NTNU

Norwegian University of
Science and Technology

Sensitive Information on Display

Using flexible de-identification for protecting patient privacy in (semi-) public hospital environments

Erlend Andreas Gjære

Master of Science in Informatics

Submission date:	July 2011
Supervisor:	Pieter Jelle Toussaint, IDI
Co-supervisor:	Maria B. Line, SINTEF IKT

Norwegian University of Science and Technology
Department of Computer and Information Science

Abstract

Background In later years, the health care work in hospitals has become increasingly fragmented, in a sense where different people and professions are required for the treatment of every single patient. As a consequence, personnel should be assisted to greater awareness of what is happening, so that they can better plan where to put in their efforts. Making information about ongoing activities more accessible to its users is hence important, but this will in turn require increased distribution of sensitive data inside the hospital. The concept of flexible de-identification has been proposed as a solution for the privacy issues raised by this, but then again new issues emerge when it comes to how useful the de-identified data are to its authorized end users, in practice.

Methods A series of six rapid field tests was executed along with a literature review on de-identification. The purpose was to explore some ideas to how de-identification could be implemented for information screens located in public and semi-public hospital environments, such as hallways, where personnel are likely to see them. The appropriateness of several techniques for de-identification was hence evaluated for being used in real-time visualizations, in contrast to previous known applications of the concept. This input was in turn used to design a high-fidelity prototype for use in a series of four experiments in a usability laboratory. The experiments involved role-play sessions, where nurses from a university hospital used the prototype in a simulation of realistic ward work. In a focused interview directly afterwards, they each assessed the usefulness of having a system available in such locations, considering that the information was de-identified. Moreover, the nurses evaluated six alternative approaches to de-identification of the sensitive information, and ranked them with respect to which, if any, would be best suited for use in their regular work environment.

Results The experiments indicate that users appreciate being notified via large screens when new information is available, but disagree on what is the preferred level of de-identification. Some would emphasize the legislative requirements and privacy issues raised, while others would put their own utility needs first. As a response to this, an interactive prototype was designed to demonstrate how users can be given interactive control over how identifiable the displayed information is. This idea of giving users flexible control over what is seen on a screen, depending on how they assess the context for access, is grounded in a framework for evaluation that considers the quality requirements of identification utility, legislation and usability.

Conclusion Useful applications of non-interactive de-identification to screens in public environments, are effectively disqualified by the legislative requirements regulating how personal health information can be disclosed. The de-identification can however be useful for enabling an intermediate security level, which can be accessed as long as there is a authorized user present. Appropriate techniques for achieving such de-identification, are found to be suppression of variables, coding, masking and generalization. With this overall approach, users may gradually authorize themselves until the required utility is reached, and hence be able to access useful information in public places. The information depth available must also be accordingly limited, so that the increased risk of abuse is mitigated. The result is possibly a security mechanism that is both legal to implement, it serves the utility needs of personnel, and it is more usable in practice than existing time-demanding login routines. Finally, these ideas have been included in the design of an interactive prototype, which still remains to see tested in practice.

Keywords:

De-identification, patient privacy, usability, data utility, digital whiteboards, hospital collaboration, usable authentication

Contents

Preface

This thesis presents my final year of studies before being awarded the Master of Science degree in Informatics, at the Norwegian University of Science and Technology (NTNU), Department of Computer and Information Science (IDI).

The work has been conducted at NTNU, the Norwegian Research Centre for Electronic Patient Records (NSEP), and Trondheim University Hospital, throughout the period of August 2010 to the end of June 2011.

Acknowledgements

I would like to thank my main supervisor Pieter J. Toussaint and co-supervisor Maria B. Line for the time and guidance you both have given me throughout the last year. I have learned a lot from you, and also had a good time doing so. The support I've had through the collaboration with Ph.D. candidate (and M.D.) Børge Lillebo, has also had an invaluable impact on my work. Thank you for your positive attitude towards my involvement in the project, and for taking on security questions with an open mind.

My appreciation goes further to NSEP and SINTEF ICT for granting me access to all the resources I've needed. Thanks to Arild Faxvaag, the leader of NSEP, for also giving useful input on issues I've worked with. I also would like to thank Inger Anne Tøndel at SINTEF ICT for both your highly constructive criticisms, and for the efforts on the MIE paper — I look forward to working more with you.

Although they are kept anonymous (or de-identified), I'd like to thank the 10 clinicians from Trondheim University Hospital who participated in the experiments. Moreover, an applause is well earned by Hanna Marie Volle for proofreading my manuscript. Credits also to Ole Andreas Alsos and Anita Das for doing pilot tests in the lab, and Terje Røsand who facilitated the video recordings. In addition, I am glad to have been working next to (and eating soup with) Andreas D. Landmark, Tobias B. Iversen and Leendert Wienhofen — and everybody else at NSEP. I'll be back...

Last, but indeed not least, I want to express my gratitute to Ragnhild, Håkon and Olive for being the best family I could ever dream of. Not to forget my in-laws Ildrid, Knut Meinert, Kathrine, Kristine and Ingebrigt for your support when time was short, and the rest of my family for your unconditional encouragement. I love you all!

And finally, praises to the Almighty enabler[1] of bits, bytes and pixels, for giving me such magnificent materials to play with.

Trondheim, June 2011

....

Erlend Andreas Gjære

[1] Credits also to Ralph Hartley, Werner Buchholz, and Paul Nipkow, respectively

List of Figures

List of Tables

List of Abbreviations

COSTT	Co-Operation Support Through Transparency
CSCW	Computer Supported Collaborative Work
DSRP	Design Science Research Process
EPR	Electronic Patient Record
GUI	Graphical User Interface
HCI	Human-Computer Interaction
HIPAA	Health Insurance Portability and Accountability Act
IPS	Indoor Positioning System
IS	Information System
ICT	Information and Communication Technology
NFC	Near Field Communicaiton
NSEP	Norwegian Research Centre for Electronic Patient Records
NTNU	Norwegian University of Science and Technology
PHI	Patient Health Information
PII	Person Identifying Information
PIN	Personal Identification Number
PoCCS	Peri-operative Communication and Coordination System
RFID	Radio-Frequency Identification
SDG	Single Display Groupware
SSN	Social Security Number

"If you think encryption solves your security problems then you don't understand encryption and you don't understand security."

Roger Needham
PROFESSOR OF COMPUTER SYSTEMS,
UNIVERSITY OF CAMBRIDGE

Chapter 1

Introduction

1.1 Background and motivation

Medical work that takes place in modern hospitals today, is well-known to often happen in a fast pace and may even appear chaotic (Sandberg, Ganous & Steiner, 2003). Personell are frequently on the move between locations, or switching between parallel tasks and interactions with people and systems. Interruptions, e.g. in door openings, the hallway and on the way from one task to another, moreover cause instant changes in the priorities of personnel (Bardram, 2005). While health care processes at the same time become more distributed and involve increasingly more specialized people in the care of every single patient, it is equally important to support the collaboration between them. Providing such support through ICT may however be difficult due to the degree of unpredictability involved (Lillebo, Seim & Faxvaag, 2010; Melby & Toussaint, 2009). A common observation from the peri-operative domain (before, during and after surgery) is nevertheless how operation personnel may "peep into the coordination room, check the status, drawing their own conclusion with regard to their own work, and silently carry on" (Bardram, Hansen & Soegaard, 2006*b*). An important approach to coping with the problems described above, is hence to increase the personnel's awareness of currently ongoing activities, so that they are better able to self-coordinate their own efforts. Creating such awareness systems will however depend on providing status information to these workers, and to provide it right where they are without causing even more interruptions.

An example of such as system is the Cetrea Surgical system, a peri-operative com-

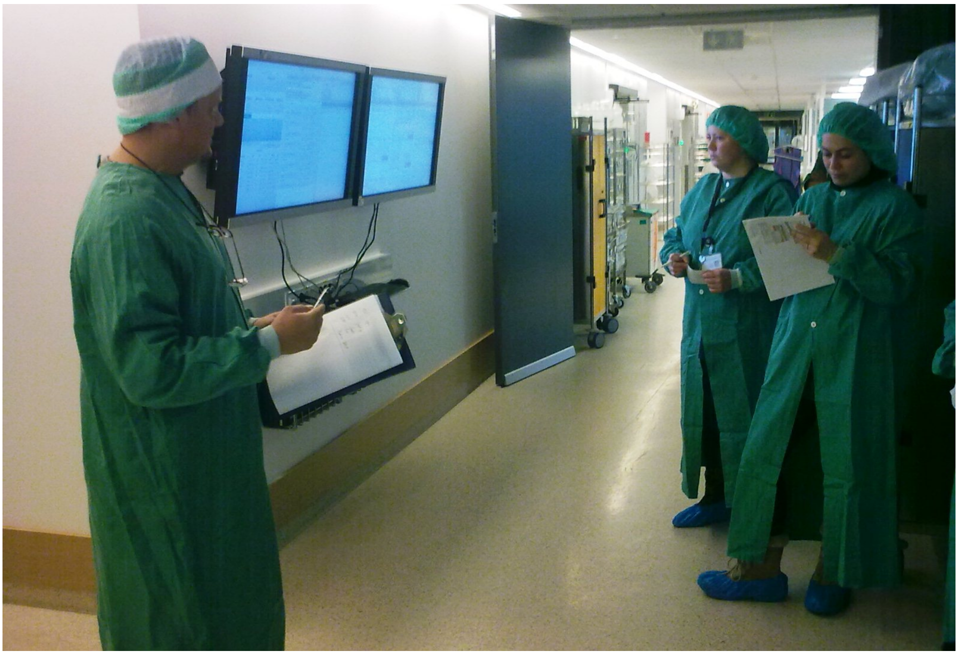

Figure 1.1: Cetrea Surgical screens installed in the hallway outside operation theatres.

munication and coordination system (PoCCS) that connects the emergency reception, patient wards, operation rooms, operation coordinators, and the sterile supplies department through one graphical user interface (Bardram, Hansen & Soegaard, 2006*a*). Via large touch-enabled screens the clinicians may communicate progress information about patients, chat, and adjust their future operation plans. When the author participated in a risk assessment session during a pilot installation launch of the system, an issue was however identified in that the screen's placement on the wall in the operation corridor (see figure 1.1), was crucial in order for all intended actors to utilize it easily. On the other hand, the screen could at the same time be seen and its contents read by everyone having physical access to the area, which included both nurses, physicians, surgeons and coordinators, as well as cleaners, security guards and even patients being transported in and out of surgery.

Whiteboards, in both "old-fashioned" and digital versions, are also well-known implementations for providing status update information to personnel. Due to their size and visibility, the information becomes very available to those who need it, while their usefulness also depend on being situated in places where clinicians often come by (Bardram et al., 2006*b*). The purpose of public boards and displays is hence to make information more publicly available, yet the sensitivity of the information displayed would normally vouch for a more private approach (O'Neill, Woodgate & Kostakos, 2004). The contradiction that must be solved here is then how to make private information publicly available, without revealing any sensitive information to outsiders like patients, visitors or other hospital workers who have no reason for accessing it.

The traditional paradigm of personal computing may therefore prove a bad match to such hospital work in many aspects, including security-wise (Rogers & Rodden, 2002). Large screens like those in figure 1.1 are not used in a private sphere like a normal PC. As commented above, everyone that pass by could catch a glimpse of the information being displayed, given that a single user has provided the credentials required to access sensitive data. Nevertheless, it is the user who has been authorized by the system, who is responsible for preventing unauthorized redistribution of the accessed information. In public and semi-public environments like this it may unfortunately be impossible for the user to actually be in such control of the displayed information, e.g. if curious visitors or cleaners happen to stop by. Therefore, this problem becomes an important design issue to solve for such systems.

Another issue is that the nomadic work, faced by e.g. clinicians whose area of action comprise several departments, is badly supported in electronic patient record (EPR) systems that tie up every user session to a single computer. In practice this requires users to spend much time only logging on and off when they change between tasks and locations, as observed by e.g. Fuglseth (2008). For the case of coordination systems in particular, users cannot be expected to spend much extra effort on becoming and staying informed. Many systems in use in hospitals today nevertheless require the users to deal with relatively time-demanding routines before getting access. In addition, users cannot simply return to computer work which was left aside just a few moments earlier, but must often start all over again when interruption occurs, unless the computer is being locked and others cannot use it in the meantime (Heckle & Lutters, 2011).

Other challenges are in addition added by the users' needs of sometimes working individually, yet suddenly change to collaborative work in a group when necessary, and then perhaps changing back again. In itself, this is a significant design issue for CSCW research (Ishii, 1990). When the group also consists of users having different system privileges and restrictions on their access rights, it becomes difficult to provide a correctly balanced view of the information, so that both the usefulness of the information is preserved, as well as the privacy of the patients involved (Faxvaag, Røstad, Tøndel, Seim & Toussaint, 2009). Sometimes the access rights of system users may not even be reflected by their needs for information in the real world, adding to the problem's complexity (Heckle & Lutters, 2011). A key challenge is therefore to establish group access control mechanisms in systems that are designed for supporting the creation of awareness on current activities in a hospital (Iachello & Hong, 2007).

1.2 Co-Operation Support Through Transparency

The Cetrea Surgical system was, as briefly introduced above, pilot tested at Trondheim University Hospital from January to June 2011 as part of a research project called Co-Operation Support Through Transparency (COSTT). The project has been funded by the Research Council of Norway's VERDIKT ("Core Competence and Value Creation in ICT") program, and run by the Norwegian Research Centre for Electronic Patient Records (NSEP). COSTT is a multi disciplinary research community at the Norwegian University of Science and Technology (NTNU) in Trondheim, and the pilot with Cetrea Surgical was used for facilitating a range of qualitative studies in the project.

The COSTT project's main goal is to support health care personnel in self-coordinating their work. Instead of an increased focus on controlling the flow of work and planning "perfectly" in advance, the involved actors should be provided an overview over process progress and status and thus become better fit to both plan and re-adjust their efforts. This is expressed through two of the project's overall objectives: (1) "To enable flexible, 'just-in-time' coordination of work in a highly collaborative and dynamic work environment" and (2) "To achieve this by creating a shared work space that gives all the actors involved in the collaboration real-time insight into the work process, e.g. its progress and possible deviations from the expected course.". The system targeted in COSTT is not intended to be another full-powered EPR system for manually entering large amounts of data. Instead, it will gather events created by e.g. sensors and data entries in other systems, and enrich those with a context, i.e. relations to patients, locations and personnel, which in turn makes the events interpretable. The information will be most suited for use in real time, despite that one could imagine statistical data being produced from this as well (Bardram & Hansen, 2010).

Eventually, the interpreted information will need to be presented in some kind of visualization and made available to the hospital workers for them to make use of it. This is also in line with the idea behind Cetrea Surgical, which is to “get information stored in clinical booking and scheduling systems ’back on the wall’ by designing a large interactive display technology [...]” (Bardram et al., 2006*b*). During the first week of pilot testing, it was however uncovered that it was hard for ward nurses to know when someone tried to reach them via the chat function, since due to privacy issues at the ward, the particular screen had to be placed alone inside a small room so that checking the screens for new progress information had to be actively pursued by the nurses there. A nurse said eagerly: “We should rather have had this screen hanging outside on the wall, so that we could more easily see whether something happens that requires our attention!”.

1.2.1 Introducing de-identification in COSTT

In previous work related to the COSTT project, a concept of “flexible de-identification” is proposed by Faxvaag et al. (2009) to mitigate the threats to privacy when implementing a system that will expose clinical data on large wall-mounted screens. The question asked is how sensitive data can be removed from a visualization, when the remaining data still must convey meaning for those who are authorized to access it. The authors examplify a justified scenario where the visualization still serves its purpose without disclosing full identification of the patients and personnel involved in the visualized cases. Then, a number of approaches on how to implement the required security measures are suggested (quote):

1. To reduce the level of granularity by de-coupling actor and role (e.g., replacing name of actor with name of role that the actor enacts)
2. To de-identify by abstracting (e.g., replacing ’patient with a tumour in ileum’ with ’patient with neoplastic disease’),
3. To de-identify by replacing direct identifiers with pseudonyms, and
4. By logging of individuals’ as well as teams’ use of information visualizations.

These approaches are still left for exploration in the context of a real-time system, both in terms of how well they can serve the users’ needs for utility, and how well they protect the privacy of those who are exposed. The suggested approaches above will hence form the main entry point for this study.

1.3 Theoretical framework

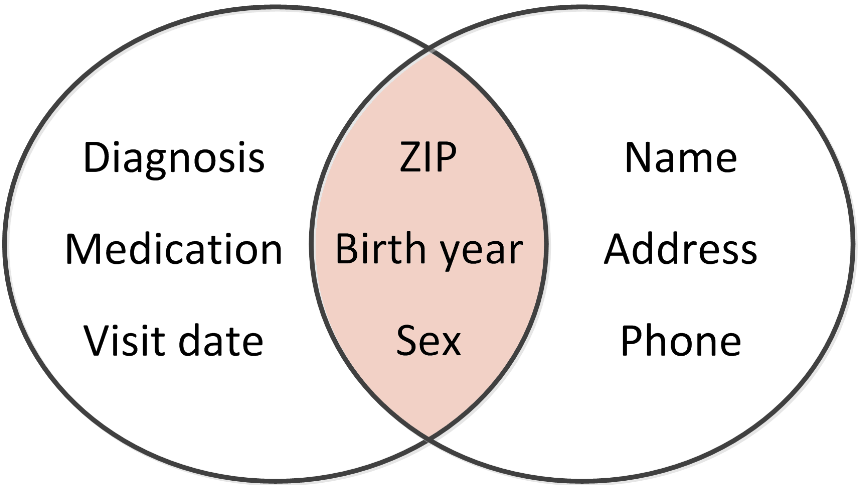

Figure 1.2: De-identified data set (to the left), intersected with a publicly available identified data set.

De-identification is normally associated with a "scrubbing" process that is applied to data sets which contain sensitive patient information. Such processing is normally mandatory, and takes place before the data are disclosed to a third party, e.g. from a hospital to a medical research group. In those cases, it is also usually an explicit objective to avoid the re-identification of any individuals in the disclosed data, by removing or altering any information that can be used to trace sensitive information back to its origin. Later, in chapter 2, it is however discussed how such re-identification still can be done in many cases, using data linking attacks where common fields, i.e. the intersection of the de-identified data set and another data set in which individuals are identified, are matched and reducing the subset of individuals to one individual, which has then become re-identified. This is also briefly illustrated in figure 1.2, where the de-identified data set (left circle) contains medical information such as diagnosis for patients, as well as their postal code, birth year and sex. At the same time, all of these three last variables are found in publicly available databases in Norway today, such as online searchable phone catalogs (right circle), where the variables are attached to an identity with name, address and phone number. If the subset of sensitive records having a particular combination of these three variables is narrowed down to one, and the attacker knows that a certain individual is contained in the set, then these three variables alone will certainly provide access to the sensitive record in question.

In this project however, de-identification will rather be explored with a goal *facilitating* re-identification, but only for those users who are actually authorized to access the identified information. The theoretical foundation which may enable such a contradiction, is based on the linking of data sets — however, as opposed to the previously mentioned attacks, the identified data set that must be correctly associated in order to re-identify the de-identified individuals, is not publicly available. Instead, it will possibly reside in the memory of those who are indeed authorized to access the meaning of the de-identified data set, i.e. nurses, physicians and others who need such information in order to provide care.

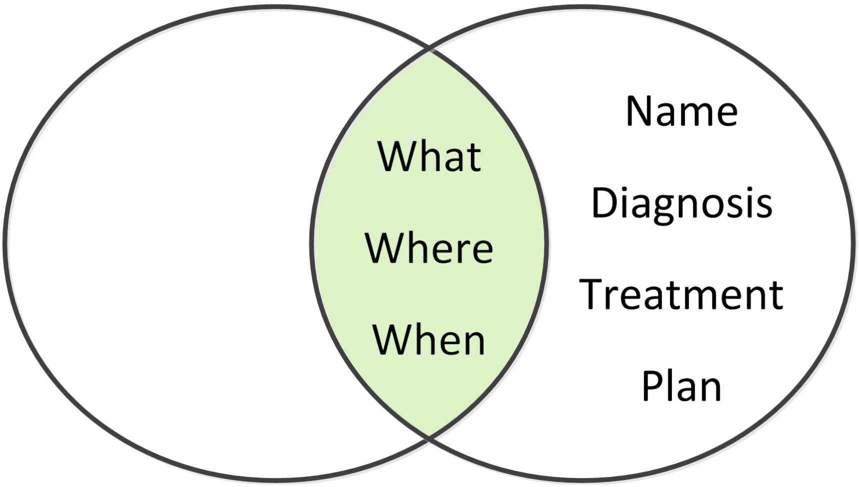

Figure 1.3: Aggressively de-identified information about a hospital event (to the left), intersected with knowledge held by its authorized users (the identified data set).

The essence of this theory is illustrated in figure 1.3 where the de-identified data set (left circle) now contains three example variables that could each be important for a typical status event message that being useful for coordination purposes, while leaving out the “who” variable to make it (possibly) de-identified. The identified “data set” (right circle), residing in the memory of those who may also have knowledge about the clinical state of the patient, will correspondingly contain some information about diagnosis and treatment. Although no-one can foretell the future with certainty, these persons might also be able to know or infer some events that can possibly happen to the patient in question. Therefore are the “where”, “what” and “when” variables also contained in this data set, making an intersection equal to figure 1.2 and hence enabling the linking “attack”.

A relevant example could be that a patient, who is currently undergoing an operation, most likely will be transferred to the post-operative unit in a foreseeable scope of time, unless something unexpected happens. The nurse who will make plans for getting the patient back, may benefit from knowing that things go according to plan, and further infer when to expect a message that the patient is ready for being taken back to the ward. At the same time, a nurse from another ward, or another patient, would probably not have such information about this particular patient, and could therefore not draw the same conclusions.

The main question nevertheless remains what level of information detail is required for personal health information (PHI) to still be meaningful to those who are authorized to access it.

1.4 Problem description

The sections above have outlined several challenges involved in digital support for collaborative work in hospitals. In order to focus the work to be done in this thesis, a more condensed problem description defines on a high level what should be investigated further:

> How can both the information utility for users be maintained and the privacy for patients be protected, when potentially sensitive personal health information (PHI) is used for supporting collaboration in a hospital setting, and hence being increasingly distributed in public and semi-public environments?

1.4.1 Research questions

As already introduced above, the use of de-identification could potentially replace possibly less appropriate access control mechanisms used today. It is however unknown to which extent de-identification may fulfill the utility needs of users, since essential information is removed for acheiving sufficient protection. The problem description above is hence broken down to two research questions, which each will be answered in chapter 5:

RQ1 Which methods are the most appropriate for de-identification in real-time oriented visualizations containing patient health information?

RQ2 How do de-identified visualizations perform compared to what medical personnel need and what current practices are?

The research questions each address the implementation of de-identification in real-time visualizations, both with respect to how usable they are in practice, and also how well they protect the privacy of patient. The first will evaluate both already known techniques and uses of de-identification, as well as new design ideas that may arise along with the project. The goal is to limit the number of possible approaches, so that it becomes feasible to investigate the remaining ones more thoroughly. Then, the second question will serve to evaluate these candidate designs against the reality of needs raised in a hospital today. This second evaluation will hence take particular interest in how such implementations are evaluated by their potential users, calling for user involvement to be done.

1.4.2 Additional considerations

For the work in this thesis, it is important to make a distinction between access control, and what can be called disclosure control. While the former deals with users' access rights to the information, the latter is concerned with the occurence itself of information disclosure, which represents a slightly broader scope. Wanted disclosure can indeed happen and be controlled through the use of access control. Unwanted disclosure on the other hand, includes security breaches in organizational infrastructure, computer systems and devices, as well as disloyal employees, but does not require a fault in the access control mechanism to occur. Unwanted disclosure can also happen through initially wanted (or authorized) disclosure, since information that should not be disclosed can become mixed with or embedded within information that is intentionally disclosed (Ohno-Machado, Silveira & Vinterbo, 2004).

Disclosure control hence aims at controlling information that is released outside protected channels also, where de-identification can be introduced as a potentially useful tool. Since access control provides protection against direct disclosures, but at the same time does not address disclosure that is a product of inferences drawn from the released data set, disclosure control is a topic outside traditional work on access control (Sweeney, 2001). This distinction is also present in a misuse case diagram for the envisioned COSTT system in figure 1.4, where some preliminary de-identification technique suggestions are included as well.

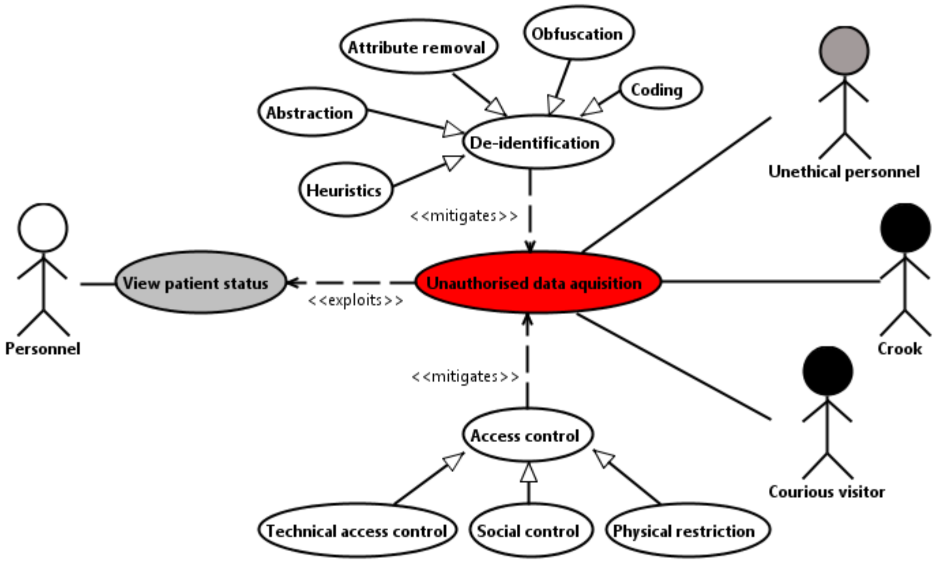

Figure 1.4: High level misuse case diagram for the targeted COSTT system.

For the question of what level of information detail can still protect the privacy of patients, the answer will most likely depend on what protection is required in different locations and contexts. The hospital has several kinds of places where information screens could potentially be placed, including ward hallways, waiting rooms and other areas being open to the public in general. Other locations are rather semi-public due to access control restrictions, being limited to an uncontrolled variety of subsequently authorized actors. Such areas are for instance operation rooms and corridors, normally accessed by e.g. surgeons, anesthesiologists and operation nurses working there, but also cleaners and security personnel. Although these areas are commonly protected with physical access control mechanisms such as door locks, hospital personnel can tell many stories about people being where they are not supposed to be, both intentionally and unintentionally. Areas protected by doors requiring ID-card and PIN-code are hence not impregnable fortresses, but exposed to e.g. people following someone else through.

Another challenge could also be to decide which thresholds can be used to determine when the information is de-identified or not, so that an implementation would be in accordance to the law. For all systems that manage PHI it is very important to properly protect the privacy of patients, and if the targeted COSTT system is not able to meet the criteria set by national legislation and privacy authorities, it may be rejected by these, independently of how well it works functionally-wise.

1.5 Ethics

When experimenting with placing sensitive patient health information (PHI) in public areas, the privacy oriented crowd might feel immediate concerns rising. As long as there is no formal evidence for a de-identification method that ot protects all information displayed at all times in every context possible, there will exist a risk for systematically sharing sensitive person data with the world. Taking a step down would be to only pursue screens placed behind a desk or in a restricted meeting room, which would surely mitigate that risk effectively. Experimenting with public exposure is however chosen to test how useful de-identification can be as a method in itself, for use in real-time coordination of patient treatment. Therefore, all other access control restrictions must be kept out of the picture while exploring this matter. Then finally, if de-identification proves not to be sufficiently powerful by itself, one must consider additional or alternative measures to protecting the sensitive information properly.

On the other hand, when introducing de-identification and hence making patient identification ambiguous, severe objections from a patient *safety* point of view may arise. If clinicians are to impose actions on a wrong patient because of a system that is designed to make this a both possible and plausible scenario, then the topic should very well be left unexplored. Having this in mind, it is still interesting to explore a potentially useful dimension within the field of information security in CSCW. It will however not be recommended to proceed with any solution in ways that might lead to fatal side effects.

It is also important to note that although this project deals extensively with sensitive personal health information, no real such information is ever encountered. All patient cases that were used in the experiments and in this report are made up, including the patient list in appendix H.

When initiating the qualitative study, a permission was granted by the privacy commissioning authorities at Norwegian Social Science Data Services (NSD), so that we could talk to health workers in the hospital. The Regional Committee for Medical Research Ethics (REK) had already approved of the COSTT project as a whole, including the types of experiments that are described in chapters 3 and 4.

All participants in the experiments have been anonymized in both the final report and papers. Those who participated in the lab experiment signed an informed consent form, and were also given an information sheet stating what data was gathered, and how it would be used — see appendix G.

1.6 Structure

This first chapter has already introduced the project and its challenges, while chapter 2 will continue with an overview of available relevant literature that contributes to the study. Chapter 3 then describes how the research was framed, before motivating and describing in detail the methods that were used for data collection. The execution of and results from the design research work is then described in chapter 4, before results from the experiments are analyzed in chapter 5, framed within the research questions. A discussion of these results, how they were obtained and how they relate to previous findings, is subsequently found in chapter 6, along with a prototype that is suggested for further work in the field. Finally, the thesis is concluded in chapter 7.

1.7 A note on collaboration

The whole period of experiments work, which is described in detail in chapter 4, was done in mutual collaboration with PhD candidate Børge Lillebo. He was at the time working at the Norwegian University of Science and Technology (NTNU) on his doctoral thesis in Medical Technology, and additionally holds a doctoral degree for physicians (M.D.), with previous experience as an intern and resident at surgical departments in three different hospitals.

Lillebo's research challenges addressed visualizations for improved coordination among personnel in a medical context, and almost at the same time as my work started out, he initiated a study on real-time visualizations of information about patient care activities. Since his visualizations would require an edge towards information security, and in addition I needed a system like the one he pursued to design, for testing purposes, we were able to merge our goals, and hence our prototypes, interviews and experiments, into a common effort.

Refer to appendix D for a more detailed specification of which are his contributions towards my work.

"In science read the newest works,
in literature read the oldest."

Edward G. Bulwer-Lytton
British politician and poet

Chapter 2

Literature review

This chapter presents research and other relevant literature that consider de-identification and related ideas for protecting patient health information (PHI). A range of different techniques for de-identification are introduced, and each is briefly assessed towards appropriateness for the intended use in this thesis. While de-identification is the main topic, its targeted application also builds on an understanding of also additional topics addressed, and a closer look at how sensitive information could be taken into shared spaces in practice. This chapter also considers legislative requirements that are considered relevant for sharing PHI in the targeted contexts.

2.1 Privacy

The purpose of exploring de-identification is not make life harder for hospital workers, forcing them to spend energy interpreting "vaguely" expressed messages. On the contrary, few people in the world will have a more internalized perception of what this research is trying to protect, namely what is known as privacy. Ever since Hippocrates invented the art of medicine, the confidentiality of patients' health information has been a principle of respect that scholars must learn and practitioners must exhibit in order to maintain the trust of the people who confide all sorts of conditions to them (Moskop, Marco, Larkin, Geiderman & Derse, 2005). Today's medical oaths echo this principle, for instance the World Medical Association's Declaration of Geneva which contains the statement "I will respect the secrets which are confided in me, even after the patient has died." (Reich, 1995).

Privacy is a concept that needs to be understood before designing solutions for it, especially when the systems are to be placed in public spaces where users often coexist in a commong space with people they do not want to interact with (Little, Briggs & Coventry, 2005). Here we have two aspects to privacy that should be considered, where the first of physical privacy concerns the spatial properties of where interaction takes place, e.g. if a screen is turned away from the public, and the second of psychological privacy that describes every individual's need to control access to oneself. One of the first models of privacy however suggests four types of privacy: solitude (being free from observation by others), intimacy (small group seclusion), anonymity (freedom from surveillance in public places) and reserve (limited disclosure of information to others) (Westin, 1967). The last type has the function of protecting communication, which limits the sharing of information to only trusted others. Westin's model has been further developed by Pedersen (1999), who re-defines the anonymity type to "being seen but not identified or identifiable by others" — an interesting definition in the context of exploring of de-identification.

Within the field of computer science, the term of privacy is often related to protection of data against various risks or during transmission (Little et al., 2005), or a statement of what data is being collected during an online session ("privacy policies"). However, in some cases another perspective is entered, such as the "ability of an individual (or organization) to decide whether, when, and to whom personal (or organizational) information is released" (Saltzer & Schroeder, 1975). This approach suggests both a certain amount of flexibility in the privacy mechanisms, and that privacy is related with an active relation from the information holder to the information itself. This will all still be heavily dependent upon the context in which it occurs (O'Neill et al., 2004).

"Public" and "private" are either not absolute and static measures (Boyle & Greenberg, 2000), but as for the case with privacy, their meaning will depend on the context. Such meaning may also change over time, in conjunction with changes in the society elsewhere. When new technology is introduced, patterns of both use and social norms may change in addition, making what is considered "acceptable" behavior subject to these changes, too (Bellotti & Sellen, 1993). This however makes developers of new technology responsible for both preserving and changing policies for privacy, since technology is not neutral when this topic is considered. Rather, it can have large impact on the extent to which people have control over their personal data.

2.2 Legislation

In Norway, rules and regulations on the duty of secrecy and how colleagues are allowed to share PHI, are mainly found in the act relating to the Processing of Personal Data (Norwegian statute, 2000). In addition, for handling personal health data in particular, the act on Personal Health Data Filing Systems and Processing of Personal Health Data (Norwegian statute, 2001) is essential. This second statute also implements the EU personal data protection directive for the health domain, and altogether, these laws provide a number of useful definitions for the work to follow. Those which are considered most relevant, have been quoted below:

Personal data "Any information and assessments that may be linked to a natural person."

Sensitive personal data "Information relating to [...] health [...]."

Personal health data "Any information subject to the duty of secrecy [...] and other information and assessments regarding health matters or that are significant for health matters, that may be linked to a natural person."

Processing of personal health data "Any use of personal health data for a specific purpose, such as collection, recording, alignment, storage and disclosure or a combination of such uses."

Data subject "The person to whom personal health data may be linked."

Consent "Any freely given, specific and informed declaration by the data subject to the effect that he or she agrees to the processing of personal health data relating to him or her."

De-identified personal health data "Personal health data from which the name, personal identity number and other characteristics serving to identify a person have been removed, so that the data can no longer be linked to a natural person, and where the identity can only be traced through alignment with the same data that were previously removed."

Anonymous data "Data from which the name, personal identity number and other characteristics serving to identify a person have been removed, so that the data can no longer be linked to a natural person."

Pseudonymous health data “Personal health data in which the identity has been encrypted or otherwise concealed, but nonetheless individualized so that it is possible to follow each person through the health system without his identity being revealed.”

As deduced from the first two definitions, personal health information (PHI) is regarded as sensitive whenever it can be connected to a specific physical individual. However, if the health information is no longer possible to connect with a specific individual, it will also not be regarded as sensitive anymore, and the need to regulate access and disclosure becomes superflous (Andresen, 2010). In the list of definitions, there are already mentioned three useful approaches to such de-identification, i.e. anonymization, de-identification and pseudonymisation. However, only the first of these qualifies to be exempted from privacy regulations. The table in figure 2.1 moreover reflects how these approaches are ordered from more to less strain on the privacy of patients, from the left to the right, in the bottom row. Finally, the middle row represents the granularity of identification, illustrating that fully identified data and pseudonymized data will have the same accuracy.

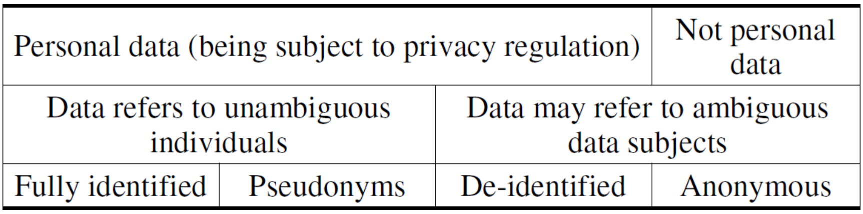

<table>
<tr><td colspan="3">Personal data (being subject to privacy regulation)</td><td>Not personal data</td></tr>
<tr><td colspan="2">Data refers to unambiguous individuals</td><td colspan="2">Data may refer to ambiguous data subjects</td></tr>
<tr><td>Fully identified</td><td>Pseudonyms</td><td>De-identified</td><td>Anonymous</td></tr>
</table>

Figure 2.1: Outline of levels of patient identification (Andresen, 2009).

For the case of anonymity, the definition mentions name and personal identity number as identifying information that must be removed in order to protect the individual from ever becoming re-identified through the disclosed data. In addition, all other characteristics that hint towards such re-identification shall be treated in the same way, however not giving any criteria to help evaluating what a sufficient threshold of anonymity will have to be in practice. Andresen (2010) comments that it is hard to give general answers to what is sufficient for anonymity, but that the person or organization responsible for disclosing anonymized health data must take great caution when doing so. In practice this requires any hints that may reveal the patient’s identity to be removed, and deliberate action may be required to sacrifice their accuracy as well, e.g. by generalizing values such as ages into age groups (Andresen, 2009). This is because anonymized data may be published and known for all future, and hence do not require additional data protection. Anonymization is hence only useful in situations where re-identification of the patient is never needed.

While the purpose of de-identification is also to prevent re-identification of individuals, it is a concept open to doing precisely this as well. For such re-identification to be allowed, the definition however requires a dependency towards re-supplying the data which has been removed in the de-identification process. This threshold implies that it should not be possible to gain access to individual identities by guessing, yet there is still no general rule for what will be an acceptable level of de-identification. In preporatory works for the EU Data Protection Directive (1995), which provides a common minimum standard to be implemented in the members' national laws, it is stated to require "an unreasonable amount of time and manpower" to consider an individual being unidentifiable in a data set. In Canada, there are statutory tests for degrees of identifiability, ranging from "reasonably forseeable" to "readily ascertainable" or "obvious" (El Emam & Kosseim, 2009). These thresholds however do not depend on being applied from a particular perspective, e.g. a motivated intruder vs. an average layperson, nor do they reference expressly the level of resources, time or effort needed for re-identification.

For situations where protected access to identities in a data set is a necessity, the last option of pseudonymization can be applied. The motivation for enabling such access is to enable linking of information to individuals, for instance if a hospital wants to know whether four given operations have been done to four different patients, or if the same patient was involved all four times (Andresen, 2009). Pseudonymization is based on encryption of all identifying information, protecting the individual's identity against anybody except those with access to the pseudonym's meaning, still referring to that unique individual. This level of uniqueness is not guaranteed nor desired with both de-identification and anonymization. Pseudonymous data will however have the same legal status as de-identified data, as shown in figure 2.1.

2.2.1 HIPAA requirements to de-identification

The Safe Harbor policy found in paragraph 164.514 of the Administrative Simplification Regulations promulgated under the Health Insurance Portability and Accountability Act (HIPAA) in the USA (1996), provides two options for data sets that shall satisfy a legal characteristic of being de-identified.

The first possibility is that the data is purged of a specified list containing 18 categories of potential identifiers that are related to either the patient or relatives, household members and employers, and any additional information that could make it possible to re-identify the individual in question. This would include personal names, addresses and telephone numbers, yet the final category is commonly interpreted to comprise clinicians who were working with the patient, as well as the names of hospitals, clinics and wards (Uzuner, Luo & Szolovits, 2007). When this process of de-identification is completed, a public-use data set is generated with no restrictions attached to its use. This is however criticised by Sweeney (2001) for incorrectly believing that de-identified data equals anonymous members of the data set, an issue taken further in section 2.6 below.

