



Stockholm
University

The Informed Gaze

On the Implications of ICT-Based Surveillance

Baki Cakici

Stockholm University
Department of Computer and Systems Sciences
ISBN 978-91-7447-669-9
ISSN 1101-8526
DSV Report Series 13-006

Swedish Institute of Computer Science
ISRN SICS-D-66-SE
ISSN 1101-1335
SICS Dissertation Series 66

Cover photo: Joel Höglund

Abstract

Information and communication technologies are not value-neutral. I examine two domains, public health surveillance and sustainability, in five papers covering: (i) the design and development of a software package for computer-assisted outbreak detection; (ii) a workflow for using simulation models to provide policy advice and a list of challenges for its practice; (iii) an analysis of design documents from three smart home projects presenting intersecting visions of sustainability; (iv) an analysis of EU-financed projects dealing with sustainability and ICT; (v) an analysis of the consequences of design choices when creating surveillance technologies. My contributions include three empirical studies of surveillance discourses where I identify the forms of action that are privileged and the values that are embedded into them. In these discourses, the presence of ICT entails increased surveillance, privileging technological expertise, and prioritising centralised forms of knowledge.

Contents

Acknowledgements	i
PART I	1
1 Introduction	3
1.1 Aim	5
1.2 Scope	6
1.3 Contributions	12
1.4 Disposition	14
2 Theory	17
2.1 Borders	17
2.2 Situating Knowledge	22
2.3 Discourse Theory	25
2.4 Studying ICTs	28
2.5 Studying Surveillance	32
3 Methodology	39
3.1 Design Documents	39
3.2 Discourse Analysis	41
3.3 A Collection of Papers	44
4 Conclusion	51
4.1 Summary of Papers	51
4.2 Answers	53
4.3 Conclusion	56
References	59

PART II	68
5 CASE: A Framework for Computer Supported Outbreak Detection	71
5.1 Background	72
5.2 Implementation	73
5.3 Results and Discussion	78
5.4 Conclusions	81
References	83
6 A Workflow for Software Development Within Computational Epidemiology	85
6.1 Introduction	86
6.2 Workflow	89
6.3 Conclusion	97
References	102
7 Sustainability Through Surveillance: ICT Discourses in Design Documents	105
7.1 Introduction	106
7.2 Material	111
7.3 Analysis	112
7.4 Conclusion	123
References	126
8 Changing Behaviour to Save Energy: ICT-Based Surveillance for a Low-Carbon Economy in the Seventh Framework Programme	129
8.1 Introduction	130
8.2 Methodology	132
8.3 Analysis	134
8.4 Discussion	139
8.5 Conclusion	142
References	145
9 Detecting the Visible: The Discursive Construction of Health Threats in Syndromic Surveillance System Design	149
9.1 Introduction	150
9.2 Background	151

Contents

9.3	Analysis	155
9.4	Conclusion	165
	References	167

Acknowledgements

While working on this thesis, I have moved between many institutions, and in each I have benefited from the kindness of friends and colleagues.

I thank my main supervisor, Magnus Boman, whose unwavering trust and intellectual integrity have guided me throughout this journey, providing me with the freedom to find my own way, and helping me continue even when I lost all hope. I am grateful to my co-supervisor Nina Wormbs for her insightful and generous support; in our every conversation I found new ideas. I also thank my co-supervisor Markus Bylund for making my stay at SICS possible, and for always valuing my ideas.

I thank my former colleagues Anette Hulth, Sharon Kühlmann-Berenzon, Maria Grünewald, Louise Magnesten and Valerie Decraene at the Swedish Institute for Communicable Disease Control where my journey began. I also thank Christian Schulte and Marianne Hellmin for their support and advice during my stay at KTH, as well as Åsa Smedberg, Fátima Ferreira, Oliver Popov, and Love Ekenberg at Stockholm University, DSV. I thank Karin Fohlstedt, Lotta Jörsäter, and Vicki Knopf for their kind support at SICS, and Joel Höglund for the cover photo he graciously provided. I also thank Jakob Axelsson, Jaana Nyfjord, and the Software and Systems Engineering Lab for giving me a home in the last two years of my research. I am thankful for having shared the burden of work with Oscar Täckström, Anni Järvelin, Pedro Ferreira, Olof Görnerup, Fabienne Cap, and Alina Huldtgren. I am especially grateful to Pedro Sanches for the academic companionship and stimulating conversations during my stay at SICS.

I thank the participants of the Theory and Method in Historical Research course, and especially Adam Netzén, for the engaging discussions. My thanks also to Sabine Höhler for her guidance during

the course, and to Daniel Pargman for giving me the chance to hold a lecture on my own research at KTH. I thank Martin French for his support and advice during the later stages of my work. I also thank the unit of Technology and Social Change (Tema T) at the Department of Thematic Studies for hosting me during my stay at Linköping University.

I thank all members of Watch This Space for keeping my mind off work twice a week every week. I also thank Therese Hellström, Markus Buschle, May Yee, Kim Nevelsteen, Tatjana Nordström, Jan Nordström and Ceren Kuşcuoğlu for their friendship. I am grateful to Justus Minx for his enduring fellowship in what was once a strange land. I thank my mother İclal Çakıcı for her love and support which have been fundamental in motivating me to pursue a career in academia.

Finally, I thank Hanna Sjögren; without you this work would not have been possible. You have taught me to question injustice, and you have continually inspired me to demand a better world. Your kindness sustains me, and I remain eternally grateful.

Stockholm, September 2013

PART I

Chapter 1

Introduction

Information and communication technologies, abbreviated as ICTs, are not value-neutral. They are marked by power struggles that shape their production, and they are affected by the materialities which sustain them. These digital technological systems connect a broad spectrum of theories, practices, and discourses. The initial term, information, holds the key to understanding the implications of the abbreviation ICT.

Two conceptualisations of information are common within computer science. The first is Claude Shannon's information theory (Shannon and Weaver 1948), which employs a definition that treats information as a quantity expressed by signals, stored as states, and always existing in relation to communication. As it provides methods for quantifying information, it renders possible the quantification of everything else that can be translated into information.

The second conceptualisation locates information in the larger process of constructing *wisdom*. Often referred to as the DIKW pyramid (data, information, knowledge, wisdom), it proposes a narrowing stack of dependencies, with data at the bottom and wisdom at the top (Rowley 2007). In the pyramid, data represent a large collection of signals and signs that are not useful on their own except as building blocks of information, and information appears as an intermediate step in the journey from data to knowledge. By connecting information to knowledge and wisdom, the act of knowing itself is associated with any activity that involves information. Information and data become units of rationalisation, representing the infrastructure

on which wisdom can be built. I am critical to these conceptualisations as will be evident from the below analysis.

Although the DIKW pyramid places data at the very bottom as the basic units of wisdom, without any trace of their pasts, innumerable standards and infrastructures support the creation, transmission, and maintenance of data. Similarly, for information there are very few traces of the work required to construct the information, and to keep it stable, although there is an emphasis on its construction from data (or from signal, in the case of information theory).

Above all, both conceptualisations allow information to be abstracted away from its source while remaining stable; in other words, information stands alone, independent of its past, future, medium, or producer. These interpretations of information disconnect it from social and political consequences. They implicitly assume that information can stand independently, unmarked by power struggles that have shaped its production, and unaffected by the materiality that sustains them. Additionally, both conceptualisations posit a type of information free from subjects involved in its production and consumption as social and political entities. For example, although senders and receivers are necessary for information to be transmitted in information theory, the characteristics that would distinguish those senders and receivers from others are non-existent. They are disembodied, and their differences are abstracted away (French 2009).

Within the ICTs I have designed, developed, and analysed, information appears mostly as a collection of *stable facts* (Latour and Woolgar 1979) that describe the world. In outbreak detection, information tells us what we know about now, so that we can know what to do next. In simulation, it tells us what the world is, to allow us to create it within the machine. In the smart home, it defines resource consumption, so that we may change it. I emphasise that these are *claims* about information. What is claimed and what happens are two stories among many, and by analysing ICT discourses I demonstrate that such stories also have implications for the types of knowledge produced by them.

ICT is put to use in a wide variety of contexts, and in those contexts it is most often described as an unproblematic good, a technology that benefits everyone. To challenge this representation, I turn to the field of surveillance studies. The digital technological systems that I examine are all used in observing, analysing, and constructing the

outside world. They are used to gather and classify observations regularly, and the classifications are used to generate information about the observed subjects. Surveillance is linked to sustaining public health, providing welfare services, and preserving and protecting the environment. In all three cases, ICTs are strongly tied to the practice of surveillance. In short, they act as surveillance systems, and I discuss this aspect in further detail in the following chapter.

I aim my critique at the tacit assumption that the development and usage of ICTs are always beneficial to society. When a technological system is described as ICT, it is easier to portray it as part of the neutral background, or the way things are. Describing it as a surveillance system, on the other hand, makes it more noticeable, as surveillance draws forth associations to social concerns and political implications. If they were to benefit *everyone* equally, then there would be a lesser need to problematise. However, given the impossibility of anything benefiting everyone equally, and indeed, the impossibility of even defining everyone, in my analyses I find it important to ask: who benefits from the design, development, and use of surveillance technologies, and who suffers its costs?

1.1 Aim

My aim is to identify how ICT and surveillance are linked in different discourses and to describe the implications of this connection. I investigate how the link between the two is motivated, and what knowledge and truth claims are made in relation to it. I identify the subject positions, practices, values, and the forms of knowledge that are prioritised in ICT discourses. The research questions for my empirical investigations are:

1. What is the role of surveillance in the discourses of ICT design and development for sustainability and for public health?
2. How are surveillance subjects positioned in relation to ICT-based surveillance in these two discourses?
3. What are the implications of these positionings?

1.2 Scope

I investigate ICT-based surveillance systems two domains, public health surveillance and sustainability, through five papers included in this thesis. Within public health, I analyse methods of computer supported outbreak detection. Within sustainability, I examine texts that propose systems that deal with energy consumption, especially in connection to smart homes where inhabitants are monitored using ICTs. The two initial papers are a result of my participation in the design and development of two systems: one for Swedish national outbreak detection, and another for providing suggestions to policy makers. In the final three papers, I analyse texts that describe the design, development, funding policy and prescribed use of various technologies in the two domains. In this section, I describe the two domains in further detail to provide a background for the papers.

Computer-Supported Outbreak Detection

Within public health, outbreak detection refers to the monitoring of the spread of communicable diseases within a population. The primary goal is to identify the increase in the number of individuals who suffer from a communicable disease as quickly as possible, because it becomes increasingly difficult to control the spread of diseases as the proportion of the population suffering from the disease rises (Lombardo and Ross 2007). This proportion is called prevalence, and it is one of the two most commonly used measures of disease frequency when reasoning about the spread of disease. The second one is incidence, which is defined as the probability of an individual becoming infected with a disease. Although it is more accurately expressed as a rate, in practice it is also used to refer to the number of newly diagnosed cases within a given time period.

In these definitions, there is a tension between the diagnosed cases and the *real* number of people with the diseases. If infected individuals remain undetected by the health authority, the *known* prevalence of the disease remains zero, but the real prevalence is higher. The tension is due to the difficulty in distinguishing the series of constructions that lead to an expression of prevalence. Although prevalence is a simple numerical expression of the state of a communicable disease within a population, the chain of knowledge that concludes with

its expression is highly complex and not easily visible. A prevalence statement requires the tools and the trained staff to obtain and store samples from the population, tests to determine the presence of the disease, statistical methods to derive the appropriate numbers from the tests, as well as a health authority that can coordinate all of these activities.

Using the two epidemiological terms, outbreak detection can be described as the practice of monitoring prevalence, and its goal can be rephrased as determining unexpected increases in incidence. It is performed regularly by epidemiologists, and it grows progressively more difficult as more diseases and larger populations begin to be monitored (Dato, Wagner, and Fapohunda 2004). The regularity of the task, the quantitative means of accomplishing it, and the growing computational demands make outbreak detection an ideal problem for ICT-based solutions (Morse 2007; Hersh 2009). Since the detection itself relies on statistical methods, and the quantitative data is easily stored and communicated digitally, computer support is highly compatible with outbreak detection (Buckeridge et al. 2008; Hulth et al. 2010; Pelecanos, Ryan, and Gatton 2010). The mobilisation of resources to design and develop the systems, to make them function in offices of epidemiologists, and to educate everyone involved in their use and maintenance, however, requires much effort (French 2009).

Computer-supported outbreak detection relies on the existence of a collection of records, called a case database, where reports of diagnosed or suspected cases of communicable diseases are stored. These databases collect various details that relate to the diagnosis, most important of which are date and location. Using these two, and given population size, prevalence and incidence can be calculated for different regions. Each field can also contain sub-categories specifying properties of the disease that might be useful in future investigations. For example, whether the disease was communicated within the national borders, or if the infected individual travelled from outside the borders after being infected. Case databases collect the details of individual cases of diseases together with date and location. Classification and sorting are already at work as the cases themselves are categorised into administrative regions, age groups, and disease types (Krieger 1992). Moving from case databases to computer-supported outbreak detection, the cases are transformed into prevalences and incidences, and the increases in incidence can be compared with calculations from

previous periods (weeks, seasons, or years) to detect outbreaks.

The first paper included in this thesis describes the design of a computer-supported outbreak detection system where I worked as a designer and developer. It marks the initial stages of my involvement in knowledge production using ICT-based surveillance. It demonstrates the complexity of constructing a system to navigate the multiple and nested disease classifications. It also serves as an example of an ICT discourse where concerns such as efficiency and timeliness are valued over considerations of the effects of such systems on the organisation of public health practice or the implications of their use for the prioritisation of certain types of knowledge.

Syndromic Surveillance

The fifth paper in this thesis deals with a sub-domain of computer assisted outbreak detection called syndromic surveillance. In that paper, we analyse reports from an EU-funded project for developing a syndromic surveillance system. In both their structure and in their assumed audience, the reports resemble my first paper on outbreak detection. Many such systems exist, and in my licentiate thesis, I provide a more detailed overview of these systems (Cakici 2011a).

Syndromic surveillance is a type of computer-supported outbreak detection where a large number of data sources in addition to case databases are used to decrease the delay between the event and its detection by the health authorities (Buckeridge et al. 2005; Das et al. 2005). The earlier applications of syndromic surveillance were characterised by the use of health-related data that precede diagnosis, and the continual monitoring of disease indicators to detect outbreaks of communicable diseases earlier than traditional methods (Mostashari and Hartman 2003; Buehler et al. 2004; Burkom et al. 2004; Henning 2004). More recently, the definition of syndromic surveillance has broadened to include the monitoring of non-communicable diseases and other health conditions such as heat-related illnesses, injuries caused by tornadoes, or respiratory illness after wildfires (Buehler et al. 2009).

Two more terms are necessary to better understand the significance of syndromic surveillance. The first, sensitivity, is the probability that an outbreak is detected, given that it is occurring, or has occurred (Kleinman and Abrams 2006). This measure can also be stated

as the answer to the question: out of all the outbreaks that have occurred in a given time period, how many have we detected? It is useful in expressing the outbreaks that a system fails to detect. The second, specificity, is the probability of not detecting an outbreak given that no outbreak is occurring (or has occurred). As a measure, it is often used to answer the question: how many of the detected outbreaks were false positives? Clearly, reasoning using sensitivity and specificity presupposes that an outbreak actually occurs or does not occur. Outbreak, however, is an ambiguous term, and to detect an outbreak is also to contribute to the construction of an event as an outbreak. For the current discussion, I use a shortcut and reason about outbreaks as they are reasoned about in outbreak detection literature (as a detectable event that exists independent of the act of detection), but in the fifth paper, we locate and analyse a similar construction mechanism at work in the term health threat with a more critical perspective.

An ideal outbreak detection method should detect all outbreaks and give no false positives, that is, it should have 100% sensitivity and 100% specificity. In practice, however, these two terms trade off against one another; more sensitive methods also create more false positives, and less sensitive methods miss more outbreaks (Bravata et al. 2004; Buckeridge 2007). The issue is further complicated by the context-dependent definition of outbreaks. While some diseases with low prevalence may have outbreaks which consist of a handful of identified cases, the same number of cases may be considered insignificant for diseases with high prevalence (Cooper et al. 2006; Buehler et al. 2008).

Syndromic surveillance aims to both reduce the delay in detection, and to increase the sensitivity of outbreak detection during the early phases of an outbreak without affecting specificity adversely. To accomplish this task, syndromic surveillance systems use data sources collected for purposes other than outbreak detection such as ambulance dispatch logs, emergency hotline calls, emergency room admissions, and over-the-counter medicine sales. It is a broader public health surveillance practice that takes advantage of the large number of ICTs used for public health administration. In the fifth paper, we analyse reports from a syndromic surveillance design and development project, and we discuss some of the implications of using ICTs for public health surveillance.

Outbreak Simulation

The second paper in the thesis discusses issues in simulating outbreaks, primarily with the aim to provide advice to policy makers on outbreak preparedness and response. In an outbreak simulation, a software model of a population is infected with a communicable disease to better understand how it spreads within that population. It is not necessarily a software-bound method, but many contemporary models require long and repetitive mathematical computations unfeasible to perform without a computer.

Outbreak simulations can be divided into two groups: compartmental models, and data-driven models. Compartmental models assume the existence of a homogeneous population where individuals can occupy a particular set of states, and transition between them. Its most basic form is an SIR-model, named after the three states: susceptible, infected, and recovered/resistant/removed (Anderson and May 1991). In this model, the whole population begins in the susceptible compartment, and parts or all of it move first to the infected compartment over time depending on the infectiousness of the disease being simulated, and then to the recovered/resistant/removed after spending a certain amount of time in the infected compartment. This type of model often produces clear graphs of the progress of the simulated disease, and remains widely used as an estimate of communicable disease spread, both despite and because of its simplicity.

