infrastructures which acquire, represent, and distribute raw sensed data and inferred data is also capable of handling imperfect data and can reveal some accuracy measurement for data. Data is always only valid within some level of accuracy. Given that data from a monitor is imperfect, it becomes essential to take the accuracy into account when deciding which actions to take. For example, the location system at the operating ward was accurate within rooms and could hence be used to display a user’s icon in the room the person is located in on the AwareMedia interface. It could, however, not be used to automatically log in a person on a computer. Similarly, self-reported status was distributed to all clients since this was considered accurate (i.e. we trust people to set their status correctly).

6.2 Confidence in Action

Consider the situation where you are absolutely certain that the inferred ‘context’ is accurate, i.e. you have accurate data upon which you can base your decisions. Given this situation, is it then possible to decide with certainty which actions the user want to execute? We argue that this is only the case for very simple actions such as switching on a light or opening a door. But even these examples are not as simple as they seem. There may be exceptions to the normal rules. When I enter a room, it may be because I want to change the light bulb and it could be catastrophic if the context aware system were to switch on the electricity while I was doing that.

One may then ask the question, if it is not possible to sense that I am about to change the light bulb and adapt the actions accordingly? And yes, it may be, but it would always be possible to come up with another example which would ruin any decision making scheme. The reason, we argue, is that human context, defined in the subjective realm of the individual, is inaccessible to sensors and excessively complex. We may create a model representing context but we will never be able to get the *appropriate* information to consistently deduce the correct actions to execute.

The trouble is that we, in order to be aware of a context, must establish what that context *means*, i.e. what is the meaning of a person entering the room? This ‘meaning’ arises in the subjective realm of a person engaged in a meaningful activity, which makes it inaccessible to our sensors. It is not just a body passing through a doorframe. It is perhaps a person entering the room with the aim of doing maintenance, i.e. changing the light bulbs. Or a kid entering the room to play hide-and-seek, in which case an automatic light would ruin his hide-out. The person entering the room is in other words not just a human body, equipped with a visual system requiring light. In that case, our mapping would be a simple mapping from meaning (person entering) to action (turn

on light). Instead, it is a person engaged in some meaningful activity, and that *meaning* exists in the subjective realm.

However, even though it is impossible to infer a person's intent from sensor data, it is still possible to make a qualified guess with some probability to what a person may need in the situation. For example, if a patient enters an operating room, there is a high probability that he or she is going to have a surgery, whereas a doctor can enter an operating room for wide variety of reasons. The point being that this difference in knowledge about what a person wants to achieve needs to be considered together with the quality of sensor data when building context awareness systems.

6.3 Consequence of Action

Independent of the accuracy of context data, and the confidence in a specific action, the actual triggering of this action still depends on the consequence of doing it. Some actions are less intrusive than others. For example, displaying the position of clinicians in AwareMedia, is far less intrusive than automatically ending anesthesia for a patient when the surgeon leaves the room, because in most cases the surgeon only leaves the room when the operation is finished. The same distinguishing goes for opening a door to the local supermarket, against opening the door to the local bank vault. In the former case it is desirable to open the door too often, whereas in the latter case the opposite is desirable.

It is therefore necessary to consider the consequences or intrusiveness of an action, when determining whether to execute a given action. The intrusiveness is tightly related to the *type of action*, but also to *the context* in which the action is executed. As we see it, there are four main types of actions:

Tagging actions focus on attaching meta-information to some data based on context information for later use. The user can then at a later point use this information to e.g. categorize captured images.

Displaying actions focus on organizing and displaying context awareness information to the user to enable decisions based on context aware information (context-mediated awareness)

Helping actions focus on preparing some actions that the computer thinks the user is going to need, but it is still the user who initialize the actions.

Executing actions focus on executing specific actions automatically without involving or asking the user in the process.

Overall the intrusiveness increase downwards in this list, but there are of course exceptions to this, which means that we also have to consider the context the action is executed in. Displaying a patient record in an operation room will seldom cause any harm, but displaying the same record in a public place can have serious privacy and legal consequences.