A second possibility allows an individual with knowledge and expertise on de-identification, for instance an expert statistician, to determine that the information contained in the data set is subject to very small risk of being used, 'alone or in combination with other reasonably available information, by an anticipated recipient to identify an individual who is a subject of the information'. This alternative approach may however allow more 'gray' room than the first approach, and should be chosen when the level of risk towards the recipient is clear (El Emam & Kosseim, 2009).

2.3 Secondary uses of patient health information (PHI)

Patient health information is primarily required and intended for the treatment of patients, such as information that is held for each patient in an EPR system, and used for e.g. giving the right medicine or amputating the right leg (or the left). These pieces of information may however also have relevance and value outside the immediate process of providing care. Such relevance may lead to so-called secondary uses that serve a variety of functions, yet somewhat remote from the patients' expectations to what their personal information is being used to.

The occurencies of secondary uses found in literature, has that in common that it happens outside the health care organization. This may not be very surprising either, since that what often happens inside a hospital is the primary use of the information. As a consequence, what is found of research involving de-identification is normally targeted at these kinds of uses as well. Using PHI for coordination purposes, may nevertheless be defined as secondary use when for example a nurse's awareness of progress in a particular ongoing surgery will allow him/her to plan other activities ahead, before returning to care for that particular patient. While this simple scenario represents the targeted context of use better than the background for existing literature which considers de-identification, a brief presentation of other secondary uses is still provided below, for the purpose of informing the discussion of legal aspects to de-identification.

Providing data for both academic and commercial research on public health and the effects of various treatment methods is a well-known example of secondary use, mentioned by e.g. Andresen (2009), Behlen & Johnson (1999) and Appari & Johnson (2010). De-identification works as a crucial enabler of making clinical data available to for instance epidemiological investigations, and collection of data on drug interactions and side-effects (Wellner, Huyck, Mardis, Aberdeen, Morgan, Peshkin, Yeh, Hitzeman & Hirschman, 2007). In fact, researchers within the field of automated de-identification do need to access large quantities of medical records for testing and further developing such tools themselves (Yeniterzi et al., 2006). Other purposes of use may include decision support on a national level and reimbursement settlements (Andresen, 2009), and policy making (El Emam, 2008). It is however likely that decision support on a local basis commonly occurs too, based on data the organization already possesses.

The most comprehensive type of secondary use is however so-called registers, although essentially being enablers of many other secondary uses. A register is not only a database, but also comprise an operating organization, responsibilities and duties for others to report to the register, and restrictions for use defined by a legal authority (Andresen, 2009). All information disclosed for secondary use is normally de-identified in a way that makes re-identification impossible, however for some registers it may be important for its function to be able to associate de-identified data from different data sets with one another. An example may be The Norwegian Prescription Database ('Reseptregisteret', in Norwegian), from which various stakeholders demanded statistics based on prescriptions and dispatch from pharmacies to individual patients, although not intending to re-identify any of the particular patients and physicians. Andresen (2009) further describes how all pharmacies report prescription data electronically every month to a central data collecting point, which in turn sends the data to a trusted third party. By using pseudonyms, the identities of both patients and physicians are encrypted, so that when the register owner, the Norwegian Institute of Public Health, receives the data, it can be linked to existing statistics, however not re-identified without the encryption keys that are held by the trusted party. The pseudonyms both ensure that the register can fulfill its assigned purpose, while the privacy of patients and physicians is protected. Another example, is if participants in a clinical study are undergoing genetic tests and it may become necessary to inform the patient or their physician if these additional tests reveal a risk to the participants or their families (El Emam & Fineberg, 2009).

Andresen (2009) subsequently describes the political debate around the introduction of pseudonymous registers in Norway, especially when the Ministry of Health in 2006 suggested making the Abortion Register a pseudonymous register, but in 2007 was decided by the Government to be a de-identified register instead. The Data Inspectorate had argued that people knowing about this register might be influenced in their decisions on whether to have an abortion or not, and so the register would not only reflect but possibly affect the health care activities. Andresen comments that the confidentiality of a pseudonymous register is based on trust in today's society, while the possibility of reversing pseudonyms might arise some time in the future if privacy values are not held intact.

Appari & Johnson (2010) nevertheless argue that little research has examined the patients' perception of their health records being shared. An exception is Campbell, Thomson, Slater, Coward, Wyatt, & Sweeney (2007) who found that about 28% to 35% of the 166 UK respondents considered themselves netural to their PHI being used anonymized by physicians for other purposes — including both age, sex, ethnicity, reason for treatment, medical history, personal habits impacting health, type of treatment obtained and side effects of treatment. Only about 1 in 8 of the survey's respondents expected to be asked for permission if their physicians used PHI for a wide variety of purposes, which included sharing how the treatment is working with other physicians in the hospital and writing research articles about diseases and treatments.

2.4 Hiding information from outsiders

A problem that commonly arise when several individuals are viewing the same computer screen, is that potentially private information, such as personal health information (PHI), may be disclosed to someone who is not authorized to see it. This may happen both when many people are using the same computer, when projecting the data on a larger screen, and if someone happens to see someone else's screen over their shoulder. Whenever the possibility of sensitive patient health information (PHI) disclosure is present, the users today often have few options available when it comes to controlling exactly what information to show on the screen. Either the system and hence PHI is open and visible for everyone being close enough to read, or it has to be closed down or minimized in order to conceal it from unauthorized insight. Some professionals who work in crowded areas such as the metro, however use physical filters on their laptop screens or cellular phones in order to limit the angle of insight (Iachello & Hong, 2007; Gass et al., 2007).

Privacy-augmented displays is on the other hand a concept for alterations of screen contents that do not require any additional physical equipment, and also distinguishes private data from public very clearly. This includes an idea of automatic blinders, being an approach where black squares are placed over particular private areas of a graphical user interface such as a web browser (Tarasewich et al., 2005). The automatic blinders, which in theory can find sensitive data such as personal financial details from a banking web site and hide them by covering these areas, can be controlled by the user with special mouse gestures. For instance, if the pointer is hovered on top of a blinder, the blinder can disappear and reveal the data behind it — and when the pointer is removed, the blinder re-appears where it was.

The technique involving blinders is also applied by Huang, Chu, Lien, Hsiao & Kao (2008), however an interesting addition is made with a generic button for hiding person identifying information (PII) being displayed, that is implemented for widespread technologies such as HTML for web, PDF documents and XML. The component is suggested for protecting patients who collect their own copies of personal health records, by equipping them with a tool for protecting themselves against spilling sensitive information after it has been acquired. It does not require any additional hardware, and can provide privacy preservation for people during reading.

A similar approach is to use the power of pixels to code information (Grimes & Tarasewich, 2005), so that only the user who has defined the coding can understand the meaning of it. For the example of financial use, a coding technique for numbers could simply be to assign a color and shape to each number, so that e.g. a green circle and two red circles could represent $100, however another user would use a black square and two blue circles to tell the same thing. Experimental results suggest that these techniques actually seem quite usable in everyday tasks. Obviously, such personalization would strengthen the coding approach to privacy, which is normally considered weak in practice, when codes are normally shared within a group. On the other hand, a group would not benefit as much from the visualization if only one of the users can read it, as a maximum. It could even lead to dangerous misunderstandings if one user reads one thing, while another reads something else.

Another way to code the data is to depict data in a graphical and artful way, very much like optical cryptography, however being easier to read by the recipient who knows what pattern to look for (Stasko, McColgin, Miller, Plaue & Pousman, 2005). Stock quotes could for instance be hidden in a cloud hovering over a landscape, although this technique would probably be less secure than the personalized approach above.

Coding also occurs in plain text, and is very common among professionals in health care. Whiteboards are used in hospitals for both reference to patients, as well as asynchronous messaging between colleagues, and normally contain a lot of encoded information. An example is information such as “v*A*3 Vgh Prn 09:00#12” that refers to a particular type of medication, where the next dose (dose number 12) should be given at 09:00 (Bjørn & Hertzum, 2010). It can however be discussed how secure this kind of coding is towards outsiders, as health care professionals do happen to be patients themselves sometimes, and could possibly decode the message with ease.

Single Display Groupware (SDG) is a research topic challenging the standard model of one user working on one computer, by allowing each individual user personal access to a system although the screen is available to the whole group (Shoemaker & Inkpen, 2001). The main idea is that if a user can supply a private input to the system, the system will return a private output specifically for that person, and at the same time that there is a public output presented to all users who can see the screen. The authors claim that mixed shared/public displays could provide opportunities for enhanced collaboration, supporting both shared data and individual exploration of the data. The implementation however requires the users to wear special 3D glasses that make only every other video frame visible to each of them, but the accompanied user study provided little negative feedback. Later, the original concept has been commercialized by Sony, and was introduced in 2011 at the E3 expo, however advertised as the PlayStation TV replacing the traditional split-screen approach to multiplayer games on a shared screen (Macrae, 2011).

A very simple strategy for limiting the possibility of unintended PHI disclosure would be to limit the size of the screen. A study involving 60 participants concluded that the users’ perceptions of privacy were strengthened by reducing the screen size where personal data was displayed, for screen sizes of 12”, 15” and 17” (Little, Briggs

& Coventry, 2005). Moreover, all screens were tested with partitions on the sides, which increased the perceived privacy additionally. At the same time, the perceived clarity of the displayed information was not determined by the screen size, because font and contrast was adjusted to make the information stand out anyway. Hence, the authors conclude that the 12" screen is best suited for private transactions in a public environment, although they recognize that not too much information can be presented there at once. This may also pose a downside for application in collaborative work, where people both should be able to collaborate on shared information, as well as being able to get a beneficial overview.

Finally, another approach could be to put restrictions to the information that is displayed, and still communicate something useful to the recipients who are able to make any use of it in such limited form. This way of limiting a data set by removing individually identifiable information, has already been introduced as the concept of de-identification. As described in chapter 1, this thesis aims at exploring whether such de-identification may be useful to users in the context of health care, more precisely in hospitals. The next sections will therefore look further into the concept of de-identification, its current uses and possible issues.

2.5 De-identification

The exact term of "de-identification" is probably younger than the concept in particular itself. When Behlen & Johnson (1999) first enclose the term in quotes, it is in conjunction with a short presentation of techniques which are used for the then-called purpose of "scrubbing" patient data, which again has been referred to as *disclosure control of microdata* since the topic appeared in the field of statistics sometime in the 1970's. The background for such "scrubbing", is identified as a recurring idea of removing personal identifying information (PII) from a database, so that the "clean" database can be made legally available for research. At the same time, there will not be disclosed any sensitive information, since no part of the disclosed information would (ideally) be possible to link back to the individual it concerns. The data could also be used for purposes other than research, which are outlined in section 2.3. The term de-identification is not only used within the context of health care either — see for instance Tarasewich et al. (2005) — although this is mainly where its applications are found.

De-identification is an approach targeted at realizing the benefits of sharing data, while privacy concerns are minimized. Exactly how the privacy of individuals is protected, can vary, but in general it is common to simply remove information for ensuring protection. An example could be to delete all patient names from a list of currently hospitalized patients and their diagnoses, so that only the diagnoses are kept on the list. Information removal must however be balanced against the wish or need to supply information that can still be useful to the recipient, i.e. keeping as much exact relevant information as possible (Ohno-Machado, Silveira & Vinterbo, 2004). Behlen & Johnson (1999) nevertheless argue that complete scrubbing is not feasible while still maintaining the usefulness of data, when applied to a whole database of patient records. Even if it was, the authors say it would not be ethically appropriate, due to the risk of individuals being re-identified. Behlen & Johnson and also Andresen (2009) therefore argue that de-identified data sets always must be protected with access control and other conventional information security measures at the point of query, in order to both provide sufficient protection and maximize utility.

Later, this assertion has been challenged, and today much patient data is made available to research through the use of de-identification (Benitez & Malin, 2010). This is accomplished by using more formal methods of verifying the protection of data such as *k-anonymity* (Sweeney, 2002), which is explained briefly in section 2.6 below. On the other hand, when considering how the de-identification should be done, it is beneficial to divide the data into two categories, each of whose presence will affect which techniques to use in the process (Benitez & Malin, 2010; El Emam & Fineberg, 2009).

First are the identifying variables, such as a personal name, social security number — and also telephone numbers, full address and email addresses — that can all be used to directly identify an individual. If the data set contains such variables, these should be suppressed (i.e. removed), which is the most common approach to de-identification in practice. This can either be solved by removing these specific variables from all records in the data set, or by removing the full records in which the directly identifiable information is contained (Sweeney, 2001). Alternatively, the contents of these variables must be either pseudonymized (Andresen, 2009), coded or randomized (El Emam & Fineberg, 2009). The technique of pseudonymization is also known as reversible coding, and can be used for cases when it may be necessary to retain the unique identities in the dataset, without disclosing them. When used in practice, pseudonyms act as direct identifiers of the individuals (Kobsa & Schreck, 2003), yet their association to particular individuals is hidden.

Second are the quasi-identifiers, which are variables that indirectly can be used to identify an individual, for instance ethnicity, sex, postal code or other location information, diagnosis information, socio-economic information, and dates such as date of birth, death, admission, discharge, or specimen collection (El Emam, 2008; El Emam & Fineberg, 2009; El Emam, 2010). Such variables can be publicly known, but are not always very obvious either. An example is that of profession, which in combination with sex could very well identify a female judge within a whole region. Profession-specific insurance company, graduation date, and eventual mammogram tests, are all additional examples that may be used to infer quasi-identifiers, respectively profession, age and sex (El Emam, 2008). The quasi-identifiers can yet be treated in a number of ways, applying one or several techniques for de-identification which are included in list below.

A distinction can also be made between de-identification of PHI in a structured format, for instance fields in a database (El Emam & Fineberg, 2009), versus unstructured free text records such as clinical notes, discharge summaries and radiology reports. While the identifying variables can be easily defined in structured data, it is more challenging to de-identify text that may contain hidden PII through acronyms and references used by the author (Gardner & Xiong, 2008), yet there is obviously much information found in free text records that can be of great value for several purposes, just like structured data. For coordination purposes, one can also imagine pieces of unstructured data to be useful if sufficiently protected. Therefore, it may also be relevant to look at automated de-identification of free text, yet Meystre, Friedlin, South, Shen & Samore (2010) only found 18 out of 200 publications on de-identification being relevant in this area. The remaining 182 publications were either focused on manual de-identification (e.g. El Emam (2008)), de-identification of structured data (e.g. McQuaid, Zheng, Melville & Green (2009)), or de-identification of images (e.g. Gross et al. (2006)). Although the architectures and methods for detecting potential PHI vary for the automated de-identification approaches, what lies beyond the sheer detection of PHI is not very relevant in this study, since what is detected will simply be removed (Uzuner, Luo & Szolovits, 2007; Friedlin & McDonald, 2008; Yeniterzi, Aberdeen, Bayer, Wellner, Hirschman & Malin, 2006).

Being aware that there are key differences between existing de-identification applications, and the application of de-identification required in the COSTT, the section below aims at presenting an overview of what are known techniques for de-identification in general. This is done in order to give them all a fair chance in this new area of application.

2.5.1 Techniques for de-identification

Since the terminology used in literature is not always consistent for de-identification techniques, some of them are grouped together in the list below. The grouped terms are closely related and could therefore be considered together. It is also possible that others will make use of other terms, that still refer to the these and very similar techniques.

Variable suppression Removal of entire columns from the data set, e.g. removal of the patient name and SSN for all records. This is the by far most commonly used approach, and mentioned by many sources including Ohno-Machado et al. (2004), Friedlin & McDonald (2008) and Sweeney (2001). The obvious downside is that it impacts the data set's utility value whenever a variable is suppressed.

Record suppression Removal of complete records from the data set. The technique is very effective for preserving privacy, but leaves no utility left whatsoever, for the information that was cotained in suppressed records (El Emam, 2008; Sweeney, 2001).

Coding, masking This technique involves data to be coded with making its representation less accessible without knowing the codes (Bjørn & Hertzum, 2010). The codes can either be reversible, or irreversible (El Emam & Fineberg, 2009), with pseudonymization as the most common concept of reversible coding. Coding often occurs as textual representation, but another possbility is also to use colors and shapes (Tarasewich et al., 2005).

Generalization, aggregation The accuracy of variables are intentionally sacrificed in order to make them less unique in the population Sweeney (2001). The technique can for instance be applied to ages, postal code areas and similar, only disclosing which age interval group the individual belongs to, e.g. 10-20 for the 14-year-old, or “above 80” for the 107-year-old (El Emam, 2008; Behlen & Johnson, 1999; Wang et al., 2004; Fung et al., 2005). It can also be applied to less numeric attributes, such as by replacing a diagnosis of “tumour in ileum” with “neoplastic disease” (Faxvaag et al., 2009). The aggregation technique can also be used for abstracted facts as such, but rather for disclosing generalities about a population, than about individuals.

Randomization, permutation, swapping, substitution While aiming at preserving as much detail as possible, the data elements are replaced with alternative fake substitutes that are randomly selected (Behlen & Johnson, 1999; Sweeney, 2001). The substitution techniques vary in sophistication, but some can for instance take geographical distribution of names into account (El Emam & Fineberg, 2009). The resulting data set should hence appear realistic, and not distort the statistical outcome of the data set (Szarvas et al., 2007; Gardner & Xiong, 2008).

Obfuscation, noise-introduction Deliberate action is here taken to make the disclosed data wrong, or artificial data is added to change the overall composition of the data set, making it harder to know what is true or not (El Emam & Fineberg, 2009; Behlen & Johnson, 1999; Xiao & Tao, 2006).

Sub-sampling This technique selects only a random subset of records from the original data set, so that it cannot be known whether or not a particular record is to be found in the final set (Behlen & Johnson, 1999).

Pixelation, blurring These are techniques for de-identification of graphical elements, including images and live video. The more sophisticated k-Same(-select) algorithm can also blend facial features from several persons into a new de-identified face (Gross et al., 2006).

Heuristics An approach based on removing uniqueness or rareness within a population, using a range of heuristics — some more formal than others. The acid test used is whether or not record linking attacks can be done using public registries (El Emam, 2008; El Emam & Fineberg, 2009), such as that briefly described in chapter 1.

Analytics A manual process where at first all quasi-identifiers in the data set are determined. Then, a threshold for risk of disclosure is determined, before the quasi-identifiers are evaluated towards this risk limit. If the risk is at an acceptable level, the disclosure can happen. If not, there must be applied other de-identification techniques until the risk falls below the threshold (El Emam & Fineberg, 2009).

2.5.2 Selection of appropriate techniques

Based on the goals for this thesis, it can already be revealed that some of the techniques are not appropriate for further consideration. Nevertheless, it is useful to define some requirements to be used in the evaluation, and judge them all by these. In the paragraphs to follow, these requirements will be introduced one at a time, while eliminating the techniques that do not fit in. The surviving alternatives will on the other hand be included in the project's next stages.

Since Norwegian legislation already has defined three categories of de-identification approaches, it should be a criteria that the selected techniques will fit inside one of these. While anonymization relies on variable suppression alone, and pseudonymization relies on reversible coding, de-identification can however make use of the rest of the alternatives presented. One particular exception is however the analytics technique, which can rather be directly applied as the overall approach to de-identification in this thesis. Another is the heuristics technique, a closely related analysis process, that uses formal evidence to achieving uniqueness

The next requirement is that the technique can be used to represent one specific individual, so that this person can in theory be uniquely re-identified from a given population, or to represent a whole population. Although both sub-sampling and record suppression would remove an arbitrary number of records, the requirement is not to preserve all original data, but rather to facilitate accurate use of what is left. Sub-sampling will however be based on random removal, so that one can not know which records are missing, and is hence not appropriate for such structured use.

The above requirement also implies the next criteria, which is that patient data must not be distorted in a way that points it directly towards the wrong individuals. Although de-identification is used for intentionally creating ambiguousness, it is important to not have introduced false facts when authorized re-identification is achieved. This requirement will hence eliminate both permutation/swapping/substitution, obfuscation/noise introduction and randomization. Substitution could on the other hand also be used to hide e.g. rare conditions without revealing that it is rare, and therefore it could still be a plausible alternative to record suppression in these cases.

Making this initial selection, reveals that the following techniques will also be included in the next stages: Variable suppression, record suppression, coding/masking and generalization/aggregation, as well as the pixelation/blurring alternative.

2.6 Risks for re-identification

Since it is decided to use the analytics approach to de-identification as a general approach, it is worth knowing something about the strength of de-identification, so that risk can be properly assessed. Consider the linking attack illustrated in figure 1.2, along with a number of individuals who are each instantiated with their own set of variables. If all these variables are considered as independent facts about the individual, and two of them are linked together, this linkage will identify a subpopulation which is found in the intersection of these variables. For each additional fact that is linked towards this subpopulation, the intersection will become smaller, until finally it only contains a single individual, which has now become precisely identified. Behlen & Johnson (1999) demonstrated that the combination of birth data and residence ZIP code, would uniquely identify one of the authors among 135 000 patients that were registered at a particular hospital.

The re-identification that occurred here, which may also be known in literature as an identity disclosure, is based on the attacker being able to make a likely match between a de-identified record and a corresponding record in a dataset where the identity is known (Benitez & Malin, 2010). Obviously, the removal of identifying variables and disclosure of only a very small number of quasi-identifiers will preserve the anonymity of individuals in a given data set. Unfortunately, the removal of all identifying variables alone will probably not provide sufficient protection (El Emam, 2010). This very same point is demonstrated by Sweeney (2001), who successfully re-identified medical information for the governor of Massachusetts, by linking a set of publicly available hospital discharge records that were believed to be de-identified, with the voter list for his residential area. Sweeney found that six people in the records had the governor's particular birth date, only three of them were men, and from these there was only match for the governor's five-digit zip code. Moreover, a number of cases where re-identification has been successfully accomplished from publicly available data is described by El Emam & Dankar (2008).

What made such re-identification possible, and indeed revealed the reason for a mysterious hospital stay of the governor's, was the linking of two data sets, in both of which the attacker could believe that the particular person would be. By combining the available facts, it may be possible to identify a sufficiently small subpopulation to infer the existence of an individual there, and hence put his privacy at risk (Ohno-Machado et al., 2004). It can however be argued that not only those who are uniquely identifiable are at danger of privacy loss. Re-identification can happen just as long as there is a certain level of individuality in the data set. An attacker who is able to isolate the only two males born in 1935 from a given ZIP code area, will for instance be correct 50% of the time with a random assignment procedure (Benitez & Malin, 2010).

The risk involved in such re-identification of public data sets which have been legally de-identified, has been more formally assessed by Golle (2006), who checked the uniqueness of simple demographics in the US population. According to this comprehensive study, 63% of the population can be uniquely identified by their combination of full date of birth, sex and ZIP code. According to Sweeney (2001), 18% of the US population could be likewise identified if the population's county of residence was given instead of the ZIP code. Removing the 18 attributes that are named in the HIPAA Safe Harbor policy and will render a data set legally de-identified, is however estimated by the US National Committee on Vital and Health Statistics (NCVHS) to result in a 0.04 percent chance of re-identification in the US (El Emam & Fineberg, 2009).

Another risk is found through the discovery of information about individuals without the need for identifying their specific record in a data set, so-called attribute-disclosure (El Emam & Fineberg, 2009). An example to illustrate this, can be that if all females in the age between 50-55 in the data set are diagnosed with breast cancer, then it will not matter how large the subpopulation of this age interval is, as long as you can know that a female of the same age is found in the complete data set — her breast cancer diagnosis will be revealed anyway. Moreover, an individual's membership in a particular database can alone be sensitive information, even if the released data does not contain any medical information, and could cause identity-disclosure as well (El Emam, 2010). For instance, a database of HIV treatment receivers will have such sensitivity, yet the challenge remains to actually verify an individual's membership here. Nevertheless, the more a patient stands out with respect to history of treatment, or changes in personal profile due to much moving around, a simple list of such dates and changes could reveal the patient's identity. In general, the more encounters, the more unique the pattern becomes, and the risk of re-identification increases (El Emam & Fineberg, 2009). Finally, patients who are suffering very rare conditions, such as progeria which is associated with unnaturally fast aging, would very likely identify this individual patient even if coupled with very little additional information (Gardner & Xiong, 2008).

The above mentioned examples of re-identification can always be assumed to happen if the attacker is in possession of some background information on the individuals (El Emam, 2010). Another issue is found when knowledge about the database, beyond its contents, is added to the equation — so-called meta-knowledge (Ohno-Machado et al., 2004). If the attacker knows what measures were taken to protect the data set, for example that no entire records were suppressed, or that it was incurred an absolute minimum of information loss in order to meet legislative requirements, this could add certainty to the findings of an attacker.

2.6.1 Managing risk

An important aspect of the risk assessment is exactly who the information should be protected from. If the purpose is making a publicly available, searchable database, the risk of re-identification would naturally be considered very high. For research purposes, it is however usual for the trusted recipient to sign a data-sharing agreement and consent to audits (El Emam, 2008), which will lower the risk. The sensitivity of the data may also vary, and highly sensitive information such as HIV treatment should always be protected by aggressive de-identification, compared to data whose disclosure would be less harmful. Also, if a consent is given from the patients that considers secondary uses, then appropriate usage of data in relation to a such consent may additionally lower the risk of invading the patients's privacy. The risk may finally also be affected both negatively and positively by the size of the disclosed data set, as a larger set would allow more redundancy and hence anonymity for each unique tuple (Sweeney, 2002), while a smaller set would perhaps include fewer "special" cases that stand out from the others.

Threats to re-identification can therefore be divided into three categories (El Emam, 2010). For many cases, a commonly disinterested general public may first not seem to be a significant threat. It is nevertheless necessary to assess this risk of motivated intruders who try to re-identify a particular patient or patients, and are also in possession of additional background knowledge on these (El Emam, 2008). Victims for such activity may thus very likely be neighbors and famous persons, or relatives of the attacker. Second, there is the risk of journalists attempting to re-identify individuals in a data-set, but who are not certain whether or not the person sought for is actually included in the set. Finally, there is a risk for the kind of attackers who try to re-identify as many people as possible, for instance if a pharmaceutical company who manufactures medicine for diabetes patients tried to identify diabetes patients from a data set, in order to target their marketing directly at those who are successfully re-identified.

While the potential threats are obviously real, the options for mitigating them are quite general and thus fewer in number. When de-identification is done manually by a data custodian or privacy officer on the health care organization's behalf, it is often based on the use of de-identification policies, since the data custodian would not necessarily be able to assess properly the risks involved in releasing a given data set, without assistance from an expert (El Emam & Kosseim, 2009) — moreover many such policies are applied without any awareness of the risk for re-identification (Benitez & Malin, 2010). Whenever the the prepared data set can be compared with demographic data, statistical methods can be used for quantifying the threats to re-identification (Appari & Johnson, 2010), but such relevant data for comparison are not always available.

As a response to this, a new principle was required for protecting data, that does not require access to externally available data as such. Sweeney (2001) and Sweeney (2002) presented such a privacy principle to be used in the process of de-identification, that is called *k-anonymity*. The principle is that in a disclosed data set, there should be at least k tuples (individuals) having the exactly same instantiation of all disclosed variables, so that for $k > 1$ not one individual can be precisely identified after publication. Thus, the principle is used for judging whether a set of data provides sufficient protection for it to be published. The k-anonymity principle is considered superior to naïve de-identification approaches where certain identifiers, like the HIPAA Safe Harbour policy allows, are simply removed (Agrawal & Johnson, 2006). While k-anonymity, effectively mitigates the risk of such linking attacks, it has been proved to still be prone to e.g. attribute disclosure, and thus subsequent enhancements have been suggested, e.g. by Truta & Vinay (2006) and Kifer & Gehrke (2006), to maintain confidentiality of patients in the cases not covered.

However, the data owners must always satisfy the two objectives well-known to be in opposition, namely the privacy of individuals and the utility value of released data — also referred to as disclosure risk and information loss (Appari & Johnson, 2010). Evidence suggest that applying the HIPAA-required removal of 18 specific variables will result in significant information loss (El Emam & Fineberg, 2009). When a set of data contains several quasi-identifiers, and an expert should find out how to make the risk as little as possible as well as minimize information loss, it may not be a straightforward task to choose which fields should be included in the de-identification process, in which order they should be processed, and by which technique they should be dealt with in order to surpass the required minimum level of k. Furthermore, El Emam (2008) argues that de-identification is a risk-management exercise, in which the objective is not necessarily remove the risk of re-identification entirely, but rather to choose a risk level that the custodian is willing to take for a particular disclosure.

What could possibly be a help in making these decisions, is a a narrative and graphical tool prototyped by McQuaid et al. (2009) for visualizing risk when preparing a data set for disclosure, and hence for finding the most privacy-preserving "scrubbing" of the data. This is technology for supporting a manual process which today is not efficient enough to release the potential value held by being able to analyzing more health data, yet it could possibly allow the data custodian to see what exact de-identification procedure from a number of candidates that preserves the most usable data, while still fulfilling the requirements to privacy protection.

2.7 Summary

This chapter has provided a review of how sensitive data such as PHI could be managed, both in theory and in practice. Techniques of known de-identification techniques have moreover been presented and briefly assessed towards their appropriateness for use in real-time visualizations in a hospital. Although this initial collection of appropriate candidates has resulted in several alternatives to proceed with, the leading selection framework is defined by Norwegian legislation. This has already defined three main approaches that can be used, namely anonymization, de-identification and pseudonymisation.

For a technique to qualify for further use, it must hence fit within one of these categories, and also be used to represent individuals or populations without obstructing the correctness of the data. The selected techniques are therefore variable suppression, record suppression, coding/masking and generalization/aggregation, as well as the pixelation/blurring alternative. In addition, the further development of de-identification solutions should both assess the risk for unwanted information disclosure, and then apply the most appropriate techniques for protecting privacy in each particular implementation.

"The most difficult part of requirements gathering is not the act of recording what the user wants, it is the exploratory development activity of helping users figure out what they want."

Steve McConnell
Former Editor in Chief of
IEEE Software

Chapter 3

Methodology

This chapter presents the methodological foundation for the research done throughout this thesis, along with the process model that was used to organize its activities. A design research strategy was chosen for framing the project as a whole, which encourages use of both qualitative and quantitative data collection methods. The methods used in this project are all qualitative however, and include a literature review, rapid field tests, role playing in a usability lab and focused interviews. These decisions are all motivated in the sections below, and the research activities are further described from a theoretical point of view.

3.1 Research approach

The problem of how de-identification should be designed in order to support utility, is naturally a design problem. Trying to assess the utility value of particular visualizations, as stated by the problem description in chapter 1, indicates that we are also talking about research work in the category of design science. For this reason, the *design science research process* (DSRP) model, as described by Peffers, Tuunanen, Rothenberger & Chatterjee (2006), was chosen as a framework for assisting the research process as a whole. Its academic approach, in combination with allowing the research to be focused on the users and a problem that needs to be solved for them, also contribute towards its appropriateness for the project.

The DSRP model, as shown in figure 3.1, is a generic tool made not only for guid-

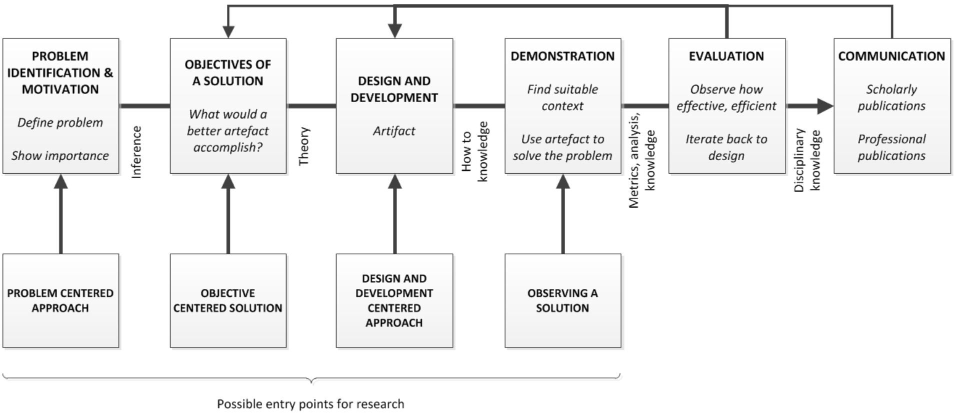
PROBLEM IDENTIFICATION & MOTIVATION
Define problem
Show importance
Inference
OBJECTIVES OF A SOLUTION
What would a better artefact accomplish?
Theory
DESIGN AND DEVELOPMENT
Artifact
How to knowledge
DEMONSTRATION
Find suitable context
Use artefact to solve the problem
Metrics, analysis, knowledge
EVALUATION
Observe how effective, efficient
Iterate back to design
Disciplinary knowledge
COMMUNICATION
Scholarly publications
Professional publications
PROBLEM CENTERED APPROACH
OBJECTIVE CENTERED SOLUTION
DESIGN AND DEVELOPMENT CENTERED APPROACH
OBSERVING A SOLUTION
Possible entry points for research

Figure 3.1: The design science research process (DSRP) model (Peffers et al., 2006)

ing researchers in the process of research itself, but also for suggesting output from different stages of design science research. It describes six main stages of a research process, stages which have already been previously described by other process models with some, or all, of these steps included. These other attempts however lacked a standardized terminology, as well as proper distinction of exactly where common activities usually (should) take place. The DSRP model further expresses four possible entry points for research, i.e. the model's first four stages, and defines what kind of research approach will be appropriate when starting out at each of these entry points.

An instantiation of the DSRP model for the work on this thesis, is presented in figure 3.2. As seen from the model, the project enters at the very first stage, which makes the approach *problem centered*. The project's problem description has in turn suggested a candidate for solving the problem that is identified in the first box, which is the yet unexplored appliance of de-identification in real-time hospital systems. The importance of finding such a solution is grounded in the need for preserving both utility and privacy. In the second box, the objectives are thoroughly described, and these together comprise the superior requirements for a working implementation, both in terms of purpose and context. The design and testing activities performed in the third and fourth box are discussed in the next sections, and this is the part where data are collected and hence users are involved. The main difference between the testing of prototypes in the design and development stage and the demonstration stage is that in the latter of the two, the prototype is contextualized and thus demonstrated in a relevant setting, instead of just being a sketch designed to stimulate the creativity of users.

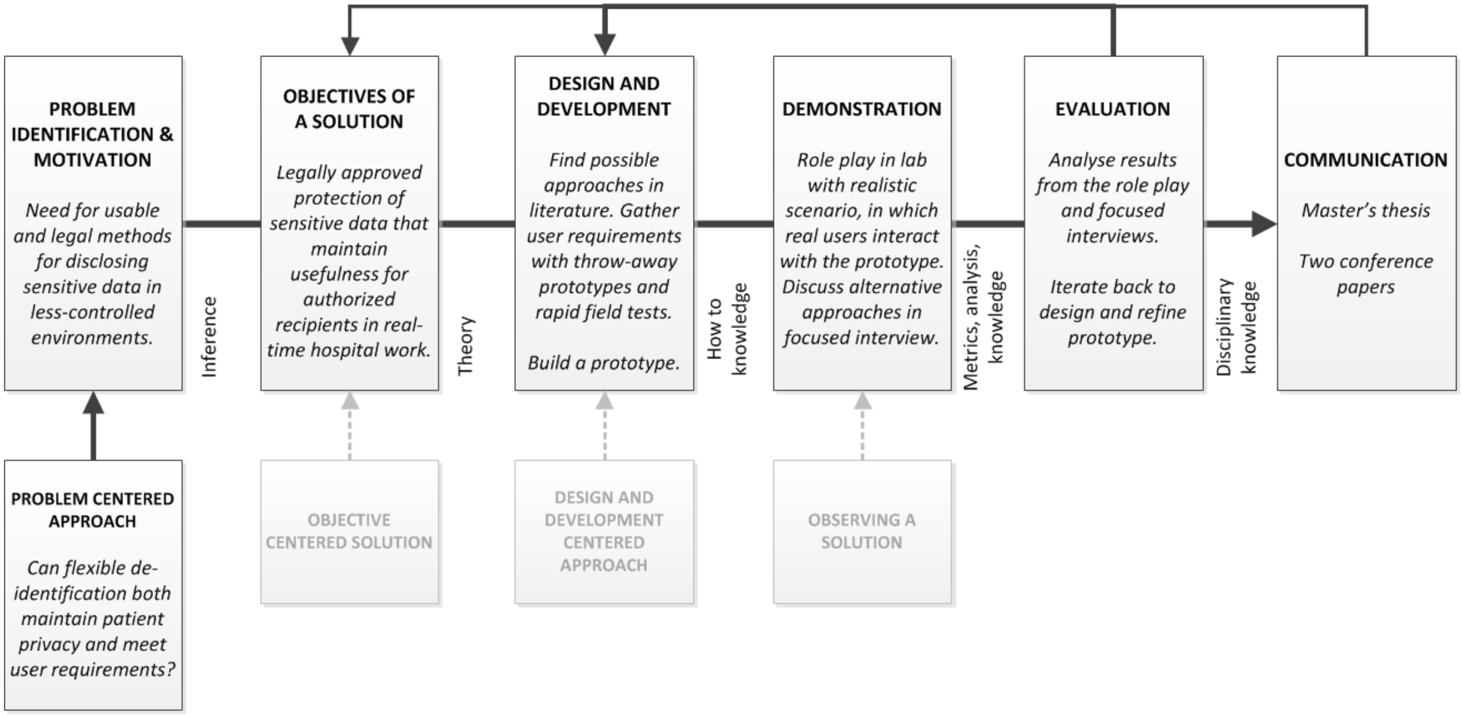
PROBLEM IDENTIFICATION & MOTIVATION
Need for usable and legal methods for disclosing sensitive data in less-controlled environments.
Inference
OBJECTIVES OF A SOLUTION
Legally approved protection of sensitive data that maintain usefulness for authorized recipients in real-time hospital work.
Theory
DESIGN AND DEVELOPMENT
Find possible approaches in literature. Gather user requirements with throw-away prototypes and rapid field tests.
Build a prototype.
How to knowledge
DEMONSTRATION
Role play in lab with realistic scenario, in which real users interact with the prototype. Discuss alternative approaches in focused interview.
Metrics, analysis, knowledge
EVALUATION
Analyse results from the role play and focused interviews.
Iterate back to design and refine prototype.
Disciplinary knowledge
COMMUNICATION
Master's thesis
Two conference papers
PROBLEM CENTERED APPROACH
Can flexible de-identification both maintain patient privacy and meet user requirements?
OBJECTIVE CENTERED SOLUTION
DESIGN AND DEVELOPMENT CENTERED APPROACH
OBSERVING A SOLUTION

Figure 3.2: This project's instantiation of the DSRP model, entering a problem centered approach

In the fifth stage, after collecting the data, the analysis takes place and an evaluation of the remaining candidate approaches is made. Instead of simply drawing some conclusions there and communicating these as the project's final stage, the process iterates back to the design and development stage in order to refine the prototype according to the test results. The improved prototype should then again be demonstrated and evaluated, before the results are ready for presentation to an academic audience. In addition to this thesis, the research has at the time of writing resulted in two[1] papers submitted and accepted for academic conferences, which are found in the attachments A and B.