Data-driven models, on the other hand, require complex software to execute, and rely on having access to a wealth of surveillance data about the population to be simulated (Ferguson et al. 2005). While the types of simulations vary based on the available data, in the second paper we are primarily concerned with simulations that use population data, that is, data gathered using different surveillance methods on parts or the whole of the population. Simulation models that use population data can be further categorised by the heterogeneity of the populations they represent. The least heterogeneous are macro-level simulations, where the individual members of the population are differentiated by few features, and the most heterogeneous are the micro-level simulations where the properties of each individual may differ from the rest (Eubank et al. 2004). The amount of assumptions that must be made to construct the models varies correspondingly, as microsimulations require making more assumptions compared to the

macro-level models.

The second paper highlights our experiences from working with a simulation model for influenza outbreaks designed to provide advice to policy makers. We detail some of the difficulties of representing assumptions articulated by the policy makers in software, as well as the limitations of regional population data when dealing with policies pertaining to broader populations.

Sustainability

The second domain covered in this thesis is sustainability. The third and the fourth paper contain analyses of ICT discourses where surveillance systems are used to contribute to goals such as environmental friendliness and lowering carbon emissions. Although I have chosen the label sustainability to position papers three and four which both focus on ICT design and development projects, the term itself refers to a much broader set of issues. It originates from an earlier term, sustainable development, defined in the Bruntland Report as “development that meets the needs of the present without compromising the ability of future generations to meet their own needs” (UN 1987). Sustainability covers similar issues, but I interpret the term as not indicating a particular focus on development, but emphasising the ability to sustain needs. While a discussion of sustainability itself remains outside the scope of my investigations, I recognise that ecological issues are, and will continue to be, significant concerns for contemporary society. For example, Bradley (2009) illustrates how strategies for sustainability are underpinned by middle-class norms, and uses a discursive approach to justice to demonstrate the consequences of various official and everyday sustainability discourses. Shove (2004) argues that policies designed to promote sustainable consumption use very narrow models of human behaviour, and may even legitimise ultimately unsustainable patterns of consumption. Similarly, in the two papers dealing with sustainability, I have been motivated by the importance of highlighting the social and political assumptions connected to particular visions of sustainability.

In the texts I have examined for papers three and four, sustainability appears as a goal of ICT-based surveillance systems. The systems discussed in the texts are labelled as *smart*, which sometimes appears as a synonym for sustainable, and sometimes works to emphasise a par-

ticular technology as more desirable. Additionally, the word “smart” attaches contemporary ICTs to other forms of technologies, including other, older ICTs. Some examples of these are smartphones, smart TVs, smart homes, smart grids, smart cars, etc. In all of these examples, a previously available technology is enhanced using data processing features, and the proliferation of data also brings an increase in the surveillance performed by and around these technologies.

One digital technological system in particular appears regularly in papers three and four. Labelled the *smart home*, it refers to a collection of technologies that use ICT to monitor, detect, and control a wide variety of features assumed to be found in homes. For example, a home can be described as smart if it includes technologies to adjust indoor temperature according to certain factors such as the outdoor temperature, the number of people currently present, or the current price of electricity. These technologies are not necessarily located inside the home, but the possibility of manipulating properties of the indoor space allows the word smart to be associated with the interiors of residential spaces.

In the third paper, I analyse the discourses of three different ICT design and development projects that deal with smart homes and sustainability. In these residential spaces, novel surveillance systems are proposed to aid the inhabitants in behaving sustainably, which in the texts stands for reducing carbon emissions or providing new ways to schedule energy consumption within the home. In the fourth paper, we analyse EU policy documents from the Seventh Framework Programme and a collection of projects financed by it which describe similar systems. In our analysis we focus on the assumption that behaviour can be changed using ICT-based surveillance, and problematise this view of technology as a value-neutral vehicle for providing information to drive behavioural change.

1.3 Contributions

The five papers included in this thesis focus on two domains, public health surveillance and sustainability. The first three papers were all published, and have been included here as published by the three respective journals, with only type-setting adjusted, to adhere to the thesis format. In all papers except the third, I have worked in collabo-

ration with other researchers, and while describing the contributions I use the first-person plural pronoun for those papers. The individual contributions of each paper included in the thesis are as follows:

1. We provide a software package for computer-assisted outbreak detection designed to ease the process of connecting a data source containing reported cases to various detection methods which require different ways of formatting the incoming data.
2. We present a workflow for using simulation models to provide policy advice regarding communicable diseases. We argue that it is challenging to use, verify, and validate regional and sensitive data sets, and we emphasise the importance of ensuring that the assumptions built into the model represent the wishes of the policy makers.
3. I argue that design documents for smart homes present intersecting visions of sustainability entailing the wide-spread use of ICT. In the documents, the inhabitants are made individually responsible for living sustainably, and surveillance is positioned as integral to this future with the help of ICT.
4. We argue that EU policy documents and project descriptions within the Energy challenge of the ICT category of the Seventh Framework Programme use models of social change that have been widely criticised as unlikely to lead to substantial changes in resource consumption. We show that these texts discuss only the potential positive effects of technological surveillance, but neither acknowledge nor require the handling of the potential negative effects of surveillance within everyday interactions.
5. We argue that reports from a European Commission co-funded syndromic surveillance project called SIDARTHa construct the concept of a health threat as a sudden, unexpected event with the potential to cause severe harm, and one that requires a public health response aided by surveillance. We further argue that syndromic surveillance discourse privileges expertise in developing, maintaining, and using software within public health practice, and it prioritises standardised and transportable knowledge over local and context-dependent knowledge.

The five papers contribute to different fields of research with overlapping interests, and to bring them together, I outline my theoretical framework and a longer discussion concerning methodology in the following chapters. I also detail additional considerations which motivated the papers, but were not included in them due to their limited scope. While working with and studying ICTs, I have encountered a universalising discourse, one that does not place as much value in local contexts and local differences, instead aspiring to create a smooth, homogeneous body of knowledge in whatever domain it is applied to. I have observed it as a hammer-and-nail issue: if the only available tool is ICT, every problem starts looking like a mix of factors that can be sorted and classified, that is, problems do not exist out there for ICTs to solve, but they are formulated to fit the available ICT (Wihlborg 2000). In the last three papers, I have worked to identify specific examples of this universalising move, and to show what is valued discursively in different domains.

In addition to the individual contributions of the papers, I have had an overarching goal which I have been able to approach from different directions in the last three papers. For that reason, I consider my overall contributions to this goal to be located within those papers: ICTs are not value-neutral tools that reflect reality; they privilege some forms of action, and they limit others. They are also imbued with values, and different subjects benefit or suffer from their use differently. My contributions are three empirical studies of surveillance discourses where I identify the forms of action that are privileged and the values that are embedded into them. In these discourses, introducing ICT entails increasing surveillance, privileging technological expertise, and prioritising centralised forms of knowledge.

1.4 Disposition

In the following chapter, I introduce the theoretical framework for my investigations. I position myself in relation to different theories I have drawn from feminist theory, surveillance studies, and science and technology studies. In the third chapter, I discuss my methodological concerns with reference to the theoretical framework. I motivate why I chose to study design documents to identify social and political implications of ICT-based surveillance, and why discourse analysis is a

suitable method for analysing design documents. In the final chapter of part one, I provide a summary of the five papers, and I state the answers to my research questions.

The second part includes the five papers in chronological order. In the earlier papers, my concern is primarily technical, and I aim to understand and discuss different issues in order to be able to construct better, more efficient systems. Beginning with the third paper, my focus shifts to understanding the discourses in which these systems are described and constructed. The first three papers have already appeared in academic journals, while the last two have not been published yet.

Chapter 2

Theory

To explain how my contributions fit together, I outline my theoretical framework in this chapter. I begin with a personal account based on my research experience to introduce issues relevant to my investigation. I then define my epistemological position, and describe my interpretations of discourse and its constituents. I continue by introducing theories concerning technologies, and by highlighting the importance of classifications within them. I conclude with a discussion of theories I have used to understand the different components of ICT-based surveillance.

2.1 Borders

During my research, working between computer science and other disciplines within the social sciences and the humanities, such as science and technology studies (STS) or surveillance studies, required that I modify my vocabulary when presenting the same work at different academic institutions. I often described my work as the critical analysis of surveillance systems, but for audiences not familiar with surveillance studies, I tended to substitute *surveillance* with *information*, and then explained how the routine collection and processing of information could be connected to theories of surveillance. The terms I chose to modify or leave unchanged led me to question my own assumptions about different disciplines, and the facts I assumed to be accepted within them. When I added a slide explaining the term information systems to a non-computer science audience, I encoun-

tered within my thoughts previously unquestioned *facts* that I have carried over from my natural science education. More importantly, I was forced to realise that my classifications of audiences as “computer scientists” or “STS researchers” fell short of describing the heterogeneity of any group I presented my work to. My computer science affiliation, and my previous degrees from technical universities, allowed others to position me as a system designer in the projects I was involved in, providing me with a different kind of access to the produced material (often in the form of project proposals and system specifications). The proximity made it easier for me to communicate with other designers, but also led me to continuously question the limits of my own reflexivity, and hinted at the existence of questions that may never occur to me due to the familiarity of the material. For example, in my research on ICTs for sustainability that I discuss in papers three and four, I tried to identify how the designers, in their descriptions of the systems, formulated their problem using the vocabulary of computer science, and how they described the problem as a purely technical matter, separate from the social and the political.

The possibility of being mobile between different disciplines was highly productive for my own research. At the same time, however, I found my own physical mobility to be constrained regularly by my Turkish citizenship, and the bureaucratic demands that it placed on me while working in Sweden as a non-EU national. Although travel within Europe is relatively free of paperwork for its own citizens, for the non-citizen these issues are much more complex. Decisions from the Migration Board take anywhere from six months to a year, and those periods required that I either not leave Sweden, or not return until the Board had reached a decision. My movement outside the country was tightly coupled to the delays in the bureaucratic process, requiring me to plan my attendance to conferences never as a certainty, but as a possibility that the decision would arrive in time. For countries that require lengthy visa procedures prior to the visit the complications were multiplied. For example, towards the end of my third year, I gave up on trying to attend a conference in the United Kingdom solely due to the complexity of the visa process, which I had been informed could require one or more interviews in person.

My most immediate strategy to deal with these various mobility limitations was to reinterpret the situation as “who is to say whether borders keep me in, or keep everyone else out”, referring to com-

monly advertised image of Sweden as a comfortable space. However, this reinterpretation did not introduce sufficient change; conference participation and visiting researcher positions that often seem an essential part of academic conduct became simply impossible to accomplish for long stretches. I wrote this section originally as an extended abstract for a workshop where I was inclined to add:

In a strange temporal shift, as I write this abstract, I wait for yet another extension, hoping that it will arrive in time to allow me to travel to Copenhagen, but if this document has reached you, chances are that I have already received my extension.

The decision arrived exactly eight months after submitting my application to the Migration Board, fortunately in time for the workshop, allowing me to cross the borders of the Kingdom of Sweden without risk of detainment. While I have now acquired a Swedish citizenship, and the challenge of European border crossings has eased significantly, my personal experience of borders, residence permits, passport checks and endless waiting has remained as a vivid reminder of a particularly unjust form of bureaucracy made highly efficient with the help of ICT, and one that countless people continue to experience while crossing borders every day. It has also informed and motivated my research into the intersections of ICTs and surveillance.

The massive effort required to uphold a nation state border illustrates that it is extremely fragile, that it would disappear if it were not upheld. Even in its fragility, however, it is a force to be reckoned with, having immense potential to cause frustration, disturbance and suffering. Inspired by my encounters with the state apparatus continually reconstructing its own borders through forms and applications, I looked for the work of upholding borders in other domains. Academic disciplines, in close resemblance of nation states, also rely on the idea of borders with the goal of creating a more homogeneous mix within than without. Although this clean demarcation serves as a shortcut for identity, as in “a computer scientist” or “a Swedish citizen”, it is constantly challenged by the intersectionality of the identities that threaten to spill outside. In facets of identity, the border is not a discrete line dividing one from another, but a continuum where change occurs only gradually. When the border is one that surrounds

a nation state, enforced by the threat of violence, its supporting network extends into the paperwork that needs to be filled before receiving the right to cross the border, the regulations that need to be followed in order to file that application, and the authorities that need to be convinced that the applicant is a legitimate subject worthy of their expertise. All play a part in imagining the borders that divide nation states. The academic discipline also supports itself with less violent but otherwise similar structures in the form of credentials and affiliations that assert the authority to publish, the norms that must be followed to create the suitable scientific article, and the subject that must be shaped to be accepted by peers as a worthy producer of knowledge. Moving closer to my own research topic, ICT-based surveillance plays a major role in the maintenance of both types of borders. Through passports, identification cards, no-fly lists, and databases of movement, the nation states reaffirm the existence of their borders. The reaffirmation of the borders of an academic discipline through surveillance is more elusive, but the most recent tool of choice seems to be the broadening usage of impact factors for journals, and ranking algorithms such as the h-index for individual researchers, where publications in a single discipline tend to translate to scores higher than their interdisciplinary counterparts, with some exceptions for a few well-known interdisciplinary journals.

Another parallel between the academic and the residence-seeking subject is in the act of applying: grants for the former, and permits for the latter. While my experience with the grant-receiving-subject is limited, I am intimately familiar with the permit-receiving one. After participating in the residence permit process several times over the past few years, I have gradually come to recognise the art of becoming the ideal subject for a residence permit application with its own collection of taboos and dangerous topics. The performance of this role leaves material traces visible throughout the application process (Does the form contain spelling errors? Is it filled out on a computer? Is the form bent, creased, or folded?), but its most intense activity is during residency interviews, and to a larger but less embodied extent, in the free-text fields in application forms that allow deviation from the standard template for a few lines.

In my performance of the applying subject, I found myself trying to relate to different construction processes that transform the diverse into the similar and vice versa. While my residence application

painted me as suitably similar to reside, my academic persona continues to gain from a certain difference marking me as the one that possesses valuable experience of the Other. I pose the questions to myself, and leave them open: What is it that makes me different, and in what ways does my difference affect the institutions I participate in? Additionally, am I different enough, and are my differences representative enough of the Other to justify the ensuing privilege of being granted a position? What role does my gender, race, and ethnicity play in the perception of myself as a subject?

An academic career provided me sufficient disguise to craft my own permit-eligible persona for the authorities, which is one facet of the academic privilege that, like many other types of privilege, tends to become invisible for those who possess it. Privileges of an academic career include participating in knowledge production, greater social mobility, and working further from the production of capital while benefiting from the fact that we live in a capitalist society, among others too numerous to list. I find it important to begin with a discussion of existing privilege, as the act of occupying a margin seems dangerously easy to romanticise. Haraway notes:

But here lies a serious danger of romanticizing and/or appropriating the vision of the less powerful while claiming to see from their positions. To see from below is neither easily learned nor unproblematic, even if 'we' 'naturally' inhabit the great underground terrain of subjugated knowledges. (Haraway 1990, 191)

While speaking from the margin is certainly a potent resistance strategy when trying to establish a position to speak from, it may become counter-productive for analysis. That is, even while speaking about margins of interdisciplinarity, or of citizenship, my discussion is grounded by the intersections of an academic privilege and a male privilege (Wennerås and Wold 1997), as well as other types of privilege which are invisible to me simply because I possess them. Prior privilege will always ground the interpretations power and equality, but it is essential to acknowledge its presence, and to try to contextualise it if I am to understand the limits of the knowledge I produce. Having said that, there is much to be gained from the experience of occupying the margins. Both in the possibilities of interdisciplinary mobility, and in

the constraints of the lack of geographic mobility, I find inspiration and questions waiting to be asked. I use these questions to construct a stable narrative expressed partially as a thesis, but also as the identity of a researcher who works at the intersection between computer science, STS, and surveillance studies.

2.2 Situating Knowledge

I begin with Harding's (1995) account of *strong objectivity* to locate my own role in knowledge production, and to understand my own role in the intersections of various power relations. Harding argues that not every subject can access all knowledge in all contexts, pointing out that "[s]tandpoint theories argue that what we *do* in our social relations both enables and limits (it does not determine) what we can know" (Harding 1995, 341). As an example, the statements I choose to include in this thesis can travel far from their spatial and temporal origins on academic networks, but at the same time they are produced and distributed in a culture of peer-review that is grounded in conferences and journals which exclude participation outside of a limited and privileged group of academics. The conventions of the academic writing format also exclude statements that are not of interest to the academic community. These are not necessarily shortcomings, but reflecting on them helps making tacit assumptions more visible. Connecting this position to my earlier discussion on privilege, the notion of strong objectivity pushes me to reflect on my current subject positions. Specifically, since I am socialised as a man and others identify me as such, I may not be able to know and understand the oppressive effects of patriarchy, especially because I am more likely to benefit from those structures, regardless of my awareness. This in no way means that I should not try to know more, but simply that I first need to acknowledge both the existence of the epistemological gap, and the potential impossibility to bridging it, before I attempt to cross it.

Harding states the strong objectivity program draws on feminist standpoint epistemology, which has its own limitations (331). While I do not provide a discussion of standpoint epistemologies themselves, I find it important to acknowledge its multiple and contested definitions as a critical feminist theory first emerging in the 1970s and 1980s by scholars concerned with the relation between knowledge produc-

tion and practices of power (Harding 2004, 1). Longino argues that the two claims of some standpoint theories, knowledge as a socially situated, value-laden construct, and that marginalised standpoints offer “epistemically superior” positions, are in conflict with one another (Longino 1993). In response, Harding states that “[s]tandpoint theory is not arguing that there is some kind of essential, universal woman’s life from which feminists (male and female) should start their thought” (Harding 1995, 344). In strong objectivity, the emphasis is not on the essential properties of a marginal subjectivity, but on how that subjectivity is positioned in relation to knowledge production with a necessarily a limited perspective. Continuing with Harding’s definition, “[s]tandpoint theories, in contrast to empiricist epistemologies, *begin* from the recognition of social inequality” (341). This stance allows theories to be constructed not on a neutral space shared equally by all, but on a contested ground where the positions of different knowers influence the knowledge produced by them.