6.4 Using the Model

The model proposes that a successful context awareness system should execute actions based on knowledge about the accuracy of sensor data, about the confidence in the

action to take, and about the consequence of performing the action. However, the model also propose that if sensor data is inaccurate, or if there is a high degree of uncertainty about which action to execute, this is not an argument for abandoning context awareness. Instead, the context aware system should then facilitate other types of action with less severe consequences, like tagging or displaying.

In retrospect, we would argue that in our design and deployment of context-aware technologies in the hospital, we had somewhat inaccurate sensor data, it was in some cases hard to decide what action to take based on the input, and the consequences of executing improper actions in a hospital might have been far-reaching. Therefore, we chose to let the system organize and display different views on the context aware information to the users, which provided less intrusiveness than executing actions. Moreover, we discovered that *juxtaposing* inaccurate and even conflicting sensor information to the user in many cases provided useful knowledge to the users. For example, a person may be scheduled for a surgery, his status is set to 'in surgery', but the location tracking system reports that he is in the outpatient clinic. This ambiguity in context would be hard to automatically deduce anything from, but the coordinating nurse would typically suspect that the operation has ended early or there might be a break in the surgery.

The main point being that there is a huge possibility for building context-awareness systems in the middle ground between automating actions and doing nothing at all, but a successful system needs to balance the action used with the sensor quality, the confidence in the action, and the action's consequences in a given context.

7 Conclusion

Based on an ongoing deployment of context-aware technologies in an operating ward we have found evidence for the usefulness of such technologies in this setting. The logs from the context-awareness infrastructure revealed that the use of these technologies has been constant over a four months period. In our qualitative research, clinicians report on the benefit of having access to view context information on colleagues and operating rooms. The display of location, status, and current activity significantly helped them improve their coordination of work. Moreover, the study has shown that privacy concerns about revealing such context information was not a major issue.

The types of context-aware computing designed and deployed in the hospital ended up mainly displaying context information instead of performing automated action. Based on an analysis of the deployed technology, we propose a design rule which states that the design of a context-awareness actions depend upon the accuracy of the sensed context information, the degree to which you know which action to perform in a certain situation, and the consequence of performing this action. Given this model, we argue that actions should only be executed automatically in very limited and well-defined situations. Our experiences show that the mapping between context information and appropriate actions is highly challenging, both from a practical and empirical point of view, as well as from a more theoretical stance.

Acknowledgments

This research has been supported by the Competence Centre ISIS Katrinebjerg. We thank the surgical staff at Horsens Sygehus who enthusiastically participated in this project. Christian Jonikeit has implemented a substantial part of AwareMedia.