3.1.1 Data collection

Deriving information from the real world can be done either qualitatively or quantitatively. The COSTT project which this work is done as a part of, is designed as an ethnographic study, which mainly includes qualitative, but also may utilize quantitative data collection methods. The research questions raised in chapter 1 might, due to the prospected assessment of appropriateness (RQ1) and performance (RQ2), also here indicate a qualitative design, which is also the chosen approach for the work on this thesis. While quantitative data can make it easier to analyze the interrelations between variables and responses, e.g. age and computer skills, the main strength with qualitative data is that it can provide more depth to what lies beyond the informant's response (Tjora, 2010). Exactly how the data collection would happen, was however not entirely specified before starting this part of the work. Both the research questions and suggested methods were hence more or less tentative and refined during the project's start-up period. An early idea was nevertheless to design some kind of experiments, to test to what extent de-identified views are considered usable by clinicians.

While qualitative research may be just as valid, or invalid, as quantitative research, some rules are good to follow in order to establish trustworthy work. Robson (2002, p. 166) therefore suggests eight characteristics of "good" qualitative research design, several of which have been pursued as central values in this project. The project hence starts with a single idea that the researcher seeks to understand, an idea that has already been outlined in the problem description section of chapter 1. Then, the use of multiple data collection techniques is encouraged, i.e. what is later referred to as triangulation, while each approach is being summarized and details are properly given

[1] A third paper which is based on parallell related work, has also been included in attachment C.

on how the data were collected.

Triangulation as such is a concept held forward by Robson (2002) as a valuable strategy for dealing with threats to validity, and it is being widely used in many research disciplines. Possible threats to validity may include incompleteness in the recording of data, failure in interpretation of the collected data, and not least biases introduced by respondents and the researcher. Focus should therefore be placed on the participants' views, while the researcher acts merely as an instrument of data collection. The involvement of multiple sources of data collection methods and/or observers, altogether enables triangulation as a strategy for mitigating such threats as those mentioned here. This project has attempted to adhere to this strategy, utilizing both field tests and lab experiments in order to build a richer and more precise understanding of the potential users, with respect to both what they say while being in their normal work environments, and also when they participate in a controlled lab experiment. These particular data collection methods have also been done in collaboration with another scientist (see appendix D for details), which has provided both duplicate field notes and valuable discussions on how the feedback and results are to be interpreted later on.

3.2 Selection of data collection methods

As a natural starting point, known techniques for de-identification of sensitive data would be revealed in a literature review on the general topic of de-identification. This was also a continuation of initial reconnaissance being done before finally choosing a topic and problem area to write about. Since the use of de-identification in this thesis was found to differentiate quite clearly from its normal application, some techniques would naturally be more appropriate than others. The remaining alternatives however set some premises for which approaches could be used to proceed the data collection, and should be considered along with implications of the research questions asked.

The next phase was therefore found to require involvement of potential users, both in order to get a preliminary assessment of the presumably relevant de-identification techniques, but even more importantly to search for new and previously undiscovered approaches that could be utilized in this new context of use. In order to give the users an impression of how the techniques could look like and hopefully trigger their own creativity, it was decided to create some simple paper prototypes. These prototypes could then be brought into short interviews with health personnel, in order to get concrete feedback from many professionals, but without taking up too much of their precious time. In addition to reviewing the sketches, there would ideally be room for checking which de-identification techniques are already in use in hospitals today, e.g. on analogue whiteboards. The data collection method chosen here was therefore decided to be *rapid field testing*, and is described in more detail later in this chapter.

The rapid field test results successfully refined the list of possible approaches to de-identification on large screens. Another prototype with higher fidelity was then built to become a vehicle for testing these remaining techniques. Assessing their individual qualities would however require a more thorough introduction to how a de-identified screen would work in practice, and hence the idea of an experiment in a usability lab was launched. By introducing “real” patients in the new prototype, and framing the experiment with a realistic scenario and corresponding role-play in which new data appeared that had to be interpreted, it could be possible to test to what extent the de-identified data were adding to the test person’s understanding of the situation. It was also decided to do a focused interview with each test person directly after the role-play was finished, very much like a de-briefing is normally found after usability tests. In addition, this would be suitable for thoroughly assessing several other alternatives, so that the test would not be limited to the prototyped approach only .

After analyzing the results from the lab tests, it was clear that none of the approaches presented there were optimal. It is undoubtedly important to listen to the users, and therefore another iteration was initiated, in which the goal was to build a more interactive prototype with new functionality as proposed from the results analysis. Due to time constraints, the demonstration and evaluation of the improved prototype could not be done as an equivalent lab experiment. This part has therefore been replaced with a discussion of the results, yet still in light of the improved prototype, which is all found in chapter 6.

3.3 Literature review

Already before a topic for the thesis was finally decided, some pieces of relevant literature was reviewed for developing an idea of what parts that may have been missing in research, and to develop the theoretical framework that the design research work could be based on. The concept of *flexible* de-identification had already been introduced by Faxvaag et al. (2009), on the background described in chapter 1, but the paper contained few references to other research on its term of origin, i.e. *de-identification*.

After settling on the topic, the second and main stage of reviewing the literature was initiated. In order to gain insight into known methods and contexts for usage of de-identification as a method for protecting information in general, a search was initiated towards existing research relevant to the chosen topic. Although much of this work could be done early in the project, many of the relevant references were relatively recent of age, bearing witness of some activity in the field, and thus it was necessary to pay attention to also new research being published after the main review phase was ended, all the way towards the thesis being finally submitted.

The objectives of doing the literature review have been picked out in accordance to Oates (2006, p. 72), who names ten objectives for a successful literature study. The most important ones for this study has been to provide an overview of known de-identification techniques, place the prospected work in the context of these previous contributions, point out gaps that have not yet been properly covered — such as de-identification in real-time visualizations — and finally identify research methods that will be used for adding new knowledge to the field.

3.3.1 Resources

Due to the high availability of research found on the Internet, searches were first performed in the online libraries IEEE Xplore, ScienceDirect, ACM — digital library, SpringerLink and PubMed. Google Scholar was also utilized for finding publications that were not accessible through these libraries.

For the searches in online databases, the following key terms were used:

> { *"de-identification", "deidentification", "re-identification", "health care privacy"* }

Moreover, all issues of the journal IEEE Security and Privacy were examined for relevant articles, as well as all issues of the Journal of the American Medical Informatics Association (JAMIA) back to 2004.

Apart from what was found in the keyword searches, some particular conferences were used as a source for papers, due to their nature of being up-to-date, more than any journal could be. Although the topic is not one that relies heavily upon groundbreaking technology, it apparently has received more focus in later years than before. The conferences that were browsed through for relevant contributions, include the *Symposium on Usable Privacy and Security* (SOUPS), and *ICST Conference on Electronic Healthcare for the 21st century* (eHealth).

Books have seemingly not addressed very much the relatively narrow topic which the literature review is concerned with. This may be because it is probably not a taught topic anywhere, and there is seldom a need for a theoretical introduction to a process of de-identification, which is specified for those who use it in practice. The topic is also relatively new.

In addition to the searches in journals and conference proceedings, most of which are refereed, it was also performed a search on the above mentioned keywords in the Google Patents beta service. This was done in order to explore further practical implementations of de-identification, and for which uses they were aimed at. It is possible that not all of the patent applications found has been published as research results later on, and could then add to the list of known methods and applications.

3.3.2 Review process

During the search for literature, publications were chosen to follow up on from their relevance in either title, abstract or keywords. Relevance was based on either their use of the term *de-identification*, whether they discussed issues with security and usability, or if they in other aspects had direct relevance to the problem domain of creating process transparency in hospital work. Some of the articles were read more in depth before making a decision on their relevancy, yet the introduction and conclusion sections would often reveal whether they considered anything of interest or not.

After building a collection of potentially relevant literature, copies were obtained either from the online libraries mentioned above, or sometimes via the author's university's web pages, as PDF-files. The copies were then imported into the Mendeley Desktop software, which allows bibliographic information and other meta data to be attached to each document, along with the researcher's personal annotations and comments.

An assessment was at the same time done on the quality of what was found, but as a general rule, all papers from academic conferences and all articles from academic journals were included in the review as long as they concerned the topic in question.

Arguing that some conferences or journals have lower quality, this choice was made due to the purpose of the review, which was merely intended as a starting point for a practical assessment of *possible* approaches to de-identification. The worst thing that could happen would be to include an idea that might not be possible to set into life. Moreover, as we later will see, the interviews brought forward new ideas too, which were welcomed into the continued design process. This also expresses clearly the main "topic specific" goal of the literature review, which has affected what were considered relevant from those sources of information.

In the cases where new de-identification techniques were discovered after the main review phase was initially concluded, these new approaches could be included in the continued prototype development described in chapter 4. Where applicable, these contributions have also been discussed in the literature review chapter along with the other related work.

3.4 Design and testing

The initial selection of appropriate de-identification techniques was based on a few obvious and high-level requirements, and without involving any potential users. When it subsequently was aimed at evaluating performance and utility value, there had to be established some more precise requirements for the design work that these evaluations would be based on. An ultimate goal would be to present PHI that is not meaningful to unauthorized outsiders, but at the same time is sufficient for providing authorized users with information of utility value comparable to that of a fully identified data set. The elicitaiton of such requirements is commonly known as requirements engineering, not being a certain phase or stage in the software engineering process, but rather a concept comprising activities for doing so:

> Requirements engineering is the branch of software engineering concerned with the real-world goals for, functions of, and constraints on software systems. It is also concerned with the relationship of these factors to precise specifications of software behavior [...]. (Zave, 1995, p. 315)

However, the purpose of doing requirements engineering in this project was not to produce a formal requirements document upon which a complete software is to be developed. Instead, the requirements should add to the understanding of how de-identification may be used for real-time hospital systems. Whatever de-identification approach that works best is still highly dependent on both the system and the context in which the system operates, and in order to test the performance of a system's de-identification "feature", it would hence be fruitful to provide a correspondingly realistic application. An already useful visualization could be carefully de-identified for its purpose, and then be used to find out how the information is interpreted by the test users. Creating prototypes is hence a useful strategy when trying to provide such a realistic frame for gathering the requirements. Not only are they good for empirically testing design ideas and hypotheses against reality, but also for establishing an interactive process with the potential end users of a system.

3.4.1 Prototyping

Prototypes have the advantage of being more flexible to ambitious ideas and rapid adjustments than working software, as they do not have to be operative at all. A prototype's lifespan does not have to exceed the test it is made for, and it must not be neither complete nor detailed, pretty, or scalable. Therefore, it can be easily thrown away if it appears to be a dead end, or perhaps it may be gradually modified until finally becoming a working system in the end. The main idea is nevertheless that the prototypes should not affect the decisions made further because they are too valuable to throw away, which is an important point made by Lauesen (2005), p. 44: "The more effort developers have spent making a prototype, the less they are willing to replace it".

The value of a prototype will hence increase when it grows closer to a working product, and so prototypes are divided into low-fidelity (lo-fi) and high-fidelity (hi-fi) categories. The lo-fi prototypes are very cheap and suitable for testing several diverse pathways early in the process, while hi-fi prototyping becomes more justifiable when the end product is better known and attention is paid to details. This is also adopted for testing the de-identification approaches, where at the first stage the original problem required many de-identification techniques to be visualized, and thus several designs existed in parallel that were rather thrown away than refined any further. Through the development of several throw-away prototypes at first, a broader understanding of the original problem could be gained (Oates, 2006), which would be of value at later stages when new suggestions might arise from the participants' feedback. These prototypes were therefore simple models of how information could be organized on the screen, drawn using design software such as Adobe Photoshop and InkScape. Then, when the data collection moved on from the field and into a usability lab, the prototyping was instead concerned with only a single prototype. This hi-fi-version was then developed further throughout the rest of the project, as it would both support interactivity and also be generic enough to suit the remaining tests well.

In prototype theory, one also separates between horizontal and vertical designs (Preece et al., 2002). The concept of horizontal prototypes is aimed at showing as much breadth of the targeted system as possible, i.e. including a majority of functions and views. The vertical prototypes are more depth-oriented on one particular concept or function, aiming at acquiring a deeper understanding of how it is understood by users. Since the de-identification functionality of the prototypes does not require very much interaction, at least initially, the prototypes were mainly designed vertically in this respect.

3.4.2 Rapid field tests

The need for qualitatively testing several candidate designs, raised certain feasibility requirements when choosing a methodology. Aiming to get a first assessment of these prototypes and their appropriateness in the prospected context, it would thus be beneficial to avoid spending very much time on recruiting test persons. In addition it would be a good idea to not waste more than strictly necessary of the test persons' time, since they are already busy with providing care to patients who need them more. It would also be need for additional recruitment to a less open-ended test later on, so that the greater effort should be put in when the collected data will have greater depth.

Rapid field testing, a simplified version of a field testing method described by Shneiderman (1997), was hence chosen as the technique for qualitative data collection at this stage. The technique is based on what is more commonly known as *hallway usability testing*, and could in many ways also be considered as an equivalent to it. Hallway testing is a methodology known to be mainly used in the industry, appropriate for quick assessment of a product at different stages. It has seemingly not been widely used in academic research work, but examples from academic use are found in publications by Danao (2010) and Ismail, Osman & Wahab (2009).

The Wikipedia article[2] for usability testing (2011) has nevertheless described hallway testing as follows:

[2] There was unfortunately not found a more reliable source to the term, but in correlation with other sources found on the Internet, it appears to reflect the method's essence.

> Hallway testing (or Hall Intercept Testing) is a general methodology of usability testing. Rather than using an in-house, trained group of testers, just five to six random people, indicative of a cross-section of end users, are brought in to test the product, or service. The name of the technique refers to the fact that the testers should be random people who pass by in the hallway.

The origin of the hallway test methodology appears to be unknown, however from the description and references above it seems to build on a publication where the ideal number of test users is estimated to being five, on the basis of a cost-benefit analysis (Nielsen & Landauer, 1993). This trade-off approach has clearly appeal to the industry, where usability testing has traditionally often been seen as a cost and thus been omitted, while Nielsen in later years has built up a famous business on usability services for the same industry.

However, in order to fit this project's purposes the original method was tweaked slightly, to a crossing of prototype assessment and interviews, where quick recruitment and short time consumption is central, such as described above. Moreover, the main purpose of the testing is not to discover design flaws, but get serious feedback on candidate ideas and how they could work. In that sense, the field tests could also be looked upon as unstructured interviews, catalyzed by a prototype shown to the informants upon which they reflect. The rapid field tests were hence used for quick access to informants, and getting quick responses on which ideas could be worth pursuing from a user's point of view. In practice, this involved first approaching random hospital workers, and then asking them for instant thoughts on a mock up or prototype. Normally, this method would restrict the time spent on each person to just a few minutes, and they should each only be asked about one prototype. However, if the participant had the time and was eager to talk, it was in addition involved an additional prototype during the session.

3.4.3 Usability lab experiment

Researching a field where the distance between improved utility and brutal failure is short, disqualifies the researcher in making crucial assumptions about the users. Instead, the users should be included in the process of making the software product usable. This involves testing certain product aspects with respect to what is defined as "usability", e.g. in the international standard ISO 9241-11 (1998):

> The extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use.

Usability experiments are a common way to test design solutions in software engineering, which is broadly described by e.g. Preece et al. (2002) and Lauesen (2005). The alternative designs that were still being considered at this stage, hence relied on further quality input from users in order to be properly assessed. The most realistic way to test the success or failure of de-identification would for instance be to set up a large screen in the hospital corridor, integrate it with software already being in place, and de-identify the events that are raised from the system. This could in turn allow something useful to be derived from is use, such as a success rate for correct assumptions. Alternatively, the Cetrea Surgical software could have been utilized further during the COSTT pilot study, by giving it a public view that contained de-identified information. It could then be observed to find cases when there was no need for the user to proceed with logging in.

Two significant matters however spoke against setting up a field test like this, one being that still having several alternative designs would make it hard to test them all, without confusing users who normally depend on a system in their work. The other would be the major ethical downside to using real patient data for a test where it is too soon to guarantee that the de-identification is sufficient, and thus risking illegal disclosure of sensitive data. At the same time, it can be questioned whether much value would be added from a field test compared to a lab experiment at this stage (Devik, 2009). A full-scale field study should rather be used for investigating how a product is actually being used (Preece et al., 2002), and that would be more appropriate when a mature implementation is ready for verification. Causing minimal intrusion in the hospital environment could hence be a guiding value when doing research on clinical systems, a topic also discussed in the paper in appendix B.

Doing a controlled lab-experiment would instead be better suited for testing the prototype. An initial idea was to use metrics normally found in regular usability testing, in order to present the data in a quantitative fashion. This would involve tasks which the test users would have to accomplish on their own, and then for instance count the number of errors made, or the time taken to accomplish each task. Examples of such tasks could be providing the user with a sheet of patient cases, and then show them the prototype with the de-identified events, and ask them to associate each event with a patient. The main reasons for not doing the experiment like this, was largely that either it would not be feasible to recruit a sufficient number of professionals as test persons, and that using non-professionals for the test would wreck the opportunity for talking qualitatively about the application and how de-identification could work for them in practice, afterwards. In addition, it would be easy to collect the raw data using high-quality video recordings, which would not be feasible in the field.

The interview guide used can be reviewed in appendix F, however it is only provided in its original Norwegian edition.

Role-play

Sometimes it is easier to show than to explain. Using role-plays for staging the lab experiment data collection was hence motivated from the need of contextualizing the system for the test users, more than what the rapid field tests could do. Role-plays tend to feel very natural, at least for participants who are either playing themselves or a role familiar to them (Seland, 2010). The role-play was for this reason used in the experiment, while at the same being based on a realistic scenario that could simulate a regular work situation.

In further comparison with a real field test, the role-play provides greater control over the situation, and one can even "freeze" a particular moment and play it over again if desired. When used in very early stages of user centered software development, the participants can even sketch improvised prototypes during the play, but prototypes can also be developed beforehand, as was done for this experiment. Instead, all relevant views of the prototype were printed on paper so that the test persons could draw on them during the interview session afterwards.

Focused interview

The concept of focused interviews is described by Tjora (2010), being highly encouraged in cases when the traditional in-depth interview would be done, only because that is the "normal" thing to do in qualitative research. When instead a probe or an artifact, such as the prototype used in the lab experiment, could be introduced to the informant at an early stage, this could facilitate a shorter but rather more focused talk about a particular issue or topic. It will for example not be necessary to spend much time on a warm-up session which has no particular meaning, except for getting to the point. This way, the length of an interview could be reduced from the “normal” hour-long session, to a mere 20-30 minutes, depending on the number of questions. Since the interview was to be done as an extended de-briefing session from what is normally used in usability-testing, it was important to limit the total duration of the role-play and interview altogether.

What is important when executing focused interviews, is that instead of using the interview itself to narrow down the topic discussed, the topic should be narrowed down much more before the interview starts. As a consequence, it may not be as easy to introduce alternative themes during the focused interview, and the interview guide should hence be well-focused and naturally structured so that it will be easy to follow by the informant. In relation to the role-play session, this was however no big issue to accomplish, since the informant had already been thoroughly introduced to both the problem domain and a possible solution then. Assessing the alternative solutions presented would not introduce any new topics, but rather elaborate on what was already experienced up front.

3.5 Analysis

Although a prototype of an application’s de-identified view was created for use in the lab experiment, this was not used alone for finding the optimal alternative for the tested system. Rather, it was used for the continued requirements engineering towards a more generally applicable approach to de-identification in real-time hospital visualizations.

Having collected data mainly as video recordings from the lab experiments, these raw data nevertheless had to be prepared before initiating the analysis. First it was necessary to transcribe all of the recordings, and this was done using a simple template created with ordinary word processing software. Screenshots from the different views of the prototype were also inserted in the text, to keep track of exactly what the words were spoken in relation to. Bold font markings were moreover used to highlight the words of the test leader(s).

The analysis then proceeded with reading through all the transcripts and marking out all segments that were directly relevant the research questions. These included comments on de-identification, the idea of having screens available public and semi-public environments, and the assessments of the de-identification approaches presented. A qualitative framework was then designed for the further analysis, in which quotes are categorized, so that main patterns and themes could be identified.

An idea of quantifying the participants' preferences towards de-identification techniques was abandoned, partially due to the low number of participants, and also because there was a qualitative depth behind the ratings that was more interesting to investigate further. After analyzing the results from the experiment, there was however qualitative data available for launching the refined design of a more interactive prototype, compared to the statical version used in the lab experiment.

3.6 Summary

This chapter has described the methodology that was relevant to this thesis, including the research process model that the project has followed throughout. After a literature review that revealed possible techniques for de-identification, it was chosen to proceed with prototyping and rapid field tests, which would be a combination well suited for quick assessments of the feasibility of several de-identification techniques. While these tests caused little impact on the participating health care workers, they also provided a limited depth of feedback. Instead, the breadth of insight gained was used to elicit a high-fidelity prototype for use in controlled lab experiements, also without causing the massive impact on a hospital environment. Now, a role-play session was also used to properly contextualize the prototype to them. Then directly afterward, the participants could provide their in-depth feedback through focused interviews. Finally, the raw data contained in video recordings would be categorized in the results analysis, revealing a need for further refinement of the prototype. Due to time constraints, this prototype has only been designed, and not tested.

"If you don't run experiments before you start designing a new system, your entire system will be an experiment!"

Mike Williams
Co-inventor of the Erlang programming language

Chapter 4

Experiments

This chapter describes how the design research moved beyond the literature study, by exploring the relevant de-identification approaches found in chapter 2, and also searching for new techniques that might be interesting towards the targeted context of use. Initial user requirements for accessing such a system through a de-identified interface are established through a series of rapid field tests, and the input gathered is analyzed and used for building a relevant prototype where the remaining candidate approaches can be tested. This high-fidelity prototype is finally tested through a series of usability lab experiments, where users first participate in a role-play and then provide their in-depth assessments through a focused interview.

4.1 Rapid field tests

In order to gain a broader understanding of the information needs of health care personnel, and especially their existing approaches to identifying patients, a series of six rapid field tests were arranged during November-December 2010. The tests were conducted at Trondheim University Hospital, and the test persons represented a range of professions and functions at the hospital — more details on this are provided in sections below.

Early paper-based prototypes were used to investigate whether the clinicians believed they would be able to tell patients' identities apart with the different identification approaches, and to see how these were related to current practices. The identification approach with the “highest” level of identification used initials and birth year of the patients, but the main focus was put on “less” identified approaches that aimed to identify the patients by e.g. their locations or relationships to health care personnel, possibly in combination with relevant health care activities. The tests were hence also used for assessing the feasibility of some new ideas to information disclosure through de-identification.

4.1.1 The paper prototypes

A total of four[1] different prototypes were explored in the field tests, all being graphical user interfaces for an application targeted at supporting hospital workers in self-coordinating their work. The prototypes included one prototype designed by the author, which is seen in figure 4.1, in addition to three prototypes designed by Ph.D. candidate Lillebo. The prototypes were designed to contain message examples related to the treatment progress of patients, e.g. “CT images taken”, “Blood test results ready: Sodium, Potassium” and “Patient has been scheduled for surgery”. The contents of the messages were either expressed using clear text, graphical elements, or a combination of both text and graphics.

The prototypes mainly differentiated on how information was organized and how the patients were identified. Two of the prototypes were also modified slightly in-between interviews, due to the feedback given. Although the prototype shown in figure 4.2 may appear more than a “slight” modification of the prototype in in figure 4.1, it is in fact only focused down to one of the many components in the original version. The headings with clinician initials represent a simple inversion of the approach where all patients are listed with reference to their responsible clinician, and instead it lists all patients relevant to each particular clinician. Finally, it was discovered and added an internally used coding technique that could make distinguishing the patients more feasible.

Some status messages were also added to prototypes during the process, but the refinement of such messages in particular is rather subject to Ph.D. candidate Lillebo’s research, and as a consequence, the prototypes designed by Lillebo are not included in this thesis. More on the information needs of healthcare workers in the perioperative domain can however be found in one of his articles (Lillebo, Seim & Faxvaag, 2011).

[1] This number is lower than the number given in the HFEHI 2011 paper (appendix B), because de-identification was not applicable to not all of the tested prototypes. See appendix D for further details.

The prototyping process started with an easily understandable approach to de-identification, since it was unknown to what extent clinicians are used to this concept. Two of the prototypes therefore included direct identifiers to support re-identification, having masked these identifiers in different ways. The main approach was here to disclose only the initials and either the birth year or age of the patients, which would at least be more secure than using full identity disclosure. Some new approaches were also explored in these prototypes, which are listed below, and were not found earlier in the literature review. Since the prototyping took place also after the field testing had begun, the two first of these approaches were both inspired by feedback given:

Treatment The treatment plan and history of a patient could be used by clinicians to re-identify patients they know, as also suggested in chapter 1. This could for instance be tested for patients undergoing an operation, and are hence expected to be transferred first recovery and then back to the ward. Revealing details on diagnosis and medical condition, such as what the patient is operated for, will however increase the damage potential in case the protection is broken.

Responsibility Relating the patient to a responsible clinician makes it possible for each particular clinician to quickly identify all patients that are relevant to them. Then, it will become easier for the clinician to distinguish between each patient case, since the data set population is narrowed down to a minimum. A risk involved is however that the population becomes too small, so that outsiders who can figure out relationships between clinician and patients could be able to guess patient identities more correctly.

Location Clinicians who operate in a limited area, e.g. a ward, may be aware of the current location of patients. Room locations are possibly stable identifiers, and should therefore be investigated further. The downside is that the room locations may be more easily sorted out by outsiders, especially by those who already have visitor access to the patient.

Non-presence While displaying current events is a straightforward approach to giving status information, it could also be possible to provide meaningful information by *not* displaying certain things. The fact that zero patients from a ward are undergoing surgery, would for instance implicate that patient Jack Jones is not undergoing surgery either. The possible utility value found in such approaches should be explored furher.

Not all prototypes were developed at the same time, but first two of the prototypes were made, and then one and finally one more, as more insight was gained. The prototype designed by the author was designed as the last one of the four, exploring some of these new ideas for de-identification techniques. More information on the design is therefore provided in the next section.

Description of the author's prototype

The prototype in figure 4.1 was made capable of demonstrating several approaches to de-identification at the same time. On the very top of the left column ("Blodprøver"), above the headline, a number is for instance given on the total number of blood tests analyzed at the hospital within the last 24 hours, as well as a trend indicator telling something about response time for all tests. This represents a technique where information is accumulated from a set, i.e. using the accumulation technique from section 2.5.1. The same approach is used in the topmost section of the center column, however from much more limited patient sets. Whenever an inpatient is discharged, the number of available beds would increment, and vice-versa when new patients arrive. This information can however not be used to re-identify individuals, but may still have a utility value.

Below the accumulative information to the left, updates are given on blood tests for a limited part of the hospital, e.g. a ward. The approach used is based on combinations of ordered blood tests, as well as the responsible physician (identified with initials) and the time for dispatching the blood to the laboratory. The idea was that physicians could recognize their own blood sample combinations supported by time. This is because different blood test indicators are known to be important for checking different medical conditions, and thus they are important to follow for being ready. Moreover, nurses could perhaps know which physician had ordered which blood tests, and therefore also infer which samples are ready for checking. The connection between patient and sample combination is hence established when the clinician adds their knowledge of which medical conditions require which specific tests. The matching can thus be done from sample combination to medical condition, and then add this, along with time and responsible clinician, to a particular patient — combining the approaches of both responsibility and treatment. This suggestion however relies on a certain minimum of variety in blood sample combinations, and that tests are being dispatched separately for each patient at different times.

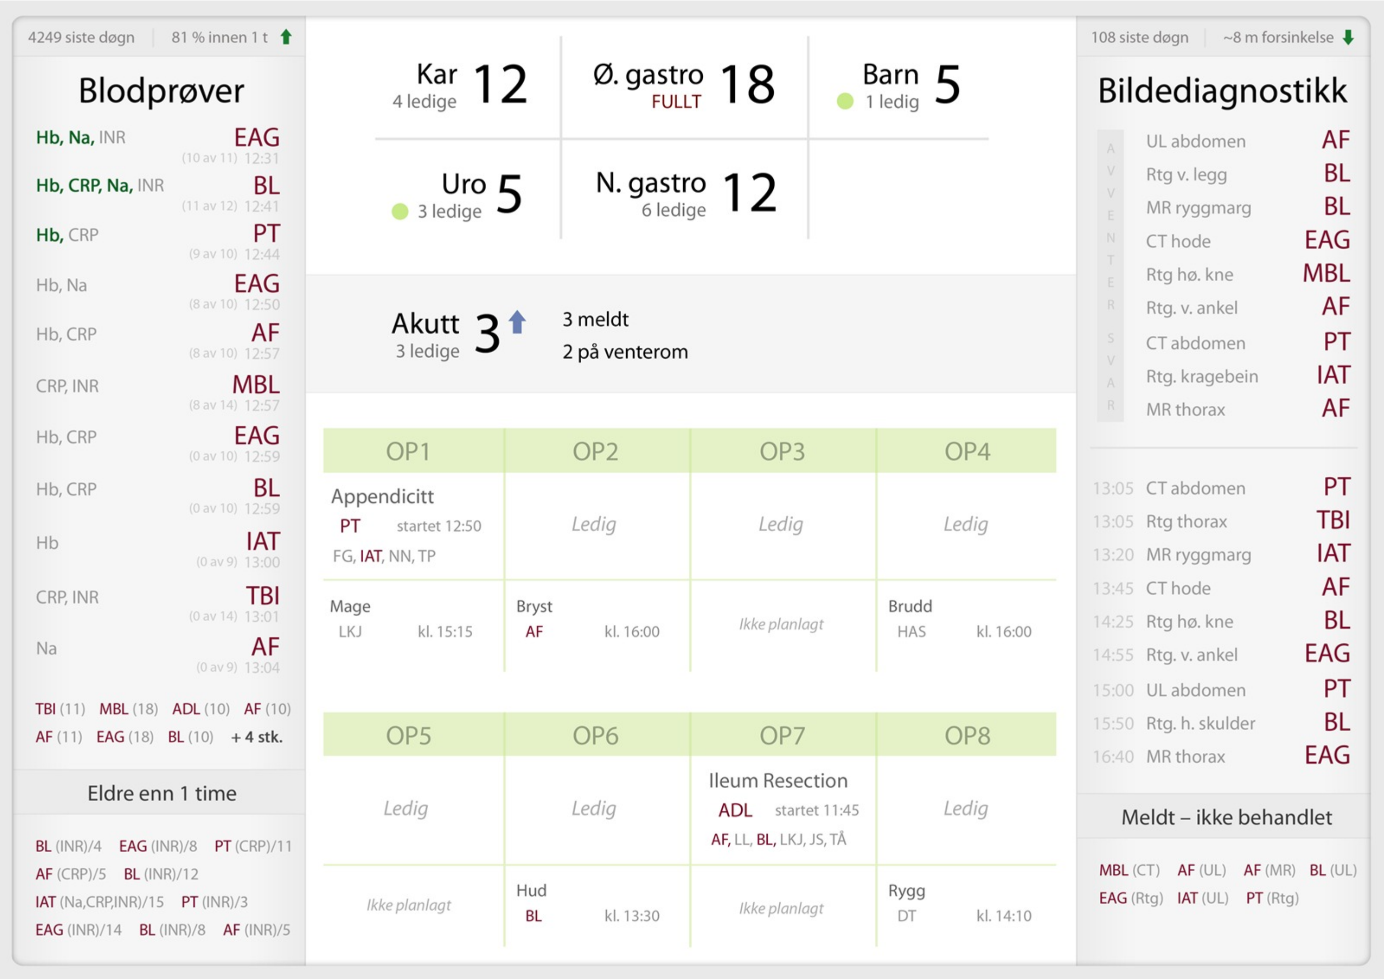

4249 siste døgn | 81 % innen 1 t

Blodprøver

Hb, Na, INR	(10 av 11) 12:31	EAG
Hb, CRP, Na, INR	(11 av 12) 12:41	BL
Hb, CRP	(9 av 10) 12:44	PT
Hb, Na	(8 av 10) 12:50	EAG
Hb, CRP	(8 av 10) 12:57	AF
CRP, INR	(8 av 14) 12:57	MBL
Hb, CRP	(0 av 10) 12:59	EAG
Hb, CRP	(0 av 10) 12:59	BL
Hb	(0 av 9) 13:00	IAT
CRP, INR	(0 av 14) 13:01	TBI
Na	(0 av 9) 13:04	AF

TBI (11) MBL (18) ADL (10) AF (10)
AF (11) EAG (18) BL (10) + 4 stk.

Eldre enn 1 time

BL (INR)/4 EAG (INR)/8 PT (CRP)/11
AF (CRP)/5 BL (INR)/12
IAT (Na,CRP,INR)/15 PT (INR)/3
EAG (INR)/14 BL (INR)/8 AF (INR)/5

Kar 12 4 ledige	Ø. gastro 18 FULLT	Barn 5 ● 1 ledig
Uro 5 ● 3 ledige	N. gastro 12 6 ledige	

Akutt 3
3 ledige — 3 meldt, 2 på venterom

OP1	OP2	OP3	OP4
Appendicitt PT startet 12:50 FG, IAT, NN, TP	Ledig	Ledig	Ledig
Mage LKJ kl. 15:15	Bryst AF kl. 16:00	Ikke planlagt	Brudd HAS kl. 16:00

OP5	OP6	OP7	OP8
Ledig	Ledig	Ileum Resection ADL startet 11:45 AF, LL, BL, LKJ, JS, TÅ	Ledig
Ikke planlagt	Hud BL kl. 13:30	Ikke planlagt	Rygg DT kl. 14:10

108 siste døgn | ~8 m forsinkelse

Bildediagnostikk

AVVENTER SVAR

	UL abdomen	AF
	Rtg v. legg	BL
	MR ryggmarg	BL
	CT hode	EAG
	Rtg hø. kne	MBL
	Rtg. v. ankel	AF
	CT abdomen	PT
	Rtg. kragebein	IAT
	MR thorax	AF
13:05	CT abdomen	PT
13:05	Rtg thorax	TBI
13:20	MR ryggmarg	IAT
13:45	CT hode	AF
14:25	Rtg hø. kne	BL
14:55	Rtg. v. ankel	EAG
15:00	UL abdomen	PT
15:50	Rtg. h. skulder	BL
16:40	MR thorax	EAG

Meldt – ikke behandlet

MBL (CT) AF (UL) AF (MR) BL (UL)
EAG (Rtg) IAT (UL) PT (Rtg)

Figure 4.1: The prototype (in Norwegian only) designed by the author, for testing several de-identification approaches.

In the green grid located in the bottom center, the hospital's operation theaters are represented each in one block. In the operation theater called OP1, there is currently an ongoing operation, started at 12:50, where the operator is identified as "PT". The anaesthesia nurses and operation nurses being present, are also identified by their initials. The important thing here, compared to the blood tests, are that in addition to the responsible people, the operation theater name can give a precise identification of the patient, provided that the user knows where the patient is located. Another hint is given in the provided operation type, which for OP1 is an appendicitis removal. This can both say something about the expected length of the surgery to someone who are not interested in the patient's identity, or if the user knows only the surgery type and not the room, this could be the key to re-identification — i.e. combining the approaches of both location, treatment and responsibility. There is also given an indication of what the next operation is going to be, but here it is demonstrated how a medical description may be generalized like Faxvaag et al. (2009) suggested, by providing a more abstract type ("mage", i.e. gastric) instead of a precise surgery description (e.g. appendicitis, but also many other operations).

Finally, in the right column, the patients are identified in very much the same way as for blood tests, however the link between description of radiology examination and medical condition may be a lot stronger than that from blood sample combination to medical condition. Additionally, there is normally a time scheduled for all imaging examinations, which can be both useful for re-identification of results as they arrive, as well as for planning ahead. Another concept introduced by this column, is the "non-presence" of results, which means that they are ready: If the clinician for instance knows that the CT of a patient's thorax (chest) is scheduled at 10:30, and that there is no entry visible for CT thorax at 12 o'clock, then the results are ready — otherwise it would still be present in the "awaiting results" area. The advantage here is that there is no retrospective display of patients' treatment history, and hence it would be very hard to abuse due to the very limited time frame results are visible in.

In figure 4.2, the two most interesting information boxes were presented in a less distracting layout, having the pending results grouped below a physician (initials *TROL*) and nurse (*HAPE*), and more clinicians could be added. In addition to the information already given about the imaging results in figure 4.1, there is also added a colored dot next to the time of two of them, indicating the degree of emergency which is already supplied by the referring physician today: Red means "As soon as possible", while orange means "within 6 hours". These dots represent a coding technique which

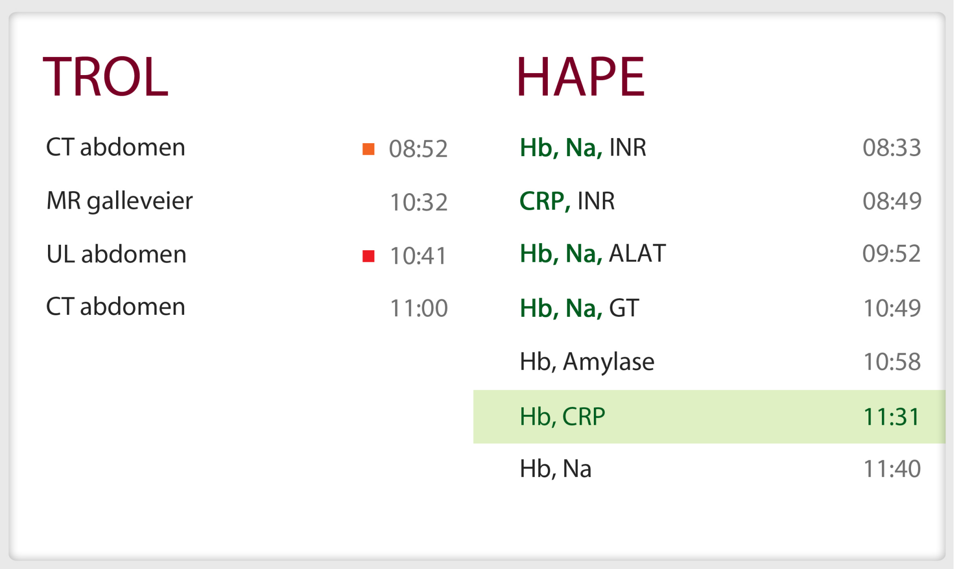

TROL

CT abdomen	🟧	08:52
MR galleveier		10:32
UL abdomen	🟥	10:41
CT abdomen		11:00

HAPE

Hb, Na, INR	08:33
CRP, INR	08:49
Hb, Na, ALAT	09:52
Hb, Na, GT	10:49
Hb, Amylase	10:58
Hb, CRP	11:31
Hb, Na	11:40

Figure 4.2: A more focused version of the prototype in figure 4.1

is currently used at the hospital where the tests were done.

Description of the other prototypes

Although not included with graphical representations in this thesis, the prototypes designed by Ph.D. candidate Lillebo were used for demonstrating some de-identification approaches not found in the author's prototype. Since they each contribute towards the results referred to in section 4.1.3, it is useful to have described them briefly with textual means, compared to not having any reference available.

The first prototype would integrate message streams from four different sources, and list each of them in their own box. One box would show a list of the last radiology examinations, and another would show a list of journal notes written in the EPR. The third would list all ongoing operations, and then finally the last containing blood test results. In this prototype, the patients would be identified with their initials and either birth year (or age) — as they both were tested towards the preference of the users. This version was shown to a total of three informants in the rapid field tests.

The second prototype was built around the physical layout of wards in a department, containing blocks of squares, each representing a patient (room). This represents the location approach described previously, as well as being a coding technique.

The idea was also that the user could zoom in on the ward of interest, but this was not implemented in the prototype. This prototype was shown to two participants during testing.