The notion of strong objectivity allows me to describe the place where I know from, and to locate the knower as a historically and culturally specific subject. Given this position, I also view knowledge not as a universal constant, but as a bounded construct that is tied to its local context. To detail this position further, and to describe how such knowledge can be used, I turn to the notion of *situated knowledge*. Haraway (1990) defines it as a “doctrine of embodied objectivity” (188), and at the same time as an alternative to relativism as “partial, locatable, critical knowledges” (191). Situated knowledges cast doubt on a scientific project that aims to understand *everything*. I use the notion to acknowledge the incompleteness of knowledge, and to be wary of universalising truth claims, especially those that necessitate reductionism to argue for validity, as situated knowledge emphasises the locality of knowledge production, and necessitates the reflexivity of the knower. As Haraway notes:

Feminists don’t need a doctrine of objectivity that promises transcendence, a story that loses track of its mediations just where someone might be held responsible for something, and unlimited instrumental power. (187)

There is a clear tension between what this position entails, and the truth claims that underlie much of the knowledge production in

the disciplines of natural science and technology. Through strong objectivity and situated knowledge, I construct a position to question these truth claims. Although Harding and Haraway have debated the compatibility of their programs in detail in their earlier texts (cf. Harding's response to Haraway in Harding 1986, 193-194, and Haraway's commentary on strong objectivity in Haraway 1996 438-439), I interpret an emphasis on the subjects involved in knowledge production through strong objectivity, and the limitations of the produced knowledge through situated knowledge as compatible positions in my own research. Having said that, both concepts have implications for method choice, on what can and cannot be said through academic inquiry (of texts, in my case), which I discuss in the following chapter on methodology.

Since I have performed my research across several disciplines, first as part of a health authority, then within two different academic departments, as well as at an applied research institute, I find it essential to ground my concerns about knowledge production, and the role of the researcher in theories. One final part of conducting research across several disciplines, and working in domains which involve the *public* (primarily as surveillance subjects in both public health and in smart homes) from a critical perspective is the issue of speaking for others. During my studies, I was occasionally asked not only to identify problems, but to provide more *correct* ways to solve them as well, and in these more normative forms (such as project reports), I have tried to highlight the importance of the voices of those that are not involved in crafting the solutions. If the answer to the question, "Can we speak for others?", is not taken for granted as affirmative, normative statements about what should be done for others become much more difficult.

I have found Spivak's (1988) analysis of the subaltern to be informative when grappling with speaking for and about others, where she emphasises the importance of acknowledging the potential impossibility of hearing and interpreting some statements. As Spivak suggests, in attempts to speak for the Other, privilege can be considered as a loss, a condition that makes the subject unable to understand others. As privilege itself is implicated in creating the marginalised Other, those who possess it may not be able to detach themselves from the forces that sustain the divisions. To give up the privilege is not a solution either, as it is often not possible to begin with, but even

when possible it may entail giving up the possibility of being heard, thus rendering the attempt at speaking for others useless. Recognising these conflicts, I use the notion of *privilege as loss* (Spivak 1988) to acknowledge the existence of the subject positions and discourses that cannot be seen, heard, or interpreted from an academic standpoint alone.

I use the theories outlined above as their combination enables a productive meta perspective on my work, and helps me to acknowledge and handle the fluctuating demands of the different disciplines I attempt to bridge with this thesis. In the course of my work, I have interpreted texts that describe different types of ICTs used for surveillance, while drawing from my own experiences as a system developer responsible for designing such technologies. In this attempt, I have encountered and occupied many different subject positions: computer scientist, philosophy of science teacher, STS scholar, Turkish citizen, Middle Eastern immigrant, Swedish citizen, to name a few. I have tried to incorporate this attempt at understanding whom I speak for in which context into my research.

2.3 Discourse Theory

Embracing situated knowledges requires questioning the link between truth and knowledge further, and to do that, I turn to discourse theory and introduce parts that relate to my work. Although my use of discourse theory, and of discourse analysis which I discuss in the following chapter, does not concern itself with the validity of truth claims, investigating the discursive construction of any reality remains productive for analysis, and necessary for the formulation of situated knowledges. In the thesis I have followed Foucault's theories on discourse. In a comment on the usage of the term discourse between his own works, Foucault provides three different definitions (Foucault 1972, 80, cited in Mills 2004, 6):

1. the general domain of all statements
2. an individualisable group of statements
3. a regulated practice that accounts for a number of statements

In the first and the broadest definition, discourse is singular, and every statement belongs to it. The second definition implies that state-

ments within a particular discourse share a common property that allows them to be separated from others, and to be treated as a group. The third definition combines statements with regulation and practice. This is the most productive definition for my work, because it allows me to focus on rules in discourses, and to investigate how these rules are expressed in texts. Additionally, the third definition provides both a non-universal theory for understanding textual exchange, and a way of locating it in practice. By using the third definition, it becomes possible to ask: What is considered to be true in a given text?

The role of truth is essential to discourse analysis, and Foucault's theories on discourse provide additional tools for their investigation. Three concepts are essential to understanding the effects of discourses: power, knowledge, and truth. The first, power, refers to a force immanent in all relations, inherently productive, and possible to resist, or rather, that constitutes resistance (Foucault 1978, 92–96). It is not only repressive, but also constructive. Power connects the social body, and it is a precondition of knowledge.

The second element, knowledge, is tied to power. It arises from power struggles (Mills 2004, 14), implying that wherever we encounter knowledge, we also encounter the traces of a power struggle which establish, or have previously established, some objects as belonging to the body of knowledge. Knowledge and power are intrinsically linked, and Foucault uses the term *power/knowledge* to emphasise their entanglement.

The final element, truth, is a representation within a discourse, and it is only accessible as part of a discourse. Jorgensen and Phillips (2002, 12) identify two distinct conceptualisations of truth within Foucault's theories: as a system of procedures for the production, regulation, and diffusion of statements during his earlier work (the archaeological phase), and as embedded in and produced by systems of power in his later work (the genealogical phase). Regardless of definition, however, the analysis of truth does not concern itself with whether a statement is true, or whether it reflects reality. Rather, it seeks to establish how knowledge is positioned as occupying the position of truth.

Although I deal with what I consider to be primarily theoretical concerns in this chapter, and I revisit issues relating to method in the next chapter, the two are intrinsically bound. Said (1978) argues that truth is itself a representation, and describes its methodological impli-

cations:

[...T]he real issue is whether indeed there can be a true representation of anything, or whether any and all representations, because they *are* representations, are embedded first in the language and then in the culture, institutions, and political ambience of the representer. If the latter alternative is the correct one (as I believe it is), then we must be prepared to accept the fact that a representation is *eo ipso* implicated, intertwined, embedded, interwoven with a great many other things besides the “truth,” which is itself a representation. What this must lead us to methodologically is to view representations (or misrepresentations—the distinction is at best a matter of degree) as inhabiting a common field of play defined for them, not by some inherent common subject matter alone, but by some common history, tradition, universe of discourse. (272–273)

The three elements of discourse theory allow the investigation of how different texts construct their knowledge and present truth claims. In discourses, the elements are associated with different subject positions, and the positions are configured in relation to knowledge and truth claims. While the positions and their configurations may vary endlessly depending on the discourses, two specific positions, the reader and the author, hold greater significance regardless of the text. The former is expected to both interpret the text and to simultaneously occupy different subject positions, for example as an expert, a student, a citizen, an observer, etc. The latter, on the other hand, is assigned the responsibility for having produced the text in the first place. In this context, strong objectivity and situated knowledge have allowed me to acknowledge the effects of my own position as a reader performing discourse analysis. I return to the issue of authorship in the following chapter on material and methods to discuss the importance of different authorship configurations for documents written by committees, texts with a large number of authors, or texts credited to institutions instead of human authors. Through strong objectivity, situated knowledges, and discourse theory, I have described

the basis for my epistemological stance. In the following sections, I turn to issues that relate closely to ICTs as the objects of my study.

2.4 Studying ICTs

To understand ICTs as systems with social and political consequences, I rely on texts and theories from the field of STS, where many scholars have long been concerned with theorising the technology-society relationship (Bijker and Law 1992; MacKenzie and Wajcman 1999). In the domains I have studied as well as in the offices where I have conducted my research, the word technology is almost always used to describe digital technological systems. Computers are considered to be technologies, while other common artifacts such as pencils, sheets of paper, tables, books or even bridges are not included in the set of technological objects. Using an STS perspective, I have been able to understand technology much more broadly as “things that people have made” (MacKenzie and Wajcman 1999, xiv), which has allowed me to locate ICT as a historically contingent set of relations with origins in a long history of constructing, using, and modifying tools.

In the previous chapter, I began by stating that ICTs are not value-neutral (Forsythe 2001). Although this is considered common knowledge by contemporary STS scholarship, it remains a fairly radical statement in most computer science discourses. While engaged in the well-known debate on the politics of artifacts, Winner states:

No idea is more provocative in controversies about technology and society than the notion that technical things have political qualities. At issue is the claim that the machines, structures, and systems of modern material culture can be accurately judged not only for their contributions to efficiency and productivity and their positive and negative environmental side effects, but also for the ways in which they can embody specific forms of power and authority. (Winner 1980, 19)

The ensuing debate within STS has ventured into questioning the facticity of the examples (Joerges 1999), and the role of narratives in constructing knowledge (Woolgar and Cooper 1999), while in a

computer science context the original argument retains its provocative power. Especially in discourses that tend toward technological determinism, stating that technologies are not driven merely by their own internal logic is controversial enough.

Faced with this conflict, where a statement about the political qualities of technologies would be considered novel in the discipline where I received most of my academic training, but common knowledge in other fields I intend to contribute to, I chose to focus on investigating the specific qualities that would allow me to state the social and political implications of ICTs. I reasoned that the former (computer and systems sciences) would benefit from discussions of ICT beyond efficiency and productivity, while the latter (STS and surveillance studies) would have room for empirical studies demonstrating specific implications. Specifically, I have consulted STS literature while investigating two issues in my analyses of ICTs: the role of classifications, and the designer-user narrative.

Searching and sorting are fundamental operations for computer science, and algorithms for performing them are usually the first to be discussed in most introductory computer science textbooks. Both activities involve forms of classification where different elements are placed in different categories. These operations are integral to ICTs as well. For example, classification is at work when an outbreak detection system separates some reports of communicable disease belonging to an outbreak from those that are not, or when a smart energy meter identifies excessive electricity consumption.

Bowker and Star (2000) state that “[t]o classify is human” (1), and that “[each] category valorizes some point of view and silences another” (5). As an activity both humans and ICTs engage in regularly, and one that often involves surveillance, I focus on classifications as constructs with the power to make things appear value-neutral and *natural*. If subjects fit into a classification smoothly, they can easily appear as part of the background. Conversely, those that do not fit may be burdened with the shortcomings of the classification even when they are not responsible for the ill-fitting categories. This possibility marks classifications as important sites when investigating the social and political implications of ICTs. In a discussion of the race classification systems under apartheid in South Africa, Bowker and Star note:

The advantages are those whose place in a set of classification systems is a powerful one and for whom powerful sets of classifications of knowledge appear natural. For these people the infrastructures that together support and construct their identities operate particularly smoothly (though never fully so). For others, the fitting process of being able to use the infrastructures takes a terrible toll. To 'act naturally,' they have to reclassify and be reclassified socially. (Bowker and Star 2000, 225)

By using classifications and sorting subjects, privilege can be assigned or withdrawn. Subjects can also be made invisible in a classification if there are no categories that can contain them. The categories that the system can see necessarily exclude other categories deemed to be unrelated to the problem at hand by their designers. While necessary and not inherently bad, classification involves ethical choices about what is silenced, and what is valorised (5-6). Their effects also influence the practice surrounding the classified objects. For example, in their discussion of the International Classification of Diseases (ICD), Bowker and Star state that the discoveries compatible with ICD are much more likely to occur than a different set of discoveries, as the established data collection routines are themselves based on the ICD, which increases the cost of data collection for alternative efforts (82). Pargman and Palme (2009) also discuss the implications of classifying language representation schemes for ICTs in their study of the ASCII standard, and Lee (2009) describes how particular visions of about knowledge, education, and learning are inscribed into technical standards.

Surveillance systems also rely heavily on classification, to identify, sort, and exclude their subjects (Burrows and Gane 2006; Graham 2005). Murakami Wood et al. (2006) argue that ICTs, as automated systems of surveillance, control, and enforcement, divide citizens into separate socio-technical realms of the premium and the marginal while their agency remains largely invisible. Similarly, Ball (2005) notes that "[s]urveillance involves the mobilization of information categories that order by hidden criteria" (105). Just as functioning infrastructure fades into the background (Edwards 2003), ICT-based surveillance also tends to become less visible for those that benefit from it. This invisibility further strengthens the undisturbed, smooth

functioning of the technologies that measure and record the *outside world* through their sensors. Graham states:

Digital surveillance systems tend to be developed, designed, and deployed in ways that hide the social judgements that such systems perpetuate. (Graham and Wood 2003, 242)

When it is difficult to observe systems under development or in action, it is also difficult to link them to desirable or undesirable outcomes. At the same time, they do not always function smoothly, and they can still be overwhelmed by the diversities of their local contexts (Graham and Wood 2003). ICT-based surveillance does not always function as intended, and even when it does, it is challenging to identify whose intent it is. Often, the term *designer* is used to denote a subject position where intent can be located. However, the agency of the designers or who can be considered in the group of designers is by no means trivial.

The designers are only one of many groups of actors that play a role in creating the technologies, and even the act of categorising designers or users as distinct groups tends to obscure the continually shifting boundaries of design and use (Suchman 1994a, 2002; Woolgar 1991). While designs can have aims that are articulated explicitly, and attempts can be made to prescribe their usage both discursively and materially, the eventual uses of the technologies cannot be inferred solely from aims or the designers' intentions (Latour 1988, 22–29). Within STS, many scholars have demonstrated that technologies are often used in ways that are not intended or foreseen by the designers, and that no single, essential use exists for any technology (Oudshoorn and Pinch 2005).

Both in my earlier work in developing technologies, and in my later work in analysing texts that describe ICTs, the categories *designer* and *user* recur. Initially, I interpreted the work I was involved in as a design process, providing me with the subject position of a designer. From that position, it seemed easy to label those who would use the system as users. This framing has been criticised as placing the concerns of technology producers above the concerns of those not directly involved in the production process (Dourish and Mainwaring 2012). When analysing the discourses of design documents, I have returned to the act of labelling others as users as an important claim

that aims to establish what the user category should include, that is, which subjects are considered to be users, and what sorts of agency are ascribed to them (Suchman 1994b).

2.5 Studying Surveillance

I view ICTs not only as solutions to specific problems, but also as expressions of a large collection of intersecting positions on how knowledge should be organised and distributed in societies. To develop this view, I turn towards the field of surveillance studies. I use theories and findings from the field to frame my discussion of the social and political implications of using ICT in knowledge production. As I have stated in the previous chapter, I aim my critique at the tacit assumption that the development and usage of ICTs are always beneficial to society. Describing an ICT as a surveillance system emphasises associations to social concerns and political implications. Surveillance practices have been shown to negatively affect those who are already underprivileged, whether they are ICT-based (Gilliom 2001; Rooney 2010; Bozbeyoğlu 2011; Maki 2011) or not (Goffman 2009; Finn 2011). However, it is also important to recognise that surveillance is not simply an oppressive force to be resisted at every turn. Depending on the context, it may also entertain (Albrechtslund and Dubbeld 2002) and empower (Koskela 2002; Shilton 2010), although this potential is strongly tied to the ability to control interpretive choices (Ottinger 2010).

Surveillance Society

Within the field, a commonly cited definition of surveillance is “the focused, systematic and routine attention to personal details for the purposes of influence, management, protection or direction” (Lyon 2007, 14). Although I interpret Lyon’s definition to include non-technological forms of watching and sorting as surveillance, in this thesis I am primarily concerned with surveillance performed using ICTs. Following the definition, it is not difficult to classify the vast majority of ICTs as surveillance systems. They are highly compatible with surveillance applications, and correspondingly, contemporary surveillance is commonly performed using ICTs as they are especially suited to performing routine tasks systematically.

While stating that most ICTs are surveillance systems explains the connection between surveillance studies and my analyses of ICT design and development, it does not lead to analytical insights, since stating that all members of a set are also members of another does not automatically provide the means to distinguish the differences within the set. As Ball and Haggerty (2005, 133) point out: “Merely labelling different sociotechnical relationships as ‘surveillance’ does little to enlighten us as to the dynamics of the control, resistance, emergence and development of surveillance practices.” (133)

To help direct the analysis of the similarities and the differences of surveillance systems, Lyon provides five “common threads”: rationalisation, technology, sorting, knowledgeability, and urgency (Lyon 2007, 26–27). I use examples from my studies of sustainability technologies to clarify the implications of these five threads on the study of ICTs. Paper three and paper four provide further analyses of these technologies.

The first thread, rationalisation, appears as a demand to quantify an ongoing process. For example, in a system for monitoring energy distribution and consumption, rationalisation drives towards systems to count the duration, amount, and price of consumption for both providing and consuming energy. The quantification drive is also helpful in understanding the second thread, technology. Computers are highly useful for recording and monitoring, and as these activities are performed using computers with increasing frequency, the meaning of the activity itself also changes. Technology does not simply enable an activity, it changes the process. It is a highly complex concept where discourse, practice, and materiality intersect, and an analysis of technology necessitates addressing questions of agency (Johnson 1988). As I have discussed earlier in this chapter, the field of STS provides a wealth of theories and empirical studies to tackle these questions.

Sorting, the third thread, and the differential treatment that results from it are highly important consequences of living with surveillance systems. How categories are constructed, and how subjects are placed in those categories (Bowker and Star 2000) are especially important to investigate when the object of study is ICTs. In these digital technologies, categories are programmed into the system, either at the level of hardware or software. This additional conversion allows the classification to function, but the criteria for deciding membership in

a particular category are made further inaccessible especially if the program itself is executed as compiled code that is nearly impossible for humans to decipher. As an example, sorting appears in multiple forms in ICTs designed for a smart city district: the sorting of waste as an activity to be performed by the inhabitants for recycling, the sorting of activities as energy wasting and energy saving (driving alone in a private car versus using the public transport), and the sorting of inhabitants as leading sustainable lifestyles or non-sustainable lifestyles.