References

1. M. C. Abhaya Asthana and P. Krzyzanowski. An indoor wireless system for personalized shopping assistance. In *Proceedings of IEEE Workshop on Mobile Computing Systems and Applications*, pages 69–74. IEEE Computer Society Press, 1994.
2. G. D. Abowd, C. G. Atkeson, J. Hong, S. Long, R. Kooper, and M. Pinkerton. Cyberguide: a mobile context-aware tour guide. *Wirel. Netw.*, 3(5):421–433, 1997.
3. J. E. Bardram. Applications of ContextAware Computing in Hospital Work – Examples and Design Principles. In *Proceedings of the 2004 ACM Symposium on Applied Computing*, pages 1574–1579. ACM Press, 2004.
4. J. E. Bardram. The Java Context Awareness Framework (JCAF) – A Service Infrastructure and Programming Framework for Context-Aware Applications. In H. Gellersen, R. Want, and A. Schmidt, editors, *Proceedings of the 3rd International Conference on Pervasive Computing (Pervasive 2005)*, volume 3468 of *Lecture Notes in Computer Science*, pages 98–115. Munich, Germany, May 2005. Springer Verlag.
5. J. E. Bardram and T. R. Hansen. The AWARE architecture: supporting context-mediated social awareness in mobile cooperation. In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 192–201. ACM Press, 2004.
6. L. Barkhuus and P. Dourish. Everyday encounters with context-aware computing in a campus environment. In *UbiComp 2004: Ubiquitous Computing: 6th International Conference*, pages 232–249. Berlin, 2004. Springer Verlag.
7. G. Borriello and L. Holmquist, editors. *Proceedings of UbiComp 2002: Ubiquitous Computing : 4th International Conference*, Göteborg, Sweden, Sept. 2002. Springer Verlag.
8. J. Burrell, G. K. Gay, K. Kubo, and N. Farina. Context-aware computing: A test case. In Borriello and Holmquist [7], pages 1–15.
9. K. Cheverst, N. Davies, K. Mitchell, and A. Friday. Experiences of developing and deploying a context-aware tourist guide: the guide project. In *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 20–31. New York, NY, USA, 2000. ACM Press.
10. A. Dey, G. D. Abowd, and D. Salber. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction*, 16:97–166, 2001.
11. A. K. Dey, M. Futakawa, D. Salber, and G. D. Abowd. The conference assistant: Combining context-awareness with wearable computing. In *Proceedings of the 3rd International Symposium on Wearable Computers (ISWC '99)*, pages 21–28. IEEE Computer Society Press, 1999.
12. A. K. Dey and J. Mankoff. Designing mediation for context-aware applications. *ACM Trans. Comput.-Hum. Interact.*, 12(1):53–80, 2005.
13. P. Dourish. What we talk about when we talk about context. *Personal and Ubiquitous Computing*, 8(1):19–30, 2004.
14. M. Fleck, M. Frid, T. Kindberg, E. O'Brien-Strain, R. Rajani, and M. Spasojevic. Rememberer: A tool for capturing museum visits. In Borriello and Holmquist [7], pages 48–55.

15. W. G. Griswold, P. Shanahan, S. W. Brown, R. Boyer, M. Ratto, R. B. Shapiro, and T. M. Truong. ActiveCampus: experiments in community-oriented ubiquitous computing. *IEEE Computer*, 37(10):73–81, 2004.
16. A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster. The anatomy of a context-aware application. *Wireless Networks*, 8(2/3):187–197, 2002.
17. J. Kjeldskov and M. Skov. Supporting work activities in healthcare by mobile electronic patient records. In *Proceedings of the 6th Asia-Pacific Conference on Human-Computer Interaction, APCHI 2004*, 2004.
18. S. Mitchell, M. D. Spiteri, J. Bates, and G. Coulouris. Context-aware multimedia computing in the intelligent hospital. In *EW 9: Proceedings of the 9th workshop on ACM SIGOPS European workshop*, pages 13–18, New York, NY, USA, 2000. ACM Press.
19. T. Moran and P. Dourish. Introduction to this special issue on context-aware computing. *Human-Computer Interaction*, 16:87–95, 2001.
20. M. Munoz, M. Rodriguez, J. Favela, A. Martinez-Garcia, and V. Gonzalez. Context-aware mobile communication in hospitals. *IEEE Computer*, 36(9):38–46, 2003.
21. R. Oppermann and M. Specht. A context-sensitive nomadic exhibition guide. In *Proceedings of Second International Symposium on Handheld and Ubiquitous Computing, HUC 2000*, pages 127–142, Berlin, 2000. Springer Verlag.
22. J. Pascoe, N. Ryan, and D. Morse. Using while moving: Hci issues in fieldwork environments. *ACM Trans. Comput.-Hum. Interact.*, 7(3):417–437, 2000.
23. R. Want, B. N. Schilit, N. I. Adams, R. Gold, K. Petersen, D. Goldberg, J. R. Ellis, and M. Weiser. An overview of the parctab ubiquitous computing environment. *IEEE Personal Communications*, 2(6):28–43, 1995.
24. H. Yan and T. Selker. Context-aware office assistant. In *IUI '00: Proceedings of the 5th international conference on Intelligent user interfaces*, pages 276–279, New York, NY, USA, 2000. ACM Press.