The third prototype had organized the messages vertically below a heading containing identifying information on the patient they concerned. There were a total of eight patients included, having each their own column that were arranged horizontally in a landscape layout. The heading was designed to possibly contain a variety of identifiers, so that it could either be de-identified or fully identified depending on the user's preference. A tested approach here was to use room numbers, being an approach to pseudonymization. This prototype was also evaluated by two participants.

Participants

The participating clinicians included one senior physician and two ward nurses from the Department of Gastrointestinal Surgery, one junior physician and one nurse from the Department of Emergency, one ward nurse from the Department of Breast and Endocrine Surgery, and one charge nurse from a ward at the Department of Orthopedic Surgery. Their ages ranged from 25 to 55, and all had been in their current positions for some while. Their identities were not collected, but for the sake of their true anonymity, both the ward nurses and the charge nurse will be referenced to as nurses when being quoted in this thesis, while both the junior physician and the senior physician will be referred to only as "physician".

All informants were randomly recruited in their respective departments during working hours. The recruitment process was integrated with the testing, and it was hence necessary to walk a bit around in the surgical departments and see if anyone showed up with some time available for the test — right there and then. Since Ph.D. candidate Lillebo is an anesthesiologist and has been working in the hospital's emergency department, he could wear his regular working clothes and ID-card during testing. While they already expected him to understand, the informants did not have to explain the, for them, basic stuff, for instance what the different systems they use are used for. This may also have contributed towards minimizing the field testing's impact on their primary duties.

4.1.2 Test design

When informants were found and agreed to participate, the test was initiated. After introducing himself and the author as a master student, Ph.D. candidate Lillebo would briefly describe the COSTT project and the idea behind the application he was designing, while one of the prototypes already would have been given to the informants from the start. They would then be asked to comment on the different prototypes that were presented to them, however only one or two prototypes were presented to each informant, and only one at a time.

A total of six different approaches to (de-)identification of patients were explored in the tests, which are found in the overview in table 4.1 below. The informants were each asked about which of the suggested messages would be useful in their work, and whether something that could potentially be useful had been missed out on. Then, it was investigated whether the patients' identities could be told apart with the different identification methods suggested, and how these related to current practices and communication of patient information today. The security and privacy relevant output that was generated from the rapid field tests can thus be mainly divided into three categories:

1. Could the prototype's approach to patient identification be usable to the informant?
2. How are patients being identified, de-identified and re-identified by clinicians in a hospital today?
3. What data are disclosed (on whiteboards, in conversations, etc.) in areas where outsiders may still gain unauthorized access to the information?

When the idea behind the prototype had been communicated, and an interactive talk about its layout was established, it was possible to follow up on the comments with questions that regarded privacy and methods for identification. The test was normally ended when both parties had emptied their desires to explore the prototype, however the tests never lasted more than 30 minutes. In one case the clinician was needed elsewhere and the session was ended quickly. This did however not occur before the review of a second prototype had already started.

The interviews were recorded with handwritten field notes, and written out directly afterwards, both by the author and Ph.D. candidate Lillebo. As a quality assurance, the transcriptions from both were then finally compared and agreed on.

Approach	**Example**	**Summary of responses**
Initials and birth year of patient	JD59	Will normally provide fairly good accuracy. Patients having the same birth year and initials (or last name) do however occur. Clinicians still found this convenient as they are used to working with basis in the patients and their name/age (various combinations of name and birth date are used for identification today).
Room number/location	(Plotted on a map of wards)	Patients move around (this may leave room lists temporarily inconsistent) or they can even be placed in the corridors. Room numbers are commonly used for reference today, but in combination with other identifiers, e.g. name, diagnosis or sex. It seems hard to remember the patients' exact locations.
Initials of responsible physician (first two letters of both first name and last name)	DAJO	Patients are not followed up by only one physician, and physicians attend many patients at each ward. Nurses will not necessarily know the name of the physician providing care for each of their patients at a specific time.
Blood test indicators, time and responsible nurse	Hb, Na, INR - 10:41 (HAPE)	Blood tests are ordered as standardized batches, so important indicators, if any (e.g. INR may decide whether to operate or not), do not stand out. Tests for several patients are often ordered at the same time, and by the same nurse, too.
Radiology type, level of urgency, time and referring physician	CT abdomen (red) - 11:00 (DAJO)	Some results (MR) take days to arrive, and often 20-30 patients with abdominal pains arrive daily. Hence, a list of pending results may become overloaded and hard to interpret.
Operation room number, surgery type, scheduled time and surgeon initials	OP3: Appendicitis - 11:00 (PT)	Nurses rarely know exactly what room an operation will take place in. But as it is uncommon to have several patients from the same ward undergoing surgery at the same time, they may still be able to deduce which operation to follow.

Table 4.1: De-identification approaches explored in the interviews, based on the paper-based prototypes.

4.1.3 Results

The following results of the interviews are also summarized in table 4.1. In general, clinicians were positive to the idea of integrating status updates from several systems. Most were still reluctant to the immediate thought of placing any patient information more publicly available than workstations or personal devices. Though the approach where patients are simply identified by initials and birth year stood out as the most convenient option, the main impression is that health care personnel have varying needs for patient identification, depending on their role and the context where identification should happen.

Combinations of person identifiers that are independent of the context, e.g. name, birth date, social security number (SSN) or even a picture, were all suggested both as preferred means of identification, and the most reliable identificators. At this stage however, our prototype contained only the patients' initials and birth year (or age):

> *Physician*: "When I see initials and age here, it's looking all right. [...] We only handle a handful of patients at a time, so duplicates have never been an issue."

Although using initials and age proves sufficient in some situations, a nurses explains that they often require more precise identifiers for their patients. The nurses are used to having precise identification available for all their patients, and carry and use patient lists actively that both contain room number, full name, birth date and hospitalization date for each of the patients they have responsibility for.

> *Nurse*: "We often have patients here [at the ward] with the same last name and birth year. We use full names and at least part of the birth date."

This will however apply mainly to situations where it is crucial for patient safety to be certain of the patient's identity, e.g. when giving medicine or preparing for surgery. On whiteboards that are at least partly shielded from public insight, it is common to use initials and birth year. Such whiteboards also seem widely accepted for the general reminders of who is who, and keeping track of where they are.

The second approach was to use patient rooms as a key to identities. The prototype had events placed over a map with square boxes representing the rooms in a ward. The idea was that if clinicians have a particular desire to follow up on a patient, they will be able to tell exactly where the patient is located. The interviewed clinicians were however all negative to the idea, which would require them to remember patient locations in general. They argued that their attention already has to be focused on more important things, and did not trust their own memory when faced with a location indicator on a screen:

> *Physician*: “I wouldn’t dare trusting just a box!”

Similarly, using room numbers as the only identifier was neither considered an option either. In addition to the personnel’s memory limitations, patients often change rooms during their stay, while new patients arrive. When this happens, updates to the room list may lag significantly behind, risking in the meantime that status updates are associated with the wrong patient. Room numbers are however commonly disclosed in combination with other identifiers, e.g. on room list whiteboards.

On the other hand, it was discovered that clinicians commonly used patients’ diagnosis or treatment history as de-identification in conversations between colleagues, both in combination with sex or even alone. It was said that remembering the diagnosis and history of treatment is easier than remembering both the name and precise birth year. An advantage that consequently appears when clinicians talk together about patients in public spaces, is that they can always start on an abstract level and go further and further into details, until adequate identification is achieved. This interaction could represent a real-world approach to flexible re-identification:

> *Nurse*: “When discussing patients in the hallways here, we try to refer to them by using their diagnosis.”

> *Physician*: “It is often easier to remember ’he with ileum, who needs another surgery in three days’.”

Ultimately, these results were used to generate a prototype of an application that could be used for testing the remaining techniques more in-depth.

4.2 Usability lab experiment

The combined results from testing in the field had provided enough feedback to build a prototype of an application which would take advantage of taking its messages to more public areas. The prototype can include a variety of such important event update messages, and has the flexibility to be presented in both a fully identified and also in a de-identified manner. During January 2011 this prototype was tested in an experiment in the usability lab at the Norwegian Research Centre for Electronic Patient Records (NSEP), involving four nurses from Trondheim University hospital as test persons.

The experiment was set up as a role-play following a realistic scenario, aiming to simulate a setting that the test persons are used to, and thus being familiar for nurses working at that specific hospital. After getting to know how a single de-identification alternative would work in practice, they could each give feedback to five additional alternatives in a focused interview afterwards. The experiment provided many qualitative considerations for each of the presented alternatives, in addition to a final ranking of the highest valued ones.

4.2.1 The prototype

For the purpose of rapid field tests, there was no need for realism in the event messages that were included in the prototypes. For the lab experiments this realism was however a vital ingredient, and the prototype was thus populated with a realistic chain of events, all corresponding to the scenario and patient list being used for the role-play — a detailed description of the scenario is found in appendix E, while the patient list is attached in appendix H. The messages were also carefully chosen to show the participants users a variety of possibilities, and to support them in becoming quickly confident with the system. At the same time, the scenario developed towards the end with several new events arriving at the same time, making the updates more challenging to handle.

In order to test the relationship between de-identified disclosure of information and full disclosure, two different screen views of the same application were also created — the *desktop prototype* where all details are revealed (figure 4.3), and the *hallway prototype* where all messages had been de-identified (figure 4.4). A typical scenario would thus be the hallway prototype indicating “X-ray description available”, while the desktop prototype specified this as “Radiologist’s X-ray description available for patient Odd Hansen”. While the desktop prototype also provided more functionality than the simple list of event messages, it was designed by Ph.D. candidate Lillebo and hence only an excerpt containing this particular list is included in the thesis for reference. A complete view of the desktop prototype can nevertheless be found in the paper in appendix B.

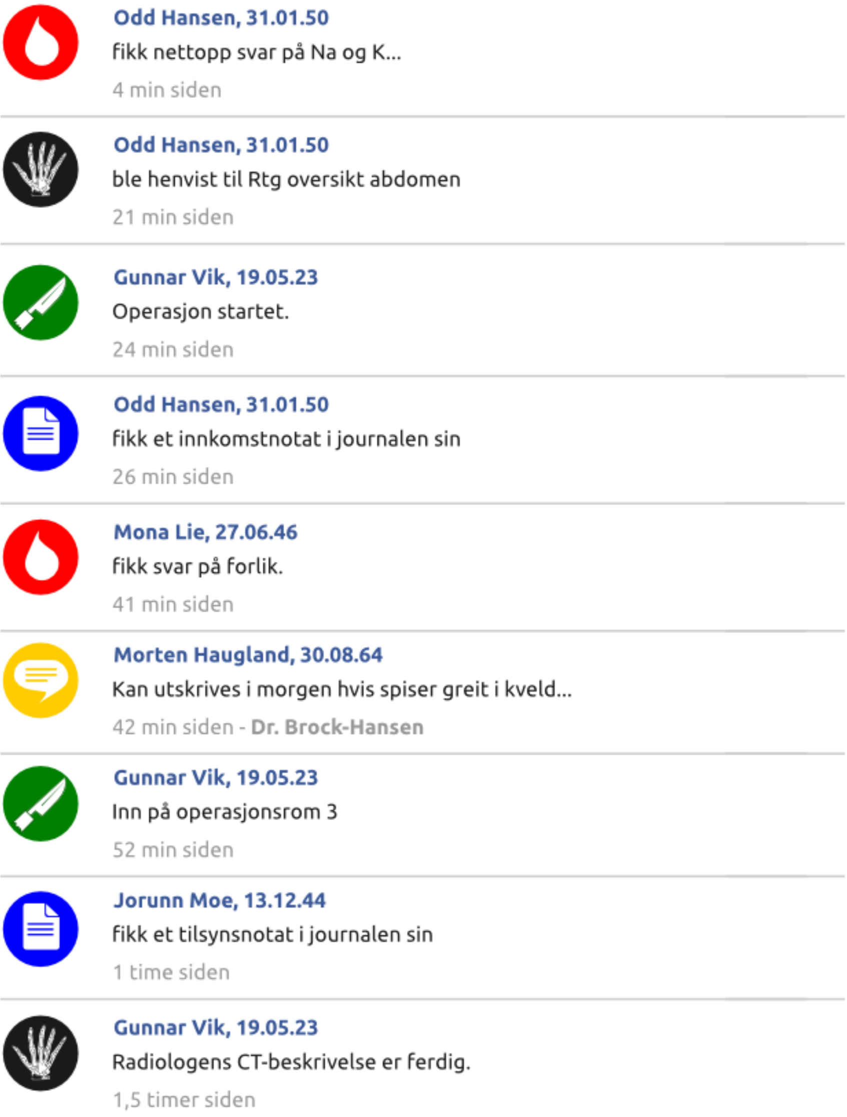

Odd Hansen, 31.01.50
fikk nettopp svar på Na og K...
4 min siden

Odd Hansen, 31.01.50
ble henvist til Rtg oversikt abdomen
21 min siden

Gunnar Vik, 19.05.23
Operasjon startet.
24 min siden

Odd Hansen, 31.01.50
fikk et innkomstnotat i journalen sin
26 min siden

Mona Lie, 27.06.46
fikk svar på forlik.
41 min siden

Morten Haugland, 30.08.64
Kan utskrives i morgen hvis spiser greit i kveld...
42 min siden - **Dr. Brock-Hansen**

Gunnar Vik, 19.05.23
Inn på operasjonsrom 3
52 min siden

Jorunn Moe, 13.12.44
fikk et tilsynsnotat i journalen sin
1 time siden

Gunnar Vik, 19.05.23
Radiologens CT-beskrivelse er ferdig.
1,5 timer siden

Figure 4.3: Excerpt of the desktop prototype developed by Ph.D. candidate Lillebo (available in Norwegian only). Used with permission.

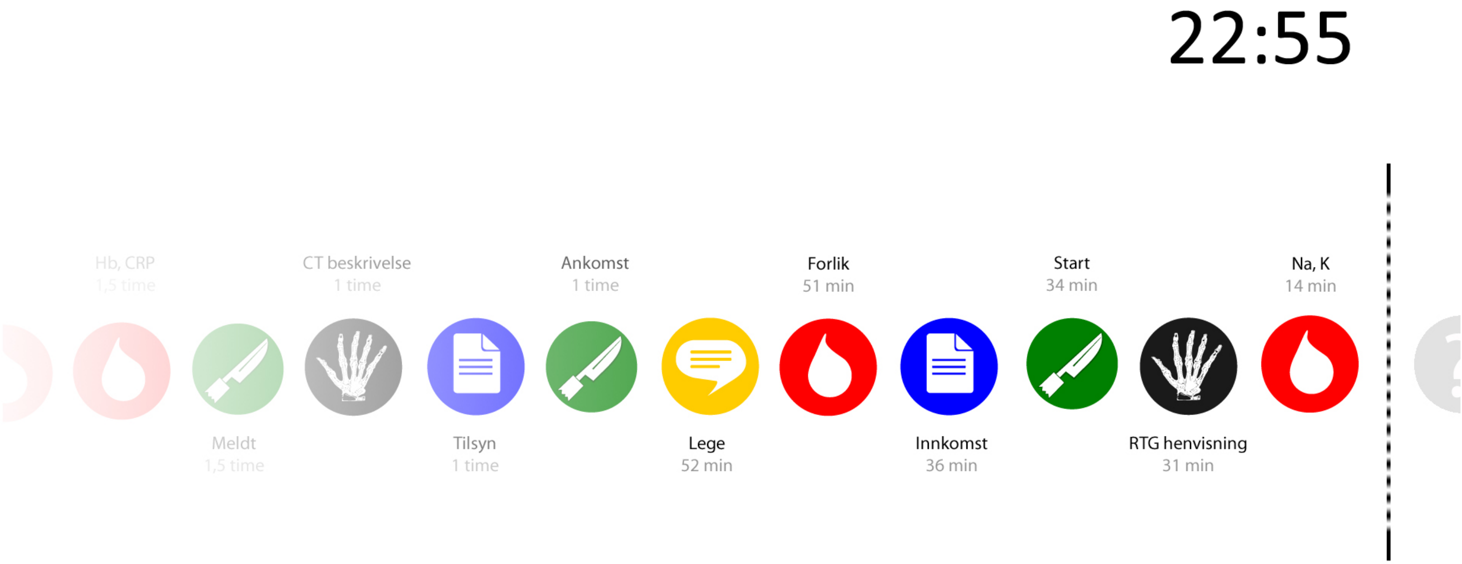

Figure 4.4: The de-identified version of the prototype, now horizontally oriented and optimized for large wall-mounted screens.

Apart from the colored graphical icons, the hallway prototype was still more or less designed by the author, as it had to be directed horizontally instead of vertically, and therefore express more clearly to the users in which direction it should be read. Moreover, space and readability did become an issue here, as opposed to the desktop view, since the text labels would be hard to distinguish and read if they were all placed on the same side of the icon, or alternatively in a vertical direction. A gradient fade in from the left was chosen to bring out the new events arriving from the right where a dotted line represents the current time, and on the future side of the line a new event is hinted in gray with a quotation mark. For readability and extra space, the text labels are positioned above every second event icon, and below for the rest. A large clock indicating the current time is also included, for increasing the test users' awareness of time passing by during the role-play.

De-identification alternatives

The usability lab experiment would naturally play the most crucial part in exploring de-identification for the targeted domain. It was therefore accordingly important to carefully select de-identification techniques for the test, in contrast to with the rather preliminary rapid field tests. In the light of these field tests, it would hence be important to evaluate the collection of possible approaches found in the literature, as well as those suggested earlier in this chapter. This evaluation of techniques is described in the following paragraphs.

Variable suppression could in this context rather be considered a question of which variables to *include* in the visualization, since the organization would already possess the complete data set. This would also be the most important part of making the contradiction described in chapter 1 possible, namely to disclose information while not disclosing it to those who are not authorized recipients. Since the information requirements for the application in general had been estalished, it was therefore a matter of deciding which of these variables to include.

Record suppression would on the other hand not be well suited for use in the prototype. Since usability requirements make it important to have conformity between the timeline of events in both the identified and the de-identified prototype, it would be problematic to completely remove individual events from only one of them. It would hence be better to use additional variable suppresion, so that the utility value of these particular records are protected.

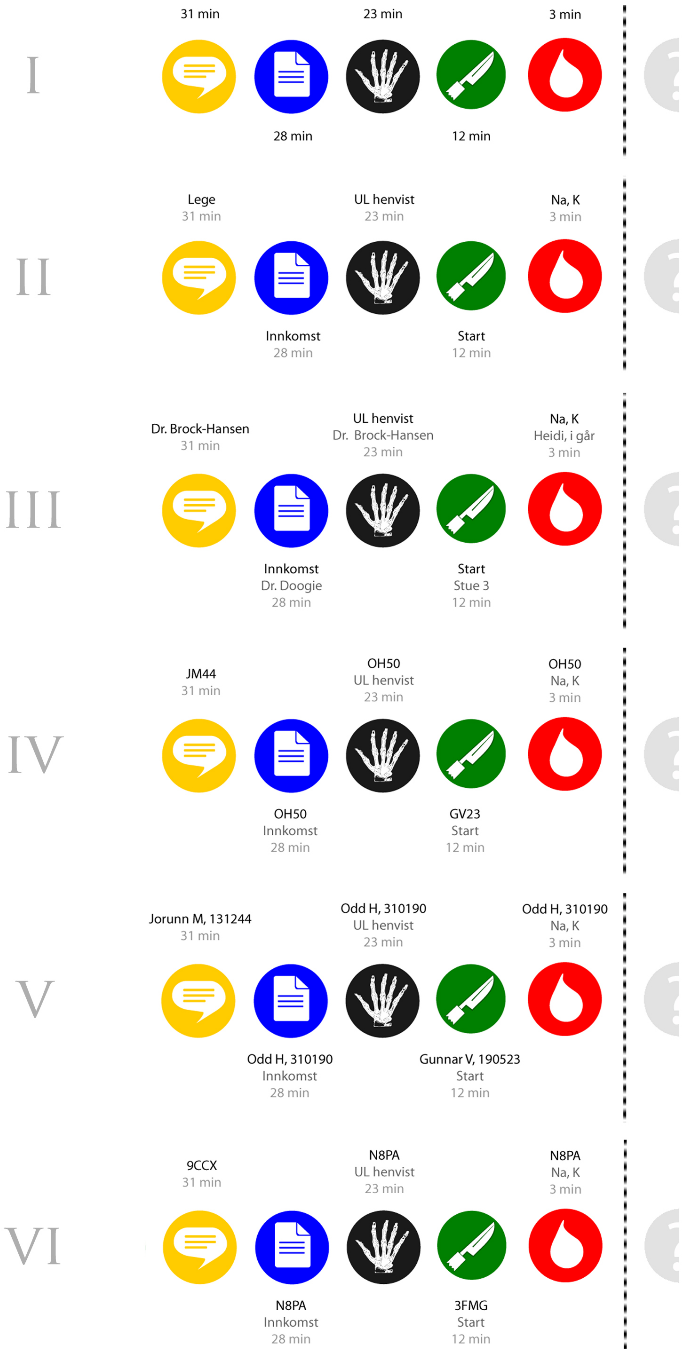

Figure 4.5: Overview of the de-identification candidates presented in the interviews. Modified for presentation (only one event from each category is included).

Coding and masking would also be important techniques, especially the masking approach used for names and birth dates in the field tests. A possible use of coding in the prototype could for instance be pseudonymization, which is reversible. The event icons themselves can also be considered as coding, yet not adding considerably strong protection. Irreversible coding would on the other hand not be useful for coordination purposes.

The generalization and aggregation techniques were both briefly tested in the rapid field tests, while their applications did not yield much feedback. For generalization this could however be since its usage appeared rather covertly, e.g. in the list of radiology examinations where it was said which main type of examination was done, but not specifically what body parts or organs were examined. Hence, this technique will also be useful for not being *too* specific in the radiology event messages, so that outsiders can guess which events concern the patient they know. While generalization for patient ages also could be possible, it is not taken any further for this application. The reason is that the rapid field tests revealed the clinicians' clear preferences towards using either birth year or full birth date, and resistance against expressions that would appear in a confusing manner.

Finally, the pixelation and blurring techniques may not seem entirely suitable for use in the prototyped application. Although one of the field test informants suggested using pictures for identification of patients, there are no images available to de-identify here. It would moreover be pointless to scramble any of the other graphics, since clinicians would not be better fit than other to interpret the hidden meaning.

These considerations have altogether resulted in a collection of six different alternative approaches. While one is pseudonymized and one appear to be anonymized, the remaining three represent different levels of information granularity in a de-identified approach, aiming at its legal definition. It can still be argued, however, that these alternatives are indeed de-identified, since two of them even include direct identifiers — although being masked. All alternatives do nevertheless utilize the treatment approach, as suggested above the rapid field tests section. In addition, the responsibility approach is represented in alternative III, along with what is left of the location approach to use — the operation theater where a patient undergoes surgery. The non-presence approach was on the other hand not found to be useful in this application, and quite possibly not in general either.

It was nevertheless early decided to only use one de-identification alternative consistently through the role-play (alternative II, see description below), in order to make the prototype easier for the users to understand and follow. Therefore, it was necessary to de-identify the initial timeline with the other five alternatives as well. A large paper sheet was hence created where all six alternatives were laid out, i.e. six different versions of the view in figure 4.4, so that the test persons could compare the alternatives altogether. The overview containing all six alternatives can be seen in figure 4.5.

Alternative I is considered anonymized, which is beyond de-identified, as it contains no information except for the type of event (the icon) and the time it was generated. It is included for a minimal representation of what an indisputable interpretation of the law may allow.

Alternative II can possibly be considered as de-identified, as it simply adds a description of the event to give more info about what is new. The idea is that users who expect something to happen for patients and know their history of treatment, in general may be able to know who the patient is when that expected event is happening, or when a foreseeable next step is taken in treatment. This event description line is also included in all alternatives below.

Alternative III still only adds more information, and still nothing that can be used to identify a person in general, but instead more related to the patient's history of treatment. The detail that is added is either the name of the clinician who has ordered the imaging diagnostics, the clinician who has taken the blood tests, the clinician who wrote an EPR note or a comment, or the location where the operation takes place.

Alternative IV has identified patients with their initials and birth year. This is by no means a bulletproof method for neither identification nor de-identification, but clinicians are used to this from whiteboards etc., and was the preferred approach from the rapid field tests.

Alternative V adds the patient's first name, as well as the first letter of the last name and the full date of birth, while revealing the first name will also reveal the patient's sex. The alternative is given for providing a close to unambiguous identification alternative of the patient, while not displaying last names that could be revealing for e.g. famous persons.

Alternative VI has a rather cryptic-looking approach, because the four character code does not represent anything other than a random code. The code is nevertheless suggested to be included on the patient list, so that it will be brought along anywhere by the clinicians. Then, a code lookup can be done whenever required, by those who can access it and thus are authorized recipients of the information.

4.2.2 Experiment design

It had been a bit challenging to interview people about de-identification in the rapid field tests, because the informants were not given a chance to properly imagine the de-identified event messages in action — they were just plain messages without any patients to associate them with. In the usability lab, it was however possible to simulate a realistic scenario in which the de-identified screen could become useful, and the layout of the lab made it possible for the test persons to get a good understanding of how the system can be used in practice during their regular work day.

Since the experiment did not involve any sequences where the actual usability of the application's interface was tested with specific tasks, the prototype did not have to implement any interactive functionality. What was mainly left could compare to a very simplified "Wizard of Oz"-setup, where the screen views only consisted of a sequence of images that are to be manually skipped forward in, as time goes by in the scenario. The sequences did not pose other changes to the user interface than the addition of new events to the timeline, and there was also nothing the test user could do to change the line of events.

The role-play scenario was driven forward by having these updates appearing in the prototypes due to skipping in the sequence. The stage was however first set with an important routine that every nurse at the local hospital is well familiar with. The *handover* (termed "rapporten" by the Norwegian professionals), is a meeting which takes place three times every day at each ward, i.e. whenever a team of nurses replaces another. The nurses going off duty will brief the ones who replace them, on each of the patients for whom they will have responsibility. At this stage the test person was given a patient list ("pasientoversikten"), which is a real document brought to these meetings (see appendix H for the fictive document we handed out to support this routine). The list describes the patients who are being handed over, and the incoming nurse will take notes here regarding each patient during the meeting, as witnessed by Munkvold, Ellingsen & Monteiro (2007). An artifact similar to the patient list is also described in Bardram (2005).

In our experiment there were eight patients on the list, providing a full bed cluster.

A bed cluster is a subset of a ward, being a group of up to eight patients, that a single ward nurse in reality will be responsible for during the night shift. During the other shifts the nurses will however have responsibility for a smaller group of patients than this. This exact group size was nevertheless chosen since it both represents a real-world grouping, as well as possibly being could be a suitable population to both infer and protect identities from.

Participants

The patient cases used in the scenario were all inpatients for gastric surgery, since this is the largest surgical department at the hospital having two large wards. The role play was thus designed for ward nurses to play, and so the four test users were all ward nurses with experience from work in the department for gastrological surgery, although one was at the time working in the emergency department.

Choosing nurses to participate in the experiment was still not only motivated by their familiarity with the scenario used. By making use of their professional experience, it was possible to check their professional understanding of the information that was provided through the two views, and how this would enable them to act in a real situation. This explains why they were not recruited to test the usability of any interactive functionality yet, but rather make it possible to gather more in-depth data on how this kind of application could work in the field. The same test could therefore not have been accomplished with non-professionals, although many people very well could say something about the relationship between identified and de-identified messages in general.

Location and equipment

In order to record data from the experiment, a usability laboratory was made available to the experiment by NSEP. The test set-up was early aimed at taking advantage of the laboratory's size and flexibility, which geared the experiment design towards simulating a real ward like the ones being familiar to our test persons. Figure 4.6 shows how we could start the session by sitting around the table, before moving over to the workstation computer at the desk (marked in blue). The desk is faced with the hallway in front, to which the rooms in the bed cluster are attached. Having a patient room on each side, so that the role play could involve moving in front of the desk area, added realism to the pretended ward work, instead of only having a single room on the one side. This way, the test persons could properly grasp the point of having the hallway prototype (marked with green) placed where it was, as they could effortlessly check it while passing through in the corridor.

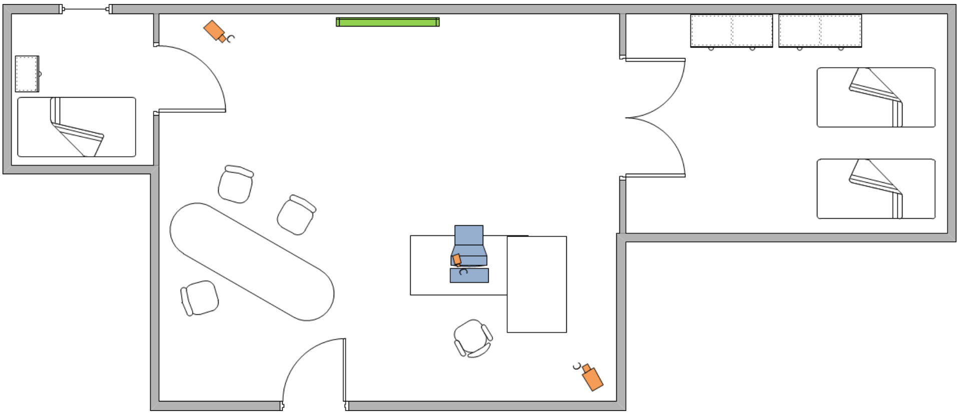

Figure 4.6: Map showing the approximate layout of the usability lab, as it was used in the experiments.

The lab was already equipped with three video cameras (marked in orange), two of which were movable at all times by the lab technician, and a sensitive microphone that could capture speech from anywhere in the room. All this recording equipment was controlled from a separate control room by a lab engineer, and the video output was recorded on a computer there, see figure 4.7. The video mix also contained a live screen capture stream from the workstation computer used in the experiment, where the test users' mouse movements could be seen. This way, it would be easier to know what they pointed at when speaking, and hence improving the quality of the transcriptions written down later.

Preparations

Two pilot tests were run before the recorded results were gathered. Initially, a fellow researcher with high expertise in the human computer interaction (HCI) field, served as test user for the first pilot. He has also worked with health informatics issues, and had some insight into both the handover meeting and the daily doings of a ward nurse. The first pilot test was run without recording video and audio, and primarily aimed at quality-assuring the methodological side to the experiment, first with the role-play, and then with the focused interview afterwards. It could hence just barely be called

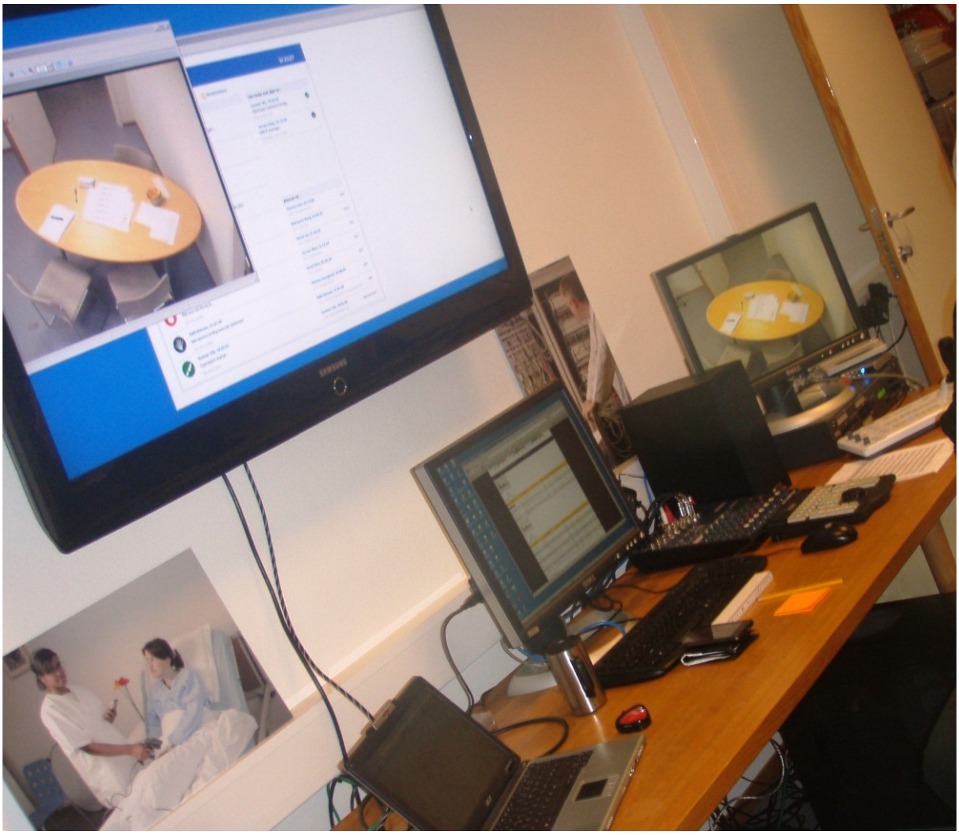

Figure 4.7: Equipment in the control room for the usability lab. The large screen shows a mix of the camera input and the active screen contents viewed by the user.

a real pilot. Some adjustments were still made based on this session, including an increase from three to four rounds of status updates when passing by in the corridor. This was done in order to not cut off the test users from the system before they got a good feeling of mastery, and thus would be satisfied when the experiment then proceeded to the interview stage. Additionally, the transition between the handover and the first interaction with the workstation had to be improved, from giving the user a task immediately, to giving the user some time to get to know the user interface first.

The second pilot test was executed with a test user who had previously worked as a nurse, but at the time was working as a researcher within the field of health informatics. She was familiar with the context of the nurse handover, and knew how she would have organized her own working day in the scenario we presented to her.

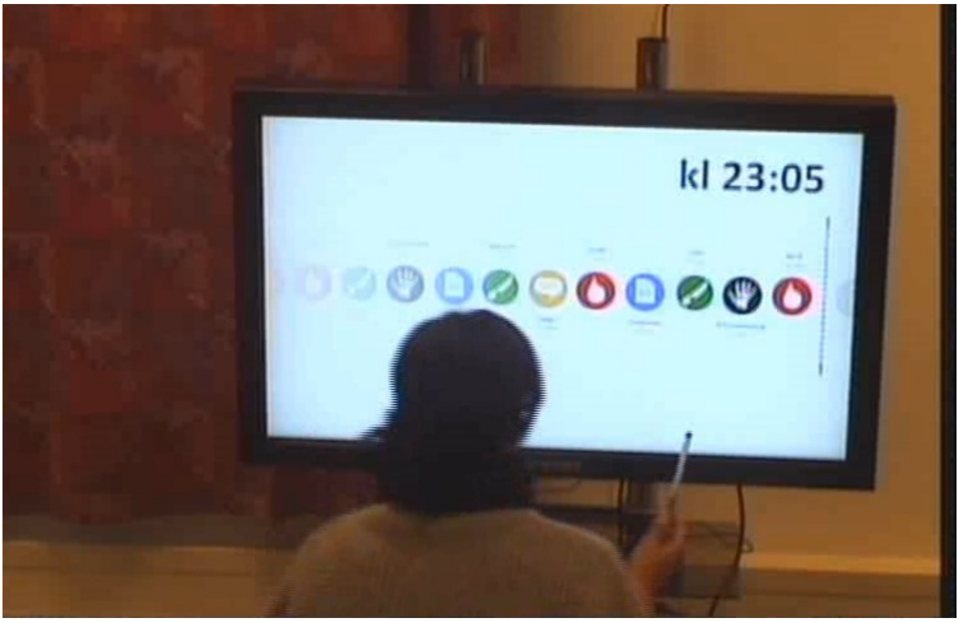

Figure 4.8: The hallway prototype in action on the large screen.

The test became a realistic check on the final test set-up, as is was in every aspect equal to what the recorded results are based on. This test also had better a flow in transitions between its different stages now. A full video and audio recording was performed so that also the lab engineer was prepared to follow the subsequent experiment runs properly.

Procedure

To begin with, the test persons were welcomed and offered some fruit and coffee along with a letter describing the experiment, as well as signing an informed, written consent to their participation. An introduction to ourselves and the project followed, with an explanation of what kind of data would be gathered and how the session would be recorded. The participant was also informed about the possibility to end the experiment at any time. The layout of the test lab and important equipment was then presented, and the test person was given instructions on how to “think aloud” during the role-play phase of the experiment. The application prototype that were subject to the test was briefly introduced, with a note on its limitations as being a non-functional prototype. Finally, if the participant had no further questions, the role-play would be initiated with the fictive handover meeting leading in on the night shift.

After the handover was completed, and appropriate notes had been taken by the test person, the nurse was guided over to the desktop workstation, which displayed the initial screen from the desktop prototype. After exploring the interface for some while, the test person was instructed to pretend, due to the fact that the patient rooms were empty, to be visiting patient rooms and carrying out regular ward work, and that time meanwhile went by. When returning from the patient room, the test person was told to take a look at the hallway prototype when passing it in the corridor, and tell whether or not something new had happened lately. When doing this for the first time, no such new events had occurred since the desktop prototype was used. It was however necessary to allow the test persons some time here to familiarize themselves with the new layout of the update messages, as this was now different compared to the desktop prototype, although the messages were entirely the same.

More pretended ward work was then done for another 10 minutes of pretended time, in reality a few seconds, before returning to the hallway prototype. The screen now indicated that new patient events had occurred, and from now on this was also the case every time when passing by the screen. The participants were at this stage asked whether something new had happened, and if they could tell who it had happened to. After the first round this was however not always necessary, as the participants completed both tasks without being prompted for it. After four such information updates separated by simulated ward work, approximately 1.5 hours of simulated time, the role-play was ended. In reality, 20 minutes of role-play plus ten minutes of introduction, had passed at this point.

Through the subsequent interview the nurses were then asked to validate their information needs and imagine how the prototype could influence their work, while exploring whether the hallway prototype in particular would be understood and valued in real work. The feasibility of each given de-identification alternative was assessed and the most valued candidates ranked from the users' individual perspective.

The experiment was finally ended within one hour after the test person's arrival.

4.2.3 Results

As already mentioned, the de-identification alternative used in the experiment was alternative II (figure 4.3). Until stated otherwise, the participant citations are also concerned with this alternative in particular.

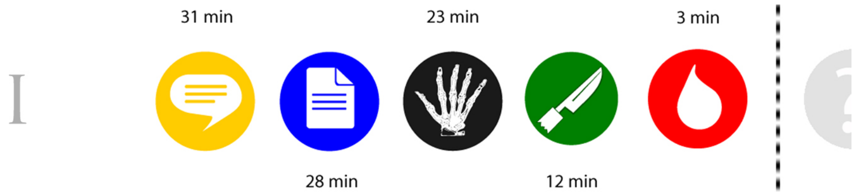

Positive	Negative
• Surgery events would have been interpreted correctly [TP2] • Would have checked all events no matter the level of detail anyway [TP1]	• Imaging events etc. would easily become indistinguishable [TP1,TP2] • A lot of “checking” would have to be done [TP1] • Would not recognize blood test results [TP2] • Users would face a steep learning curve [TP3]

Table 4.2: Feedback on events de-identified with alternative I

The first observation made, is based on the theory presented in chapter 1, that giving the users an idea of what is happening, but not explicitly who it is happening to, requires them to already have an understanding of what might happen in the near future. After doing the lab experiments, it appears that ward nurses do in fact have a certain understanding as such, since they were all able to interpret some events correctly in the test:

> *Nurse 1*: “Yes, now someone has taken an X-Ray scan, it is probably he who was waiting for that overview, Hansen. And I will go and check the [radiologist’s] description, because I can see that the description is ready too. And the remaining of Hansen’s blood tests are also ready...”

> *Nurse 3*: “Yes! Now I see that there has arrived a result for troponine. And the haemoglobine to him [Hansen].”