The fourth thread, knowledgeability, refers to the knowledge or awareness of the surveillance subjects about the ongoing surveillance. Lyon points out that surveillance works best with the cooperation of its subjects, but “what they do know and what they do with what they know makes a difference” (Lyon 2007, 27). Some surveillance systems require or encourage the active cooperation of users, for example smart energy meters which offer different consumption and billing plans. Others also require the explicit agreement of the user, but once in use they are designed to fade into the background and achieving a smooth integration with daily life, for example wearable sensors monitoring bodily signs. Questions of knowledgeability point to political consequences of surveillance with regards to knowledge distribution, but they also emphasise the importance of the subjective experience of surveillance: How much are the surveillance subjects informed about the ongoing monitoring while their traces are recorded by the systems? Are they able to influence the monitoring process or the classifications it relies on? Do they resist surveillance, and if they do, how do they do it?

The final thread, urgency, is particularly important for ICTs where the speed of searching, storing, retrieving, and analysing are common concerns. With ICT, records can be transmitted rapidly from place to place. However, urgency does not necessarily imply short time scales. For smart electricity meters which motivate their use as sustainable activities, there are two time scales where urgency is visible: the everyday time scale, where systems monitor consumption and generate reports regularly (daily, weekly, monthly, etc.) and urgency appears as the wish to transfer the reports rapidly; and a longer time scale extending into the future where consumption methods must be made sustainable urgently to prevent environmental catastrophes (within the next 50 years, within two centuries, etc.).

In the definition of surveillance, Lyon includes the phrase “atten-

tion to personal details". However, systems which do not deal with personal details can still be included in the broad category of surveillance systems. For example, environmental monitoring for detecting temperature spikes, heat waves, or pollution do not directly deal with personal details, but certainly fit the rest of the definition, and their use also has social and political consequences.

The five threads of surveillance as defined by Lyon provide different perspectives to analyse and understand surveillance systems. Through them, I have shown the parallels between ICTs and surveillance systems, and discussed how these elements can be applied to understand the motivations for constructing surveillance technologies. The five threads explain what I refer to when I use the term ICT-based surveillance.

Elements of Surveillance

While the five threads are highly useful in connecting ICT to surveillance and suggesting research directions, an earlier framework, elements of surveillance (Ball 2002) provides tools for identifying the social and political implications of surveillance. In the framework, surveillance comprises four interconnected elements: re-presentation, meaning, manipulation, and intermediation.

Re-presentation refers to the material aspects of surveillance subjects captured by surveillance technologies. These technologies *re-present* the gathered traces, and they allow observers to interpret them differently. Through technology, these traces are moved between different contexts, taking on different meanings. Much like the technology thread suggested later by Lyon (which I have described above), re-presentation ties knowledge produced by surveillance to specific technological capabilities to present that knowledge in different forms and contexts. In the third paper where I discuss smart home technologies, this aspect of surveillance is especially evident in connecting ICTs to sustainability. The traces of everyday life at home are re-presented as consumption stories, which are then used to interpret whether the surveillance subjects live their lives sustainably. Descriptions of the technical capabilities of the wide variety of sensors present in the smart home illustrate the facets of re-presentation in the surveillance practice. Movement, temperature and acts of consumption are routinely observed with the help of sensors placed in the smart home.

The second element, meaning, refers to interpretations of surveillance, and Ball details its three common meanings, as knowledge, as information, and as protection from threat, although it is not limited to these three. According to Ball, knowledge is “held by the people within the surveilled domain, enabl[ing] it to function as a socio-technical network” (Ball 2002, 581). It is a type of knowledge held both by those who perform the surveillance and those who are watched by it. Another common meaning, information, is one of the end-products of surveillance for those operating the systems, and it allows classifications of the surveillance subjects to be made. In the smart home case from paper three, the collected data are given meaning by situating them in the process of creating sustainable city districts, life styles, and travel patterns. With the help of surveillance technologies, the city district is transformed from a polluting space of excessive consumption to a clean space of responsible living.

The third element, manipulation, “refers to the inevitability of power issues under surveillance, not the least because surveillance practice captures and creates different versions of life as lived by surveilled subjects” (583). This element also highlights the subject positions and the forms of agency available to those under surveillance. Ball’s definition of manipulation suggests that different interpretations of technologically mediated events can be used by different parties to further their own interests, connecting analyses of surveillance to subject positions of both the watched and those who watch.

The final element, intermediation, “binds (or unbinds) networks of individual actors or institutions” (Mansell 2002, 4, cited in Ball 2002, 584). Intermediaries mobilise meanings to sustain forms of surveillance. For example, demonstrators which provide narratives about how fictional inhabitants would live sustainably with the help of ICT-based surveillance (paper three), or project reports that describe systems to monitor emergency data sources (paper five), act as intermediaries in both sustaining and resisting surveillance. When the usefulness of a particular technology is questioned, such reports can be used to argue for their benefits, or for their undesirable effects (paper four).

Ball poses several questions to understand surveillance that arise from this framework, and two of the questions are applicable to the last three papers in this thesis:

- “What kind of processes determine the categories used by indi-

viduals who design software which sorts through surveillance-generated information?” (587)

- “What kind of power relations are created under different alignments of the elements of surveillance?” (587)

The research questions I have stated in the previous chapter have been based on the above. Referring back to my earlier discussion of power/knowledge and truth, searching for answers to these questions also leads to the need to identify the specific forms that truth and knowledge take in surveillance domains.

In this chapter I have described the theoretical framework that I use in my investigation of ICT that spans different domains as well as multiple academic disciplines. This particular configuration of theories provides me with the tools to analyse ICT-based surveillance critically, and to both situate and question the truth claims that arise during their production. In the next chapter on methodology, I discuss the challenges of combining practices and findings from different disciplines while working towards this goal, and I describe my research process to contextualise the five papers.

Chapter 3

Methodology

The theoretical framework I have presented in the previous chapter is tightly connected to methodological issues. In this chapter, I discuss my methodological concerns with reference to my theoretical framework. I motivate why I chose to study design documents to identify social and political implications of ICT-based surveillance, and I discuss why discourse analysis is a suitable method for analysing design documents. Given my theoretical assumptions, deduction is not possible when analysing texts, as every statement of analysis acts on the analysed texts. Hence, what is analysed does not remain constant, it shifts as the analysis proceeds. For that reason, my use of discourse analysis is necessarily inductive. As I have mentioned in the first chapter, I consider my overall contributions to be located primarily in the last three papers, with the first two papers serving as the background that shaped my understanding of ICTs and their implications. That said, the first two papers constitute an important and necessary backdrop also for my more recent research. Hence, this chapter focuses on the methodological issues relating to my overall contributions as represented in the last three papers, although the discussion remains relevant for the first two papers as well.

3.1 Design Documents

In ICT design and development, texts can take on different roles, as research funding applications, project descriptions, system specifications, policy guidelines, user manuals, scientific publications, etc. In

these roles, they influence the shaping and distribution of material and other resources (Smith 2001). Texts describe what is worthy of observation and contain valuable pointers when trying to critically analyse the consequences of living with ICTs. In the projects I have analysed, the design process is characterised by the occurrence of many meetings, the production of a large number of reports, and development of deliverables including the technologies. I use the term design documents to refer to all the texts written during this lengthy process. These documents express a collective position of a large group of experts including system designers, project managers, developers, funding agencies, etc., who may have different or even conflicting goals individually. The things worthy of monitoring and analysing in data sources for syndromic surveillance, or in smart homes, are discussed, selected, and written down in this phase. Star (1999) argues:

In information infrastructure, every conceivable form of variation in practice, culture, and norm is inscribed at the deepest level of design. Some are malleable, changeable, and programmable—if you have the knowledge, time, and other resources to do so. Others—such as a fixed-choice category set—present barriers to users that may only be changed by a full-scale social movement. (389)

These documents contain many indicators of the designers' intentions, describing what is assumed to be true, and the relationships between knowledge production and the different subject positions. However, they may just as easily mislead if the goal is to understand how the actual implementation functions. For this reason, I focus on the discourses formed partially by these documents and their implications. During the design process, these documents are written to describe the systems fully, to ensure that they can be developed in the future. Many choices that become invisible later when the system is operational are clearly described in the design documents.

Differing from interviewing designers, another potential method, studying design documents makes it possible to describe the system as it is being constructed. As these documents are also addressed at other *experts*, they make it possible to capture the assumptions shared between different actors involved in the project. They are mostly written for an audience that is considered to be part of the design process,

and they tell stories of how ICTs come to be. Due to these assumptions, documents produced during an ICT design and development project allow interpretations of the expected implications of the technologies being designed. Bowker and Star (2000) state that “software is frozen organizational and policy discourse” (135). In design documents, the specifics of how software comes to freeze at particular points and not at others can be gleaned from the design documents. Graham and Wood (2003) emphasise the need to “get inside” the systems when studying them from a critical perspective, and add:

This might mean switching the focus of research to the social and the political assumptions that software producers embed (unconsciously or consciously) into their algorithms years before and thousands of miles away from the site of application. (242)

What a system is designed for and what it eventually comes to be used for may differ greatly, but design discourses engage in truth claims and position subjects, and the implications of these can still be analysed critically before the systems are constructed. Analysing ICTs at the design stage opens new possibilities for voicing concerns about their implications before they become operational. As Winner (1980) observes, following the implementation of socially demanding technologies, those who cannot accept the difficult requirements tend to be dismissed as “dreamers and fools” (38). Discussing design documents allows concerns about how these systems might affect their surroundings when they are operational to be raised before they are implemented.

3.2 Discourse Analysis

Referring back to the definition of discourse I have discussed in the previous chapter, “a regulated practice that accounts for a number of statements”, discourse analysis is an attempt to understand textual exchanges. It involves identifying what are represented as truths in texts, and how those truths are constructed. Within the analysis, justifications that the texts provide, the issues they problematise, as well as the issues they leave out, are all essential in understanding the construction of truths in discourses. Beginning with the question of what is

considered to be true in a given text, other related issues such as the positionings of different subjects and objects, or the possibilities made available or unavailable by them can be discussed.

Discourse analysis is a suitable method for locating the possible implications of ICT-based surveillance in design documents, because it provides the tools required to investigate the roles assigned to different groups in these texts, and how those roles relate to knowledge production. In other words, it is possible to identify particular configurations of power relations in texts using discourse analysis. These configurations can then be used to interpret which groups are given access to knowledge and the means to issue truth claims, and which groups are considered to be outside knowledge production. Since design documents contain not only technical details about how systems should be implemented but also motivations about why they should be implemented, their contents necessarily declare what is considered valuable and hence worth doing. Discourse analysis allows me to relate these value assignments to existing power relations, and to interpret their possible implications.

My theoretical framework is compatible with discourse analysis only if I position myself in relation to the discourses I analyse. I acknowledge that there is no neutral ground detached from all discourses from which I can provide my interpretations. Every subject position, including my own, is deeply embedded in different discourses simultaneously, and the act of analysis does not free the text from all discourses, but merely places it in a different one. As Mills (2004) explains:

[T]he process of finding a position for oneself within discourse is never fully achieved, but is rather one of constantly evaluating and considering one's position and, inevitably, constantly shifting one's perception of one's position and the wider discourse as a whole. (87)

While performing discourse analysis, I position myself as a reader who will first interpret a discourse, and following the analysis, participate in the production of another kind of discourse. While I identify implicit assumptions in the documents I analyse, I make implicit assumptions myself. Just like any interpretation, my work constructs a narrative based on the analysed texts, and my analyses construct and contribute to other discourses.

Limitations

As I mentioned in the previous chapter, the texts I have analysed have different authorship configurations: some are written by committees and have a large number of authors, others are credited to institutions. Additionally, they are often written by a majority of non-native English speakers. For these reasons, I decided that focusing on language features such as grammar and wording (which is necessary in a different method such as Critical Discourse Analysis, for example (Fairclough 1989)), would not be as productive for answering my research questions, and that I needed to interpret the texts with more flexibility, focusing less on factors such as syntax and word choice. In my analyses, my concern is not identifying the *actual* intent of the author(s), but engaging the implications of the objects being described. The intentions of the authors are necessarily discursive constructions themselves (hence the impossibility of declaring a singular, actual intent), and these texts circulate and affect the world around them independently of the intentions of their authors. Although the subject position of the author serves different and contested purposes (Foucault 1979), in my interpretations I have focused on the generated discourse. That is not to say that the author position in general is irrelevant to my analyses of implications. My point is simply that the intentions of the authors do not occupy an epistemologically privileged position as evidence in explaining why certain statements are made, but they do play their roles in making truth claims as elements of the discourses I analyse.

I have argued that discourse analysis is suitable for identifying the possible social and political implications of design, development, and use of ICT-based surveillance in design documents. However, a consequence of focusing on discourse and on texts to understand surveillance is that the experiences of surveillance subjects are not visible in my analyses. For example, the experience of being watched by a video camera differs from that of touching a fingerprint scanner, but providing a fingerprint to unlock a personal computer also differs from providing it for a criminal investigation even when the underlying ICT is identical. Additionally, the same surveillance system may have different effects on different embodied subjects depending on their positions in different power relations (Conrad 2009). Equally important, the effects of the same surveillance system may be *experienced* differently by

different subjects. As Ball argues, “we lose something analytically if we focus on just, for example, a discursive notion of the subject.” (Ball 2009, 652) Although my investigations concern discursive representations of surveillance, I acknowledge that subjects are not constituted solely within discourses. In my research, this translates to grounding my own claims in the specific discourses of specific design documents, and recognising that these claims are shaped by my partial perspective. Although knowledge and truth claims are discursive acts, they are influenced by the materiality of the contexts they occur in, and embodied experience remains an essential part of any context. Embodiment matters, because as humans we have access to a vast collection of subjective experiences of being surveilled, and insights derived from those experiences matter to studies of surveillance.

3.3 A Collection of Papers

In this section, I describe the five papers included in the thesis. I reflect on issues and insights which did not fit into the papers, and I provide additional details about where and how the papers were written, as well as a more detailed description of my method. As I have mentioned at the beginning of this chapter, my contributions are primarily located in the last three papers, and therefore my theoretical and methodological concerns have primarily dealt with those three. To highlight the role the first two papers have played in my research, I discuss them in further detail below.

In the last three papers, I followed a similar procedure for analysing each document: I began by reading through the text and marking all sentences that made knowledge or truth claims related to the overall argument, stated hypotheses, or presented research questions. I also marked sentences which defined subject positions, either as subjects acting in some capacity, or as being acted upon. I continued by marking sentences which informed the reader about the existence of entities outside of the text, such as external objects. In this category I included any statements that refer to uncertainty, because such statements also provide insights into what is considered to exist only partially outside the text. After the marking process was complete, I reviewed the marked sentences for each paragraph, section, and chapter, and tried to group them according to common themes based on either the top-

ics they discuss, or the frequency of the words they use (Ryan and Bernard 2003). I examined these themes in relation to claims about knowledge, truth, and subject positions. In my analysis, I repeatedly asked the question “what is considered true in this statement?”, and answered it by referring to the themes, the collection of marked sentences, as well as the unmarked sentences referenced by the marked ones.

Paper One

In September 2007, I was employed part-time by the Swedish Institute for Communicable Disease Control (Smittskyddsinstitutet, abbreviated SMI) to work as a developer for a software model simulating the spread of pandemic influenza. After about four months, my employment became full-time, and my tasks were split between programming the simulation model, and developing a system for computer-supported outbreak detection. The first paper describes the design and development of that system.

During this time, I worked with Paul Saretok in programming the first prototype system, and Anette Hulth led the project and communicated SMI’s requirements as well as the wishes of the epidemiologists who were to be the future users of the system. As the project progressed, we were also joined by Maria Grünwald who worked on the statistical detection methods used by the system, and Kenneth Hebing, who initially aided us in connecting the system to the existing databases at the institute, and took over all software development and maintenance duties when the project was complete. The system has since been developed further and remains in regular use at the institute.

The primary goal with developing a system to detect outbreaks automatically was to aid the epidemiologists in handling large sets of data. Our aim within the project was to create a tool that would be able to provide an up-to-date overview of existing data daily, and to automate the statistical detection of outbreaks. We hypothesised that if the system could accomplish these goals, the epidemiologists would have more time to commit to other tasks that are not possible to automate, such as the investigation of events that may or may not be classified as outbreaks, following up on earlier investigations,

and contacting local health authorities for additional information in ongoing investigations, to name a few.

My role throughout this project was to program the software. Although we already had access to SmiNet, the Swedish communicable disease case database (Rolfhamre et al. 2006), which included classifications of diseases and their subtypes, attempting to transfer that information into an easily accessible interface was a complex challenge. Many of the diseases had a large number of subtypes with different transmission vectors, which meant that they were impossible to aggregate. In the end, we allowed subtypes to inherit the settings from their parent diseases, but also allowed them to be configured independently, meaning that different statistical methods and different parameters could be used when performing detection for a particular subtype. During development we received regular feedback from the epidemiologists about how they performed their work, what indicators they looked at, and how much information they would want from the system. It quickly became clear that almost every disease required different parameters, and these depended on a wide variety of factors impossible to generalise. To support this style of detection, Maria Grünewald worked with epidemiologists individually to choose suitable statistical methods, and to come up with parameters that would allow the detection software to produce the output in line with the epidemiologists' wishes. After the initial software development stage was completed, we released the source code with a GPLv3 license (SMI 2013b). At this point, Kenneth Hebing took over the further development and maintenance of the software, and I started working on the publication describing our work. This period also coincided with my admission to the Royal Institute of Technology (KTH) to begin my studies for a licentiate degree.

Although I had already examined many similar systems used in outbreak detection while developing CASE, performing a literature review for the first paper also allowed me to systematically investigate the field of syndromic surveillance, where many different data sources were combined to form indicators for early outbreak detection. The research in the field promised earlier detection, lower costs, and highly complex software, all of which appealed to me at the time. Part of this work informed the first paper, but also provided the material that I used later in my licentiate thesis (Cakici 2011a). In 2012, researchers from SMI published an evaluation of how the epidemi-

ologists perceived the system (Kling, Grünewald, and Hulth 2012), which is also described in an SMI report (SMI 2012).