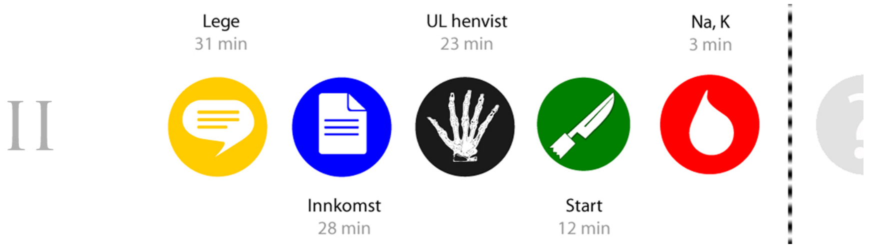

Positive	**Negative**
• Privacy for patients is well protected [TP1] • Only one patient would undergo surgery at each time [TP1] • You often known what you're waiting to happen [TP4] • Clean layout, not overloaded with text [TP3]	• Must know the patient to know what to expect [TP4] • Comment events always have to be checked [TP1]

Table 4.3: Feedback on events de-identified with alternative II

> *Nurse 4*: "Mmmm.. I see that he [Hansen] has been scheduled [for surgery] 5 minutes ago."

In the interview afterwards, the test persons could each elaborate on how they re-identified the patients concerned in the de-identified messages:

> *EAG*: "Did you then feel confident that you knew which patient it had was concerned with?"
> *Nurse 3*: "Yes, I probably was. It was mostly the one patient that things happened to... But then I knew that I was expecting the troponine, and I knew that I was waiting for the X-Ray description also."

> *Nurse 1*: "If we're talking about eight patients, like in our bed clusters, I would have understood who the event was concerned with. [...] We

> never have two in surgery at the same time, so the only thing would be the comment events, which would be hard to figure out.”

> *Nurse 2*: “If patients from the whole department showed up, I wouldn’t have a chance to follow up. But as long as we’re limited to 8 patients...”

Since the scenario was staged, the test persons were also asked to confirm whether the chain of events presented was realistic compared to a normal day at work. The response was that this could in theory be true, and the scenario was even more difficult to the test persons than reality, being introduced to eight new patients at once. On the other hand, the patient cases were considered a bit too diverse for a normal mix, which often contains some patients having the same diagnoses:

> *Nurse 2*: “Often we have several with the same problem, but often they do not arrive at the same time and do not follow the same trajectory, so I think I could manage with that [alternative II]. Perhaps not on the morning shift, because then we take all the blood samples early in the morning. And if it then was only said ’sodium and potassium’ for a patient, I would have had difficulties with knowing which patient it is, because we test that on almost all. At least liver and bile samples, and those kinds.”

There were still differences between the different event types used in the prototype. While radiology messages (black) and operation events (green) would be feasible to follow with alternative II, it as nurse 2 stated be harder with blood tests (red). These events are the most difficult to present in a de-identified view, in line with what was already found in the rapid field tests. It is however pointed out that blood tests from the daily ward rounds are often routine tests for following trends, and not always important to be updated quickly on. Hence it could be important enough to know when these results start popping in, so that one can be prepared to sit down for a while with them later.

The individual blood test indicators are not always important to follow up either, but are ordered in groups where only some indicators are interesting. The patient Odd Hansen in the scenario was important to follow up on, and his symptoms could possibly indicate a heart attack. This possibility would be settled by the troponine indicator, and hence the nurses above paid attention to the arrival of this in particular. It is therefore relevant to the nurses not only to know that blood test results have

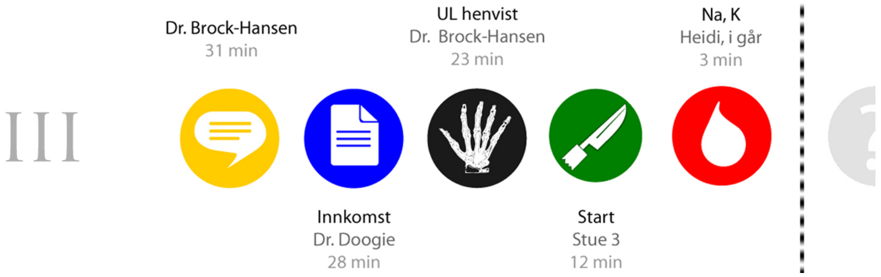

Positive	**Negative**
• "All information is good information" [TP3]	• Uninteresting and hard to define [TP1,TP3,TP4]
• Imaging diagnostics events are indeed related to the referring physician [TP3]	• Many physicians come by, we don't always know what they do [TP3,TP4]
• There might be exceptions where this adds value.. [TP4]	• Same physician may check up on all of the ward's patients [TP2]
	• Who took the blood tests is not known the next day [TP4]

Table 4.4: Feedback on events de-identified with alternative III

arrived in general, but also to know which indicators are ready. If not, they would have to check up on each and every indicator immediately on arrival, since it could be the important one arriving:

> *Nurse 1*: "If it isn't very interesting with sodium or CRP, you possibly wait until the CRP has arrived too, then, before you check it [sodium]."

Journal events (blue) were unfortunately not tested in the scenario. It is however likely that these would have been hard to contextualize with the de-identified alternative II, since journal notes could have been written for anyone. An exception might still be the arrival note, which is often written by a physician after the patient has arrived at the ward. Nurses would hence wait for this to be ready, so that they can read up on the patient's problem and condition. As for the comment event (yellow)

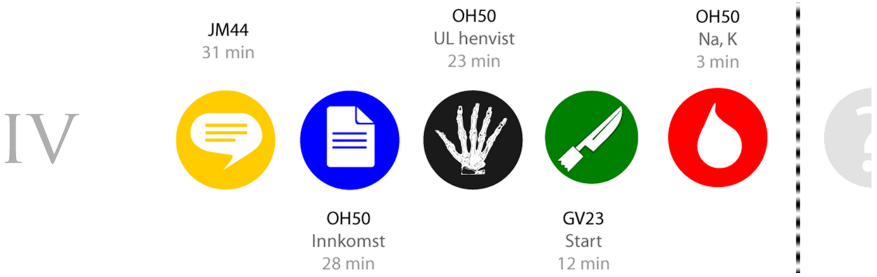

Positive	**Negative**
• Initials and birth year helps a lot [TP1,TP2] • Identification is acheived [TP4] • Easy to identify known patients [TP3] • The screen could be placed somewhere less public.. [TP1] • Event description can be removed if identity is given [TP1] • Similar to the patient sheet [TP3]	• Duty of secrecy hard to maintain when patients see or talk to each other [TP2,TP4] • Patients reckognize themselves [TP3] • Not sure how legal it would be [TP3] • An HIV test would be too sensitive here [TP1,TP3] • The combination of the knife and the initials look scary! [TP2]

Table 4.5: Feedback on events de-identified with alternative IV

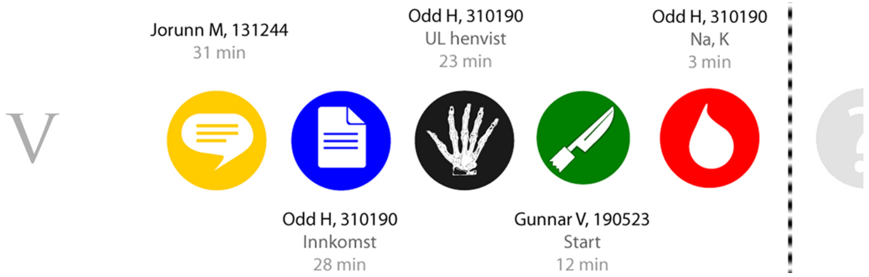

Positive	**Negative**
• Would correspond to patient lists [TP2]	• No, that's not possible [TP1,TP2,TP3] • A violation of the duty of secrecy [TP4]

Table 4.6: Feedback on events de-identified with alternative V

was however impossible to relate to any context:

> *Nurse 3*: "But when a comment event appeared, I was of course not certain who it concerned."

All test persons likeways argued that they in real work would most likely understand what patient the de-identified event concerned, by knowing the context of their patients. This is especially true if that patient is more important to follow up on, like Odd Hansen was in our scenario. At the same time, they always knew when something that was expected to happen, had not yet occured:

> *BL*: "Has anything new happened, compared to what was on the desktop?"
> *Nurse 3*: "Sodium and potassium.. Nope, that was the first entry there, too."

> *Nurse 1*: "So, now I've spent 10 minutes here, and some blood test results arrived 14 minutes ago. [...] Then it has not happened anything new."

As for the other de-identification alternatives, the feedback was mixed. The anonymous alternative I (figure 4.2) would for instance reduce the utility value of the screen considerably. Although still being capable of indicating whether or not something new has happened lately, it would be hard to relate just about any message to an individual:

> *Nurse 2*: "I would've figured out the operation events, though. And I'd see when blood tests arrived, but I wouldn't have a clue which one I would be expecting, because normally we wait for more than one. And for radiology, it could be both MRCP, ultrasound, CT and X-Ray overview... So it would have been more difficult if I received two or three patients during the shift or earlier the same day, and they all would have some examinations. It would have been hard."

For the different de-identification alternatives, it was obvious that increasing the identifying information displayed will allow more precise identification to be achieved. It could however be discussed whether the extra information supplied in alternative III is actually useful to the clinicians. Radiology examinations are only weakly related to physicians, and blood samples are not especially related to a single clinician either. At least for the nurses, they do not know everybody who works in all care activities the patient is subject to:

> *Nurse 3*: "We don't even know everybody who are there, who take orders and such [...] There are a lot of physicians, so it is hard to [...] know it just by looking here."

The suggested use of pseudonyms in alternative VI (figure 4.7) was criticised for requiring the users to remember something that is not already kept in their memory. In addition, one of the nurses pointed out the risk involved in losing a patient list to an outsider, that would immediately reveal everything. On the other hand, two of the participants also commented that the codes used would only be a supplement to the de-identified alternative II, where the description of event was already disclosed. And if the patient was especially important to follow up on, the code could be memorized and hence be used for these particular cases.

> *Nurse 2*: "If I'm most worried for Odd Hansen, I could have taken note of the code and be on alert for when it appears on the hallway screen."

The participants were all conscious with respect to the duty of secrecy, clearly rejecting alternative V (figure 4.6). Their personal interpretations of what would actually be required from a visualization in the hallways was not entirely equal. One

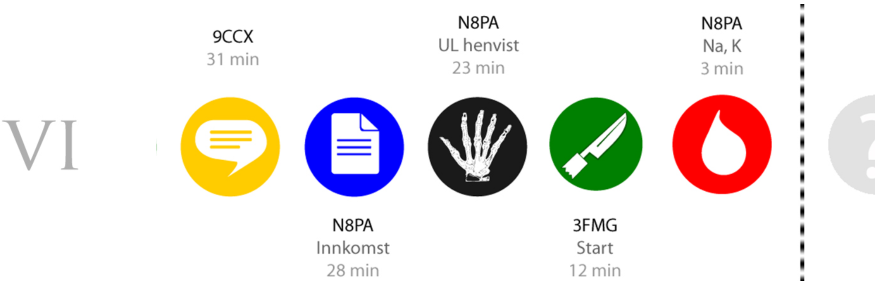

Positive	**Negative**
• Compared to alternative II it only adds something (the code) [TP1,TP3] • You can remember the code for particular patients [TP2] • Precise identification [TP2,TP4] • The code can always be brought along on the patient sheet [TP3,TP4] • Identity can be verified without additional checking [TP3] • Logical codes may be easier to remember [TP4] • Easier to ignore irrelevant messages with access to patient identities [TP4]	• Codes are hard to remember [TP1,TP3,TP4] • Code is not used elsewhere [TP1] • Does it really save any time? [TP2]

Table 4.7: Feedback on events de-identified with alternative VI

feared that even if the information was presented without a name tag, the patients could be identified by others, for instance by observing other clinical activities at the ward. Another was positive to the idea of using initials, as in alternative IV (figure 4.5), but was slightly worried about the association between patients and the rather expressive icons:

> *Nurse 2*: "I don't know with respect to patients who pass by and see that thing with the knife [operation events]. That is maybe the most frightening [laughs]. 'Now they're cutting an 'IN' born in 1944.' [...] It would of course help us more, but whether or not I can fulfil my duty of secrecy when patients perhaps know each other.. That's when it becomes hard, 'cause then they know the patients' initials."

Moreover, one believed that initials could be used in a public space, since the risk of anyone wanting to re-identify patients was not considered high compared to the increased utility value they would add to the visualization:

> *EAG*: "Do you think it would be okay if something like this [IV] was displayed in the hallway?"
> *Nurse 1*: "[...] I could not imagine that this would be very interesting for others who are not interested. We already have that whiteboard, and that contains first names, but its location is maybe slightly hidden after all..? [...] I would have known immediately who it concerned, [...] that it was concerned with a patient under my responsibility."

Finally, the use of initials was also considered safe in itself by a third of the participants:

> *Nurse 3*: "I think that it's impossible for other patients to figure who it is, so we're not breaking any duty of secrecy, I don't think so, no."

Nevertheless, all four test users declared to be very positive to the tested prototypes, both the desktop version as well as the de-identified view. A central point of the feedback was how they imagined themselves saving a lot of time spent logging repeatedly in and out of a variety of clinical information systems to check for new information:

> *Nurse 3*: “I would’ve surely used it a lot, rather than checking DocuLive [the EPR] and radiology to check, I would have rather used this, and then enter the DocuLive. [...] it tells whether results are ready or not, and that is the big issue because we are looking all the time and waiting all the time, very often. Especially with such emergency patients where time is crucial, we keep logging in and wait for results.”

> *Nurse 4*: “This would be super to have, so we don’t have to spend time checking out things that may not be worth checking out.”

	TP1	**TP2**	**TP3**	**TP4**
I				
II	2	2		2
III			3	*
IV	1	3	1	
V				
VI	*	1	2	1

Table 4.8: Rankings given to the de-identification alternatives.
The highest score is 1, while * means ranked third, but not voted for.

Their current authentication and authorization methods were moreover considered very ineffective, including an ID-card that must be inserted into the computer, and usernames and passwords for both the operating system and information systems. The participants thus appreciated having a screen in the corridor, allowing them to check for updates while passing by, instead of having to sit down and log in on the workstation each time:

> *Nurse 4*: “I immediately think that the hallway screen is very nice. [...] the workstation computers require us to use the card to log on, so it is good to have this in the hallway for checking whether something new happens that is worth logging on the workstation for.”

> *Nurse 1*: “[...] if you’re going to the medicine room to get something or fetch mail from the pipeline, you must take out the card, so I don’t want the hassle with logging back on before I think that ’now there’s soon

> something I can do', so then this would have been amazing. [...] And you know that you don't have to spend a lot of time interting your card and waiting for the login, and then logging back in there, and then again."
>
> *Nurse 2*: "I might have had to check DocuLive several times already now only when I could pass by and see it there."

Another advantage with having a large screen, is that it can be viewed from a distance. This makes it easier to pick up updates quickly, when other work in the ward keeps the clinicians from sitting down by the computer:

> *BL*: "Can you imagine approximately how often you would use such a system during a regular shift?"
> *Nurse 1*: "Hard to tell exactly, but I guess that if I'd had this screen available, I would have taken a quick look at it each time I passed by. And if it is located in the bed cluster, it will be something to pay attention to, and you'll always notice when a new symbol or update appears."

4.3 Summary

Through a series of rapid field tests, using simple throw-away prototypes, it was gained further insight in visualization alternatives and information needs of the end users. This was in turn used to build a high-fidelity prototype which was populated with realistic patient data. The prototype was moreover separated into two views; the main view intended for interactive use on a desktop computer, and the hallway view which contained a timeline with possibly de-identified activity updates.

The prototype was then installed in a usability lab, which was set up to simulate a hospital bed cluster. A scenario was written for a role-play to correspond with the prototypes, so that the test persons could better contextualize the new prototype, and thus be more fit for the focused interview directly afterwards. For the role-play, the hallway view only displayed a fixed level of information granularity, which included no patient identifying data. Later in the interview however, the four test persons were each introduced to a set of six different alternatives to de-identification, that could potentially also be used for the same purpose. The test persons were asked to assess each of these de-identification alternatives, and then finally suggest a ranking of the most appropriate candidates, in their opinion. These overall rankings are summarized in table 4.8.

In general, the participants were also positive to be given immediate access to updates from information systems, and being able to be notified through a large screen in a location where they frequently pass by. It was moreover confirmed that the ward nurses often have an idea of what is happening next to particular patients, which resulted in all participants being able to interpret some de-identified events correctly in the experiment.

"Humans do not always err. But they do when the things they use are badly conceived and designed."

The Design of Everyday Things
DONALD NORMAN

Chapter 5

Results analysis

In this chapter, all the data collected in the experiments will be analyzed in light of the findings from the literature review in chapter 2. First, a framework for the following analysis is however presented, which has also served both as background and title for the paper attached in appendix A.

5.1 A framework for analyzing the data

De-identification of PHI is definitely associated with a conflict between the privacy of patients and the utility that the data may have to those who will use it. The secondary use of de-identified PHI for research purposes is an example where increased richness of data will increase the value of the data sets, however the rules for exactly what data may be included in such data sets may effectively reduce the possible purposes of use the data may have. When de-identification is taken further into a setting where the data is used in real-time, it also has to present something of particular value to the users of a specified system in a specific context. There are hence three perspectives that altogether will define what an acceptable de-identification solution can be.

The first perspective is what the users will *need* in order to make use of data that are de-identified. This involves the amount of identifying information that is necessary for making the information presented meaningful to them, and hence provide utility. A physician responsible for a particular patient may for instance read that the results of a patient's radiology examination is ready for further assessment. When there is no

Figure 5.1: The best solution is found somewhere inside the triangle.

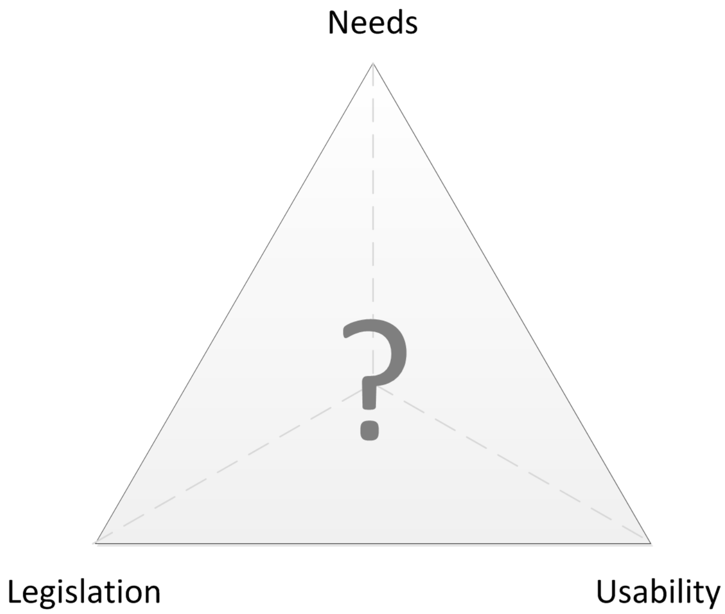

indication of which patient the results belong to, the message will however have very limited utility value. Desired indications may have to be as specific as e.g. the patient's initials and birth year, or as vague as only a known limitation of which patients may be subject to the message, e.g. including only patients from bed cluster A in ward 1 at department X, or that the patient's responsible physician is the one who reads the message. It is however necessary for the utility value of information to raise above a certain threshold for the system to be perceived as useful, which makes this perspective essential. Being able to correctly identify the patient whom a message concerns, is hence a deciding factor in this respect.

The second perspective may however be in direct conflict with the needs of the users, as it involves both national and international *laws and regulations* concerning how sensitive data may be used. As described in chapter 2, there are for instance 18 data elements expressly specified by the HIPAA that should be removed from a data set in order for it to be deemed as de-identified. For use with real-time visualizations this list may however be both stricter than necessary, or it may not take the particular risks of such uses into account, and other interpretations of the law may decide what is legal to show and not. The question is also not always in general what particular variables are legal to show, but what combinations of elements will maintain the privacy of patients, facing a calculated risk that is considered smaller than the overall achieved utility value.

The third perspective has taken into account the cases where the users have to perform re-identification of some or all of the de-identified data, if it cannot be deducted directly from the disclosed data set. This will generates requirements towards the *usability* of such visualizations, since it will affect the time and effort taken in order to make the system fulfill a user's needs. The problem hence specifies how well the users are supported in the re-identification process, as it is likely to be done subject to time constraints, especially in dynamic work environments like hospitals are. The system's utility value may also be destroyed if forcing users to spend a lot of time in either repeating authentication processes, or repeatedly misinterpreting the disclosed information, although it serves both the users' perceived needs for information and the patients' lawful rights to privacy.

Striking a balance between these three perspectives is hence a basic feasibility requirement for a system such as that targeted in COSTT. The possible implementation of de-identification will subsequently be affected by each of them, although it will likely have to consider tradeoffs between them all. Such a balance may therefor exist somewhere within the borders of the triangle in figure 5.1, yet it is not certain whether any of the tested approaches will be able to fulfill the requirements from all sides. The real-world requirements can moreover be different for varying systems and contexts, and the perspectives instantiated differently for another system. What is examined in this thesis, is however how well de-identification can work for PHI in a public and/or semi-public hospital environments, where one can assume that unauthorized outsiders can obtain access to reading the screen's contents.

5.2 Answers to research questions

5.2.1 Research question 1

> "Which methods are the most appropriate for de-identification in real-time oriented visualizations containing patient health information?"

After starting with a list of possible de-identification techniques in the literature review, some were excluded from the study due to lack of appropriateness for the targeted purpose of use. Another few were abandoned after the rapid field tests, because the users did not find them feasible to use in practice, but some new ideas were also included. Finally, six different alternatives were tested in the usability lab experiment, which in turn could be grouped into three main approaches. These three main approaches correspond to the three categories found in Norwegian legislation for protecting information through disclosure control.

Anonymization

The first category, which is represented by alternative I (figure 4.2), was in the lab experiment clearly pointed out to not suit the needs of users very well. The only thing it can contribute to, is that one can possibly see whether something has *not* happened lately, e.g. that no blood tests results have arrived during the previous two hours, or that only one imaging diagnostics event has occurred within the same period of time. This second example may in turn tell that the pictures have possibly been taken as planned, but that the radiologist's description is not ready yet.

From a legislative point of view, there are hence reasons to believe that this approach would be entirely legal to use, even if exposed in a public environment, and the data set's population is limited to the persons on a ward — or even a bed cluster. Since the provided information detail is so limited, it will be impossible from an outsider's perspective to figure out which events may belong to which patient, and even a patient himself could not be able to track his "own" events. None of the participants in the lab experiment had second thoughts to implementing this in the specified context either.

From a usability point of view it is however clear that personnel would have little benefit from this approach. A visualization that would anyway require them to check another system for the importance of a message, before then finally turning over to the system where the information is actually found, would require much unnecessary workload. The total amount of workload would possibly even be comparable to what users today are faced with, and would therefore make such a system unnecessary and useless.

De-identification

The second category is represented by the alternatives II, III, IV and V, yet the last one is very close to a completely identified alternative and was discarded by all four of the lab experiment test persons. For the remaining three however, the test persons did not agree just as much. In the rapid field tests, alternative IV with initials and birth year (figure 4.5) was told to be sufficiently precise for identification of patients that are in the clinician's responsibility. This was also repeated in the lab experiments, and hence it was ranked first by two. Alternative II (figure 4.3) was agreed on as not being optimal, yet it was ranked second by three (but never first), having obvious weaknesses when many similar events occur within the same scope of time, or when something that is not expected to happen do indeed trigger a message. Alternative III (figure 4.4) was hence provided as a more detailed alternative, but apparently the extra level of detail given about "where" or "by whom" was not considered to be very useful by the participants. The suggested approaches of clinician responsibility and location may hence seem to provide little help towards authorized re-identification.

When it comes to how well these alternatives protect the privacy of patients, it is hard to make a final conclusion. However, the popular alternative IV — and of course V — would not comply with Norwegian regulations for processing personal health information, at least if the screen was to be placed without any physical or computational access control. Two of the test persons had a strong opinion in the same direction too, yet one chose to rank it third if "it would be considered legal anyway". The two remaining candidates were on the other hand open to its usage in public areas, at least to the author's power of judgement. The difference between the two is not considerable, which is also confirmed by the lab experiment results, but the question posed by legislation is whether or not the information is de-identified, which is seemingly a binary construct. As discussed by El Emam (2010), it can still be argued that in reality there exists a continuum of identifiability, where some data sets can be easier to re-identify than others, which may require much time, effort, cost and skill for the same to be accomplished. The size of the data set, composition of patients in a ward, ordering of events, and so forth, can all make re-identification either harder or easier for attackers to achieve. In addition, one must assess what harm an unauthorized re-identification would cause, which depends on the information that is disclosed. An important aspect is still that what is displayed on these screens will not be searchable files such as traditional data sets disclosed to a third party, but a visualization that endures only until new messages replace the old ones. Nevertheless, is seems impossible to draw a final conclusion without knowing more about exactly where the screen is located, for instance if alternative II would be illegal in a public hallway, but approved if it was turned away from the by-passing audience behind a simple pillar.

Usability-wise, alternative IV was the clearly most popular, since it would require no further action from the users to achieve a normally sufficient re-identification. The alternatives II and III were again very much equal, both requiring additional confirmation of identities while giving a fair hint of who the message concerned. Another possibility is that users will accept it at first, but then make a considerable amount of wrong guesses, causing both frustration and time being wasted. It was nevertheless pointed out that II looked a bit “cleaner” than III, since the additional info in III was not perceived to be very useful. Another participant still stated that “all information is good information”, and ranked alternative III instead of II.

Pseudonymization

Alternative VI (figure 4.7) was a pseudonymized version similar to both IV and V, where patient names and birth year/date were replaced with pseudonym codes. This solution would enable an equally precise, or possibly even more precise, re-identification key, compared to even V . From an information need perspective, such precise identification of patients would be very beneficial, in terms of removing all ambiguousness when the code can be matched against a direct identifier possessed by authorized users.

Pseudonymization will in addition provide a legally valid de-identification, since de-identified and pseudonymous registers have the same legal status in the terms of Norwegian legislation. It is yet uncertain whether the event description can still be disclosed along with such a pseudonym, as this would possibly expose larger parts of the process if someone connects the dots between several messages concerning the same patient. Moreover, this is a possible risk in the same category as those mentioned for de-identification above, making it a question of implementation and context, rather than something that can be given an answer in general. One of the participants also commented the risk of losing the patient list, which was suggested to contain the patient’s pseudonyms for reference, so that an attacker could immediately re-identify messages regarding all patients contained on the list. The same is a shortcoming of pseudonyms also pointed out by Riedl et al. (2008).

Still, it will overall be a sub-optimal solution, as the test persons all agree that codes are hard to remember, and that such codes are not being used anywhere else in the hospital today. Implementing such an approach would therefore add to their overall workload, but having direct access to identifiers without consulting another system would at the same time most likely overrule this additional workload. As pointed out in the experiment, it will also be possible to remember the code for particular patients, whenever they are important to follow up on. Likewise, it will also not be necessary to spend time checking up on patients that are not of any interest.

5.2.2 Research question 2

> "How do de-identified visualizations perform compared to what medical personnel need and what current practices are?"

During the lab experiments, the four participating nurses got a chance to be familiar with how a de-identified visualization could work in practice. A first observation was that all of the test users were able to correctly associate at least one new event, as they all recognized several events appearing for patient Odd Hansen. The scenario was nevertheless staged and the number of test users was so low, so this will not necessarily be true in general. Not all events could be identified either, such as the nurse's comment, which was added to another patient "out of nowhere". In combination with feedback given in the rapid field tests and the interviews, it may however be true that clinicians do know their patients by their trajectory, sometimes even better than by their names. When an event is expected to happen, the event may hence also be possible to distinguish from others when it occurs, yet unexpected events do not work very well within the theory's scope unless either an unknown amount of extra detail *or* a pseudonym is given for reference.

The test users also agreed that they were exposed to an acceptibly realistic scenario in the experiment, at least for the night shift. The chain of events was considered likely to occur in reality, making the participants confident that from a selection of eight patients they would be able to recognize at least some events, if they happened in real life. How well the re-identification works, could however depend on the context in which the de-identified events operate. A situation where many similar events would happen, for instance in the mornings when all blood tests are taken at more or less the same time, the results would be returned within an equally limited time frame, and thus "flood" the list of events.

While the participants each gave their personal rankings to which of the de-identification alternatives they would prefer to have available in such a setting, the outcome did not provide a clear answer to which of the alternatives were the best to use in practice. A central point of disagreement was related to whether or not alternative IV (initials/birth year) would be feasible to present on a screen in a hallway, like the tested scenario suggested. Two of the users ranked this first, arguing that the approach using initials and birth year is already very much in use today, and hence would suit both their needs for identification and usability. One of them also considered this approach close to anonymous, and while the others had concerns about how it would affect the privacy of patients negatively, only one of them did not vote for it at all.

As already mentioned, all participants however agreed that alternative V (first name, first letter of last name, and full birth date) was too revealing to be placed in a close-to public environment such as the hallway where they work. Moreover, alternative I was never ranked either, but there were still some differences in how it was assessed by the four, with one claiming that she would indeed be happy with that, compared to having nothing at all. It was also commented that all uncertain but possibly interesting events would probably be checked on the desktop prototype anyway. In practice this would very likely add extra unnecessary workload for all events that are related to patients not concerned by the clinician, or events that prove not to be interesting after all.

A system where de-identified events are used for hinting whether something has happened, hence requires someone to interpret if the events are relevant or not. In case of uncertainty, the user may want to check a fully identified view that can confirm or refute the assumption that it may be relevant. This could lead to much frustration if more time is spent on checking identities than time is saved by not having to check the other systems as frequently as today. Although three of the participants were more or less happy with alternative II, ranking it second, it was hence never ranked in first place. Instead, alternative IV with patient initials and birth year was ranked first by two, but at the same time it was ranked third by one and rejected by the last. Alternative VI was also ranked first by two, then second by one, and with the last being somewhat positive herself, yet skeptical due to what others might think of it, she did not vote for it at all.

There is however not any quantitative weight to find in these rankings, due to the low number of test persons. Instead, from a qualitative perspective, it can be said that none of the four test persons agreed entirely in their rankings. Not even two out of the four did actually vote for the same three alternatives, even with their internal rankings left aside. Their comments associated with doing the ranking also revealed different opinions towards the weighting of priority for each of the three quality requirements.

5.3 Summary

Anonymization, de-identification and pseodonymization have each been assessed on their qualities towards supporting identifiability needs, complying with legislation and being usable in practice. None of the three categories provided optimal solutions for all three, which is also reflected by test persons' disagreements in their rankings.

While anonymization would be legal to use in public settings, it is neither user friendly nor useful towards identification. Pseudonymization is, on the other hand, the most precise means of identification, while being a possibly legal candidate too. The approach nevertheless faces a considerable downside towards usability. Finally, de-identification may be okay in all three aspects, but raises challenges towards potentially many "false positives" when re-identification is uncertain, as well as the overall need to verify identities.

When used to support hospital coordination in a public or semi-public context, a balance between needs, legislation and usability is for the reasons presented not obvious. Moreover, this is most likely not feasible at all, using a *non-interactive* approach to de-identified visualizations, such as those six that were tested.

> "To design something really well, you have to get it. You have to really grok what it's all about. It takes a passionate commitment to really thoroughly understand something, chew it up, not just quickly swallow it."
>
> *Steve Jobs*
> CO-FOUNDER OF APPLE INC.

Chapter 6

Discussion

While the results analysis concluded that de-identification could not provide an optimal balance between identification needs, legislative requirements and usability, it does not necessarily mean de-identification is not at all useful for real-time visualizations. This chapter will take a look forward to how an interactive implementation could possibly benefit such systems to a larger degree than the non-interactive approaches used in the experiments. Supplementing approaches to information disclosure in the same context will also be considered in the light of this study's outcome. It will also be discussed methodologically how the resulting data have been obtained, before finally some ideas for further work are presented.

6.1 Finding a balance

For each of the three requirement perspectives, i.e. utility, needs and legislation, adjustments can possibly be made that would enable a viable solution to the design problem attacked in this thesis. Suggestions to such adjustments are given in the three sections below, based on both results from the study, as well as other research available on each topic. These suggestions build on a fundamental insight held forward by DePaula et al. (2005), on the need to trade-off different factors against each other, in order to effectively put available resources into use — including both security, usability, availability, and so forth.

6.1.1 Identification needs

From the identification needs perspective, there are basically two strategies available for providing sufficient identification to the users of a system. The first is to use the visualization itself to reveal identities along with the messages, which is the approach that has been investigated in this study. The other is to utilize other devices for making the re-identification happen.

When exploring the first strategy with several alternatives for identification, there are identified three techniques which entirely satisfy the identification needs of personnel. Using full names and birth dates, will of course be the most precise and also usable method. This alternative is followed by a slightly less precise, but still very acceptable approach (from this perspective) with initials and birth date. Finally, pseudonyms are considered an effective technique for precise identification, but it has a significant value-loss towards usability, especially when the pseudonyms are only used in the organization for this purpose alone. Still, it is the only of the three alternatives that may be legal to use in general for a static visualization placed in close-to-public areas of a hospital. There may however be system-specific cases where some messages can actually be displayed to the public, without requiring the de-identification to be watertight. An example here is the "comment"-event (yellow) in the lab experiment prototype, where it is only told that a new comment has been written for a patient, but with no medical information attached — which could very well be displayed along with initials and birth year.

Another approach would be to introduce personal mobile devices as a support system for pulling out interesting information that appears in a non-obtrousive way on the large screens. This would implicate going back to a one-to-one user-computer ratio, but being a whole lot more flexible for collaboration than traditional personal computing. Each user in a group could for instance be allowed to pull up individually needed information. Moreover, these users can authenticate themselves individually on the device, and still make the switch between working individually and as a group (Heckle & Lutters, 2011).

During the lab experiments, one of the participants also came up with a suggestion based on the ward nurses' care relationship to the patients. In that particular ward, the patients are supposed to have one primary nurse each, who will take care of the events that have been included as messages in the prototyped system. If the messages could be linked to this particular nurse, it would not be necessary to know anything more about the patients they concern than the anonymized view would provide. The messages could rather be placed in employee-specific feeds, since the nurses would be interested in all such messages regarding "their" patients anyway. An analogy would be that each nurse had their own mailbox that could be easily checked for new mail, where all were being opened regardless of what would actually be acted upon. A similar idea was also explored in the rapid field tests using the prototype in figure 4.2, but the messages were too abstract for the concept to turn any heads. It may still prove challenging to provide such a correct link towards all responsible nurses in a dynamic environment. In addition, what will happen if nurses becomes occupied for a while, and hence cannot pay attention to their incoming messages?

6.1.2 Legislation

From the legislative perspective, it may be possible to use access control restrictions for additional protection of the information, and hence making it possible to provide more precise identification than could be used in the close-to-public setting like the ward hallway. This could be acheived by using two main strategies also, one being physical access control and the other being computational access control.

Physical access control is normally not a watertight concept in hospitals, where few areas are entirely shut down to explicitly authorized access. Still, there are a lot of places where sensitive information may be displayed with less risk than in the examplified hallway setting, including meeting rooms, offices, lunch rooms and behind the ward control desk that may be located in the hallway — only turned away from the public audience. Depending on the screen's function, a compromise between physical protection and amount of sensitive information could therefore be negotiated, such as leaving initials and birth year for patient identification if the screen is e.g. placed inside a meeting room. The placement of such screens may however have a significant function when used in CSCW, especially when designing for awareness and transparency.

Scupelli et al. (2010) has observed how a whiteboard's usefulness is limited by lack of information detail when placed in public and semi-public areas. Moreover, it should be natural to search for the information the screen contains in the near proximity of where it is located, e.g. making it improper to locate a medication schedule too far away from the medicine room. Valuable information about patient status could on the other hand be left unnoticed if it is located off the course of staff's normal work activities. Placing it where nurses and anesthesiologists are likely to pass by at the same time, can instead make collaboration feel natural and signal that the organization sees them as a team. Nevertheless, Scupelli et al. recommends whiteboards to be positioned in staff-only areas for increasing utility value of the information, although this problem might be solved differently with interactive displays.

Computational access control is normally associated with the well-known login interfaces greeting the users who need access to almost any system these days. Such user authentication is an essential mechanism that identifies the user and accordingly verifies this identity. If the identity is verified, the user may be granted privileges to access resources on the computer, according to a descriptions of such access rights, a part of the process known as authorization. This kind of login routine became important when multi-user environments were introduced in computing, with usernames and password as the far most common approach to user authentication. Authentication of users can be done using three approaches, either alone or in combination, identifying them by something they *know* (e.g. a password), *have* (e.g. a smart card) or *are* (e.g. fingerprint, iris pattern or other biometrics) (Bardram, 2005). Each of these are still in turn associated with certain risks. Passwords are commonly known to be shared in hospital environments (Vaast, 2007), smart cards can be lost or stolen, and biometric protection may be not impossible to defeat either. Therefore, they are often used in combination as two-factor solutions to protection, making penetration more difficult — at least through the "front door".

A relevant aspect could still be the amount of information that can be reached through the system, and the potential damage it may cause if abused. A system like the EPR in use at hospitals could for instance require both a smart card and a corresponding username and password in order to grant the user access. This is because such access would give the user privileges to both read and manipulate highly detailed sensitive information, which has also been collected throughout a longer period of time. For a coordination system, such information depth will not be required, and it may not be crucial to provide write access to sensitive data either. The lab experiment participants were all okay with not having direct access to e.g. test results through the coordination system, but still appreciated how they could know when to take the required measures and check the EPR for updates. The objective should therefore not be to reduce risk of re-identification in general to its lowest possible level, but to choose a total line of risk that is acceptable in terms of what is gained from the system (El Emam, 2008; Ohno-Machado et al., 2004). It should subsequently be added to the equation that the system will neither be searchable to the public crowd nor contain information that can be seen in a historical perspective, still being aware that all health information is considered sensitive information by law.

6.1.3 Usability

Due to the above introduced requirement of access control, the usability of the targeted solution is under attack. Since physical access control will be highly dependent on the overall architecture where the information is placed (Iachello & Hong, 2007), this aspect is not considered further here. User authentication is on the other hand a field where much work remains undone, in order to provide generic mechanisms that are more usable, and not only stronger (Bardram, 2005). This topic cannot simply be ignored either, as the consequence of inappropriate login mechanisms may cause circumventions that jeopardize security, and subsequently leads to more vulnerable systems (Vaast, 2007; Bardram, 2005). Thus, when making trade-offs between security and usability, it is a goal to equally weight the objectives of the system against the needs of its users Heckle & Lutters (2011), although some legislative issues still cannot be ignored when dealing with sensitive information.

Since the credibility of an authenticated user is proportional to the strength of the authentication mechanism used, there are limitations to how simple the implementation of access control can be when providing access to completely identifiable PHI. Nevertheless, if the risks of information disclosure are mitigated in other ways, like de-identification is definetely capable of — although not bulletproof — the users' burdon of authenticating themselves can be lightened equally. When de-identified information already requires the user to *know* something quite exclusive to access the information, adding something that e.g. the users *have* would make it altogether a secure two-factor access control solution. One such token could for instance be an indoor positioning system (IPS) tag that could identify clinicians passing by and hence allow the de-identified view to be activated even before the user arrives at the screen — in theory (Heckle & Lutters, 2011). A capability assessment of such positioning systems in a hospital setting, however argues that the idea of automatically logging on and off a computer may not be entirely feasible just yet (Landmark, 2009). The idea of context-aware security systems may still be a useful path to follow in the future, along with upcoming advances in technology (Covington et al., 2001).