Paper Two

As mentioned earlier, I was initially employed by SMI to work as a developer for a software model simulating the spread of pandemic influenza. The second paper, co-authored by Magnus Boman, is based on our experiences while working with this model. I worked on the simulation model between the second half of 2007 and the end of 2010, initially part-time, and as part of my licentiate work after 2009. During this period, I developed a large set of tools to aid in the analysis of the data, worked on various parts of the code, and developed a module for testing vaccination strategies within the software model. In early 2010, after roughly two years of experience with the model, I started to question the methodological limits of testing vaccination strategies in simulation models. As my final task in the project, I prepared the source code for open sourcing using GPLv3 (SMI 2013a). At the end of the year, I asked for a leave of absence from SMI to finish my graduate degree and returned to KTH full-time.

The second paper was written during the transition period from SMI to KTH, and it was published in August 2011. It arose from extended discussions I had with Magnus Boman on the issues we faced simulating influenza in software. We converted our experience into a scientific paper that could inform others in similar situations. We framed it as a workflow, and provided a set of statements as “lessons learned”. Many of the lessons were translated directly from issues we both experienced during the two years which culminated in my departure from the project. During this time, my familiarity with concepts such as reflexivity and voice in the text were limited, and revisiting the text today I find the detached position we assumed in relation to our material to be a strong indicator of the type of discourse we participated in at the time.

Although I had been interested in surveillance systems and the consequences for their subjects ever since I applied for my first residence permit to study in Sweden, my experience in this project motivated me to take into account academic perspectives on surveillance in my own research.

Following the submission of the second paper, I started working on my licentiate thesis based on the research I had performed at SMI. I successfully defended it in June 2011. The thesis includes my research on syndromic surveillance, as well as the first two papers I have also included in this thesis (Cakici 2011a). After my licentiate defence, I was informed that my department at the time would not be financing my studies past licentiate level. Although I remained at KTH for six more months, I was not able to continue with my PhD work. During this time, I received funding from the European Institute of Innovation and Technology, and I worked on a report dealing with syndromic surveillance and digital cities (Cakici 2011b). This six-month period gave me the chance to re-evaluate my thoughts on surveillance, and to start familiarising myself with the field of STS.

Paper Three

In January 2012, I was admitted to Stockholm University as a PhD student, with 50% co-funding from the Swedish Institute of Computer Science (SICS). The third paper is based on the research I performed during my first year at SICS. The projects that financed my PhD were primarily concerned with sustainability, which meant that I had to leave public health as a domain of study, although I did eventually return to it for the fifth and final paper. During this period, I retained my surveillance focus, and gradually brought in more insights from the field of science and technology studies into my research.

In the third paper, I investigated how different projects linked sustainability and surveillance with the help of ICT, and how they reasoned about the future users of the proposed technologies in their design documents. In this publication, I switched roles: Instead of writing design documents and developing systems, I started analysing them. I found design documents to be open to investigation using discourse analysis, and my training as a computer scientist provided the background to interpret the more technical implications of the systems described in the documents. At the same time, I was constantly aware that due to my familiarity with the discipline I was likely to be blind to many of the assumptions in the texts, and that I risked leaving many problematic constructions unquestioned. As I detailed in the previous chapter, I believe that everyone who engages in discourse analysis has to deal with their own relation to text, and to be able to

proceed with the analysis one solution is to remain aware of the limitations posed by (un)familiarity during the process, and to receive all criticism from a position of partial perspective.

The paper includes a section describing the three ICT design and development projects that I analysed. Through my employment at SICS, I was involved in one of the projects, Stockholm Royal Seaport (SRS) – Smart ICT. During this time, I authored a report discussing social concerns related to ICTs for the project (Cakici 2012). The issues I raised in the report were the initial inspirations for the third paper.

Paper Four

In the fourth paper I collaborated with Markus Bylund. Together, we identified some common assumptions that appear regularly in texts related to research on ICT. We chose the Seventh Framework Programme (FP7) as it is the most significant source of research funding in Europe, and correspondingly, the texts used in their calls are highly influential in shaping European ICT research discourses.

This paper differed from papers three and five because it involved analyses of policy documents and project descriptions instead of design documents. Although the form and content of project descriptions resembled design documents closely, dealing with policy documents introduced new challenges. These described different areas of research that were being funded by FP7, but they had undergone many revisions. Additionally, the documents belonged to a large set of supporting policies which remained constant throughout FP7. One of the first tasks was to separate the documents that remained unchanged, and to identify the changes in the revised texts. The paper includes a methodology section providing additional details about their revision history and how we arrived at the relevant texts.

A much broader question I considered while working on this paper was: Where can the influence of discourses on material flows regarding scientific research be located? One candidate was the texts produced by institutions that finance research, since the process of writing research applications is essential to contemporary academic practice. I felt that since applying for funding plays such a large role in research, texts that describe how and in which fields to apply for funding would contain potentially problematic assumptions of their

own. Hence, my participation in the paper was also driven by the wish to understand the role of critical research, and where it could potentially make a difference in current practices in a field as broad as ICT research.

Paper Five

In the fifth paper, I worked together with Pedro Sanches. We began with a much broader focus on health-related ICT-based surveillance, and we gradually narrowed it down to syndromic surveillance for this paper. For my part, the choice was motivated by my previous experience with computer-supported outbreak detection, and my wish to revisit the role of ICTs in public health surveillance from a more critical perspective. We started from a review of disease surveillance systems I had included in my licentiate thesis, and we originally intended to conduct a study involving multiple systems. However, the amount of documentation available for the different systems was highly varied, and we eventually settled on examining only one, SIDARTHa, as the project publicly provided many documents which described different stages of the design, development, and testing process in detail.

Compared to CASE (the system described in paper one), the scope of SIDARTHa is much larger, involving multiple member states and many health institutions. However, the algorithms used in the software, and the *technical* problems described in the design documents mirrored many of my experiences while working on the development of CASE. As I began my studies in ICT design and development for public health surveillance, I felt that it was important for me to return to the field and to view my earlier work from a different perspective, and the fifth paper allowed me to do so.

Chapter 4

Conclusion

Surveillance is often an uncomfortable topic to think about, and the discomfort it causes highlights its importance. My aim was to identify how ICT and surveillance are linked in different discourses and to describe the implications of this connection. In this chapter, I provide a summary of the five papers, I state the answers to my research questions, and I conclude with some reflections.

4.1 Summary of Papers

The five papers contribute to different fields of research with overlapping interests, and as I have mentioned, I consider my overall contributions to be located primarily in the last three papers, with the first two papers constituting an important and necessary backdrop for my more recent research. As before, I use the first-person plural pronoun when summarising the papers with multiple authors.

In the first paper, “CASE: A Framework for Computer Supported Outbreak Detection”, we describe the design and development of a system for national outbreak detection at the Swedish Institute for Communicable Disease Control. The software package we provide draws data from a database of case reports, and processes the retrieved cases using statistical methods that can be selected by the epidemiologists that use the system. If an outbreak is detected, the system sends an automatically generated email to those included in the list of recipients for that disease. The system simplifies the process of connecting a case database to multiple detection methods, all which may require

different ways of formatting the data. The software package includes input generators and output parsers for handling the data, and it is available as open source software, licensed under GNU General Public License Version 3. As I noted in the previous chapter, researchers from SMI have performed an evaluation and found that the epidemiologists consider the system to be valuable for their daily surveillance work (Kling, Grünewald, and Hulth 2012).

In the second paper, “A Workflow for Software Development Within Computational Epidemiology”, we present a workflow for using simulation models to provide policy advice regarding communicable diseases. We ground our discussions on our experiences with a spatially explicit micro-meso-macro model for the entire Swedish population built on registry data, and we collect the lessons learned into a check list intended for use by computational epidemiologists and policy makers. We argue that it is challenging to use, verify, and validate regional and sensitive data sets, and we emphasise the importance of ensuring that the assumptions built into the model represent the wishes of the policy makers. The main result is arguably the depiction of the workflow itself.

In the third paper, “Sustainability Through Surveillance: ICT Discourses in Design Documents”, I examine design documents from three different ICT design and development projects. I argue that design documents for smart homes present intersecting visions of sustainability entailing the wide-spread use of ICT. In these visions, the technologies provide ways of judging the inhabitants who are described as behaving compatibly with such technologies. The inhabitants are made individually responsible for living sustainably, and surveillance is positioned as integral to this future with the help of ICT. In my analysis of the visions from the three projects, I also identify a translation process that captures the traces of the inhabitants’ lives and classifies them according to different criteria for sustainable living. These classifications are later returned to the tapestry of everyday life to convince the users to behave differently. I conclude that in the discourses of these documents, surveillance translates the traces, and the translations exert new pressures on existing power relations.

In the fourth paper, “Changing Behaviour to Save Energy: ICT-Based Surveillance for a Low-Carbon Economy in the Seventh Framework Programme”, we analyse Seventh Framework Programme policy documents within the Energy challenge of the ICT category pub-

lished by the European Commission, and descriptions of research projects granted funding from it, to highlight the uncritical development and application of surveillance technologies to change human behaviour. We identify a belief that human behaviour can be monitored at the individual level to generate different signals, and that these signals can be used to influence individuals to behave differently. We argue that EU-financed projects dealing with sustainability and ICT use models of social change that have been widely criticised as unlikely to lead to substantial changes in resource consumption. Additionally, we show that these texts discuss only the potential positive effects of technological surveillance, but neither acknowledge nor require the handling of the potential negative effects of surveillance.

In the fifth paper, “Detecting the Visible: The Discursive Construction of Health Threats in Syndromic Surveillance System Design”, we analyse reports describing the design, development, testing, and evaluation of a European Commission co-funded syndromic surveillance project called SIDARTHa. We show that the reports construct the concept of a health threat as a sudden, unexpected event with the potential to cause severe harm, and one that requires a public health response aided by surveillance. Based on our analysis, we state that when creating surveillance technologies, design choices have consequences for what can be seen, and for what remains invisible. Finally, we argue that syndromic surveillance discourse privileges expertise in developing, maintaining, and using software within public health practice, and that it prioritises standardised and transportable knowledge over local and context-dependent knowledge. We conclude that syndromic surveillance contributes to a shift in broader public health practice, with consequences for fairness if design choices and prioritisations remain invisible and unchallenged.

4.2 Answers

In this thesis, I posed the following questions: (i) What is the role of surveillance in the discourses of ICT design and development for sustainability and for public health? (ii) How are surveillance subjects positioned in relation to ICT-based surveillance in these two discourses? (iii) What are the implications of these positionings? I detail answers to each question with pointers to the relevant papers below.

I identified the role of surveillance in the discourses of ICT design and development for sustainability in paper three and paper four, and for public health in paper five. In these discourses, surveillance is a method of knowledge production where individuals are viewed as a source of information. ICT makes the collection and processing of information possible at a large scale, and different claims are made about the monitored world based on the gathered traces. When designing systems for smart homes, the traces of the inhabitants tracked by ICTs are essential to claims of sustainability, environmental friendliness, and low carbon emissions. These claims are made stronger by surveillance because they are linked to subjects outside the system. In public health, surveillance is necessary for understanding the current health state of the population. In paper one, case databases containing records of individuals who have contracted communicable diseases are essential to the outbreak detection. Surveillance is used initially to produce knowledge about the population, and the outbreak detection system contributes to the produced knowledge with statistical detection methods and ways of linking different data sources. In paper five, surveillance is not only necessary for understanding, it is essential to constructing the notion of a health threat as a sudden, unexpected event with the potential to cause severe harm.

As mentioned above, surveillance subjects are positioned in relation to ICT-based surveillance as sources of information. Additionally, in the discourses I analyse in paper three and paper four, they are made individually responsible for changing their behaviour based on the information they receive, and living sustainably using the ICTs designed and developed for that purpose. This also entails a kind of expertise that the inhabitants must acquire to operate the technologies. The systems described in the design documents make the differences in subjects commensurable with one another, aiding the creation of new categories such as the smart home inhabitant and the sustainable city district. The subjects are sources of information, but within the discourse, the classifications used by the systems to make truth claims remain invisible to the subjects. They are not given the agency to contest the produced knowledge.

One implication of these positionings is that ICT-based surveillance appears as value-neutral, seemingly representing the world as it is, instead of actively constructing it. This position is also instrumental in reducing social and political concerns about knowledge produc-

tion to issues of efficiency and productivity. It prioritises competence in building ICTs, as well as technological advancement and economic growth aided by surveillance. In syndromic surveillance discourses, such expertise is prioritised within public health practice, and the needs of those who will use the systems are made secondary to constructing standardised and transportable knowledge. Furthermore, subjects under surveillance are represented as passive sources and receivers of information, and as paper four argues, representing them as such is unlikely to lead to the substantial changes that the designers aim at. Finally, in the domain of sustainability, ICTs are assumed to be valuable for society as a whole simply because their domain is sustainability, but on the contrary, social inequalities and discriminatory biases can also be sustained by these technologies (Alaimo 2012), necessitating critical investigations of ICT-based surveillance.

The high technical complexity of the ICTs represented in the texts suggest proportionally high costs of development. The implications of these costs are essential to understanding the effects of developing the systems, especially if the motivations include sustainability or lowering carbon emissions. As with many other infrastructural projects, it is possible that the benefits created by these systems remain local to where the systems are installed while their harmful effects are pushed away, e.g., electricity consumption is decreased in the new city district while natural resources from another region far from the new city district are exhausted in creating the computer screens that visualise electricity consumption.

To build ICTs is to encode a particular way of seeing the world into software and hardware. As Bowker and Star (2000) note, in classification “[p]olitically and socially charged agendas are often first presented as purely technical and they are difficult even to see.” (196) The categories that these systems are designed to construct and uphold come at the expense of other categories and alternative classifications that are deemed unrelated to the problem at hand by their designers. Excluding and sorting may very well be necessary for the systems to function, but the decisions about what the system should observe must be made with the constant awareness of the possibilities that are not followed when the decision is made, as those that would have benefited from the paths not followed may consider the resulting systems to be severely inadequate for their own needs.

4.3 Conclusion

I set out to locate the influence of discourses on material flows for scientific research, and my investigation led me to design and policy documents. I aimed my critique at the tacit assumption that the development and usage of ICTs are always beneficial to society. I identified this assumption as part of a universalising discourse, one that does not place as much value in local contexts and local differences, instead aspiring to create a smooth, homogeneous body of knowledge in whatever domain it is applied to.

As I state in the third paper, sometimes ICT might be the wrong answer. ICTs can be oppressive, even when that is not the designers' intention, because they embody complex power relations expressed in software and hardware. Categories embedded into ICT can be understood as attempts to control and discipline those who use the systems, although as with any other expression of power, they can be contested and resisted as well. On the other hand, they can also empower subjects by encouraging cooperation, granting a sense of safety, and allowing companionship to prosper over great distances. The productive potential of ICT-based surveillance is always present, but it is fraught with danger. Foucault asserts:

My point is not that everything is bad, but that everything is dangerous, which is not exactly the same as bad. If everything is dangerous, then we always have something to do. (Foucault 1982, 343)

ICTs are not disconnected from contexts, and they are never without values. Hence, it is not possible for them to be neutral in any sense of the word. As I have demonstrated, the discursive move to portray them as value-neutral is highly problematic. It is not that the texts and their authors occupy this position to actively avoid being tied to a particular perspective, but that they are unaware of occupying any perspective, as if viewing the world from nowhere. As wide-spread as this view may be, there are other ways. Bowker and Star (2000) argue that “a key for the future is to produce flexible classifications whose users are aware of their political and organizational dimensions and which explicitly retain traces of their construction” (326). ICTs need not be judged only on their cost-efficiency, or their feasibility.

As social and political constructs, it is possible to reason about them differently, not simply as tools for solving problems, but as historically and culturally specific methods for organising knowledge production in societies. Adopting these perspectives, it becomes possible to reach beyond a deterministic view of technological progress, and reason about alternatives. Haraway explains:

So, with many feminists, I want to argue for a doctrine and practice of objectivity that privileges contestation, deconstruction, passionate construction, webbed connections, and hope for transformation of systems of knowledge and ways of seeing. (Haraway 1990, 191–192)

Different configurations of knowledge production are always possible, and in guiding alternative ways of knowing, the essential question remains unchanged: Who benefits from ICT-based surveillance, and who suffers its costs?

References

- Alaimo, Stacy. 2012. "Sustainable This, Sustainable That: New Materialisms, Posthumanism, and Unknown Futures." *PMLA* 127 (3): 558–564.
- Albrechtslund, Anders, and Lynsey Dubbeld. 2002. "The Plays and Arts of Surveillance: Studying Surveillance as Entertainment." *Surveillance & Society* 3 (2/3).
- Anderson, Roy M., and Robert M. May. 1991. *Infectious Diseases of Humans: Dynamics and Control*. Oxford University Press. ISBN: 978-0198540403.
- Ball, Kirstie. 2002. "Elements of Surveillance: A New Framework and Future Directions." *Information, Communication & Society* 5 (4): 573–590.
- . 2005. "Organization, Surveillance and the Body: Towards a Politics of Resistance." *Organization* 12 (1): 89–108.
- . 2009. "Exposure." *Information, Communication & Society* 12 (5): 639–657.
- Ball, Kirstie, and Kevin D. Haggerty. 2005. "Editorial: Doing Surveillance Studies." *Surveillance & Society* 3 (2/3): 129–138.
- Bijker, Wiebe E., and John Law, eds. 1992. *Shaping Technology / Building Society: Studies in Sociotechnical Change*. The MIT Press. ISBN: 0262521946.
- Bowker, Geoffrey C., and Susan Leigh Star. 2000. *Sorting Things Out: Classification and Its Consequences*. The MIT Press. ISBN: 0262522950.

- Bozbeyoğlu, Alanur Çavlin. 2011. "Citizenship Rights in a Surveillance Society: The Case of the Electronic ID Card in Turkey." *Surveillance & Society* 9 (1/2): 64–79.
- Bradley, Karin. 2009. "Just Environments: Politicising Sustainable Urban Development." PhD diss., KTH, School of Architecture and the Built Environment (ABE). <http://www.diva-portal.org/smash/record.jsf?pid=diva2:209293>.
- Bravata, Dena M., Kathryn M. McDonald, Wendy M. Smith, Chara Rydzak, Herbert Szeto, David L. Buckeridge, Corinna Haberland, and Douglas K. Owens. 2004. "Systematic Review: Surveillance Systems for Early Detection of Bioterrorism-Related Diseases." *Annals of Internal Medicine* 140 (11): 910–922.
- Buckeridge, David L. 2007. "Outbreak Detection Through Automated Surveillance: A Review of the Determinants of Detection." *Journal of Biomedical Informatics* 40 (4): 370–379.
- Buckeridge, David L., Howard Burkom, Murray Campbell, William R. Hogan, and Andrew W. Moore. 2005. "Algorithms for Rapid Outbreak Detection: A Research Synthesis." *Journal of Biomedical Informatics* 38 (2): 99–113.
- Buckeridge, David L., Anna Okhmatovskaia, Samson Tu, Martin O'Connor, Csongor Nyulas, and Mark A. Musen. 2008. "Understanding Detection Performance in Public Health Surveillance: Modeling Aberrancy-Detection Algorithms." *Journal of the American Medical Informatics Association* 15 (6): 760–769.
- Buehler, J. W., R. S. Hopkins, J. M. Overhage, D. M. Sosin, and V. Tong. 2004. "Framework for Evaluating Public Health Surveillance Systems for Early Detection of Outbreaks." *MMWR Morbidity and Mortality Weekly Report* 53:1–11.
- Buehler, J. W., A. Sonricker, M. Paladini, P. Soper, and F. Mostashari. 2008. "Syndromic Surveillance Practice in the United States: Findings from a Survey of State, Territorial, and Selected Local Health Departments." *Advances in Disease Surveillance* 6 (3): 1–20.

- Buehler, J. W., E. A. Whitney, D. Smith, M. J. Prietula, S. H. Stanton, and A. P. Isakov. 2009. "Situational Uses of Syndromic Surveillance." *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science* 7 (2): 165–177.
- Burkom, Howard S., Y. Elbert, A. Feldman, and J. Lin. 2004. "Role of Data Aggregation in Biosurveillance Detection Strategies with Applications from ESSENCE." *MMWR Morbidity and Mortality Weekly Report* 53:67–73.
- Burrows, Roger, and Nicholas Gane. 2006. "Geodemographics, Software and Class." *Sociology* 40 (5): 793–812.
- Cakici, Baki. 2011a. "Disease Surveillance Systems." Licentiate thesis, KTH. Accessed May 13, 2013. <http://kth.diva-portal.org/smash/record.jsf?pid=diva2:416962>.
- . 2011b. *Innovation Radar: Public Health in Digital Cities*. Technical report TR2011-003. 978-91-87253-04-1. EIT ICT Labs.
- . 2012. *Social Concerns in Information Systems for Sustainability*. Technical report. Swedish ICT. Accessed September 4, 2013. <http://smartict.swedish-ict.se/files/2012/05/Social-Concerns-in-Information-Systems-for-Sustainability-2012-05-14.pdf>.
- Conrad, Kathryn. 2009. "Surveillance, Gender, and the Virtual Body in the Information Age." *Surveillance & Society* 6 (4): 380–387.
- Cooper, D. L., N. Q. Verlander, G. E. Smith, A. Charlett, E. Gerard, L. Willocks, and S. O'Brien. 2006. "Can Syndromic Surveillance Data Detect Local Outbreaks of Communicable Disease? A Model Using a Historical Cryptosporidiosis Outbreak." *Epidemiology and Infection* 134 (1): 13–20.
- Das, D., K. Metzger, R. Heffernan, S. Balter, D. Weiss, and F. Mostashari. 2005. "Monitoring Over-the-Counter Medication Sales for Early Detection of Disease Outbreaks—New York City." *MMWR Morbidity and Mortality Weekly Report* 54 (Suppl): 41–46.

- Dato, Virginia, Michael M. Wagner, and Abi Fapohunda. 2004. "How Outbreaks of Infectious Disease are Detected: A Review of Surveillance Systems and Outbreaks." *Public Health Reports* 119 (5): 464–471.
- Dourish, Paul, and Scott D. Mainwaring. 2012. "Ubicomp's Colonial Impulse." In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, 133–142.
- Edwards, Paul N. 2003. "Infrastructure and Modernity: Force, Time, and Social Organization in the History of Sociotechnical Systems." In *Modernity and Technology*, edited by Thomas J. Misa, Philip Brey, and Andrew Feenberg, 185–225. The MIT Press. ISBN: 9780262633109.
- Eubank, Stephen, Hasan Guclu, V. S. Anil Kumar, Madhav V. Marathe, Aravind Srinivasan, Zoltán Toroczkai, and Nan Wang. 2004. "Modelling Disease Outbreaks in Realistic Urban Social Networks." *Nature* 429 (6988): 180–184.
- Fairclough, Norman. 1989. *Language and Power*. Reprint, 2001. Pearson Education. ISBN: 9780582414839.
- Ferguson, Neil M., Derek A. T. Cummings, Simon Cauchemez, Christophe Fraser, Steven Riley, Aronrag Meeyai, Sopon Iamsirithaworn, and Donald S. Burke. 2005. "Strategies for Containing an Emerging Influenza Pandemic in Southeast Asia." *Nature* 437 (7056): 209–214.
- Finn, Rachel L. 2011. "Surveillant Staring: Race and the Everyday Surveillance of South Asian Women After 9/11." *Surveillance & Society* 8 (4): 413–426.
- Forsythe, Diana. 2001. *Studying Those Who Study Us: An Anthropologist in the World of Artificial Intelligence*. Stanford University Press. ISBN: 9780804742030.
- Foucault, Michel. 1972. *The Archaeology of Knowledge*. Reprint, 1982. Translated by Alan Sheridan. Vintage. ISBN: 0394711068.
- . 1978. *The Will to Knowledge: The History of Sexuality Volume 1*. New Ed, 1998. Translated by Robert Hurley. Penguin. ISBN: 0140268685.

- . 1979. "What is an Author?" In *The Foucault Reader*, Vintage Books Edition, 2010, edited by Paul Rabinow, translated by Josué V. Harari, 101–120. Vintage.
- . 1982. "On the Genealogy of Ethics: An Overview of Work in Progress." In *The Foucault Reader*, Vintage Books Edition, 2010, edited by Paul Rabinow, 340–372. ISBN: 0394529049.
- French, Martin. 2009. "Picturing Public Health Surveillance: Tracing the Material Dimensions of Information in Ontario's Public Health System." PhD diss., Queen's University. Accessed July 16, 2013. <http://hdl.handle.net/1974/1689>.
- Gilliom, John. 2001. *Overseers of the Poor: Surveillance, Resistance, and the Limits of Privacy*. University Of Chicago Press. ISBN: 0226293610.
- Goffman, Alice. 2009. "On the Run: Wanted Men in a Philadelphia Ghetto." *American Sociological Review* 74 (3): 339–357.
- Graham, Stephen. 2005. "Software-Sorted Geographies." *Progress in Human Geography* 29 (5): 562–580.
- Graham, Stephen, and David Wood. 2003. "Digitizing Surveillance: Categorization, Space, Inequality." *Critical Social Policy* 23 (2): 227–248.
- Haraway, Donna J. 1990. "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective." In *Simians, Cyborgs, and Women: The Reinvention of Nature*, 183–201. Routledge. ISBN: 0415903874.
- . 1996. "Modest Witness: Feminist Diffractions in Science Studies." In *The Disunity of Science: Boundaries, Contexts, and Power*, edited by Peter Galison and David Stump, 428–441. Stanford University Press.
- Harding, Sandra. 1986. *The Science Question in Feminism*. Cornell University Press. ISBN: 0801493633.
- . 1995. "'Strong Objectivity': A Response to the New Objectivity Question." *Synthese* 104 (3): 331–349.

- Harding, Sandra, ed. 2004. *The Feminist Standpoint Theory Reader: Intellectual and Political Controversies*. New York: Routledge. ISBN: 9780415945004.
- Henning, Kelly J. 2004. "What is Syndromic Surveillance?" *MMWR Morbidity and Mortality Weekly Report* 53 Suppl:7-11.
- Hersh, William. 2009. "A Stimulus to Define Informatics and Health Information Technology." *BMC Medical Informatics and Decision Making* 9 (1): 24.
- Hulth, A., N. Andrews, S. Ethelberg, J. Dreesman, D. Faensen, W. van Pelt, and J. Schnitzler. 2010. "Practical Usage of Computer-Supported Outbreak Detection in Five European Countries." *Eurosurveillance* 15 (36).
- Jorges, Bernward. 1999. "Do Politics Have Artefacts?" *Social Studies of Science* 29 (3): 411-431.
- Johnson, Jim. 1988. "Mixing Humans and Nonhumans Together: The Sociology of a Door-Closer." *Social Problems*:298-310.
- Jorgensen, Marianne W., and Louise J. Phillips. 2002. *Discourse Analysis as Theory and Method*. SAGE Publications Ltd. ISBN: 0761971122.
- Kleinman, Ken P., and Allyson M. Abrams. 2006. "Assessing Surveillance Using Sensitivity, Specificity and Timeliness." *Statistical Methods in Medical Research* 15 (5): 445-464.
- Kling, A. M., M. Grünewald, and A. Hulth. 2012. "Two Years of Computer Supported Outbreak Detection in Sweden: The User's Perspective." *Journal of Health & Medical Informatics* 3 (108): 2.
- Koskela, Hille. 2002. "Webcams, TV Shows and Mobile Phones: Empowering Exhibitionism." *Surveillance & Society* 2 (2/3).
- Krieger, Nancy. 1992. "The Making of Public Health Data: Paradigms, Politics, and Policy." *Journal of Public Health Policy*:412-427.
- Latour, Bruno. 1988. *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press. ISBN: 06747-92912.
- Latour, Bruno, and Steve Woolgar. 1979. *Laboratory Life: The Construction of Scientific Facts*. Sage Publications. ISBN: 0-8039-0993-4.

- Lee, Francis. 2009. "Letters & Bytes: Sociotechnical Studies of Distance Education." PhD diss., Linköping University, Technology and Social Change, Faculty of Arts and Sciences.
- Lombardo, Joseph S., and David Ross. 2007. "Disease Surveillance, a Public Health Priority." In *Disease Surveillance: A Public Health Informatics Approach*, edited by Joseph S. Lombardo and David L. Buckeridge, 1–40. Wiley-Interscience. ISBN: 0470068124.
- Longino, Helen. 1993. "Subjects, Power, and Knowledge: Description and Prescription in Feminist Philosophies of Science." In *Feminist Epistemologies*, edited by Linda Alcoff and Elizabeth Potter, 101–120. New York: Routledge. ISBN: 9780415904513.
- Lyon, David. 2007. *Surveillance Studies: An Overview*. Polity. ISBN: 074563592X.
- MacKenzie, Donald, and Judy Wajcman. 1999. *The Social Shaping of Technology*. Open University Press.
- Maki, Krystle. 2011. "Neoliberal Deviants and Surveillance: Welfare Recipients Under the Watchful Eye of Ontario Works." *Surveillance & Society* 9 (1/2): 47–63.
- Mansell, Robin. 2002. *Inside the Communication Revolution: Evolving Patterns of Social and Technical Interaction*. Oxford University Press. ISBN: 9780198296553.
- Mills, Sara. 2004. *Discourse*. Routledge. ISBN: 0415290139.
- Morse, Stephen S. 2007. "Global Infectious Disease Surveillance and Health Intelligence." *Health Affairs* 26 (4): 1069–1077.
- Mostashari, Farzad, and Jessica Hartman. 2003. "Syndromic Surveillance: A Local Perspective." *Journal of Urban Health* 80 (1): i1–i7. doi:10.1093/jurban/jtg042.
- Murakami Wood, David, Stephen Graham, Mimi Sheller, and John Urry. 2006. "Permeable Boundaries in the Software-sorted Society: Surveillance and Differentiations of Mobility." In *Mobile Technologies of the City*, 177–191. New York: Routledge. ISBN: 978-0415655606.

- Ottinger, Gwen. 2010. "Constructing Empowerment through Interpretations of Environmental Surveillance Data." *Surveillance & Society* 8 (2): 221–234.
- Oudshoorn, Nelly, and Trevor Pinch, eds. 2005. *How Users Matter: The Co-Construction of Users and Technology*. The MIT Press. ISBN: 0262651092.
- Pargman, Daniel, and Jacob Palme. 2009. "ASCII Imperialism." In *Standards and Their Stories: How Quantifying, Classifying, and Formalizing Practices Shape Everyday Life*, edited by Martha Lampland and Susan Leigh Star, 177–199. Cornell University Press. ISBN: 978-0801474613.
- Pelecanos, Anita, Peter Ryan, and Michelle Gatton. 2010. "Outbreak Detection Algorithms for Seasonal Disease Data: A Case Study Using Ross River Virus Disease." *BMC Medical Informatics and Decision Making* 10 (1): 74.
- Rolfhamre, P, A Jansson, M Arneborn, and K Ekdahl. 2006. "SmiNet-2: Description of an Internet-Based Surveillance System for Communicable Diseases in Sweden." *Euro Surveillance: European communicable disease bulletin* 11 (5): 103–107.
- Rooney, Tonya. 2010. "Trusting Children: How Do Surveillance Technologies Alter a Child's Experience of Trust, Risk and Responsibility?" *Surveillance & Society* 7 (3/4): 344–355.
- Rowley, Jennifer. 2007. "The Wisdom Hierarchy: Representations of the DIKW Hierarchy." *Journal of Information Science* 33 (2): 163–180.
- Ryan, Gery W, and H. Russell Bernard. 2003. "Techniques to Identify Themes." *Field Methods* 15 (1): 85–109.
- Said, Edward. 1978. *Orientalism*. Reprint, 2003. Penguin Books Limited (UK). ISBN: 9780141187426.
- Shannon, Claude Elwood, and Warren Weaver. 1948. *A Mathematical Theory of Communication*. American Telephone / Telegraph Company.
- Shilton, Katie. 2010. "Participatory Sensing: Building Empowering Surveillance." *Surveillance & Society* 8 (2): 131–150.

- Shove, Elizabeth. 2004. "Changing Human Behaviour and Lifestyle: A Challenge for Sustainable Consumption?" In *The Ecological Economics of Consumption*, edited by Lucia A. Reisch and Inge Røpke, 111–132. Cheltenham, UK; Northampton, MA: E. Elgar. ISBN: 9781843765127.
- SMI. 2013a. *MicroSim*. Smittskyddsinsitutet. Accessed May 13. <https://smisvn.smi.se/sim>.
- . 2013b. *The CASE Project*. Smittskyddsinsitutet. Accessed May 13. <https://smisvn.smi.se/case>.
- . 2012. *Datorstödd Utbrottsdetektion på Smittskyddsinstitutet: Datorprogrammet CASE – Erfarenheter och Användarnytta*. Technical report 2012-101-1. Smittskyddsinstitutet.
- Smith, Dorothy E. 2001. "Texts and the Ontology of Organizations and Institutions." *Studies in Cultures, Organizations and Societies* 7 (2): 159–198.
- Spivak, Gayatri Chakravorty. 1988. "Can the Subaltern Speak?" *Marxism and the Interpretation of Culture* 58 (1): 271–313.
- Star, Susan Leigh. 1999. "The Ethnography of Infrastructure." *American Behavioral Scientist* 43 (3): 377–391.
- Suchman, Lucy. 1994a. "Do Categories Have Politics?" *Computer Supported Cooperative Work (CSCW)* 2 (3): 177–190.
- . 1994b. "Working Relations of Technology Production and Use." *Computer Supported Cooperative Work (CSCW)* 2 (1-2): 21–39.
- . 2002. "Practice-Based Design of Information Systems: Notes from the Hyperdeveloped World." *The Information Society* 18 (2): 139–144.
- UN. 1987. *Report of the World Commission on Environment and Development: Our Common Future*. Technical report. Accessed May 15, 2013. <http://www.un-documents.net/wced-ocf.htm>.
- Wennerås, C., and A. Wold. 1997. "Nepotism and Sexism in Peer-Review." *Nature* 387 (6631): 341–343.

- Wihlborg, Elin. 2000. "En Lösning som Söker Problem: Hur och Varför Lokala IT-policyer Utvecklas i Landsbygdskommuner [A Solution in Search for Problems: How and Why Local ICT-policies are Developed in Rural Municipalities]." PhD diss., Linköping University, Technology and Social Change, Faculty of Arts and Sciences.
- Winner, Langdon. 1980. "Do Artifacts Have Politics?" *Daedalus* 109 (1): 121–136.
- Woolgar, Steve. 1991. "Configuring the User: The Case of Usability Trials." In *A Sociology of Monsters: Essays on Power, Technology and Domination*, edited by John Law, 58–99. Routledge. ISBN: 978-0415071390.
- Woolgar, Steve, and Geoff Cooper. 1999. "Do Artefacts Have Ambivalence? Moses' Bridges, Winner's Bridges and Other Urban Legends in S&TS." *Social Studies of Science* 29 (3): 433–449.

PART II

Chapter 5

CASE: A Framework for Computer Supported Outbreak Detection*

Baki Cakici^{†,‡}, Kenneth Hebing[‡], Maria Grunewald[‡], Paul Saretok[‡], and Anette Hulth[‡]

Abstract

In *computer supported outbreak detection*, a statistical method is applied to a collection of cases to detect any excess cases for a particular disease. Whether a detected aberration is a true outbreak is decided by a human expert. We present a technical framework designed and implemented at the Swedish Institute for Infectious Disease Control for computer supported outbreak detection, where a database of case reports for a large number of infectious diseases can be processed using one or more statistical methods selected by the user. Based on case information, such as diagnosis and date, different statistical algorithms for detecting outbreaks can be applied, both on the disease level and the subtype level. The parameter settings for

*Cakici, Baki, Kenneth Hebing, Maria Grunewald, Paul Saretok, and Anette Hulth. 2010. "CASE: A Framework for Computer Supported Outbreak Detection." *BMC Medical Informatics and Decision Making* 10 (1): 14.