Requiring the users to present something they know, like a password, is often considered a golden mean between security strength and usability, however in a hospital setting this approach is used so much that users spend a whole lot of time entering passwords already. All four participants in the lab experiment pointed out that the login routines required for keeping track of important updates like those included in the prototype, takes up a considerable amount of their time each day. Since the passwords cannot easily be brute-force attacked on displays in open areas, it will however be sufficient to use PIN codes that are shorter (e.g. four digits). Apart from a PIN code that may already be in use, requiring the users to remember certain keys or codes to deciphering a de-identified view appears to be a bad idea. Both in the rapid field tests when testing a prototype that required room locations to be remembered, and in the lab experiment where pseudonyms were introduced, the users were all negative to the need for keeping such things in mind, in addition to everything else. Although the pseudonyms can be brought along on a patient list like suggested, it might still be just as simple and convenient for the users to spend their time on a simplified authentication process, such as described above. Coding of information with colors or shapes, like suggested by Tarasewich et al. (2005), has not been tested extensively in this study, but chances are that the same applies to this concept also in the health care domain. Some coding thechniques are then again already used on shared whiteboards, for instance the degree of emergency added to the prototype in figure 4.2. Using these in coordination systems as well, will not be faced with the same issues as new coding techniques are. On the other hand, coded information (also known as security by obscurity) is not a very secure approach in itself, and should hence also be used with caution, like all other sensitive information.

6.2 Towards an interactive prototype

While the static view of the de-identified prototype would have been accepted by the ward nurses who tested it, the low population limit in the displayed patient group required for it to be useful enough (maximum eight patients), would require the information to be more access controlled than a close-to-public area and non-interactive screen can provide alone. If the physical placement of the screen would reduce the risk of outsiders reading its contents, it could however be possible to increase this base level of information detail accordingly. This location-orientated approach is also part of the recommendation presented in the paper included in appendix C, where a risk-based evaluation of mechanisms for group access control is presented. Still, it

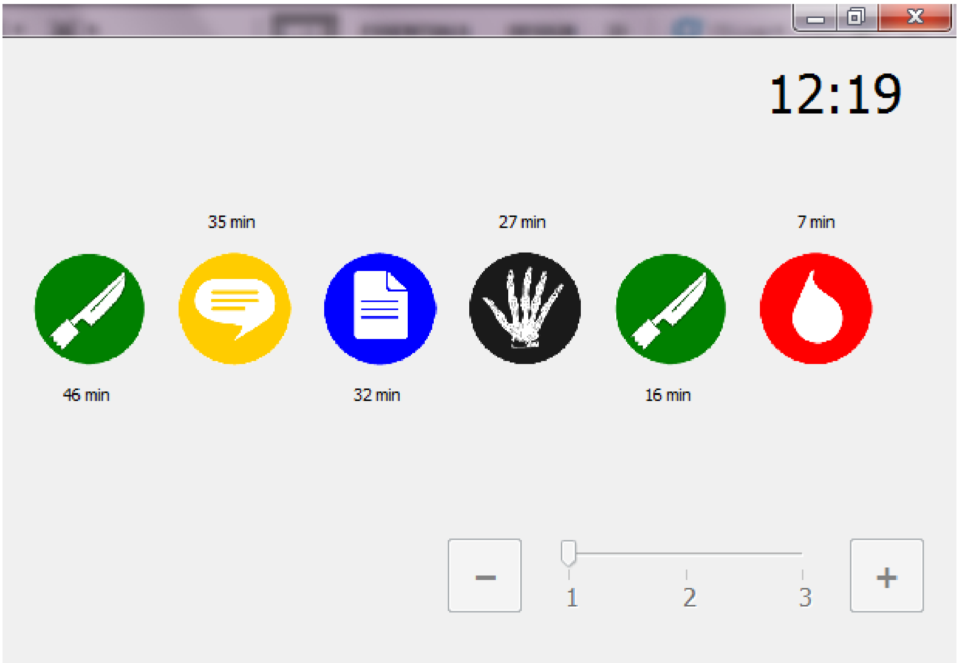

Figure 6.1: Lowest level of detail, allowed for public display (anonymized).

will not be sufficient to provide the required utility alone in open areas, and hence the recommendation involves a combination of several approaches that are referred to in this section too. Nevertheless, the anonymized view could serve well as a public layer on top of the more identifiable information, since it in practice could be disclosed to anyone without significant risk (see figure 6.1).

An important criticism of traditional security systems is their "all or nothing"-approach (DePaula et al., 2005), since there in theory could be many degrees of security available depending on the system and context. The next information level in the interactive prototype (figure 6.2) should therefore take advantage of especially usable access control mechanisms, and rather add an additional layer of protection through the use of de-identification. The conclusion of the paper in appendix C also encourages further research on methods for situation-awareness, in order to provide a login mechanism that is especially usable in the hospital setting. Although such awareness will not match the real world situation in all possible cases, it is necessary to have methods for overriding both too little access granted, as well as too much information being disclosed. Using de-identification as an intermediate level would hence protect the information if unintentionally disclosed, as well as provide a next level to which the user can proceed, if authentication becomes more certain. If producing such awareness is not feasible, the much used personal ID smart card could meanwhile be suggested, often containing an radio frequency identification (RFID) chip for near field communication, which could simply be scanned in front of a card reader for access to be granted to this level.

While many cases are likely to be solved using this kind of de-identified representation, the lab experiment clearly witnessed that often this will not be sufficient for precise re-identification. However, if the above described screen is now already made interactive, there are no reasons not to make such precise identification possible here too, and hence a third level is introduced with figure 6.3. As suggested in the paper in appendix C, making use of pop-ups and handheld devices could be beneficial extensions whenever the information is potentially disclosed to others within reading distance of such limited views. If one of the users are already (weakly) authenticated by their IPS tag or ID smart card, further access to identified information could thus be granted to this individual, either alone, or on behalf of the group, if the authentication is strengthened.

This will unfortunately require the introduction of a new authentication factor, since the de-identification part is taken out, but this does not necessarily mean a traditional username and password is required either. Biometrics could for instance solve this part, yet it is uncertain whether it is indeed as usable as often advertised (Bardram, 2005) — especially if wearing gloves or having the face partially covered, which is not uncommon in hospitals. Another approach, that would already be familiar to hospital workers who use their ID smart cards for opening doors, is the PIN code. Compared to implementations for restricted access to physical areas, this factor in combination with either the smart card or IPS tag, will not pose any greater threats while displaying information in support to coordination. The amount of available sensitive information is still very limited, and none of it can be manipulated at this stage either — as opposed to when full access to an EPR system is granted.

Although the authorization mechanism may automatically figure out which patients the authenticated user has access rights to see information on, it is not necessarily appropriate to reveal all these identities immediately when the user has been sufficiently identified. When using large screens in close-to-public environments, there is always a chance of somebody passing by, who are not authorized to view the screen's restricted contents, e.g. cleaners, patients and visitors. This motivates a restriction as

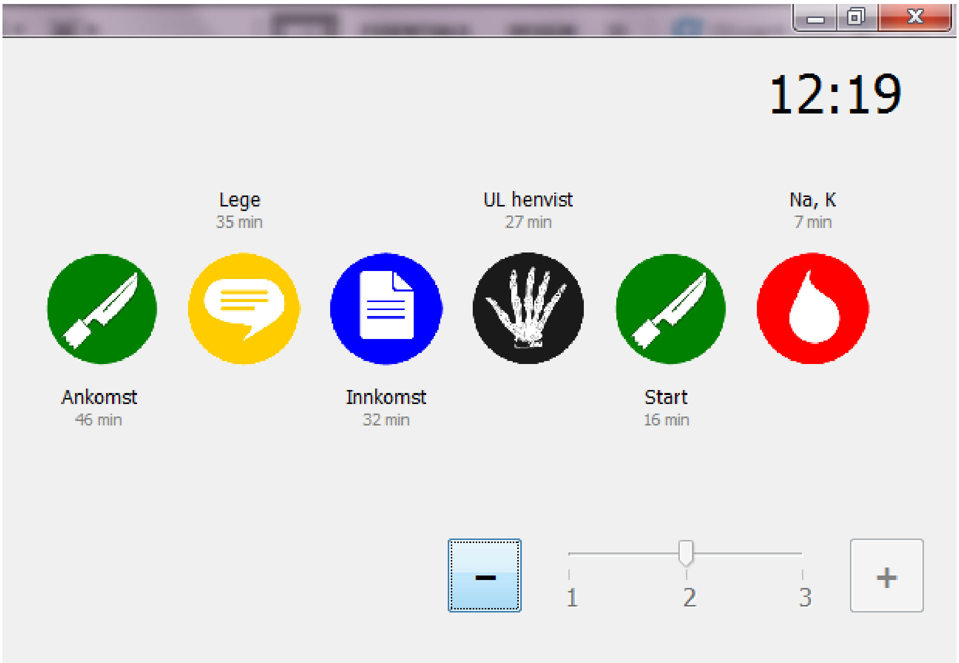

Figure 6.2: Information detail on its intermediate level (de-identified).

such, and instead, the slider in figure 6.3 indicates that third security level is reached, but only the message that the user points his finger at will provide the patient's identity in clear text. At this point, it could also be possible to re-introduce the pixelation/blurring techniques for de-identification, that were left unexplored in chapter 4. When for instance the third security level is reached, the patient names could be immediately added to the messages, indicating their availability, but being blurred so that they cannot be read. This approach is still not used in the preliminary design, since such blurred items could unnecessarily clutter the overall interface appearance.

Still, when the users are identified, it not only enables such additional information to "pop up", but also requests for even more information to be sent to a personal handheld device. For the case shown in figure 6.3, the user could simply slide his finger down to the letter icon[1] for having a message sent to e.g. his IP phone or PDA, containing the results of the chosen blood tests. If there are several concurrently authorized users, a list of these could appear for selecting the exact user(s) who should

[1] The icon is distributed under the Creative Commons licence (http://creativecommons.org/licenses/by-nd/3.0/), with credits to the author VisualPharm (http://www.visualpharm.com/)

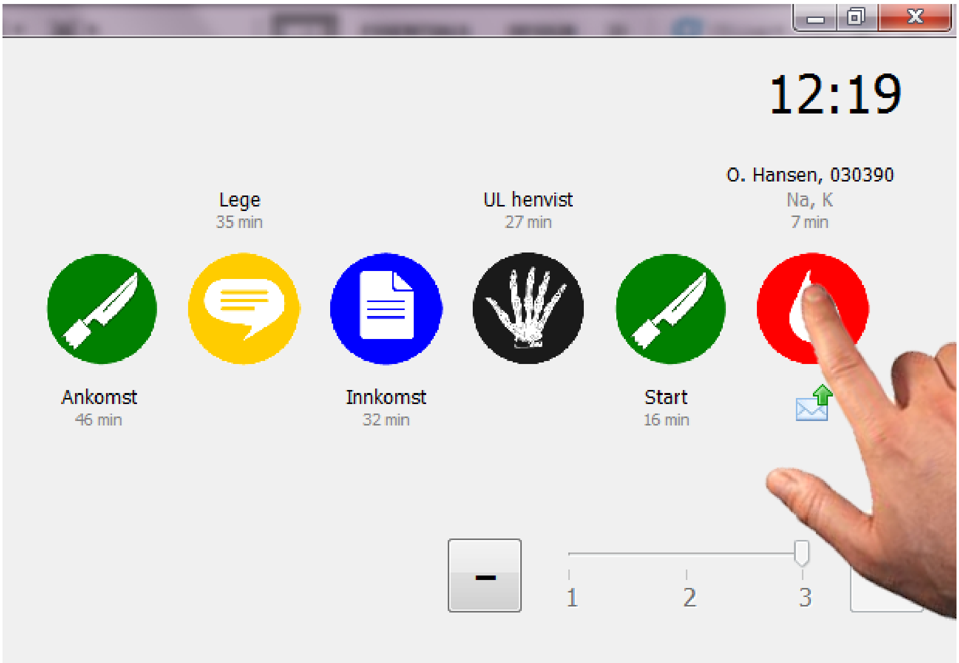

Figure 6.3: Full identification is enabled, but only when requested by the user.

receive the extended information.

A possible component in the interactive prototype's graphical user inteface (GUI) is also the slider found in its lower right corner. This slider (with "–" and "+" buttons) can be used for controlling the level of information detail shown, making it possible to also decrease this level when the context changes and such action becomes necessary. When the user is given such power to de-identify or anonymize the information according to the current context, it is both possible to work safely when outsiders are present, and then turn back to a usable, fully identified view when that becomes required and safe. In addition, the component adds transparency to the security functionality, so that the user can stay informed on what basis the current level of information detail is chosen, and which further possibilities exist.

Using a kind of "incremental" authentication, as explained above, will also enable logging of what information is accessed by whom, focusing on those information requests that require security level 3, and whenever extended information is sent to a personal device. Tracking all the simple requests for de-identified information would not only swamp the logs with all patient instances that are being checked every five minutes by each individual user — but logging the movements of employees with such a high frequency of sampling, could lead to privacy loss for the employees themselves.

6.3 Relationship to earlier findings

Iachello & Hong (2007) has pointed out that privacy is a holistic property of interactive systems, that need to consider the people who use them. The development of standard privacy-enhancing interaction technologies is hence named as one of five “grand challenges” in HCI and privacy. Studies by both Heckle & Lutters (2011) and Vaast (2007) have moreover revealed that secure systems become vulerable when the security measures do not match the real world needs of users, emphasizing the importance of user-centered security design. While Bardram et al. (2006*b*) focuses on user authentication methods in this respect, this thesis has mainly taken its approach to the visualization of information in itself. Whenever the statical approach is not sufficient, the end-user sould however be given the power to determine what is appropriate considering the current context, making it critically important to make security features visible in the user interface (Iachello & Hong, 2007). DePaula et al. (2005) also argues that users are able to understand and appreciate flexibility in systems that allow them to understand the consequences of their actions and develop new practices.

Results from the experiments indicate what was also observed by Shoemaker & Inkpen (2001) and Bellotti & Sellen (1993), that individual users have different preferences towards what information should be kept private, and what can be publicly disclosed. DePaula et al. (2005) moreover describes the gray-area between “secure” and “insecure” as a relative matter that cannot be legislated in advance and resolved simply by a system. A users’s privileges to a resource throughout its lifetime could also not always be known beforehand (Heckle & Lutters, 2011; Røstad & Edsberg, 2006), making it challenging to define an appropriate level of granularity for security systems. In line with these findings, it was already by Shoemaker & Inkpen (2001) suggested to develop methods of keeping the users informed about what information is private. This is also viable in the context of COSTT, considering the feasibility challenge of creating a system that always knows what will be both a secure and appropriate level of security. In addition, Chung et al. (2004) have suggested design patterns for ubiquitous computing, which would also apply to context-aware screens in public and semi-public hospital environments. One of the pattern categories is called “developing successful privacy”, comprising e.g. appropriate privacy feedback, privacy-sensitive architectures, partial identification, physical privacy zones and keeping personal data on personal devices — all of these being relevant to the suggested interactive prototype. The use of mobile devices for accessing more detail is also suggested by O'Neill et al. (2004) and Bardram J. E. (2003).

The interactive prototype's feature of turning de-identification on and off, is moreover giving the authorized users control over the visualization's level of sensitivity. The proof of concept by Huang et al. (2008) shares the same motivation, introducing simple "on"/"off"-switch button for runtime privacy control in documents and web-pages containing sensitive information. Although being a clear reference for the interactive prototype's sensitivity zoom control, the approach is nevertheless to display either all or nothing, which has already been criticised by DePaula et al. (2005). Rather, the system should match security to task and allow flexible security so that it becomes secure enough for the users' immediate needs. The earlier described blinder approach by Tarasewich et al. (2005) however offers precise control over which specific blinders to reveal the contents behind, by the use of mouse gestures. Iachello & Hong (2007) also predicts that the blinding technique may increasingly common in the HCI privacy landscape, however mentioning semi-public and public displays in particular, as opposed to the single-user focus of both earlier applications. The interactive prototype has hence implemented this concept with allowing the user to point a finger at specific event icons for revealing identifying information.

With the introduction of interactivity, it also becomes a requirement to authenticate users. Behlen & Johnson (1999) had already warned against using de-identification for publicly released data sets, believing that data sets cannot be entirely free from links to the individuals they concert. While this is not necessarily true, especially when de-identification is used in an application where searches and automated linking of information across databases is not possible, the required level of de-identification certainly impacts the utility value for use in real-time. The subsequent recommendation of protecting certain identifiers with access control mechanisms at the point of query, also supported by Andresen (2009), hence becomes a solution that should be adopted.

Logging in however is in this respect a time bandit, emphasized both by all lab experiment participants, as well as a study by Fuglseth (2008). 97 users sessions were here reported to require on average 1 minutes and 11 seconds for logging in, although most of them lasted between only 2 and 10 minutes. The access control mechanism should nevertheless be balanced against the risks involved in what users are in fact given access to, as opposed to a mathematical guarantee that are likely to be requested from systems in health care (DePaula et al., 2005). The users do not always seek e.g. full write access to the complete medical record, but rather just “enough” information and the easiest way to get to it (Melby & Toussaint, 2009). The concept of different levels of security depending on the strength of authentication is for instance used by the Norwegian Tax Administration for public reporting and dialogue via the portal *Altinn*. The users can here choose between a total of five security levels at login, an approach that on the other hand has been criticised for being too complex for most users to understand (Yoga, 2010). Reducing possible levels of security to three in the interactive prototype therefore makes sense.

As for the legal issues, the original PoCCS system design by Bardram et al. (2006*b*) did not comply with the law. The authors vaguely argue that this is the same practice as analogue whiteboards today are following around the world, however being aware that both national and international legislation pose requirements to personal user authentication for clinical systems. A digital whiteboard implementation by Aronsky, Jones, Lanaghan & Slovis (2008) in Massachussetts, USA, also uses a practice that is against the HIPAA privacy requirements by naming patients with full last name and first letter of the first name. On a screenshot included in the article, the system further displays the patients’ sex, age and medical record number. The solution here is to give the patients an opportunity to sign a form rejecting the use of their names on whiteboards. When this happens, or the patient is considered a security risk (e.g. has a gun shot wound), a “No info” label replaces the real name. Since both these systems are only found in semi-public areas however, it can be considered whether any of these two systems could benefit from a de-identified view as the base level of information — instead of the anonymization requirement in public areas — and then allow fully identified patients to be revealed when simple authentication is provided.

6.4 Methodological checks

In chapter 3 it was aimed at establishing a triangulation of data collection methods, in order to make the final results more confident. The possible threats mentioned include both incompleteness in the recording of data, and biases introduced by respondents. These threats have also been identified as real during the study, and correspondingly been countermeasured with the collection of additional data using other methods.

6.4.1 Literature review

The literature review laid out an important foundation for the consecutive work, and making a selection of applicable and possibly applicable data protection techniques implicated the elimination of others. When doing this elimination, there is always a possibility that something good is considered bad and left out of future explorations.

Applying de-identification to real-time visualization was however not an idea found anywhere else, so that envisioning what de-identification methods that would be appropriate was indeed a challenge, in which errors could arise. Literature that considered alternative approaches to protecting privacy in the targeted context, was therefore also included in the study, so that all relevant ideas on the field could be included in the search for the most appropriate implementation. Moreover, there were still other approaches to de-identified real-time visualizations that came up during the other research activities, that were not found anywhere in earlier literature, but nevertheless should be included in the study.

“Current state of research”-articles and literature reviews conducted by other authors were also used in order to find relevant sources that may have been overlooked in the search, including Appari & Johnson (2010), El Emam & Fineberg (2009), Iachello & Hong (2007) and Meystre et al. (2010).

6.4.2 Rapid field tests

The field tests were aimed at primarily generating requirements for information needs of clinicians, and the privacy questions were thus not the first to be asked to the participants. Each session was short in duration, focused on getting initial thoughts on the proposed system from potential users, and there was not much time to explain the concept of de-identification properly. While respondent bias is normally diminished in such sessions, the responses could also have been affected by the respondents not being used to relate to de-identified data sets. Hospital physicians who are on the contrary used to access the complete EPR of every patient they deal with, and may be expected to not approve of de-identified data due to concerns for erroneous treatment. Retrospectively, one can say that the informants who responded very negatively to the idea of de-identified views on large screens, were speaking too hastily not properly knowing their potential usefulness. This problem was therefore countered by introducing a role-play in a lab, which revealed a more positive opinion towards this kind of de-identified information screens.

6.4.3 Usability lab experiment

The lab experiment was not intended to be a conventional usability test, although it was set up in a usability laboratorium. Instead, it was designed to use a realistic scenario involving a role play in order to properly contextualize the system for the participants. The responses here on de-identified visualization were quite different from those when only talking about a de-identified prototype on piece of paper, in general much more positive, possibly indicating a success in this respect.

The experienced benefit may however be imagined, due to the nature of the role-play used. First of all, the scenario was carefully crafted so that we would be able to test all the things we intended to test. In addition, the test person was protected from being intimidated by the system, by only being given simple tasks first, and by portioning the events evenly out so that they were intuitive to follow. Moreover, the context was based on the night shift, which only requires one nurse (our test person) on duty at each bed cluster, in contrast to day and evening shifts with two nurses or more. Nevertheless we introduced a total of eight hospitalized patients in the bed cluster, which would make it full and hence be the “worst case” scenario. In order to make the scenario easier to follow, there was however mainly one patient that required special attention by the test nurse throughout the role-play. The events were still related to a total of three patients, although this could not counter the fact that re-identification of de-identified messages was easier than with a possible worst case scenario. Such a scenario could for instance be where the appearing messages regarded all eight patients, only during a short period of time. Still, the used scenario was considered quite realistic by all four test persons.

As a quantitive alternative to the qualitative approach used in the lab, it could have been possible to do six iterations of ward work and checking the hallway prototype for updates, and then change the de-identification approach for each iteration so that the user could see them all in action, trying to give a clear guess on what happened to whom in the last update. For quantitative measures, the de-identification alternatives would receive scores on how may right and wrong guesses were made using each of them, and how much time they required from the test person. However, this would require a lot more test persons for diminishing person bias, and the sequence must have been shuffled between each test so that possible bias from the sequence itself could be removed. Although it would not require real clinicians to serve as test persons, this approach was believed to become too confusing for the users, and hence excluded early from the experiment design. Doing a qualitative assessment in a focused interview was instead the feasible thing to do in this project's lab experiments, but the small number of test persons can not be sufficient for suggesting a level of de-identification that is the most convenient for nurses in general to use. Even if increased with five or ten participants, the value of the results would not be raised accordingly, since even in the group of four they all disagreed.

6.5 Evaluation

6.5.1 Generalizability of results

The review of possible de-identification techniques, and subsequent selection of those which may be applied to visualizations concerned in this thesis, can also be applied to other similar uses. There are however made some important desicions towards the intended context of use and information included, which may affect the overall generalizability of the exact selection. For instance may the use of pixelation and blurring be more appropriate in a different system, although it has not been totally rejected in this study either. The initial selection done in chapter 2, along with the ideas brought forward in chapter 4, could therefore be included altogether when, before considering contextual constraints for the particular implementation.

The idea of using de-identification for creating an intermediate view of a visualization, may be also applicable to many applications. Since this technique could serve as an enabler of more usable user authentication, it could potentially be considered for other multi-user systems that are either used collaboratively, or/and in public and semi-public areas. The particular implementation of de-identification can however be done in may ways, including variables totally different from those that are used in this study's application.

6.5.2 Lessons learned

When involving real people in both rapid field tests and lab experiments, it became apparent that data collection methods each had their own qualities exclusive to the other. While the rapid field tests enabled quick recruitment and a chance to get some very intial feedback on design ideas, it was apparent that a more in-depth security topic of de-identification was not something that they could assess very precisely, just by looking briefly on a paper sketch. Although the resistance towards e.g. the map where room locations should have been remembered by the clinician was not at all surprising, it would have been impossible to understand how e.g. de-identification alternative II (figure 4.3) could have actually been interpretable in certain situations, after only discussing it over a sheet of paper for five minutes. Instead, illustrating its use in practice through a role-play in the usability lab led to rather a positive assessment overall to this approach, that would have not been found in a more preliminary rapid field test.

This leads to another lesson, being that potential users should not always be believed uncritically when telling what they want or need from a future system. This is not because they intentionally do not tell the truth, but rather because they do not necessarily know for themselves what they really want or need. An example is that in the rapid field tests, a couple of prototypes were presented as specifically made for smart phones. This immediately seduced the informants to believe that receiving all messages on their own smart phone would be the solution to more or less everything, (unknowingly) uncritical to all unwanted interruptions this could potentially lead to.

6.5.3 Questions raised

- Under which circumstances may the proposed approaches to de-identification actually be legal?
- How could generalization be applied in a usable way, enabling additional quasi-identifiers to be publicly disclosed?
- To which extent does information access workload affect utility value?
- If a written consent is required from the patient for such secondary use; what will the visualization look like if individual do not want to appear in the system?

6.6 Further work

The suggestions for an interactive prototype altogether form a new design hypothesis that should be explored further. From the line of reasoning, the author will expect the prototype to be considered both useful and usable by future users, depending on being implemented for an application that can provide more convenient access to important information than today. In order to evaluate this hypothesis, it will hence be natural to involve potential users in a new series of experiments, testing the interactive implementation in practice.

An experiment like this would at this stage be interesting to perform inside a real hospital setting. Using an existing system that is frequently checked for updates, for instance the radiology software where the radiologist's assessments are delivered to nurses and physicians. At the local hospital, this software already has a view that contains a list of examinations during the last 24 hours. This list could therefore be used as a basis for a "Wizard of Oz"-type of interactive test, since it is not feasible to perform a complete integration for use in a usability test only. If this view is placed on a screen of its own, it could be initially displayed as an anonymous list of relevant examinations to the department, for instance. Then, when a user comes by the de-identified view is activated, and finally if identified information is requested, the screen could show this for selected items. The different information levels could however be generated using real functionality, but since there will be no real integration, the data must be added manually by someone who can follow the list in the real system, e.g. the author, if permitted.

The purpose of doing such an experiment, would be to investigate whether the less access-controlled screen has become a valuable asset to the users, or if it was used all. This could for instance be investigated with a user satisfaction survey among the users, and the risks with doing such an experiment would not be very considerable, since the users themselves can choose not to use the system if it is only in the way.

"The software isn't finished until the last user is dead."

Sidney Markowitz
SOFTWARE DEVELOPER

Chapter 7

Conclusions

7.1 Displaying sensitive information on large screens

In this study, a collection of techniques have been evaluated towards their appropriateness for use in real-time visualisations in public and semi-public hospital areas. While some techniques were found in existing literature, some others were added through a series of rapid field tests. The techniques found most appropriate were suppression of variables, coding, masking and generalization. These could all be used in the high-fidelity prototype that facilitated the controlled lab-experiments, and could all be used to support de-identification in line with national legislative requirements. In addition, it is important to analyze the prospected system and context for use in particular, before making desicions on how severe the de-identification process has to be. This will be an essential consideration to make in order to maximize the utility value of the disclosed information.

The evaluation of large screens in public environments however suggests that the information should be rather anonymized in order to align with legislative requirements. These requirements are however not explicit to the extent of exactly which attributes must be removed, in order to achieve sufficient de-identification. Although de-identification may be implemented from several variations over patient identitifiers and quasi-identifiers, only complete removal of patient references will in practice lead to anonymization. Then, there will exist no chance whatsoever to know what has happened to whom, with any degree of certainty, but displaying such abstract and anonymous data could unfortunately not provide much utility for personnel. It may be a way of indicating what has not yet happened, and can thus potentially save some wasted login time, but if personnel have too little knowledge of what has happened when the event icon appears on the screen, they might think even more often that something might have happened that requires their attention — and then much time is wasted.

The possible solutions to this would be either to physically restrict access to the screen, or to use access control mechanisms in the systems that allow the user to be authorized, and hence access less de-identified — possibly even fully identified — information on the screen. A combination of both physical and computational access control may also be possible, however it is vital for such a coordination support system to be more available to the users than existing systems are today. If an access control based solution is pursued in combination with the use of de-identification, it is hence important to choose a level of protection that is proportional to the risk of unwanted disclosure of the information, as well as the damage potential.

7.2 Trading information depth with usable authentication

Time-demanding login routines for workstation computers take up a lot of health personnel's time, especially when the login sessions are short and the purpose may be just checking whether a test result is ready yet. The authentication burden being placed on the health care worker is however the same, regardless of what is the user's purpose for login. While sometimes full access to both reading and editing complete EPR entries is required, the same authentication process is necessary when only checking for the arrival of test results. As revealed in the experiments in chapter 4, this often results in time just being wasted if the information is not yet available in the system.

Although Norwegian legislation does not separate between levels of sensitivity of such information, there is still an obvious difference between the risk of losing a complete EPR entry, and telling someone that CT images are ready for a patient on the ward having initials O.H. and birth year 1950 — without disclosing any of the images or attached comments. Users should therefore be given a way of authenticating themselves with a weaker authentication method, if preferred, and hence only getting access to less de-identified information. Using only an identity card with RFID, the user could for instance be given access to information no more explicit than being well protecting against an intruder stealing someone's ID card. Then, finally, if the user requires precise identification for the displayed information, the ID card could be supplemented with entering a PIN code, and hence reveal an identified view. While the loss of identity card may pose a certain threat to systems that manage PHI, the risk of disclosing sensitive information is accordingly mitigated as of the de-identification routine.

7.3 Empowering the authorized user

In the lab experiments, the four participants did not agree on what would be the preferred level of de-identification for a screen placed in the hallway. Being introduced to a set of six alternatives, they all based their rankings on personal assessments and weighting of several parameters, which altogether affected each alternative's overall evaluation. As a consequence, the outcome of these experiments did not suggest a valid ranking of the proposed alternatives, and neither could it be taken into account for evaluating how useful each alternative was in isolation, as the scenario was very limited. Instead, the results suggest that the users themselves should be involved in configuring the application of de-identification, so that they can add their own preferences to the rendering of a screen in ways that make it useful to them.

Since the clinician who accesses the PHI is also responsible for any re-disclosure of what is accessed, it would be beneficial to provide the users with a usable mechanism for controlling what is displayed on the large screen. If for instance a visitor passes by, it is not user-friendly to require the user to log out entirely from the session, and hence be forced to start over again. Instead, by providing the users with a simple GUI component that can control the level of identification used in the visualization by e.g. "zooming" in and out, it will be simple to take action in less-controlled environments whenever an outsider would pass by.

In such a component, de-identification can be a very useful tool for creating an intermediate view between the anonymized and the identified view. Such an intermediate view can be dynamically brought forward with ease, whenever the user needs to protect information. Moreover, it can be used as described above, when user authentication is not certain enough to provide full access. In chapter 6, an interactive prototype using this approach is illustrated, making the described concept a new design hypothesis for future work in this area.

Bibliography

Agrawal, R. & Johnson, C. (2006), 'Securing electronic health records without impeding the flow of information.', *International journal of medical informatics* **76**(5-6), 471–9.

Altinn (2011), *Altinn — Simplified electronic dialogue*. http://www.altinn.no/, *last accessed: June 27, 2011*.

Andresen, H. (2009), The Policy Debate on Pseudonymous Health Registers in Norway, *in* A. Fred, J. Filipe & H. Gamboa, eds, 'Biomedical Engineering Systems and Technologies', Vol. 25 of *Communications in Computer and Information Science*, Springer Berlin Heidelberg, pp. 413–424.

Andresen, H. (2010), Tilgang til og videreformidling av helseopplysninger, PhD thesis, University of Oslo, Faculty of Law.

Appari, A. & Johnson, M. (2010), 'Information security and privacy in healthcare: current state of research', *International journal of Internet and enterprise management* **6**(4), 279–314.

Aronsky, D., Jones, I., Lanaghan, K. & Slovis, C. M. (2008), 'Supporting patient care in the emergency department with a computerized whiteboard system.', *Journal of the American Medical Informatics Association* **15**(2), 184–194.

Bardram, J. E. (2005), 'The trouble with login: on usability and computer security in ubiquitous computing', *Personal and Ubiquitous Computing* **9**(6), 357–367.

Bardram, J. E. & Hansen, T. R. (2010), Why the plan doesn't hold: a study of situated planning, articulation and coordination work in a surgical ward, *in* 'Proceedings of the 2010 ACM conference on Computer supported cooperative work', ACM, New York, NY, USA, pp. 331–340.

Bardram, J. E., Hansen, T. R. & Soegaard, M. (2006*a*), AwareMedia: a shared interactive display supporting social, temporal, and spatial awareness in surgery, *in* 'Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work', CSCW '06, ACM, New York, NY, USA, pp. 109–118.

Bardram, J. E., Hansen, T. R. & Soegaard, M. (2006*b*), Large Interactive Displays in Hospitals - Motivation, Examples, and Challenges, *in* 'CHI'06: Proceedings of the SIGCHI conference on Human Factors in computing systems'.

Bardram J. E., Kjær T. K., N. C. (2003), Supporting Local Mobility in Healthcare by Application Roaming among Heterogeneous Devices, *in* C. L, ed., 'Proceedings of the Fifth International Conference on Human Computer Interaction with Mobile Devices and Services, volume 2795 of Lecture Notes in Computer Science', Springer Verlag, Udine, Italy, pp. 161—176.

Behlen, F. M. & Johnson, S. B. (1999), 'Multicenter Patient Records Research', *Journal of the American Medical Informatics Association* **6**(6), 435–443.

Bellotti, V. & Sellen, A. (1993), Design for privacy in ubiquitous computing environments, *in* 'Proceedings of the third conference on European Conference on Computer-Supported Cooperative Work', Kluwer Academic Publishers, pp. 77–92.

Benitez, K. & Malin, B. (2010), 'Evaluating re-identification risks with respect to the HIPAA privacy rule.', *Journal of the American Medical Informatics Association* **17**(2), 169–177.

Bjørn, P. & Hertzum, M. (2010), 'Artefactual Multiplicity: A Study of Emergency-Department Whiteboards', *Computer Supported Cooperative Work (CSCW)* **20**(1-2), 93–121.

Boyle, M. & Greenberg, S. (2000), Balancing awareness and privacy in a video media space using distortion filtration, *in* 'Proceedings of Western Graphics Symposium 2000'.

Campbell, B., Thomson, H., Slater, J., Coward, C., Wyatt, K., & Sweeney, K. (2007), 'Extracting Information from Hospital Records: What Patients Think About Consent', *Quality and Safety in Healthcare* **16**(6), 404–408.

Chung, E., Hong, J., Lin, J., Prabaker, M., Landay, J. & Liu, A. (2004), Development and evaluation of emerging design patterns for ubiquitous computing, *in* 'Proceedings of DIS 2004: Designing Interactive Systems', ACM Press, Boston, MA, USA, pp. 233–242.

COSTT (2011), *Co-operation Support Through Transparency*. http://www.costt.no/, *last accessed: May 4, 2011*.

Covington, M. J., Long, W., Srinivasan, S., Dev, A. K., Ahamad, M. & Abowd, G. D. (2001), 'Securing context-aware applications using environment roles', *Proceedings of the sixth ACM symposium on Access control models and technologies - SACMAT '01* pp. 10–20.

Danao, G. (2010), 'The SMU eLearning System: A Students Usability Assessment', *SCSIT Research Journal*.

DePaula, R., Ding, X., Dourish, P., Nies, K., Pillet, B., Redmiles, D., Ren, J., Rode, J. & Silvafilho, R. (2005), 'In the eye of the beholder: A visualization-based approach to information system security', *International Journal of Human-Computer Studies* **63**(1-2), 5–24.

Devik, K. (2009), Brukbarhetstesting av en mobil butikkløsning i laboratorium og felten, Master's thesis, Norwegian University of Science and Technology (NTNU).

eHealth (2011), *ICST Conference on Electronic Healthcare for the 21st century*. http://www.electronic-health.org/, *last accessed: March 30, 2011*.

El Emam, K. (2008), 'Heuristics for De-identifying Health Data', *IEEE Security & Privacy Magazine* **6**(4), 58–61.

El Emam, K. (2010), 'Risk-Based De-Identification of Health Data', *IEEE Security & Privacy Magazine* **8**(3), 64–67.

El Emam, K. & Dankar, F. K. (2008), 'Protecting Privacy Using k-Anonymity', *Journal of the American Medical Informatics Association* **15**(5), 627–637.

El Emam, K. & Fineberg, A. (2009), 'An overview of techniques for de-identifying personal health information', *Access to Information and Privacy Division of Health Canada*.

El Emam, K. & Kosseim, P. (2009), 'Privacy Interests in Prescription Data, Part II: Patient Privacy', *IEEE Security & Privacy Magazine* **7**(1), 75–78.

European Union (1995), *Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data*. http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CONSLEG:1995L0046:20031120:EN :PDF, *last accessed: June 24, 2011*.

Faxvaag, A., Røstad, L., Tøndel, I. A., Seim, A. R. & Toussaint, P. J. (2009), Visualizing Patient Trajectories on Wall-Mounted Boards — Information Security Challenges, *in* K.-P. Adlassnig, B. Blobel, J. Mantas & I. Masic, eds, 'Medical Informatics in a United and Healthy Europe - Proceedings of MIE 2009', IOS Press, pp. 715–719.

Friedlin, F. J. & McDonald, C. J. (2008), 'A Software Tool for Removing Patient Identifying Information from Clinical Documents', *Journal of the American Medical Informatics Association* **15**(5), 601–610.

Fuglseth, E. (2008), Brukervennlig og sikker innlogging på sykehus: dagens praksis og forslag til alternative løsninger, Master's thesis, Norwegian University of Science and Technology (NTNU).

Fung, B. C. M., Wang, K. & Yu, P. S. (2005), Top-Down Specialization for Information and Privacy Preservation, *in* 'ICDE'05', pp. 205–216.

Gardner, J. & Xiong, L. (2008), HIDE: An Integrated System for Health Information DE-identification, *in* '2008 21st IEEE International Symposium on Computer-Based Medical Systems', Ieee, pp. 254–259.

Gass, P., Walton, E., Winlow, R., Sagardoyburu, M., Stubbs, P., Kean, D., Tillin, M., Bourhill, G., Yabuta, K. & Takatani, T. (2007), 'Privacy LCD technology for cellular phones', *Sharp Tech. J.* **30**(95), 45–49.

Golle, P. (2006), Revisiting the uniqueness of simple demographics in the US population, *in* 'Proceedings of the 5th ACM workshop on Privacy in electronic society', ACM, pp. 77–80.

Google (2011*a*), *Google Patents*. http://patents.google.com/, *last accessed: May 31, 2011*.

Google (2011*b*), *Google Scholar*. http://scholar.google.com/, *last accessed: May 31, 2011*.

Grimes, A. & Tarasewich, P. (2005), Investigating Privacy-Augmented Displays for Mobile Devices, *in* 'Proc. of Human-Computer Interaction International 2005'.

Gross, R., Airoldi, E., Malin, B. & Sweeney, L. (2006), Integrating Utility into Face De-identification, *in* G. Danezis & D. Martin, eds, 'Privacy Enhancing Technologies', Vol. 3856 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, pp. 227–242.

Heckle, R. R. & Lutters, W. G. (2011), 'Tensions of network security and collaborative work practice: Understanding a single sign-on deployment in a regional hospital.', *International journal of medical informatics* pp. 1–13.

Huang, L.-C., Chu, H.-C., Lien, C.-Y., Hsiao, C.-H. & Kao, T. (2008), 'Embedding a Hiding Function in a Portable Electronic Health Record for Privacy Preservation', *Journal of Medical Systems* **34**(3), 313–320.

Iachello, G. & Hong, J. (2007), 'End-User Privacy in Human-Computer Interaction', *Foundations and Trends in Human-Computer Interaction* **1**(1), 1–137.

International Standards Organization (1998), *ISO 9241-11: Ergonomic for Office Work with Visual Display Terminals (VDTs) — Part 11: Guidance on Usability*. http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm? csnumber=16883, *last accessed: June 11, 2011*.

Ishii, H. (1990), TeamWorkStation: towards a seamless shared workspace, *in* 'Proceedings of CSCW 90', ACM Press, Los Angeles, pp. 13—26.

Ismail, M., Osman, A. & Wahab, N. (2009), Evaluation of Location-Aware Travel Guide, *in* 'Computer Technology and Development, 2009. ICCTD '09. International Conference on', pp. 237–239.

Kifer, D. & Gehrke, J. (2006), Injecting utility into anonymized datasets, *in* 'SIGMOD Conference'06', pp. 217–228.

Kobsa, A. & Schreck, J. (2003), 'Privacy through pseudonymity in user-adaptive systems', *ACM Transactions on Internet Technology* **3**(2), 149–183.

Landmark, A. D. (2009), Capability Assessment of Indoor Positioning Systems, Master's thesis, Norwegian University of Science and Technology (NTNU).

Lauesen, S. (2005), *User Interface Design: A Software Engineering Perspective*, Pearson Education, Harlow, UK.

Lillebo, B., Seim, A. R. & Faxvaag, A. (2010), Situated Coordination of Healthcare Professionals: A Field Study of Operating Room Personnel, *in* M. H. Safran, C., Reti, S., ed., 'MedInfo 2010 - Proceedings of the 13th World Congress on Medical Informatics', IOS Press.

Lillebo, B., Seim, A. R. & Faxvaag, A. (2011), Information and communication needs of healthcare workers in the perioperative domain, *in* 'Forthcoming: XIII International Conference of the European Federation for Medical Informatics', Oslo, Norway.

Little, L., Briggs, P. & Coventry, L. (2005), 'Public space systems: Designing for privacy?', *International Journal of Human-Computer Studies* **63**, 254–268.

Macrae, G. (2011), 'E3 Expo 2011: Sony unveils 3D PlayStation TV alongside PlayStation Vita', International Business Times online edition. http://uk.ibtimes.com/articles/158738/20110607/e3-expo-2011-sony-psn-hack-3d-playstation-tv-alongside-playstation-vita.htm, *last accessed: June 7, 2011*.

McQuaid, M., Zheng, K., Melville, N. & Green, L. (2009), Usable deidentification of sensitive patient care data, *in* 'Proceedings of the 5th Symposium on Usable Privacy and Security - SOUPS '09', ACM Press, New York, NY, USA.

Melby, L. & Toussaint, P. J. (2009), 'Supporting operating nurses' collaborative work: Preventing information overload and tailoring information access', *22nd IEEE International Symposium on Computer-Based Medical Systems*.

Mendeley Ltd. (2011), *Mendeley: Academic reference management software for researchers*. http://www.mendeley.com/, *last accessed: April 15, 2011*.

Meystre, S. M., Friedlin, F. J., South, B. R., Shen, S. & Samore, M. H. (2010), 'Automatic de-identification of textual documents in the electronic health record: a review of recent research.', *BMC medical research methodology* **10**(1), 70–80.

Moskop, J. C., Marco, C. A., Larkin, G. L., Geiderman, J. M. & Derse, A. R. (2005), 'From Hippocrates to HIPAA: privacy and confidentiality in emergency medicine–Part I: conceptual, moral, and legal foundations.', *Annals of emergency medicine* **45**(1), 53–9.

Munkvold, G., Ellingsen, G. & Monteiro, E. (2007), From plans to planning: the case of nurse care plans, *in* ‘Proceedings of the 2007 international ACM conference on Supporting group work’, ACM, New York, NY, USA, pp. 21–30.

Nielsen, J. & Landauer, T. K. (1993), A mathematical model of the finding of usability problems, *in* ‘Proceedings of the INTERACT ’93 and CHI ’93 conference on Human factors in computing systems’, CHI ’93, ACM, New York, NY, USA, pp. 206–213.

Norwegian statute (2000), *Act of 14 April 2000 No. 31 relating to the Processing of Personal Data, An English translation*. http://www.ub.uio.no/ujur/ulovdata/lov-20000414-031-eng.pdf, *last accessed: June 1, 2011*.

Norwegian statute (2001), *Act of 18 May 2001 No. 24 on Personal Health Data Filing Systems and the Processing of Personal Health Data*. http://www.regjeringen.no/nb/dep/hod/tema/folkehelse/Act-of-18-May-2001-No-24-on-Personal-Health-Data-Filing-Systems-and-the-Processing-of-Personal-Health-Data-Personal-Health-Data-Filing-System-Act-.html?id=224129, *last accessed: June 23, 2011*.

Oates, B. J. (2006), *Researching Information Systems and Computing*, SAGE Publications, London, UK.

Ohno-Machado, L., Silveira, P. S. P. & Vinterbo, S. (2004), ‘Protecting patient privacy by quantifiable control of disclosures in disseminated databases’, *International journal of medical informatics* **73**(7-8), 599–606.

O’Neill, E., Woodgate, D. & Kostakos, V. (2004), Easing the wait in the emergency room: building a theory of public information systems, *in* ‘Proceedings of the 5th conference on Designing interactive systems: processes, practices, methods, and techniques’, ACM, New York, NY, USA, pp. 17–25.

Pedersen, D. M. (1999), ‘Model for types of privacy by privacy functions’, *Journal of Environmental Psychology* **19**(4), 397—405.

Peffers, K., Tuunanen, T., Rothenberger, M. A. & Chatterjee, S. (2006), ‘The design science research process: A model for producing and presenting information’, *Journal of Management Information Systems* **24**(Number 3 / Winter 2007-2008), 45–77.

Preece, J., Rogers, Y. & Sharp, H., eds (2002), *Interaction Design: Beyond Human-Computer Interaction*, John Wiley and Sons, New York, USA.

Reich, W. T., ed. (1995), *Encyclopedia of Bioethics Vol. 5*, Macmillan, New York, NY, pp. 2646–2647.

Riedl, B., Grascher, V., Fenz, S. & Neubauer, T. (2008), Pseudonymization for improving the Privacy in E-Health Applications, *in* 'Proceedings of the 41st Annual Hawaii International Conference on System Sciences (HICSS 2008)', IEEE Computer Society, Los Alamitos, CA, USA, pp. 255–255.

Robson, C. (2002), *Real World Research*, Blackwell Publishing, Oxford, UK.

Rogers, Y. & Rodden, T. (2002), Designing new workspaces to provide physical and social affordances for successful interaction, *in* 'Workshop paper presented at CSCW 2002'.

Røstad, L. & Edsberg, O. (2006), A Study of Access Control Requirements for Healthcare Systems Based on Audit Trails from Access Logs, *in* 'Proceedings of the 22nd Annual Computer Security Applications Conference', IEEE Computer Society, Washington, DC, USA, pp. 175–186.

Saltzer, J. & Schroeder, M. (1975), 'The protection of information in computer systems', *Proceedings of the IEEE* **63**(9), 1278–1308.

Sandberg, W., Ganous, T. & Steiner, C. (2003), 'Setting a research agenda for perioperative systems design', *Semin Laparosc Surg* **10**, 57–70.

Scupelli, P., Xiao, Y., Fussell, S., Kiesler, S. & Gross, M. (2010), Supporting coordination in surgical suites: physical aspects of common information spaces, *in* 'Proceedings of the 28th international conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 1777–1786.

Seland, G. (2010), Role-Play Workshops as a User-Centred Design Method for Mobile IT, PhD thesis, Norwegian University of Science and Technology (NTNU).

Shneiderman, B. (1997), *Designing the user interface: Strategies for Effective Human-Computer Interaction*, Addison-Wesley Longman Publishing, Boston, MA, USA.

Shoemaker, G. B. D. & Inkpen, K. M. (2001), Single display privacyware: augmenting public displays with private information, *in* 'Proceedings of the SIGCHI conference on Human factors in computing systems', ACM, New York, NY, USA, pp. 522–529.

SOUPS (2011), *Symposium on Usable Privacy and Security*. http://cups.cs.cmu.edu/soups/, *last accessed: March 30, 2011*.

Stasko, J., McColgin, D., Miller, T., Plaue, C. & Pousman, Z. (2005), Evaluating the InfoCanvas Peripheral Awareness System: A Longitudinal, In Situ Study, Technical Report Technical Report GIT-GVU-05-08, GVU Center/Georgia Institute of Technology, Atlanta, GA, USA.

Sweeney, L. (2001), Computational Disclosure Control — A Primer on Data Privacy Protection, PhD thesis, Massachusetts Institute of Technology.

Sweeney, L. (2002), 'k-anonymity: a model for protecting privacy', *International Journal on Uncertainty Fuzziness and Knowledgebased Systems* **10**(5), 557–570.

Szarvas, G., Farkas, R. & Busa-Fekete, R. (2007), 'State-of-the-art Anonymization of Medical Records Using an Iterative Machine Learning Framework', *Journal of the American Medical Informatics Association* **14**(5), 574–580.

Tarasewich, P., Avenue, H. & Campbell, C. (2005), What Are You Looking At?, *in* 'In Proc. SOUPS 2005', number 1368.

Tjora, A. (2010), *Kvalitative forskningsmetoder i praksis*, Gyldendal Akademisk, Oslo, Norway.

Truta, T. M. & Vinay, B. (2006), Privacy Protection : p-Sensitive k-Anonymity Property, *in* 'Data Engineering Workshops, 2006. Proceedings. 22nd International Conference on', IEEE Computer Society, pp. 94–104.

United States of America (1996), *The Health Insurance Portability and Accountability Act (P.L.104-191)*. http://www.cms.gov/HIPAAGenInfo/Downloads/HIPAALaw.pdf, *last accessed: June 11, 2011*.

Uzuner, O., Luo, Y. & Szolovits, P. (2007), 'Evaluating the state-of-the-art in automatic de-identification', *Journal of the American Medical Informatics Association* **14**(5), 550–563.

Vaast, E. (2007), 'Danger is in the eye of the beholders: Social representations of Information Systems security in healthcare', *The Journal of Strategic Information Systems* **16**(2), 130–152.

Wang, K., Yu, P. S. & Chakraborty, S. (2004), Bottom-Up Generalization: A Data Mining Solution to Privacy Protection, *in* 'ICDM'04', pp. 249–256.

Wellner, B., Huyck, M., Mardis, S., Aberdeen, J., Morgan, A., Peshkin, L., Yeh, A., Hitzeman, J. & Hirschman, L. (2007), 'Rapidly retargetable approaches to de-identification in medical records', *Journal of the American Medical Informatics Association* **14**(5), 564.

Westin, A. (1967), *Privacy and Freedom*, Atheneum, New York.

Wikipedia (2011), *Usability testing*. http://en.wikipedia.org/wiki/Usability_testing #Hallway_testing, *last accessed: May 11, 2011*.

Xiao, X. & Tao, Y. (2006), Anatomy: Simple and Effective Privacy Preservation, *in* 'VLDB'06', pp. 139–150.

Yeniterzi, R., Aberdeen, J., Bayer, S., Wellner, B., Hirschman, L. & Malin, B. (2006), 'Effects of personal identifier resynthesis on clinical text de-identification.', *Journal of the American Medical Informatics Association : JAMIA* **17**(2), 159–168.

Yoga, R. (2010), 'Hvorfor må innlogging i Altinn være så vanskelig?', Brukskvalitet.no — en fagblogg om brukervennlighet. http://www.brukskvalitet.no/2010/hvorfor-ma-innlogging-i-altinn-v%c3%a6re-sa-vanskelig/#more-1207, *last accessed: June 27, 2011*.

Zave, P. (1995), Classification of research efforts in requirements engineering, *in* 'Proceedings of the Second IEEE International Symposium on Requirements Engineering', RE '95, IEEE Computer Society, Washington, DC, USA.

Appendix A

Paper accepted for MIE 2011 conference

XXIII International Conference of the European Federation for Medical Informatics, Oslo, Norway, August 28–31, 2011.

Personal Health Information on Display: Balancing Needs, Usability and Legislative Requirements

Erlend Andreas GJÆRE [a], Inger Anne TØNDEL [b, 1], Maria B. LINE [b],
Herbjørn ANDRESEN [c], and Pieter TOUSSAINT [a]
[a] *Dep. of Computer and Information Science, NTNU, Trondheim, Norway*
[b] *SINTEF ICT, Trondheim, Norway*
[c] *Dep. of Private Law, University of Oslo, Norway*

Abstract. Large wall-mounted screens placed at locations where health personnel pass by will assist in self-coordination and improve utilisation of both resources and staff at hospitals. The sensitivity level of the information visible on these screens must be adapted to a close-to-public setting, as passers-by may not have the right or need to know anything about patients being treated. We have conducted six informal interviews with health personnel in order to map what kind of information they use when identifying their patients and their next tasks. We have compared their practice and needs to legislative requirements and conclude that it is difficult, if not impossible, to fulfil all requirements from all parties.

Keywords. Personal health information, de-identification, privacy, coordination

1. Introduction

The COSTT[2] project aims at supporting coordination in the peri-operative hospital environment by visualising status information regarding current operations and patients under treatment on large wall-mounted screens. This will help the personnel predicting when their time and effort are needed, and which colleagues are available for advice or assistance. As a result, both physical resources and staff can be utilised more effectively. Research on similar computerised coordination systems implemented as electronic whiteboards are also presented by Bardram et al. [1] and Aronsky et al. [2]. In order to maximise coordination support, the screens should be placed at locations where the relevant health personnel are likely to see them, e.g. in corridors. This however makes them available to everybody present, including patients, their relatives, and personnel not directly involved in patient treatment (e.g. cleaners and technicians). Such availability has consequences for the privacy of patients and employees.

In previous work [3] we have introduced the concept of *flexible de-identification*, and described how it is possible to present patient information at various levels of

[1] Corresponding Author: Inger Anne Tøndel, SINTEF ICT, N-7465 Trondheim, Norway; e-mail: inger.a.tondel@sintef.no

[2] Co-operation support Through Transparency, http://costt.no/

details, both with regards to identifying information and the medical condition. Three perspectives have to be taken into account when developing solutions for de-identification. The first perspective is that clinical personnel require a certain amount of identifying information for the medical information presented to be meaningful and useful. The second perspective is that laws and regulations restrict the amount of patient identifying information that can be presented. The last perspective is usability. A system that requires users to log on to multiple systems in order to obtain patient information, might fulfil both the information need and requirements set by laws and regulations, but is not very usable in a dynamic work environment where clinicians work under time pressure. These three perspectives generate different demands, and designing the right level of de-identification means balancing these different demands.

The rest of the paper is organised as follows: Section 2 presents the results of unstructured interviews with personnel working in the surgical clinic at a Norwegian hospital, and Section 3 outlines the Norwegian legislative requirements. Then, Section 4 discusses how needs, usability and legislative requirements can be balanced, and Section 5 concludes the paper.

2. Interviews

In order to improve our understanding of the information needs of health care personnel, and specifically their need for identifying information, we conducted six unstructured interviews at Trondheim University Hospital, during November-December 2010. Six different identification approaches were explored (see overview in Table 1), where the one with highest identification level used initials and birth year of the patient. The less identified approaches aimed to identify the patient by his location or his relation to health care personnel, possibly in combination with the test or surgery type performed. In the interviews we wanted to gain feedback on whether the less identifying approaches still resulted in useful status information for health care workers.

The participating clinicians included one senior physician and two ward nurses from the Department of Gastrointestinal Surgery, one junior physician and one nurse from the Department of Emergency, one ward nurse from the Department of Breast and Endocrine Surgery, and one charge nurse from a ward at the Department of Orthopaedic Surgery. Their ages ranged from 25 to 55, and all had been in their position for some while. The informants were recruited randomly during work hours, and interviewed straight away in their regular work environment. They were each asked to comment on some early-stage paper-based prototypes of information visualizations, containing message examples related to the treatment progress of patients, e.g. “CT-image description is ready” and “Patient has been scheduled for surgery”. We explored in total four different prototypes, but only one or two were presented to each informant. Some status messages were added during the process, and two of the prototypes were modified slightly in-between interviews, due to feedback given. The prototypes mainly differentiated on how information was organised and how the patients were identified). We used the prototypes to investigate whether the clinicians would be able to tell patients’ identities apart with the different identification approaches, and to evaluate how these related to current practices. The feedback was recorded with handwritten field notes, and written out directly afterwards.

The results of the interviews are summarised in Table 1. Generally, clinicians were positive to the idea of integrating status updates from several systems. Most were still reluctant to the immediate thought of placing *any* patient information more publicly available than workstations or personal devices. Though the approach where patients are identified by initials and birth year stood out as the most convenient option, our main impression is that health care personnel have varying needs for patient identification, depending on their role and the context where identification should happen. We also discovered that clinicians commonly used patients' diagnosis or treatment history as de-identification in conversations between colleagues, (e.g. "he with ileus who needs another operation in three days").

Table 1. De-identification approaches explored in the interviews, based on the paper-based prototypes.

Approach	Example	Summary of responses
Initials and birth year of patient	JD59	Will normally provide fairly good accuracy. Patients having the same birth year and initials (or last name) do however occur. Clinicians still found this convenient as they are used to working with basis in the patients and their name/age (various combinations of name and birth date are used today).
Room number/location	*(Plotted on a map of wards)*	Patients move around (this may leave room lists temporarily inconsistent) or they can even be placed in the corridors. Room numbers are commonly used for reference today, but in combination with other identifiers, e.g. name, diagnosis or sex. It seems hard to remember the patients' exact locations.
Initials of responsible physician *(first two letters of both first name and last name)*	DAJO	Patients are not followed up by only one physician, and physicians attend many patients at each ward. Nurses will not necessarily know the name of the physician providing care for each of their patients at a specific time.
Blood test indicators, time and responsible nurse	Hb, Na, INR 10:41 (HAPE)	Blood tests are ordered as standardised batches, so important indicators, if any (e.g. INR may decide whether to operate or not), do not stand out. Tests for several patients are often ordered at the same time, and by the same nurse, too.
Radiology type, level of urgency, time and referring physician	CT abdomen (red) 11:00 (DAJO)	Some results (MR) take days to arrive, and often 20-30 patients with abdominal pains arrive daily. Hence, a list of pending results may become overloaded and hard to interpret.
Operation room number, surgery type, scheduled time and surgeon initials	OP3: Appendicitis 11:00 (PT)	Nurses rarely know exactly what room an operation will take place in. But as it is uncommon to have several patients from the same ward undergoing surgery at the same time, they may still be able to deduce which operation to follow.

3. Legislative Requirements

In Norway, rules and regulations on the obligation of secrecy, and the criteria for sharing or disclosing data, are mainly found in the Personal Health Data Filing System Act [5] which implements the EU personal data protection directive [4] for the health domain, and in the Health Personnel Act [6] which are national rules of conduct for health personnel. The authorisation rule for granting access to health data [5] consists mainly of two criteria. The first is a general need-to-know restriction: "Access may only be granted insofar as this is necessary for the work of the person concerned" [5]. The second criterion is that access must be "in accordance with the rules that apply regarding the duty of secrecy" [5]. The general rule on secrecy goes beyond a mere duty to "keep silent". It is a proactive duty on institutions as well as individual health personnel to "prevent others from gaining access to or knowledge of information relating to people’s health or medical condition” [6]. There are a few derogations to the secrecy rule [6], mainly the need to share information with co-operating health personnel, the duty to supply patient administrative systems with key data, and a few more rules on sharing information with a patient’s next of kin, and with students, health care assistants or data processing expertise. However, there are no general permissions for making health data available to *other* patients, or to other patients’ next of kin.

There are, in principle, two possible strategies on how to make the envisioned wall-mounted displays legitimate under data protection law. The first strategy would be to generalise or trivialise the data in ways that put the information content below the threshold of “relating to people’s health or medical condition”. An example could be to make the displayed data read something like “patient x to be present in room 101 from 9:30 to 14:00” without revealing what activities would take place there. The second strategy would be some sort of de-identification of the patient, in order to avoid that the displayed data pertains to a specific part of the definition of “personal health data” [5], namely a criterion that it “may be linked to a natural person”.

Norwegian law contains several useful concepts for de-identification [5]. These legal concepts were initially aimed at central health registers, spanning information originating from different hospitals, but they could also be relevant for de-identification purposes within a single hospital. The definition of “de-identified personal health data” has two components. First, any identifying data is removed. Second, any re-identification shall be *dependent* on re-supplying the data that was removed. This second component implies a high threshold; an acceptable level of de-identification may not be pro forma, and re-linking data to the right patient cannot be easily accomplished by guessing. An alternative is to aim for “pseudonymous health data”, which implies that identifying information is encrypted.

4. Discussion

The interviews indicate that status updates for patients under treatment are useful. Health care personnel would like to know when test results are ready, how operations proceed, etc. Making such information easily available on wall-mounted screens will however expose the information to everybody who has physical access, something that is not permitted by Norwegian legislation. As mentioned in Section 3, two main strategies are available in order to adhere to the legal restrictions: Removing all health-related information or de-identifying the information. The first strategy may work for some events, but using it as a general strategy, will probably render the system useless. The second strategy seems more appealing, as it can supply more useful information. Finding an appropriate level of de-identification that makes personnel able to identify patients yet remains a challenge.

Results from the interviews reveal that variations over name and birth date are commonly used for identification. At a ward with a limited number of patients, this close to identifies most patients. The other de-identification techniques tested in the interviews, such as using the room number or the identity of health care personnel, turned out not to be usable. Thus we need to work on alternative de-identification methods. Existing literature on de-identification of health information [7] is mainly concerned with de-identification of large datasets that are to be used for secondary purposes (e.g. research). Still we plan to look into how existing techniques such as pseudonymisation can be used for our setting. We will also investigate to what extent information will still be useful if all identifiers are removed.

If it turns out that the level of de-identification required by legislation will render the system useless, we are left with no option but to limit access to the information to authorised personnel only. This can be ensured by placing the screens at locations where only health personnel have access or by access control mechanisms on the screens, although this will exceedingly reduce the usability for coordination purposes. If such an approach is necessary, it will be important to investigate smart ways of doing access control, e.g. by providing more details on a personal handheld device, or by mechanisms that automatically detect who is present and present information based on the access rights of that group of people.

Reducing the level of identification will result in an increased risk of erroneous interpretation of information. Though this will reduce the benefits of the coordination support system, it is important to state that the system will not replace any of the medical information systems. These will still use full identification for all medical data, and thus there should be no increased risk of treatment errors.

5. Conclusion

Public display of health information poses an obvious risk to patient privacy, and thus there is a need to determine the appropriate level of identification. As the legislative requirements are in conflict with the needs of health personnel, it may be impossible to fulfil all the legislative requirements, without sacrificing usability.

Acknowledgments

We would like to thank Børge Lillebo for his work on the prototypes and cooperation on the interviews. Thanks also to our other colleagues in the COSTT project, Arild Faxvaag especially, for useful comments and discussions. This work was supported by the Norwegian Research Council’s VERDIKT program (grant no. 187854/S10).

References

[1] J.E. Bardram, T. Hansen, M. Soegaard, AwareMedia – A Shared Interactive Display Supporting Social, Temporal, and Spatial Awareness in Surgery, *Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work (CSCW '06)* (2006), 109-118.

[2] D. Aronsky, I. Jones, K. Lanaghan, C.M. Slovis, Supporting Patient Care in the Emergency Department with a Computerized Whiteboard System, *Journal of the American Medical Informatics Association* **15** (2008) , 184-193.

[3] A. Faxvaag, L. Røstad, I.A. Tøndel, A.R. Seim, P.J. Toussaint, Visualizing Patient Trajectories on Wall-Mounted Boards – Information Security Challenges, *Studies in Health Technology* **150** (2009), 750-759.

[4] Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data

[5] Act on personal health data filing systems and the processing of personal health data [Personal Health Data Filing System Act]

[6] Act of 2nd July 1999, no 64 relating to health personnel etc. [The Health Personnel Act]

[7] K. El Emam, A. Fineberg, An overview of Techniques for De-identifying Personal Health Information, Health Canada, January 2009.

Appendix B

Paper accepted for HFEHI 2011 symposium

The fifth international symposium on Human Factors Engineering in Health Informatics, Trondheim, Norway, August 26-27, 2011.

Combining Field Studies and Laboratory Experiments in the Design and Preclinical Validation of the UI of a Hospital Patient Care System

Børge Lillebo, Erlend Andreas Gjære, Arild Faxvaag

Norwegian EPR Research Centre,

Norwegian University of Science and Technology, Trondheim

Abstract

Coordination of patient care at a surgical ward is a challenge. Complex examinations and treatment protocols require multiple hospital actors to stay informed on what happens at different stages. In this study we have combined iterative prototyping and user tests to evaluate the utility value of distributing such information in real-time. Six paper prototypes used in field tests were used to elicit a high-fidelity prototype populated with realistic data, for use in a usability laboratory experiment. Here, a role-play session was first used to contextualise the prototype for the participants, while a focused interview afterwards gathered in-depth feedback on the different parts of the prototype design. Results indicate that real-time information on certain patient care activities are useful for improving coordination of care. In addition, the methodology has provided valuable feedback while causing minimal impact of hospital activities.

1. Introduction

Surgical patients are taken care of by teams of hospital actors. These actors represent various clinical and non-clinical specialities[1], all of which depend on the work of the others. The coordination of these actors' patient care activities is complicated by the fact that the actors belong to different organisational entities and work apart – thus impeding their ability to monitor the collaborative progress of patient care. Adding to this, patients' illnesses are unpredictable and may require rapid changes in care activities and priorities.

Developments in information and communication technologies (ICT) have demonstrated possible solutions to coordination challenges within the health domain [1]. A key feature of such technology is the distribution of real-time information of patient care status, enabling the alignment of actors' individual care activities.

In this study, our primary goal was to explore whether real-time visualisation of information about patient care activities has the potential to improve coordination of care from the perspective of a surgical ward nurse. By combining field studies, iterative prototyping and laboratory experiments we tried to falsify our main hypothesis: 'Real-time information of patient care activities is useful for coordination of care'. Our secondary goal was to use tangible prototypes to validate supplementary hypotheses on what pieces of information that should be included in such an information system. We considered the

[1] Operating room (OR) nurses, ward nurses, post-anaesthesia care unit nurses, surgeons, anaesthesiologists, anaesthetic nurse, OR technicians, cleaners, secretaries, coordinators, radiographers, radiologists and laboratory technicians.

prototypes as a catalyst for validating information requirements rather than trying to assess the prototypes per se – this being in contrast to conventional usability testing.

2. Methods

First we conducted field studies including both observations and interviews with healthcare professionals working in the domain. This part focused on information needs in situations where the healthcare professional was trying to decide on what to do next. In practice, this involved care events that could support perioperative patient handovers e.g. “Patient entered the OR”, “Surgery ended” and “Patient returning to ward”. Further details on this can be found in another article [2]. When analysing the data it became clear that some of the actors – including the ward nurse – were in need of information about a broader set of patient care events than those directly related to perioperative patient handovers. Events related to planning and execution of imaging, laboratory analyses and specialist referrals were also regarded as valuable candidates for facilitating coordination. Hence we expanded our set of care activities and began designing paper prototypes that would illustrate a real-time patient care activity information system.

Six different paper prototypes visualising varying information about patient care activities were presented to 13 hospital actors in simple field tests. This involved recruiting physicians and nurses at work, asking them to give their immediate opinions on to what extent the prototype would support their work. All prototypes were not presented to all actors, due to the parallel and iterative prototyping process. The feedback from the tests provided further insight in information needs, visualisation alternatives and potential effects on the work of the end users. From this we chose to invest more time in a prototype that contained information about specific events from several patients and patient care activities in a time-oriented way, i.e. with the newest event sorted on top – similar to micro blogs on the Internet (e.g. Twitter, Facebook). Other alternatives were time-line visualisations, and patient-oriented, room-oriented and activity-oriented prototypes.

Based on the feedback from the paper prototyping, we developed two complementary high-fidelity prototypes and populated them with realistic patient data. The main prototype (figure 1) consisted of three major parts: a multi-patient care event log, a current status patient list, and a futuristic “next patient care event” forecast. The other prototype (from now on referred to as the “de-identified prototype”) consisted only of a de-identified copy of the multi-patient care event log, i.e. without patient names (figure 2). It was intended to be placed in an open area where it easily could be seen, such as the ward hallway, and thus provide at-a-glance coordination support to the nurses. Since both patients and visitors are likely to also have access to such open areas, the concept of de-identification was tested as an approach to protecting the privacy of the patients whose health information is put on display. More information about how several de-identification techniques were explored with respect to both usability and legislative properties, is available in another article [3]. The quality of the user interfaces of the prototypes were of such that they could easily be mistaken for fully functional systems (in contrast to paper prototypes), when in fact they only consisted of series of images and were non-interactive.

Although the information provided by the prototypes and the scenario that we used in the experiments were fictive, the role-play session was considered realistic as it was based on the previous experience of one of the authors (BL) as a student, intern and resident at surgical departments in three different hospitals. One pilot experiment was done to ensure that laboratory equipment functioned as planned and to practice on the experimental script.

Participants were recruited from the surgical clinic of a university hospital. All participants were nurses and had several years of experience with surgical patient care (25-35 years old). The participants were recruited by one of the authors (BL) personally visiting their respective departments. A total of four nurses participated one hour each.

The experiments took place in the usability lab of the Norwegian Electronic Health Record Research Centre. Each experiment consisted of three parts: Introduction (10 min.), role-playing (15 min.), and a structured focused interview (35 min.). This experimental design was chosen to enable realistic contextualisation of the prototypes, in order to validate information needs and to get more thorough feedback from the nurses compared to what we already knew from field observations and field interviews. The lab was configured as a surgical ward with patient rooms, a corridor and a central desk. The main prototype was available on a desktop computer located at the central desk, and in the corridor – on a wall-mounted large monitor – the de-identified prototype was visible.

All experiments were videotaped and the conversations were later transcribed. Participants gave their informed, written consent before the experiment started. The experiments were approved by the Norwegian Social Science Data Services.

3. The role-play

The role-play started with a simulated meeting were one of the authors acted as the departing nurse and the test subject as the next shift nurse. During this meeting information about the eight patients at the ward was given both orally and written. The written information was given as a patient list, almost identical to the one that is actually used at the hospital the nurses were recruited from. The test subject could also ask for supplementary information about the patients, at his/her own request.

After the meeting, the nurse was presented with the main prototype for the first time. The prototype provided full access to all recent care events[2] of all the eight patients (figure 1).

Having explored the main prototype, the nurse was told to visit patient rooms to carry out regular ward work. Instead of actually doing such work (since our patient rooms were empty) the nurse was instructed to pretend that this work was done and that time went by. Furthermore, he/she was instructed to leave the patient room and have a look at the de-identified prototype in the corridor. The first time the nurse used the de-identified prototype, no new patient care events had occurred, compared to what the nurse already had seen on the main prototype. The nurse had never seen this prototype before, and was therefore given a simple explanation of the features of the de-identified view.

The nurse was also told to verify the information provided in the de-identified prototype by cross-checking with the main prototype. After this the nurse was supposed to visit a second patient room and the same story repeated: no patient, time went by, and the nurse was instructed to look at the screen in the corridor. This time the de-identified prototype contained new patient care information – “troponin and haemoglobin test results were available for one of the eight patients 1 minute ago” (figure 2).

[2] In this context “care event” denotes what we considered to be an important event for coordination of care.

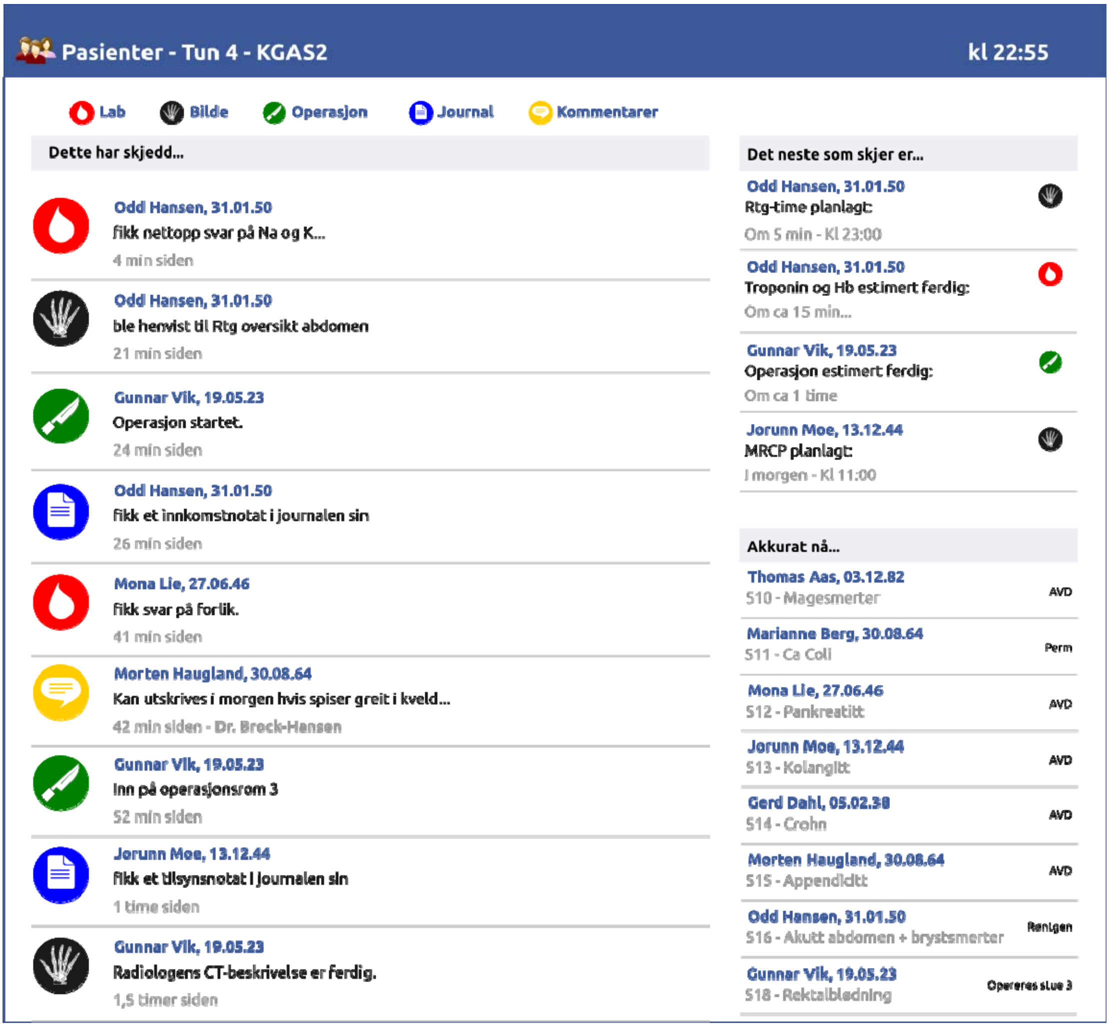

Pasienter - Tun 4 - KGAS2 kl 22:55

Lab Bilde Operasjon Journal Kommentarer

Dette har skjedd...

Odd Hansen, 31.01.50
fikk nettopp svar på Na og K...
4 min siden

Odd Hansen, 31.01.50
ble henvist til Rtg oversikt abdomen
21 min siden

Gunnar Vik, 19.05.23
Operasjon startet.
24 min siden

Odd Hansen, 31.01.50
fikk et innkomstnotat i journalen sin
26 min siden

Mona Lie, 27.06.46
fikk svar på forlik.
41 min siden

Morten Haugland, 30.08.64
Kan utskrives i morgen hvis spiser greit i kveld...
42 min siden - Dr. Breck-Hansen

Gunnar Vik, 19.05.23
Inn på operasjonsrom 3
52 min siden

Jorunn Moe, 13.12.44
fikk et tilsynsnotat i journalen sin
1 time siden

Gunnar Vik, 19.05.23
Radiologens CT-beskrivelse er ferdig.
1,5 timer siden

Det neste som skjer er...

Odd Hansen, 31.01.50
Rtg-time planlagt:
Om 5 min - Kl 23:00

Odd Hansen, 31.01.50
Troponin og Hb estimert ferdig:
Om ca 15 min...

Gunnar Vik, 19.05.23
Operasjon estimert ferdig:
Om ca 1 time

Jorunn Moe, 13.12.44
MRCP planlagt:
I morgen - Kl 11:00

Akkurat nå...

Thomas Aas, 03.12.82
510 - Magesmerter
AVD

Marianne Berg, 30.08.64
511 - Ca Coli
Perm

Mona Lie, 27.06.46
512 - Pankreatitt
AVD

Jorunn Moe, 13.12.44
513 - Kolangitt
AVD

Gerd Dahl, 05.02.38
514 - Crohn
AVD

Morten Haugland, 30.08.64
515 - Appendicitt
AVD

Odd Hansen, 31.01.50
516 - Akutt abdomen + brystsmerter
Røntgen

Gunnar Vik, 19.05.23
518 - Rektalblødning
Opereres stue 3

Fig. 1. The main prototype

The nurse could cross-check this information with the main prototype to be certain of which patient this event pertained. On request BL would also give supplementary medical information related to that particular event. This was done because clinical information such as a health records were not included in the experiment. A common work-flow in the experiment would thus be: Nurse entered patient room, time went by, the de-identified prototype indicated e.g. “Radiologist's X-ray description available”, the main prototype specified “Radiologist's X-ray description available for patient John Smith”, supplementary medical information given orally was “Air under the diaphragm [serious condition]”.

After four such cycles with new event information updates of different kinds (and approximately 1.5 hours of simulated time) the test was ended. Through the subsequent interview the nurses were asked to validate information needs and describe how the prototype could influence their work. They were also shown alternative de-identification techniques and requested to assess whether or not the alternatives would be understood under normal working conditions.

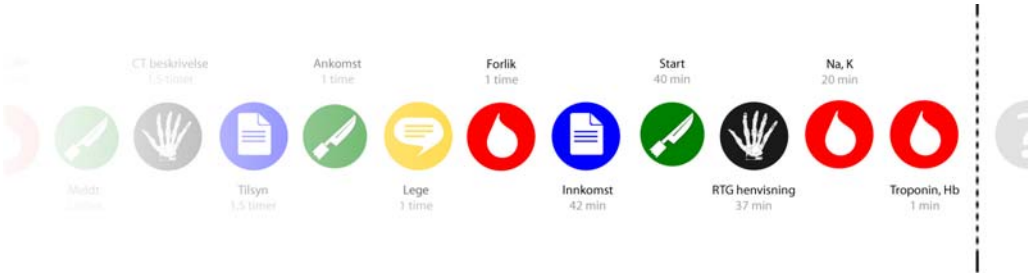

Fig. 2. Prototype with de-identified information

4. Preliminary results

All four nurses declared that the prototypes would facilitate the work they do at the ward to ensure proper patient care. They claimed that the prototypes would reduce their current dependence on phone calls and time-consuming routines for coordination of care. They imagined how they could avoid looking up various clinical information systems (CISs) repeatedly just to find out that no more information was available yet. Instead they wanted to use the prototypes to get notified when new information was available and use specific CISs only when they know that new information existed.

The nurses were also positive about the de-identified prototype, and the idea of having a de-identified notification tool that they could use without having to personally log in and out of it all the time, was regarded useful. They considered their current authorization methods (ID card and user-names and passwords in multiple clinical information systems) to be very inefficient. The nurses argued that they most likely would understand what patient a de-identified event would concern by knowing the context – especially if that patient was relatively important to follow up on. However, the experiment was too limited to say anything about this in general. More importantly, though, the nurses were concerned about maintaining the confidentiality of the patient. They feared that even if the information was presented without specific identifiers, the patients could be identified by others – for instance by observing other clinical activities at the ward.