[†]Royal Institute of Technology (KTH), 164 40 Kista, Sweden

[‡]Swedish Institute for Communicable Disease Control (SMI), 171 82 Solna, Sweden

the algorithms can be configured independently for different diagnoses using the provided graphical interface. Input generators and output parsers are also provided for all supported algorithms. If an outbreak signal is detected, an email notification is sent to the persons listed as receivers for that particular disease. The framework is available as open source software, licensed under GNU General Public License Version 3. By making the code open source, we wish to encourage others to contribute to the future development of computer supported outbreak detection systems, and in particular to the development of the CASE framework.

5.1 Background

In this paper, we describe the design and implementation of a *computer supported outbreak detection system* called CASE (named after the protagonist of the William Gibson novel *Neuromancer*), or Computer Assisted Search for Epidemics. The system is currently in use at the Swedish Institute for Infectious Disease Control (SMI) and performs daily surveillance using data obtained from *SmiNet* (Rolfhamre et al. 2006), the national notifiable disease database in Sweden.

Computer supported outbreak detection is performed in two steps:

- 1 A statistical method is automatically applied to a collection of case reports in order to detect an unusual or unexpected number of cases for a particular disease.
- 2 An investigation by a human expert (an epidemiologist) is performed to determine whether the detected irregularity denotes an actual outbreak.

The main function of a computer supported outbreak detection system is to warn for *potential* outbreaks. In some cases, the system might be able to detect outbreaks earlier than human experts. Additionally, it might detect certain outbreaks that human experts would have overlooked. However, the system does not aim to replace human experts (hence the prefix “computer supported”); it should rather be considered a complement to daily surveillance activities. To a smaller extent, the system can also aid less experienced epidemiologists in identifying outbreaks.

Systems for outbreak detection which support multiple algorithms include RODS (Tsui et al. 2003), BioSTORM (Crubezy et al. 2005) and AEGIS (Reis et al. 2007). Additionally, computer supported outbreak detection systems operating on the national level have been used previously in a number of countries, including Germany (Krause et al. 2007) and the Netherlands (Widdowson et al. 2003).

Health Care in Sweden

The health care system in Sweden is governed by 21 county councils. Each county has appointed a medical officer, who is in charge of the regional infectious disease prevention and control. Every confirmed or suspected case of a notifiable disease is reported both to the county medical officer and to SMI. At SMI, the regular national surveillance is currently performed by thirteen epidemiologists, each in charge of a number of different diseases.

All 21 county medical officers as well as the majority of the hospitals and the laboratories in Sweden are connected to the SmiNet database. The database collects clinical reports and information on laboratory verified samples. In 2008, a total of 174811 reports were submitted to SmiNet. 87 per cent of these reports were submitted electronically and those that were not submitted electronically were entered into SmiNet manually. Of the 92744 lab reports, as much as 97 per cent were submitted electronically and 62 per cent fully automatically. The reports were subsequently merged into 74367 case reports. These reports form the basis of the data used by CASE to perform outbreak detection.

5.2 Implementation

CASE is designed to be administered using a graphical interface, and can operate on all of the 63 notifiable diseases in Sweden. One or more statistical detection methods can be applied to each disease. If more than one method is activated, result reports are generated independently. By default, the data are aggregated over all disease subtypes, but the system allows detection of single subtypes as well. When an outbreak signal is generated, an alert is sent by email to all members of the notification list for that particular disease.

CASE is composed of three interconnected components for *configuration*, *extraction* and *detection*. The configuration component provides a graphical user interface for modifying detection parameters and editing the list of recipients for generated alerts. The extraction component is used to copy data from the national case database to the local database. The detection component is scheduled to run at regular intervals and automatically applies the chosen statistical methods to the currently selected diseases.

System Description

CASE is developed using Java to ensure platform-independence of all components. Currently at SMI all three components run on Ubuntu, a Linux-based operating system. The local database for CASE is MySQL and the national database, SmiNet, is Microsoft SQL Server 2005.

Figure 5.1 shows the flow of information within the framework. The extraction and detection components are scheduled to run once every 24 hours at midnight using the standard Unix scheduling service *cron*. When the extraction component is executed, it transfers data from SmiNet to the local database. The local database stores the case data and the configuration parameters for all algorithms. The configuration module can be used to view and modify the parameters. The detection component is executed automatically after all required data have been extracted from SmiNet. It applies the detection methods with the given parameters to the case data for the selected diseases, and emails notifications if any alerts are generated. Detailed logs of these processes are generated automatically.

Configuration

The configuration component is a graphical user interface that allows the administrator to mark diseases for detection, choose the detection methods to be applied to each diagnosis/subtype and manage the list of epidemiologists that will receive alerts in case a warning is generated. The settings are stored in a local database that is also accessed by the other two components. The system can be administered by multiple users who access the same local database.

Figure 5.2 shows a screenshot of the graphical user interface for the CASE administrator. The notifiable diseases are displayed in the




Figure 5.1: A flowchart demonstrating the detection process in CASE.

left column. These entries can be expanded using the arrow to display their subtypes. Parameters for the current selection are shown on the right hand side. The *Algorithms* tab lists the available methods. Parameters for the selected method can be modified by double-clicking the name of the method. The *E-mail* tab contains a list of recipients for the selected disease and/or subtype. If an alert is generated after detection, the algorithm that generated the alert is highlighted in red. The flag is automatically cleared every night before a new detection batch is executed.

Extraction

CASE uses data retrieved from SmiNet to perform outbreak detection. A case report is created in SmiNet when a clinical or a labo-

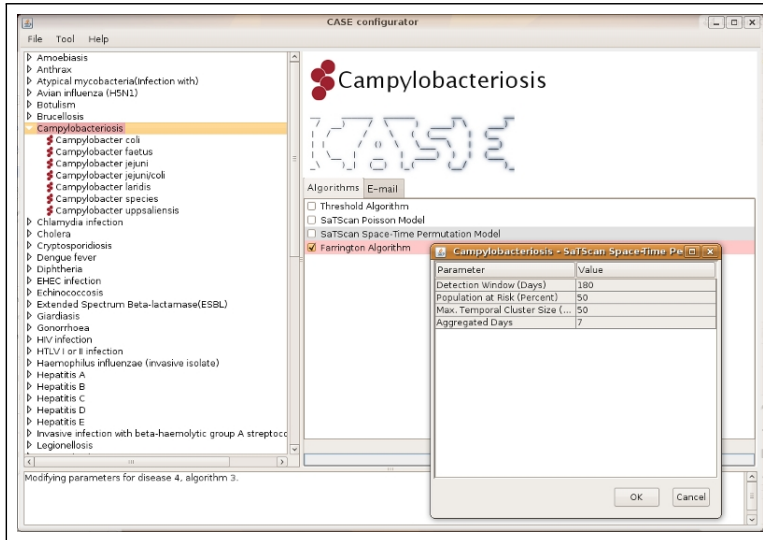


Figure 5.2: A screenshot of the graphical user interface for the CASE administrator.

ratory report is received, provided that this patient does not already exist in the database. When additional reports arrive, the original case report is automatically updated with the new information. Depending on the number of days that have elapsed since the last time a patient received a particular diagnosis, a new case report might be created for the same diagnosis and patient. For a detailed technical description of SmiNet, see Rolfhamre et al. 2006.

The extraction component populates the local database with data from the case reports stored in SmiNet. Diagnosis, lab species, date, and reporting county are copied for every case, except those with infections that are reported to have originated abroad. No information that can reveal a patient's identity is used in the outbreak detection process. There are approximately twenty dates in SmiNet for each case report, ranging from dates that are automatically generated by the system to dates entered by the clinician or the laboratory. There is, however, only one date that is available on all case reports, namely *statistics date*. This automatically set date corresponds to when a patient first appears in SmiNet with a particular diagnosis. The date that would best reflect when a patient fell ill is the date when the sample

was taken from the patient. However, many case reports do not contain this date. For example, for 2008 this date is missing in 29 per cent of the case reports. When the case information is copied from SmiNet to the local database, the extraction component fetches the statistics date as the date for the case.

Detection

CASE is developed by the Swedish Institute for Infectious Disease Control, and has a national perspective on outbreaks. Its primary role is to find outbreaks that cover more than one county, especially those with few cases in each affected county, as these might be difficult to detect for the local authorities.

The detection component uses the selected statistical method(s) on all activated diseases and sends notification emails if any alerts are raised. If there are too few data points for a detection algorithm to produce a result — which is often the case for detection on the subtype level — this information is written to the log file. The system currently supports four different statistical methods for detection: SaTScan Poisson (Kulldorff 1997), SaTScan Space-Time Permutation (Kulldorff et al. 2005), an algorithm developed by Farrington et al. (1996), and a simple threshold algorithm. The methods are briefly described below. Three of the four methods are freely available implementations, while the fourth was developed within the project and is included in CASE's source code. For the external programs, input generators and output parsers are also contained within the source code. It is possible to extend the system with additional statistical methods, although this requires a certain familiarity with the Java programming language. We are currently in the process of adding the *OutbreakP* method (Frisén, Andersson, and Schiöler 2009) to the core package.

SaTScan is a freely available spatial, temporal and space-time data analysis platform (SaTScan). Two algorithms from this application are used in CASE: *SaTScan Poisson* which uses the discrete Poisson SaTScan model to search for spatial clusters and *SaTScan Space-Time Permutation*, which searches for spatio-temporal clusters. Both models are applied to data at the county-level resolution. The population data required by SaTScan Poisson are obtained from Statistics Sweden (SCB). The SaTScan Poisson parser, developed specifically for

CASE, raises an alert if a detected cluster ends within the last week.

The third detection method was developed by, and is in regular use at the Health Protection Agency in England and Wales (Farrington et al. 1996). In CASE, we use the `surveillance` R-package implementation (Höhle 2007) of the method and we refer to it as the *Farrington algorithm*. The algorithm is used on data aggregated at the national level, to investigate if the current disease incidence exceeds that of the reference data from previous years. The CASE parser for the Farrington output ensures that an alert is sent only if an exceedance occurred during the last two weeks. The required window size is implemented as a sliding window of seven days and detection is performed daily.

The *threshold algorithm* is used to generate alerts when the number of cases for a particular disease rises above a manually defined value, with the number of cases aggregated at the national level.

For all methods, as long as an outbreak is ongoing according to the results of the statistical analysis, a new alert is raised every night. Figure 5.3 shows an alert email that is sent to the recipients of “MRSA infection”. The graph is automatically generated by the detection component and shows all computed alarms on the x-axis. The computed threshold is denoted by the blue curve (the graph in Figure 5.3 was generated using simulated data). The email also includes a brief description of the algorithm that generated the alarm.

5.3 Results and Discussion

CASE is a technical framework designed to ease the process of connecting a data source with reported cases to various statistical methods requiring different input formats. When using CASE, the user can select the methods that are best suited to the characteristics of a particular disease.

CASE can also be used as a platform for comparing different detection algorithms, although that is not its primary purpose. Since all algorithms use the same data, running multiple detection methods on the same disease regularly and comparing the successful detections and the false warnings can provide insights into the accuracy of a certain method for a given disease. Comparisons and evaluations of the statistical methods currently included in CASE can be found in, for example, Rolfhamre and Ekdahl (2005) and Aamodt, Samuelsen, and

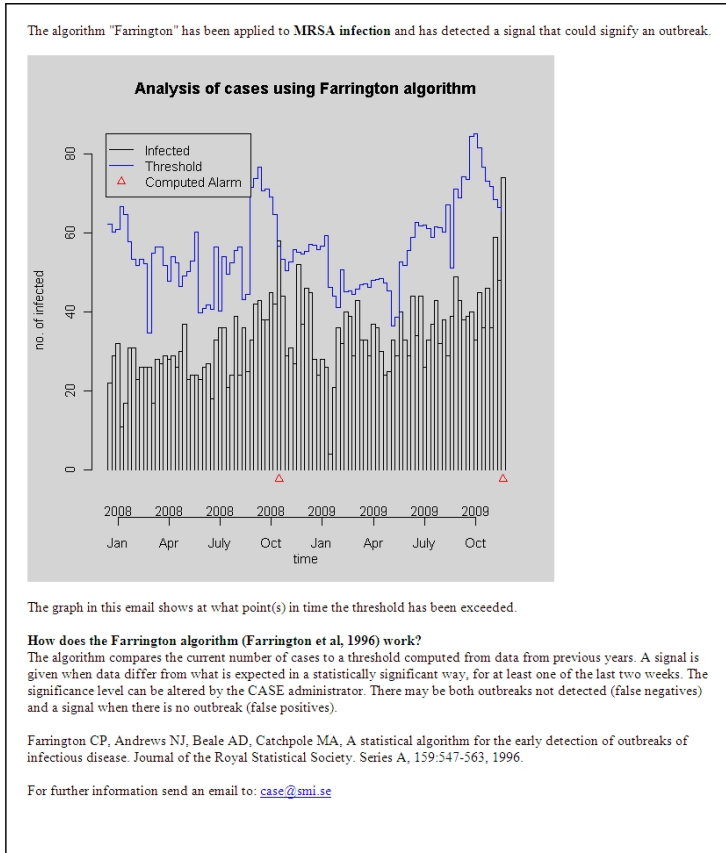


Figure 5.3: A sample email for a disease alert.

Skrondal (2006). Here, the importance of calibrating the parameters for the detection methods must be emphasized, something which is still an ongoing work at SMI.

At present, the evaluation of the system is mainly qualitative, consisting of frequent discussions between the epidemiologists and the CASE developers. There is, however, a need for more systematic evaluations of the system, including a questionnaire assessing the users' experience, in addition to quantitative evaluations of the performance of the algorithms and the parameter settings. To facilitate the quantitative evaluations, we plan to extend the functionality of CASE to incorporate an evaluation module allowing the algorithms to be run

retrospectively, with analysis carried out for each day in a specified time period. The main objective is not a general comparison of the algorithms, but an assessment of their performance in the specific context of the data they are used on. Where external data telling when actual outbreaks have occurred are available, measures such as sensitivity and specificity can be calculated. The evaluation module would provide valuable guidance in the choice of algorithms and parameter settings for the end user. Another evaluation feature we consider implementing is the possibility to run simulated data in the system.

CASE currently uses emails for notification. The advantage of this approach is that it presents information to the users in a familiar way and does not require them to learn how to operate a new interface. The disadvantage, on the other hand, is that the system becomes one-sided if the emails do not include a feedback mechanism. Regardless of the actual implementation, a system for providing feedback from the receivers of the signals is essential. Currently, users who would like to provide feedback on CASE output are instructed to email the administrator.

As expected, a relatively simple method operating on accurate and informative data produces better results than a complex method operating on noisy data. Therefore, the most important factor for creating a reliable outbreak detection system is to ensure the quality of the input data. If the input is not reliable, improving the data collection process from local medical centres is a much better investment than trying to perform automatic detection on inaccurate data. Additionally, expectations from an automated detection system must be realistic. For a computer, detecting ongoing outbreaks and seasonal regular outbreaks is possible, but predicting an outbreak at onset is currently not feasible.

CASE is designed primarily to analyze case reports and does not provide syndromic surveillance support using external data sources, unlike RODS (Tsui et al. 2003) or BioSTORM (Crubezy et al. 2005). The only requirement for the operation of CASE is access to a case database for notifiable diseases. All scripts to create and configure the intermediate local database are included in the software package. The local database is used to selectively copy and store case reports after removing all information that can reveal a patient's identity. We believe that the ease of configuration and maintenance in addition to the possibility of operating without storing highly sensitive data make CASE

a strong candidate for use in national infectious disease surveillance.

5.4 Conclusions

In this paper we have described the design and implementation of a publicly available technical framework for computer supported outbreak detection. The source code is licensed under GNU GPLv3 (FSF) and is available from <https://smisvn.smi.se/case>.

The CASE framework is designed to be a complete system for computer supported outbreak detection at the national level. We are aware that any outbreak detection system must always be adapted to a particular context, where national requirements and regulations will affect the implementation of the system. Such adaptations can easily be made within the described framework. By making the code open source, we wish to encourage others to contribute to the future development of computer supported outbreak detection systems, and in particular to the development of the CASE framework.

Availability and Requirements

The source code for CASE is licensed under GNU General Public License Version 3 (GPLv3), and is available for download from <https://smisvn.smi.se/case>. The provided documentation and the interface are written in English. The following software must be installed on the target system in order to use CASE:

- Linux or Windows operating system that can run Sun Java Runtime Environment 6.0 (or higher)
- MySQL 5.1 (or higher)
- SaTScan version 8.0.1 (or higher)
- R version 2.9.1 (or higher)
- ImageMagick 6.5.4 (or higher)

Competing Interests

The authors declare that they have no competing interests.

Authors' Contributions

AH, BC and PS designed and developed the CASE framework. BC, KH and PS implemented the framework. KH and MG worked on improving the application. AH and BC drafted the manuscript. All authors read and approved the final manuscript.

Acknowledgements

We would like to thank the epidemiologists at SMI, especially Margareta Löfdahl and Tomas Söderblom, both enthusiastic recipients of the notifications during the early stages of the project. We also thank Martin Camitz for naming the system. Finally, we would like to thank everyone who provides reports for and works with SmiNet. The project is funded by the Swedish Civil Contingencies Agency (formerly the Swedish Emergency Management Agency).