Some suggestions for improvement where also given by the participants. The de-identified view could for instance be replaced with a personal mobile device, or the multi-patient, de-identified event log in the de-identified prototype could be replaced with a nurse-specific event log including only the events subscribed to specifically by the nurse.

5. Conclusion

The results from this study indicate – as a proof of concept – that real-time information of patient care activities has the potential to improve coordination of care. The prototypes have evidently faced validity with the nurses. However, the study design has major limitations, and all results must be interpreted cautiously. It has nevertheless provided certain hints on potential effects of a novel hospital patient care system, and it has also given fruitful insight into the information needs of surgical ward nurses.

The results warrant more research on this topic to validate that the perceived usefulness will prevail, using real patient care events in a real hospital environment. Additionally, other outcome measures should be added to the study, such as number of unnecessary patient record accesses, time from new results are made available until they have been accessed by nurse and/or physician, and time until proper diagnosis has been made and therapy is initiated. In a broader perspective this study illustrates how some aspects of health information systems could be evaluated with minimal impact on hospital activities. We think more research should be done on achieving valid preimplementation evaluation of health information systems, before clinical implementation is done.

We are currently preparing to run a similar small-scale experiment in a real hospital setting with real patient care events that are supervised and updated manually by one of the researchers.

References

[1] J.E. Bardram, T.R. Hansen, and M. Soegaard, 'AwareMedia: a shared interactive display supporting social, temporal, and spatial awareness in surgery', Banff, Alberta, Canada: ACM, 2006, pp. 109-118.

[2] B. Lillebo, A.R. Seim, and A. Faxvaag, 'Information and communication needs of healthcare workers in the perioperative domain', “XXIII International Conference of the European Federation for Medical Informatics,” Oslo: 2011.

[3] E.A. Gjære, I.A. Tøndel, M.B. Line, H. Andreasen, and P. Toussaint, 'Personal Health Information on Display: Balancing Needs, Usability and Legislative Requirements', “XXIII International Conference of the European Federation for Medical Informatics,” Oslo: 2011.

Appendix C

Paper accepted for MURPBES 2011 conference

International Cross Domain Conference and Workshops on Availibility, Reliability and Security - Multidisciplinary Research and Practice for Business, Enterprise and Health Information Systems, Vienna, Austria, August 22-26, 2011.

A Risk-based Evaluation of Group Access Control Approaches in a Healthcare Setting

Maria B. Line, Inger Anne Tøndel, Erlend Andreas Gjære

SINTEF ICT, Trondheim, Norway
{maria.b.line, inger.a.tondel, erlend.andreas.gjare}@sintef.no

Abstract. This paper focuses on access control approaches usable for information sharing through large screens where several individuals are present at the same time. Access control in this setting is quite different from traditional systems where a user logs on to the system. The paper outlines a number of possible approaches to access control, and evaluates them based on criteria derived from risk analyses of a planned coordination system for the perioperative hospital environment. It concludes that future work should focus on extending the location-based approach with situation awareness, and add support for using pop-ups or handheld devices for sharing of the most sensitive information.

Keywords: Access control, privacy, health care, information security

1 Introduction

There are a number of systems available whose main purpose is to inform the public about status and status changes. Examples are screens showing incoming flights at airports, or overviews of meeting room occupancy at hotels. In these example systems the information on the screen is unlikely to be sensitive, and thus there is no need to control information visualization. But imagine such information displays being used in healthcare or in other businesses where some status information should be considered internal.

In this paper we focus on access control solutions for wall-mounted screens that show status information in a perioperative hospital environment (before, during and after surgery). This environment is characterized by multidisciplinary teams, the need to react to unanticipated events, and utilization of expensive resources. Planning and coordination are difficult but important in such a setting [1]. To improve coordination, wall-mounted screens visualizing progress and current status can be placed at waiting rooms, wards, operating rooms, recovery rooms etc. where health care personnel is likely to see them. As a result it becomes easier to understand how the patient care is progressing and adapt own behaviour.

Access control in systems communicating via large public wall-mounted displays is quite different from traditional access control where a single user logs on to a system. First, it is difficult to know at a given point of time who is able to view the information on screen. In the perioperative environment, there can be a number of health care personnel having access to the location of a screen, in addition to patients, their next-of-kin, and other types of personnel such as cleaners or technicians. Still, if we were able to know who were present, it is not a straight-forward task to determine how this knowledge should affect what information to display. Access policies are normally defined on a single-user level, while the coordination system may just as well have none or several users to consider in its decisions on what to display. Second, the information - though it may be sensitive - is displayed because it is needed for a purpose. What to display in a given situation must be based on proper trade-off decisions between privacy on one hand and efficiency and patient safety on the other. This calls for access control solutions that are dynamic and context-aware, and that fit the way of work in the perioperative domain. To further complicate matters, there is in general no time for users to login to the system, as information should be available by just by taking a quick look at the screen.

In our previous work [1] we have decided on a strategy in order to overcome these challenges, termed flexible de-identification. With *flexible* we mean that decisions on what to display should not be static but adapted to the situation and current context. *De-identification* leads to solutions that go beyond the more traditional consideration of whether to display identifying information or not. Instead we assume that information needs not be at the highest level of granularity to be useful for coordination purposes.

In this paper we focus on providing flexible access control, while de-identification is left out[1]. We do not consider access to information by single users, but instead how to determine access rights for a dynamically changing group of individuals. This group is likely to consist of personnel with different professions and different needs and rights for information, who - though they work together - will have different opinions as to what information gives meaning and is useful. We present several approaches to access control in this setting. Our main contribution is a preliminary evaluation of these approaches based on criteria derived from a risk analysis of a the COSTT system[2] - a planned coordination system for the perioperative hospital environment.

The paper is organized as follows: Section 2 outlines relevant approaches for access control. Section 3 explains the method used to perform the risk analyses and to deduce evaluation criteria. Section 4 gives the results of the risk analyses, explains the evaluation criteria and applies them in order to evaluate the different access control approaches. Section 5 discusses the validity of the evaluation result, and suggests directions for future research, before section 6 presents our concluding remarks.

[1] See Gjære et al. [2] for more information on de-identification solutions for this kind of systems.

[2] Developed by the Co-operation Support Through Transparency (COSTT) research project. http://www.costt.no

2 Relevant Approaches to Access Control

Access control related to shared disiplays have been given some attention from researchers, and proposed solutions include using special types of glasses to allow different people to see different types of information [3] and using visualizations (e.g. colors) instead of text to present the most sensitive information [4] (similar to our de-identification approach). Available research on how to determine access rights of dynamically changing groups of users is however sparse. We are only aware of one publication [5] that addresses this issue to some extent. This publication lists three different approaches that can be used in combination. The first approach is *aggregation* where the access rights of a group correspond to the sum of the access rights of the individuals in the group. As a result, larger groups are likely to get access to more resources than smaller groups. The second approach is *maximum/minimum* where the individual with the highest/lowest access rights determines access rights for the whole group. Third, the *group structure* approach computes access rights based on the structure of the group, e.g. ensures that at least two users with a certain access level are needed in order to gain access to a specific resource.

Current access control solutions developed for healthcare are mainly based on Role Based Access Control [6] where users are granted access based on their profession [7][8]. A study of an up-and-running Electronic Patient Record system at a Norwegian hospital [9] identified the need for access control solutions more tailored to the needs of health care personnel. Solutions should be better able to handle dynamic events, workflow and collaboration.

Dynamic context-aware access control solutions have been suggested in various forms, some also specifically addressing health care (e.g. Hu et al. [10]). We will not go into specifics on the different models, but rather point to parameters that can be used when creating more dynamic access control solutions in a health care setting. Hu et al. mention time, location, trust-level of authentication, relationship to patient, and specialist area. Most of these are also mentioned by Alam et al. [11], who add device type, duration, purpose, number of accesses, user consent, presence of the patient, delegation and emergency situations to the list. Risk and benefit are also factors that can be taken into account in access control decisions. Examples of access control solutions that include the concept of risk is the work of Cheng et al. [12], Dimmock et al. [13] and Diep et al. [14].

Below we present the main access control approaches considered in this paper. It is assumed that we have technology available for authenticating and locating users[3].

Location-based: Only the screen's location applies, which means that all screens have a given default view that can not be affected by persons being present. Access control to the screen is managed through physical access control; people being allowed to be at a certain location are also allowed to see all information displayed on the screen at the given location.

[3] The purpose of locating users will NOT be surveillance of all their movements and actions.

Minimum [5]: A group's access rights will correspond to the lowest level of access rights present in the group. This will clearly ensure patients' privacy, as nobody in the group will see more information than they are allowed to. However, the usefulness may be low, as the persons with higher levels of access rights will not always see all information meant for them.

Maximum [5]: A group's access rights will correspond to the highest level of access rights present in the group. This will ensure usefulness, as all health care workers will see all the information meant for them. However, the patient's privacy may be compromised.

Group structure [5]: The access level is decided by computations on the group structure. An average value is calculated based on who is present, including weighting of who is closest to the screen and considerations of how many is present with limited access rights vs. wide access rights.

Facilitator: One of the users in the group acts as a facilitator. The facilitator is authenticated, and makes decisions as to whether to include new users into the group and which information is needed/appropriate based on the users present.

Situation aware: The system is aware of the situation in which it operates, e.g. it combines information on type of patient and diagnosis with time of day and an understanding of whether this is an emergency or normal operation. Situational awareness is then used to decide access rights of the group.

Possible extensions: Pop-up window and handheld devices: All the suggested approaches can be extended with solutions that grant individual users access to more information. This can be done by utilizing small pop-up windows where e.g. surgeons can authenticate themselves and get access to more details shown in a limited part of the screen, or get the information sent to a handheld device. Getting access to information in a pop-up window limits the reading access for other people being within proximity, information is only readable for the one/those standing really close to the screen. Sending information to a handheld device further reduces the risk of confidentiality breaches.

3 Method

In this work, we use the results of two risk analyses of the planned COSTT system to identify criteria that the access control solution need to fulfil, and use these criteria to evaluate and compare the alternative access control approaches. The main motivation for using risk analysis in this respect is twofold. First, the results of the risk analyses are already available and provide valuable insight into the environment in which the access control solution will be put to use. Second, access control aims to protect (some of) the system assets by reducing risks, but may also introduce new risks. Performing a risk analysis is a good way to identify both the assets and the risks towards these assets. We recognise that using risk analysis of systems to evaluate access control policies is uncommon. Still, we uphold that the results of a risk analysis are useful for performing a preliminary evaluation of alternatives in order to decide which should be further investigated and evaluated.

Two risk analyses have been performed, and both were carried out in two stages: 1) Asset identification[4]: “What are the most valuable assets in the COSTT system? What do you want to protect?”, and 2) Risk identification and ranking: “What are the most important risks for COSTT? What are you most afraid of happening?”. In each stage the participants were given five minutes to write down their answers to the questions posed. Both stages were summarized by organizing the brainstorming results into groups that the participants agreed upon. The risk identification stage also included ranking of the risks. Each participant was given three votes they could use to prioritize risks. The risks were then ranked according to votes in total.

The COSTT project group was used as participants. Together they represent a broad spectrum of specialist areas; IT, sociology, medicine, and technology management. The first risk analysis was performed at the stage where the system itself existed only as a concept and many decisions that would affect the outcome had not been made yet. The purpose of this preliminary analysis was to get an initial sense of what are the key risks as perceived by the project team. The second risk analysis of the future COSTT system was performed 10 months after the first one. The system itself still existed only as a concept but some research, including literature studies and empirical studies, had been performed. The purpose this time was to see if the results would differ a lot from earlier, and to identify the major changes, if any.

4 Results

In this section we present the results of the risk analyses, and use these results as a basis for identifying evaluation criteria. Then we show to what extent the identified access control approaches are able to meet the criteria.

4.1 Results of the risk analyses

The findings from the first risk analysis represent a starting point and a snapshot of the project status at that point of time. The second risk analysis revealed the same results, but both broader and in more depth. One of the main differences, was the ranking of the risks related to sensitive information and access control; as the participants increased their understanding of what the COSTT system will be, they also increased their worries of sensitive information leakage, while they decreased their worry of the access control mechanisms not being strong enough. We choose to present the results from the latter only, because that is sufficient in order to cover all identified issues.

Table 1 presents all assets and risks identified. Note that the categories of assets and risks are not considered to be mutually exclusive. The participants themselves sorted the input from their brainstorming process and gave names to the categories of information. The assets mainly include types of information available in the system; both to be displayed on the screen and underlying

[4] Inspired by the asset identification method described by Jaatun and Tøndel [15].

Table 1. Identified assets and risks

Category	**Assets**
Patient information	Identification, medical data, secret relations, irrelevant health history
Employee information	Name, role, actions, personal data
Location data	Position for all tagged persons, info about rooms, movements
Aggregated/reasoned data	Efficiency of employees, process statistics, surveillance of procedures
Deviations	Unwanted incidents, info on operations, system errors in hospital
Usefulness	Utility value by using the COSTT system

Category	**Risks**
Poor quality of data (9)	Drawing wrong conclusions on what info means, inaccurate catching of events, misinterpretation of events, coordination trouble due to erroneous data, patient injury
Surveillance of employees (6)	Management monitoring efficiency of employees, wrong/incomplete statistics on employees, employees feel they are being monitored, public negative exposure of some employees
Sensitive personal information (6)	Info displayed to persons not concerned, deduction of patient having a sensitive diagnosis, unintended access to sensitive info, hacking/data theft, info is taken out from the hospital
Unintended/erroneous use (5)	Location tag theft, active bypassing of access control lists in other systems, bypassing physical access control, unhealthy changes in work processes, employees working against the system or refusing to use it, conflicts due to low efficiency
Patient (2)	Patients choosing a different hospital, theft of patient data, patient info known to public press
Access control (1)	Limitations hide important data when it should be available, bugs giving illegitimate access
Wrong focus (1)	Fussbudget, loss of efficiency, debates on prioritizing, critical questions due to insignificant errors
Relatives (0)	Creating unnecessary feelings, unhappy relatives calling frequently on health personnel
Public (0)	Negative newspaper headlines

information needed to make the system work properly. Also, parameters that can be deduced from information in the system are considered valuable assets. The risks span from concerns of the underlying sensors not being able to catch events to breaches of both patients' and employees' privacy. The numbers listed in parenthesis in the column of categories indicate the prioritizing done by the process participants. In the table, the risks are presented in prioritized order.

4.2 Identification of Evaluation Criteria

In the process of identifying evaluation criteria, we focused on the highest ranked risks. The criteria are referred to as C1, C2 etc., which constitutes a mapping to table 2 where they all are summed up.

Poor quality of data: The main risk was considered to be poor quality of data. As shown in table 1, this is mainly a concern about the underlying system not being able to catch and/or interpret events correctly. In the COSTT project, the coordination information that is to be displayed on the screens are built by capturing events in other information systems [16]. Simple events (e.g. access to the medical record of patient A by a given health care personnel) are combined into composite events (e.g. cardiology assesment of patient A has been performed), and it is these composite events that will be displayed on the screens. It is however important to be aware of the uncertainty involved in the event enrichment process. As an example, access to medical record of patient A can indicate that the health care personnel that accessed the record is performing an examination of the patient, but it can also be that the health care personnel is preparing for the examination. Thus, events in the COSTT system will be associated with a quality attribute that is a measure of the validity of the event [16]. This quality attribute should ideally influence the access control decision, as presenting information that is correct is an important part of the information security of any system (integrity). This is reflected by the data quality awareness criteria (C1).

Surveillance of employees: The next highest ranked risk was that of surveillance of employees. This covers the employees' fear of being monitored and the possibility for management to misuse registered data about their employees to measure efficiency or other statistics. Data registered for the purpose of COSTT may not give the complete and correct picture of employees' actions, which means that it should be used with high caution, if at all, for management purposes. Thus it is relevant to consider whether the solutions increase the need for surveillance, e.g. by requiring location information (C2). Employee surveillance is also related to what information is published on the screens (further addressed for the risk of sensitive personal information).

Sensitive personal information: The third highest ranked risk is that of displaying sensitive personal information in ways that makes the information available to unauthorised persons. It is important that solutions are able to maintain privacy of both patients and employees, and strive towards the ideal solution where everybody gets access to what they need - and no more. This is reflected by the privacy preserving criteria (C3).

Unintended/erroneous use: The risk related to sensitive personal information should also be considered together with the risk of unintended/erroneous use (rated fourth) and also the much lower prioritized risk related to access control. The concern that access control does not support the work flow is reflected in all these three risks. Failure in this respect can lead to active bypassing of access control due to important information not being available (C4). To meet this challenge it is important to consider dynamic and/or user controlled access control solutions that is able to fit into the way people work. It is also important to consider the effort required from users in order to use the systems in a secure manner (C5), as expectations on user involvement may require changes in work processes in itself in addition to requiring time and effort from the users. This

Table 2. Identified criteria based on risk analysis

Nr	Criteria	Explanation
C1	Data quality awareness	The ability of the solution to take the data quality into account in the access control decisions.
C2	Minimisation of employee surveillance	The need for use of employee surveillance techniques, e.g. for monitoring the location of employees.
C3	Privacy preservation	The ability to restrict sensitive/private information to those that are authorized for access.
C4	Availability ensurance	The ability to ensure that information important for safe and efficient treatment of patients are available when needed. This can e.g. be ensured by using dynamic approaches able to adapt to the situation , or to ensure that users can override the access control decision.
C5	Workload reduction	The ease of use for users. Solutions that rely on user cooperation will require some time and effort on behalf of the users.
C6	Complexity	The more complex the access control mechanism is the higher risk of mistakes that may render the access control solution vulnerable.

Table 3. Evaluation of access control approaches with respect to the selected criteria

Approach	**C1** qual.	**C2** surv.	**C3** priv.	**C4** avail.	**C5** workl.	**C6** compl.
Location-based	-	+	?	?	+	+
Minimum	-	-	+	-	+*	+
Maximum	-	-	- -	+	+*	+
Group structure	-	-	?	?	+*	-
Facilitator	-	+	-/+	+	-	+
Situation aware	?	+	?	?	+	-
Extension: Pop-up/handheld	-	-/+	+	+	-	+

influences the perceived system efficiency and is likely to have an impact on the employees' attitudes towards the system; employees working against the system or refusing to use it. Risks related to access control can also increase with complexity (C6). With complex access control solutions it is easier to make mistakes e.g. during implementation or during policy specification.

4.3 Evaluation of access control approaches

Table 3 summarises the evaluation of the suggested access control approaches with respect to the evaluation criteria. A '+' indicates a positive score while a '-' indicate a negative score. A score of '-/+' indicates that the approach is able to meet the criteria to some extent, but not fully. A '?' is used in situations where the evaluation result will depend on trade-offs made when defining access control policies. A '*' is used to illustrate that the score depends on the mechanism used to determine who is present.

The *location-based* approach is in many ways the most simple approach. Its ability to meet the needs of COSTT is however dependent on how easy it is to determine beforehand which information should be available in given locations. As several of the envisioned locations (e.g. corridors and examination rooms) are likely to be accessed by a number of different groups of users, we envision that it will be difficult to make such pre-set trade-offs that are able to meet the criteria for both privacy and availability.

The *minimum* approach is unlikely to meet the availability requirements of the users in need for most information. The same way, the *maximum* approach will probably result in too many privacy breaches, as everybody will get access to whatever information should be available to the one present with the highest access rights. Considering group structure is likely to perform better than using the minimum/maximum access rights, but is complex and its success depends on the ability to make proper trade-offs between the access rights of the highest and lowest ranked users present.

Relying on a *facilitator* seems to be a solution that could fit COSTT well, as it is able to meet the majority of criteria to some extent. The privacy achieved will be dependent on whether we can trust the facilitator to make good decisions as to what information to display in given situations. In a study of clinicians' experiences related to privacy and security of health information systems [17] Fernando and Dawson noticed that most clinicians used measures such as lowering their voices and omitting to ask relevant questions in order to protect privacy and security when residing in a shared workplace. At the same time they found that privacy and security implementations on electronic health information systems often took time from patient care, and were therefore considered to hamper patient care. Sharing of passwords was mentioned in this study, as well as in a study by Vaast [18]. In his study he also found that physicians were concerned that employees on wards were overwhelmed with work and therefore were likely to forget to close programs or patient charts. To be able to reach a conclusion as to whether the facilitator approach is adequate in the COSTT setting, more research is needed on the situation in which the COSTT system will operate when it comes to the work process and the general attitude of the employees.

Making the access control solution more *situation aware* is also likely to improve the trade-off between privacy and availability, but at the cost of complexity. This is the only approach that has the potential to meet the data quality awareness criterium.

The *pop-up/handheld extension* also seems promising, and is able to meet the majority of criteria. By implementing this extension one is able to introduce more flexibility and user involvement without sacrificing privacy. It should however not be used as the only approach.

To sum up, none of the access control approaches studied is able to meet all evaluation criteria. However, the ones that seem most promising are either using the facilitator role or using an access control solution that is situation aware. Alternatively, one of the more simple automatic approaches, e.g. the location-based approach, can be combined with the pop-up/handheld extension.

5 Discussion

The preliminary evaluation performed is based on a risk analysis of the COSTT system at an early stage and with participants from the project group. At this stage the project participants are the ones most likely to have the best understanding of how the COSTT system will work and what are the main challenges and risks. The results of the risk analysis would however be more reliable if it had included project-external representatives as well.

Basing the preliminary evaluation on the results of a risk analysis is useful in that the evaluation criteria will be risk-based and likely to reflect the top issues. The criteria derived are however high level and have not been evaluated by the intended users of the system. It is also not possible from this initial evaluation to state how the different access control approaches will perform in real life. User evaluations are needed in order to assess how the alternative approaches are able to fit the work processes of the perioperative domain. In particular we suspect that the facilitator-based approach, though getting good scores in the evaluation, will fail in this respect.

In the evaluation of the alternative access control approaches, we have studied the approaches individually and evaluated how they perform related to the identified criteria. It is however possible to combine several of the approaches into a final solution and in this way achieve a solution that better fits the needs of COSTT. To illustrate, screens in waiting rooms may have a preset access level (location-based approach) while screens at other locations may have a maximum access level that is determined by their location but where the group structure, the general situation or a facilitator determines the access level at a given point of time. It is also possible to use the extension suggested where individual employees can get access to more information by utilizing pop-up windows or handheld devices. The preliminary evaluation of the approaches suggests that future work looks into combinations of the location-based approach, the facilitator approach, the situation aware approach and the pop-up/handheld extension approach. The location-based approach is a simple one with the possibility of offering good baseline security. The facilitator approach is able to meet the majority of the criteria. The situation aware approach is the only one able to meet the data quality criterium, and has the potential to also perform well on most of the other criteria. The pop-up/handheld extension approach has good scores on both the privacy and availability criteria.

As the location-based approach is in many ways the most feasible solution, we plan to use this solution as a starting point and look into how it can be combined with risk-based access control approaches in order to add situation awareness. Making proper trade-offs between the risk of privacy breaches and the risk that information is not available is central to the success of the COSTT solution. Risk-based access control solutions can utilise knowledge of the screens' location in order to determine the probability of privacy breaches, as well as the availability requirements for an information item. Other context information, like the time of day and who is likely to be present, can influence the risk evaluation as well. This way, the combined solution will likely perform better on the criteria related to privacy (C3) and availability (C4). In addition, the quality of the information can be taken into account (criterion C1). Though the facilitator approach gets quite high scores on the criteria used in our evaluation, we believe that this solution will not be usable for this type of systems, as it requires quite a lot of interaction with the users. Instead we recommend using handheld devices in combination with large wall-mounted screens in cases where highly sensitive information is needed (to better meet criteria C3 and C4).

6 Conclusion

The main contribution of this paper is a preliminary evaluation of several approaches to access control for public screens used in a perioperative setting. The evaluation criteria used are derived from a risk analysis of the COSTT system. In the evaluation, the facilitator based and the situation aware approaches received high scores, and so did the possible extension of using pop-up windows or handheld devices to get access to additional information. Of the simpler and most feasible approaches, the location-based approach turned out to be the best candidate. As none of the approaches were able to perform well on all evaluation criteria, the results motivate to look further into combining access control approaches. As there are major usability concerns with the facilitator approach in this setting, we recommend focusing on extending the location-based approach with situation awareness, and add support for pop-ups or handheld devices.

Acknowledgments

Thanks to the COSTT project group for participating in the risk analyses. Thanks also to Torstein Nicolaysen for his work in a related MSc thesis, and to Åsmund A. Nyre, Arild Faxvaag, Andreas Seim and Warren Sandberg for reading and commenting on this paper during the writing process. This work has been performed as part of the COSTT project that is supported by the Norwegian Research Council's VERDIKT program (grant no. 187854/S10).

References

1. A. Faxvaag, L. Røstad, I. A. Tøndel, A. R. Seim, and P. J. Toussaint, "Visualizing patient trajectories on wall-mounted boards - information security challenges," in *MIE*, ser. Studies in Health Technology and Informatics, K.-P. Adlassnig, B. Blobel, J. Mantas, and I. Masic, Eds., vol. 150. IOS Press, 2009, pp. 715–719.
2. E. A. Gjære, I. A. Tøndel, M. B. Line, H. Andresen, and P. Toussaint, "Personal health information on display: Balancing needs, usability and legislative requirements," in *MIE , ser. Studies in Health Technology and Informatics (to be published)*, 2011.
3. G. B. D. Shoemaker and K. M. Inkpen, "Single display privacyware: augmenting public displays with private information," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, ser. CHI '01, 2001, pp. 522–529.

4. P. Tarasewich and C. Campbell, “What are you looking at,” in *The first Symposium on Usable Privacy and Security (SOUPS 2005)*, 2005.
5. A. Bullock and S. Benford, “An access control framework for multi-user collaborative environments,” in *GROUP ’99: Proceedings of the international ACM SIGGROUP conference on Supporting group work*, 1999, pp. 140–149.
6. ANSI, “American National Standard for Information Technology - Role Based Access Control,” 2004, ANSI INCITS 359-2004.
7. A. Appari and M. E. Johnson, “Information security and privacy in healthcare: Current state of research,” *Forthcoming: International J. Internet and Enterprise Management*, 2009.
8. A. Ferreira, R. Cruz-Correira, L. Antunes, and D. Chadwick, “Access control: how can it improve patients’ healthcare?” *Studies in Health Technology and Informatics*, vol. 127, pp. 65–76, 2007.
9. L. Røstad and O. Edsberg, “A study of access control requirements for healthcare systems based on audit trails from access logs,” in *ACSAC ’06: Proceedings of the 22nd Annual Computer Security Applications Conference*, 2006, pp. 175–186.
10. J. Hu and A. Weaver, “Dynamic, context-aware access control for distributed healthcare applications,” in *Proceedings of the First Workshop on Pervasive Security, Privacy and Trust (PSPT)*, 2004.
11. M. Alam, M. Hafner, M. Memon, and P. Hung, “Modeling and enforcing advanced access control policies in healthcare systems with SECTET,” in *1st International Workshop on Model-Based Trustworthy Health Informaton Systems (MOTHIS 07)*, 2007.
12. P.-C. Cheng, P. Fohatgi, and C. Keser, “Fuzzy mls: An experiment on quantified risk-adaptive access control,” IBM Thomas J. Watson Research Center, Tech. Rep., January 2007.
13. N. Dimmock, A. Belokosztolszki, D. Eyers, J. Bacon, and K. Moody, “Using trust and risk in role-based access control policies,” in *Proceedings of the ninth ACM symposium on Access control models and technologies*, ser. SACMAT ’04, 2004, pp. 156–162.
14. N. N. Diep, L. X. Hung, Y. Zhung, S. Lee, Y.-K. Lee, and H. Lee, “Enforcing access control using risk assessment,” *European Conference on Universal Multiservice Networks*, vol. 0, pp. 419–424, 2007.
15. M. G. Jaatun and I. A. Tøndel, “Covering your assets in software engineering,” in *Third International Conference on Availability, Reliability and Security*, 2008, pp. 1172–1179.
16. L. W. M. Wienhofen and A. D. Landmark, “Poster: Representing events in a clinical environment - a case study,” in *The 5th ACM International Conference on Distributed Event-Based Systems (DEBS 2011) (to be published)*, 2011.
17. J. Fernando and L. Dawson, “The health information system security threat lifecycle: An informatics theory,” *International Journal of Medical Informatics*, vol. 78, no. 12, pp. 815–826, 2009.
18. E. Vaast, “Danger is in the eye of the beholders: Social representations of Information Systems security in healthcare,” *The Journal of Strategic Information Systems*, vol. 16, no. 2, pp. 130–152, 2007.

Appendix D

Specification of collaboration

This appendix describes in detail the relationship between the work done by the author of this thesis, and the contributions made by Ph.D. candidate Børge Lillebo. Ph.D. candidate Børge Lillebo is from now on denoted as “BL”.

The rapid field tests

Before starting the testing phase, BL wrote the application to Norwegian Social Science Data Services (NSD), in order to get permission for accessing the informants. BL was moreover responsible for creating three of the four prototypes that were used in the rapid field tests, while the one designed by me, including its refined edition, is included in chapter 4. Recruiting informants was done on the fly when we were roaming in the hospital corridors, but BL was the one doing the talking. The interview guide for the tests was mainly written by BL, but I added two questions of my own to it, as described in chapter 4. We both recorded personal field notes, and wrote out each our own transcription document for every interview.

The lab experiment

The common interpretation of the results from the field tests led to BL’s creation of three new prototypes of a system, one of which he decided to proceed with for a test in the usability lab. Together, we then populated the prototype with realistic data, which in turn a fictive patient overview sheet was based on, as seen in appendix H. This part was done in mutual collaboration. I then designed de-identified views of this information using the graphical icons from his work, and proceeded with writing a scenario for a role play, suggested a physical layout for the lab, and wrote a document describing the experiment as a whole. BL was on the other hand responsible for recruiting test persons for the experiment. Finally, each of us wrote our own interview guide for the focused interview set up stright away after the role play, mine of which is included in appendix F. The two pilot tests we done together, and so were three of the real tests as well. I unfortunately missed out on one myself, due to illness. Transcribing the video recordings was split in half, so that we did two of them each.

Appendix E

Role-play scenario

This appendix shows how the scenario proceeded, that supported the role-play session in the lab experiments.

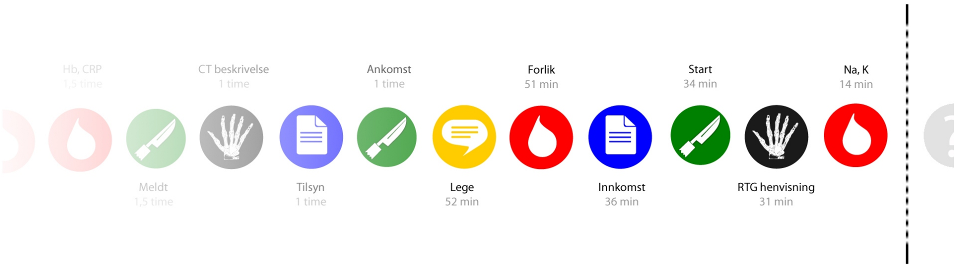

Figure E.1: The intial list of events displayed on the hallway screen (stage 1)

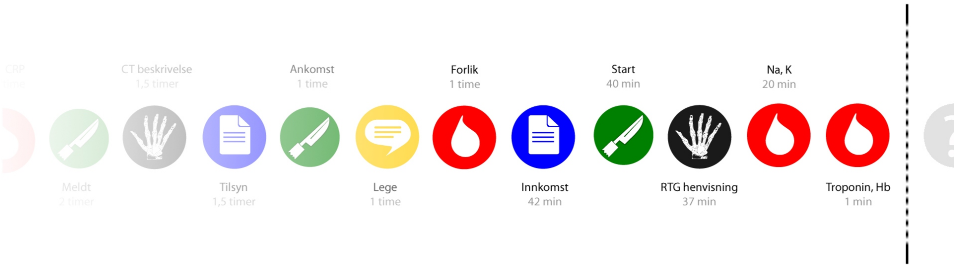

Figure E.2: The hallway prototype after 16 minutes have “passed” (stage 2)

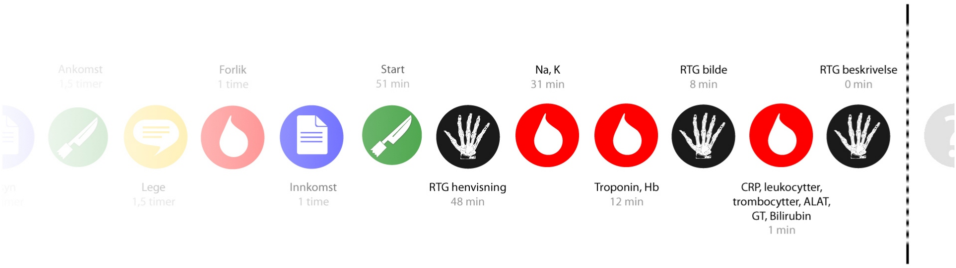

Figure E.3: The hallway prototype after another 10 minutes has “passed” (stage 3)

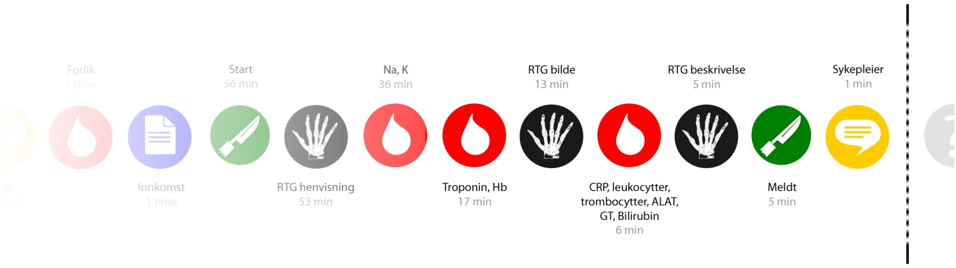

Figure E.4: The hallway prototype after another 10 minutes has “passed” (stage 4)

Appendix F

Interview guide

This appendix contains the interview guide that was used in the lab experiments.

Introduksjon

Vi er . . . og arbeider med . . .
I denne testen skal vi vurdere . . .
Vi tester ikke deg. Vi ønsker å lage et godt og nyttig verktøy, og har derfor utarbeidet et scenario som du forhåpentligvis vil kjenne deg igjen i.
Du har full anledning til å avbryte dersom du blir ukomfortabel eller ikke vil gå videre med en bestemt oppgave. Beskrivelse av videoopptak.
Beskrive utstyr i rommet: Sengetun med arbeidsstasjon og skjerm i korridoren. To pasientrom.
Tenke høyt underveis er bra.
Beskrivelse av prototypen og dens begrensninger.
Introduksjon av produktet. (Gi en overordnet punktliste over hva som skal skje.)
Er det noe du lurer på?
Begynn med oppgaver!

Rollespill

Simulert vaktrapportmøte — bruk pasientoversikten. Etter vaktrapporten: "Nå skal du få mulighet til å gjøre deg litt bedre kjent med et nytt system som kan gjøre det lettere å se hva som skjer med pasientene dine."

Følgende spørsmål stilles når testpersonen passerer foran den av-identifiserte skjermen:

- Kan du se om det har hendt noe nytt?
- Kan du si noe om hvem du tror dette har hendt med?
- Hvor sannsynlig tror du at hendelsen gjelder den du tror den gjelder?

Fokusert intervju

Legg fram oversiktsark med alternativer. Start med alternativ II som allerede er kjent, samt alternativ I som er enda mer begrenset, og fortsett så nedover. Analyser hvert enkelt informasjonsnivå:

- Gjør den nye informasjonen du får det enklere å identifisere pasienten hendelsen gjelder?
- Hvis nei: Kan du likevel ha nytte av informasjonen på dette nivået?
- Er den nye informasjonen problematisk å henge opp i en gang med tanke på taushetsplikt/personvern, sett i sammenheng med hendelsen den identifiserer?
- (Alternativ V blir sannsynligvis forkastet). Presenter da til slutt alternativ VI som et like nøyaktig identifiserende alternativ, der det er ved hjelp av pasientoversikten at man kan identifisere pasienten (legg fram pasientoversikt med påskrevne pasientkoder):
- Kunne du akseptere dette alternativet selv om du da kanskje må sjekke pasientoversikten når du skal se om en ny hendelse gjelder en av "dine" pasienter eller konkret hvem det gjelder?

Be testpersonen om å velge ett av alternativene som han/hun ville foretrekke å ha på en slik skjerm. (Eventuelt kan alternativer kombineres, dersom det er hensiktsmessig.)

Appendix G

Informants agreement

This following is the document that the was signed and hence witnessed the test persons' informed consent to participation in the lab experiment.

Samtykkeerklæring

Prosjekt: ....

Ansvarlig: ....

Denne samtykkeerklæringen gjelder for et forsøk som gjennomføres i brukbarhetslaboratoriet til Norsk senter for elektronisk pasientjournal. Forsøket dreier seg om å vise prototyper av fremtidige kliniske informasjonssystem for personer som har erfaring som gjør dem i stand til å vurdere nytteverdi av slike informasjonssystem.
Detaljer om denne studien og denne prototypen gis muntlig og skriftlig (i eget skriv).

Du bes med dette å erklære følgende før vi starter forsøket:

- Jeg deltar av fri vilje.
- Jeg forstår formålet med forsøket.
- Jeg vet at det blir gjort lyd- og bildeopptak, og at anonymiserte versjoner av dette kan bli benyttet i forskningsartikler.
- Jeg vet at jeg kan trekke meg når som helst og kreve at all data som stammer fra meg slettes (med unntak av publiserte data).

Sted/dato Underskrift

....

Appendix H

Patient list

This document is a fictional representation of the document that is in reality being used at wards in the hospital where the lab experiment participants were recruited from.

Pasientoversikten for Sykepleier – KGAS2 – Tun 4

510		**Thomas Aas, 03.12.82**	PlTyngde:	
	Diag.	Magesmerter		
	Beh.	÷		
	Rap.	Normale blodpr. Normal UL.		
511		**Marianne Berg, 12.04.76**	PlTyngde:	
	Diag.	Ca. coli		
	Beh.	Op. 20/1.		
	Rap.	UL nyrer 15/1. Perm.		
512		**Mona Lie, 27.06.46**	PlTyngde:	
	Diag.	Pankreatitt		
	Beh.	Drenasje.		
	Rap.	MCRP 19/1. Kontroller CRP+amylase daglig.		
513		**Jorunn Moe, 13.12.44**	PlTyngde:	
	Diag.	Kolangitt		
	Beh.	Blærekat. Ampicillin intravenøst.		
	Rap.	UL abd 17/1. Kvalme, oppkast.		
514		**Gerd Dahl, 05.02.38**	PlTyngde:	
	Diag.	Crohn		Opr. 16.01.2011
	Beh.	TPN fra 13/1.		
	Rap.	CT 14/1. Med.tilsyn bestilt.		
515		**Morten Haugland, 30.08.64**	PlTyngde:	
	Diag.	Appendicitt		
	Beh.	Op. 18/1.		
	Rap.	Rtg oversikt abd 18/1. FF fra kl 12, NK fra kl 21.		
516		**Odd Hanssen, 31.01.50**	PlTyngde:	
	Diag.	Akutt abdomen + brystsmerter		
	Beh.			
	Rap.	Faste inntil videre.		
518		**Gunnar Vik, 19.05.23**	PlTyngde:	
	Diag.	Rektalblødning		
	Beh.			
	Rap.	Kjent AAA. Gastroskopi 17/1		