References

- Frisén, M, E Andersson, and L Schiöler. 2009. "Robust Outbreak Surveillance of Epidemics in Sweden." *Stat Med* 3:476–493.
- Aamodt, G, S O Samuelsen, and A Skrondal. 2006. "A Simulation Study of Three Methods for Detecting Disease Clusters." *Int J Health Geogr* 5 (15).
- Crubezy, Monica, Martin O'Connor, David L Buckeridge, Zachary Pincus, and Mark A Musen. 2005. "Ontology-Centered Syndromic Surveillance for Bioterrorism." *IEEE Intell Syst* 20 (5): 26–35.
- Farrington, C P, N J Andrews, A D Beale, and M A Catchpole. 1996. "A Statistical Algorithm for the Early Detection of Outbreaks of Infectious Disease." *J Roy Stat Soc Stat Soc* 159 (3): 547–563.
- FSF. *Free Software Foundation - Licenses*. <http://www.gnu.org/licenses/gpl.html>.
- Höhle, M. 2007. "surveillance: An R Package for the Monitoring of Infectious Diseases." *Comput Stat* 22 (4): 571–582.
- Krause, G, D Altmann, D Faensen, K Porten, J Benzler, T Pfoch, A Ammon, M H Kramer, and H Claus. 2007. "SurvNet Electronic Surveillance System for Infectious Disease Outbreaks, Germany." *Emerg Infect Dis* (October).
- Kulldorff, M. 1997. "A Spatial Scan Statistic." *Commun Stat Theory Methods* 26:1481–1496.
- Kulldorff, M, J Heffernan Hartman, R. Assunção, and F Mostashari. 2005. "A Space-Time Permutation Scan Statistic for Disease Outbreak Detection." *PLoS Med* 2, no. 3 (February): e59.
- Reis, Ben Y, Chaim Kirby, Lucy E Hadden, Karen Olson, Andrew J McMurry, James B Daniel, and Kenneth D Mandl. 2007. "AEGIS: A Robust and Scalable Real-time Public Health Surveillance System." *J Am Med Inform Assoc* 14 (5): 581–588.
- Rolfhamre, P, and K Ekdahl. 2005. "An Evaluation and Comparison of Three Commonly Used Statistical Models for Automatic Detection of Outbreaks in Epidemiological Data of Communicable Disease." *Epidemiol Infect* 134 (4): 863–871.

- Rolfhamre, P, A Janson, M Arneborn, and K Ekdahl. 2006. "SmiNet-2: Description of an Internet-based Surveillance System for Communicable Diseases in Sweden." *Euro Surveill* 11 (5): 626.
- SaTScan. *SaTScan - Software for the Spatial, Temporal, and Space-time Scan Statistics*. <http://www.satscan.org>.
- SCB. *Statistics Sweden*. <http://www.scb.se>.
- Tsui, Fu-Chiang, Jeremy U Espino, Virginia M Dato, Per H Gesteland, Judith Hutman, and Michael M Wagner. 2003. "Technical Description of RODS: A Real-time Public Health Surveillance System." *J Am Med Inform Assoc* 10 (5): 399-408.
- Widdowson, Marc-Alain, Arnold Bosman, Edward van Straten, Mark Tinga, Sandra Chaves, Liesbeth van Eerden, and Wilfred van Pelt. 2003. "Automated, Laboratory-based System Using the Internet for Disease Outbreak Detection, the Netherlands." *Emerg Infect Dis* (September). <http://www.cdc.gov/ncidod/EID/vol19no9/02-0450.htm>.

Chapter 6

A Workflow for Software Development Within Computational Epidemiology*

Baki Cakici^{†,‡}, Magnus Boman^{‡,§}

Abstract

A critical investigation into computational models developed for studying the spread of communicable disease is presented. The case in point is a spatially explicit micro-meso-macro model for the entire Swedish population built on registry data, thus far used for smallpox and for influenza-like illnesses. The lessons learned from a software development project of more than 100 person months are collected into a check list. The list is intended for use by computational epidemiologists and policy makers, and the workflow incorporating these two roles is described in detail.

*Cakici, Baki, and Magnus Boman. 2011. "A Workflow for Software Development Within Computational Epidemiology." *Journal of Computational Science* 2 (3): 216–222.

[†]Swedish Institute for Communicable Disease Control (SMI), 171 82 Solna, Sweden

[‡]Royal Institute of Technology (KTH), 164 40 Kista, Sweden

[§]Swedish Institute of Computer Science (SICS), 164 29 Kista, Sweden

6.1 Introduction

Computational Epidemiology

In 1916, Ross noted that mathematical studies of epidemics were few in number in spite of the fact that “vast masses of statistics have long been awaiting proper examination” (Ross 1916, 205). In the 90 years which followed, the studies made were analytic, and the micro-level data available were largely left waiting, to leave room for systems of differential equations built on homogeneous mixing. This is remarkable not least because the modeling problem remains the same throughout history: “One (or more) infected person is introduced into a community of individuals, more or less susceptible to the disease in question. The disease spreads from the affected to the unaffected by contact infection. Each infected person runs through the course of his sickness, and finally is removed from the number of those who are sick, by recovery or by death. The chances of recovery or death vary from day to day during the course of his illness. The chances that the affected may convey infection to the unaffected are likewise dependent upon the stage of the sickness.” (Kermack and McKendrick 1927, 700). Heterogeneity is present already in this classic description, in several places; susceptibility, morbidity, and also contact patterns, if only implicitly. Only with the advent of powerful personal computers, were micro-level data given a role in the modeling of epidemics. Executable simulation models in which each individual could be modeled as an active object with its own attributes (Boman and Holm 2004), often referred to as an agent, began to appear (Eubank et al. 2004; Ferguson et al. 2005; Longini et al. 2005). A new area within computer science, computational epidemiology, has recently become established as the scientific study of all things epidemiological except the medical aspects. This area is turning into computational science (see, e.g., Balcan et al. 2010), following the example of computational biology, computational neurology, computational medicine, and several other new areas focusing on building computationally efficient executable models. This development also includes the social sciences, as in computational sociology (Epstein 1999).

Model Description

The model on which the analysis below is based has been continuously developed since 2002 by a cross-disciplinary group of researchers from the fields of medicine, statistics, mathematics, sociology and computer science. Since 2004, a team of developers have implemented various versions of a software tool, representing the computational part of the model, recently made available as open source software and licensed under GNU General Public License Version 3 (SMI 2010). In parallel with the implementation, the requirements on the model have changed many times. It began as a model for predicting the effects of a possible smallpox outbreak in Sweden (Brouwers et al. 2010), which was later transformed into a model for studying pandemic influenza, and is now a model that could be used for many different kinds of communicable disease studies (excluding vector-borne diseases, i.e., diseases with animal reservoirs). The model is a detailed representation of real situations, sometimes referred to as a tactical model, as opposed to simpler strategic models (Coelho, Cruz, and Codeco 2008). For instance, the model was recently used to study a fictitious scenario of H4N6: a new influenza virus strain that was assumed to be deadly, highly contagious, and introduced into a completely susceptible population. In all, the development project has included more than 100 person months of implementation work, and consists of more than 5000 lines of C++ code.

The parameters used to represent individuals in the model are age, sex and current status (alive or deceased). Each individual is also assigned a home, a workplace, and a department within that workplace. The movement of individuals outside of home and workplace are represented using travel status (home or in another location), emergency room visits, and hospitalizations.

Infections caused by social contact outside of work or home are classified as context infections. When the context infection process is active, there is a probability that an infectious individual will infect those that live within a fixed radius. Context contact radius defines the size of neighborhoods, mirroring the interaction of every individual with others, based on geographical proximity and the social network.

The disease affects every individual through three parameters: infectiousness, death risk, and place preference. The infectiousness parameter influences the probability that the infected individual will

infect others in the same home, workplace, or neighborhood. The death risk depends on the disease level and is expressed as a probability. Place preference is the probability distribution used when deciding where the individuals will spend their day (workplace, home, primary care, or hospital). These parameters are defined for five levels of severity: asymptomatic, mild, intermediate, severe, and critical. In addition, there are four disease profiles: asymptomatic, mild, typical, and atypical.

The model description is combined with Swedish data on workplaces, households, and individuals. Workplaces include companies, schools, healthcare, and other state institutions. For each workplace, the data indicate the total number of workers, geographical coordinates, and workplace type. The current version of the simulation platform uses data from the Swedish Total Population Register, the Swedish Employment register, and the Geographic Database of Sweden (cf. SCB 2010).

Because the model was developed with the purpose of being run with data for the country of Sweden, it has been used solely for studying outbreaks in that country. Sweden has relatively many infection clinics and good international reputation for detailed clinical reports of communicable disease. Thus, in some areas of disease control, Sweden works well as a role model. Other countries face special local problems, however, and results have sought to be generalizable, for example contributing to the complicated model of EU care-seeking behavior. Generally speaking, the project goals have included to sensitize policy makers to the scope of possible disruption due to a newly emergent disease event, and to identify a range of policy handles which can be used to respond to such an episode.

A sample case description illustrates how an experiment would be described using the executable model. The sample case simulates the effects of pandemic influenza in Sweden, without any interventions, for 300 days. The simulation is initiated with 50 infected individuals, randomly selected from the entire population. Since the data set is registry data for the entire country, any random selection procedure is uniform, i.e., an individual has a 50 in nine million chance of being initially infected. This does not mirror realistic spread, which would more typically be an airplane or a boat arriving to Sweden with one or more infected individuals on board, but in the sample case it at least provides an opportunity to discuss the complex matter of how

epidemics start. The maximum size for an office is set to 16 individuals and all workplaces with more than 16 employees are split into departments, each containing 16 or fewer members. This value is not arbitrary, but corresponds to the average size of a Swedish workplace. Context contacts—the parameter representing the average number of contacts outside the home or the workplace—is set to 15. Even if that number was recommended by the sociologists in the project, it is somewhat arbitrary, and is therefore subjected to sensitivity analyses in our sample case. Naturally, such analyses would be extensive in a real policy case; here the reason for their inclusion is chiefly pedagogical.

Disposition

A report on lessons learned from the software development project constitutes the bulk of the analysis below. It starts with a description of the workflow in a computational epidemiology project, and observations on the micro-meso-macro link follow. More detailed descriptions of what it actually means to manage and run a simulator are then provided, before discussing the scientific merits and challenges of this kind of research, and the concluding check list is presented.

6.2 Workflow

Model Development

The process of developing a model for outbreaks today often includes the development of a simulator, allowing for scenario execution and relatively swift sensitivity analyses. The simulator does not capture the entire model, but only those parts that are subject to uncertainty or those that involve stochastic parameters. The instigator is typically a policy maker (PM), knowledgeable in public health issues, and seeking to evaluate various scenarios. The PM may well have medical training, or even be an epidemiologist. The implementer of the simulator is a computational epidemiologist (CE): a modeler knowledgeable in computer science, and the social sciences, typically without much medical training. Naturally, both PM and CE could denote a team instead of a single person. A schematic workflow for develop-

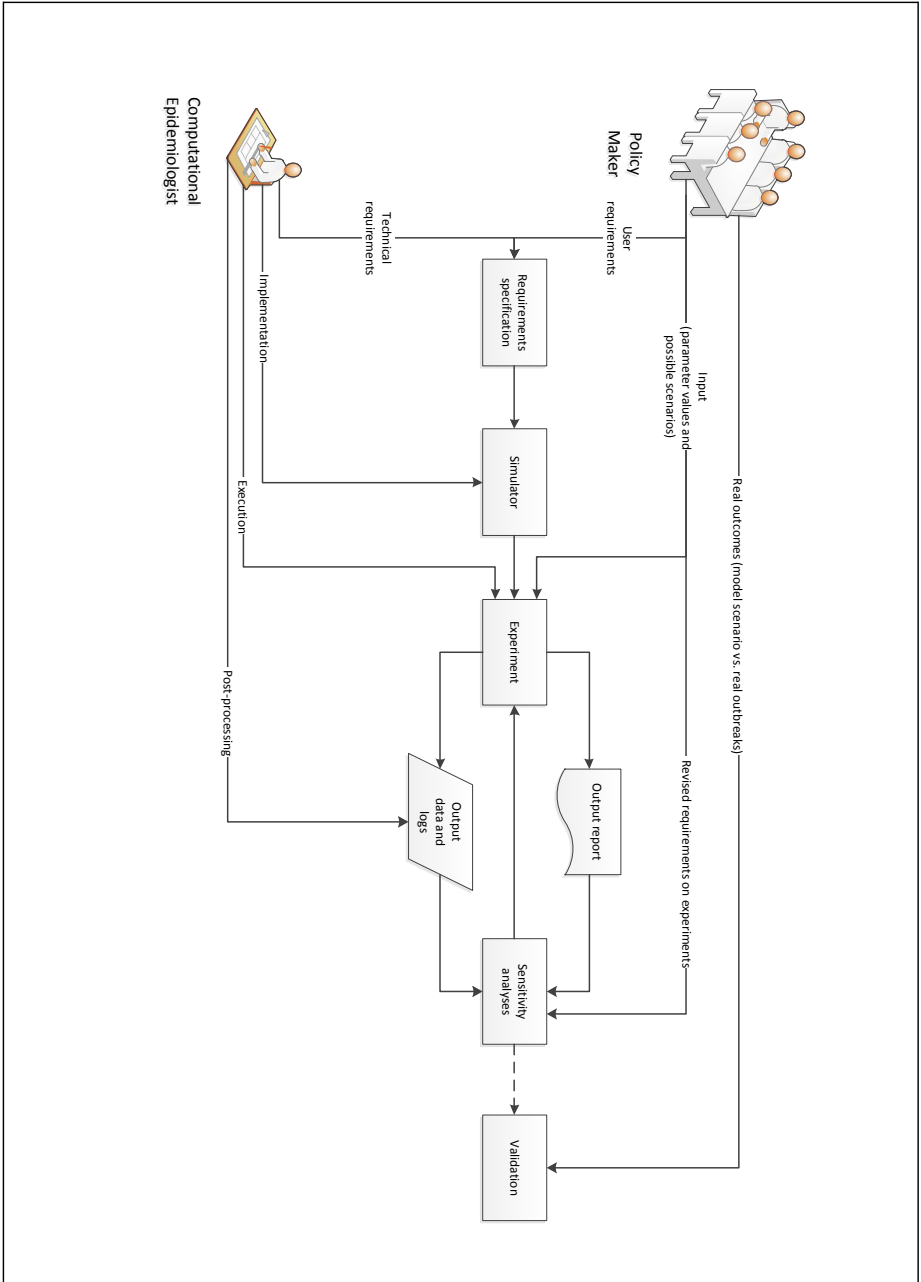


Figure 6.1: The schematic workflow of developing and running an executable model, incorporating policy makers and computational epidemiologists.

ing and using a simulator, depicting the roles of both PM and CE, is presented in Figure 6.1.

As in all development projects, work begins with a requirements specification, to which the PM contributes user requirements and the CE contributes technical expertise. From this specification, the simulator is built. It consists of a software package with two parts: a simulation engine and a world description. The latter is not the complete description of the world under study, but covers only those parts that have a bearing on the executable model. This modeling work is carried out by the CE, with considerable assistance from medical professionals. The CE implements the simulator in accordance with the specification and medical expertise. The CE will also seek to verify the accuracy of the simulator (e.g., through extensive testing, or even logical proof). The CE works in two distinct sequential steps that cannot be combined: design and implementation. Software engineers are taught not to modify their design during the implementation stage to “improve” the model, no matter how tempting this might be. If design decisions leak into the implementation stage, the software project quickly becomes impossible to maintain. What software design means in the area of computational epidemiology is the craft of knowing which parameters to vary, being aware of their mutual dependence, and how to openly declare all simplifying assumptions.

Once the simulator is complete it is given a version number, and one may proceed to experiments. For an experiment to be meaningful, the PM must envisage scenarios. The PM must also provide values for some input parameters. Each parameter in the model is important, and even slight changes to an input value might have a drastic effect on the output. The kind of model considered here is a complex system: a system which cannot be understood through understanding its parts. Before the CE can run the system, the world description must be populated with data, which typically need a significant amount of post-processing to allow for smooth use in the simulator. In addition, one must then attempt to ascertain that the resulting data set is accurate and noise-free. The data set in the here described model was sensitive with respect to personal integrity, as it consisted of registry data on the entire Swedish population of approximately nine million individuals. This sensitivity rendered many kinds of replication experiments impossible.

Once the system runs, it will produce a vast amount of output,

so experiments must be set up carefully to avoid information overload. The so-called induction trap—the lure of running too many experiments for each scenario because it is easy to produce more output, and then jumping to inductive conclusions too swiftly (Popper 1957)—must also be avoided. The output and logs of a set of runs typically do not lend themselves to straightforward reading, but require post-processing. In practice, this means turning huge text files into calculable spreadsheets, and further into graphs and diagrams. Those outputs can then be presented back to the PM, who can then call for more experiments, sensitivity analyses, or even a revision of the requirements specification. The CE in this process makes certain design choices, e.g., which output data to present and how. It is important that this process is iterative and that the PM is given the option of making informed choices, by having at least some grasp of what is realistic to do, given the constraints of computational complexity. The CE must provide technical specifications on further experiments, and the technical competence used also comes with a responsibility to inform: the PM must know what options there are, and why and how certain results were omitted or deemed irrelevant. Because the PM is typically the one responsible for acting upon results obtained, a chain of trust to the CE must be upheld. Likewise, the CE should react if the PM, for example, calls only for certain experiments to be run, or if the selection is made so as to confirm a preconceived truth, in a pseudo-scientific fashion (Lakatos 1977).

In principle, the output of the executable model can finally be validated by comparing its predictions to real outcomes of actual policy interventions for the population modeled, given that the input parameters adequately model the real population prior to those interventions. Naturally, some scenarios could be considered extreme (e.g., the introduction of an entirely new influenza virus to a population without native immunity) and are simulated precisely because they cannot be studied in the real world. In such scenarios, validation can, at best, pertain only to parts of the model. More importantly, simulations of outbreaks are difficult to validate because the simulated event is rare. Catastrophic events are characterized by low probability and disastrous consequences (see, e.g., Thom 1993), and yet the input data are collected from the normal state of the population in non-outbreak situations. Using this input, the simulator is expected to produce one possible yet highly unlikely scenario to provide re-

searchers and policy makers with more opportunities to observe and learn about the unlikely event.

Since computational epidemiology is problem-oriented and constitutes applied science, models are often pragmatic in the sense that they are adapted to their use as policy-supporting tools. Any provisos made have to be grounded in the culture of the decision making entity, such as a government or a pharmaceutical company, making alignment studies, in which models are docked for replication studies (Axtell et al. 1996) difficult.

The Micro-Meso-Macro Link

In microsimulation models of outbreaks, individuals are exposed to the disease and may infect other individuals that they come into contact with. The most primitive unit is the individual and the focus is on the activities of the individual, for the purposes of studying transmission. By contrast, macrosimulation focuses not on the individual, but on the whole society. All members (i.e., the whole, possibly stratified, population) share the same properties and move between different disease states such as susceptible, infected, and resistant.

Even if originally conceived as a pure microsimulation model, the executable model discussed here has macro-level parameters, e.g., workplace size. This parameter governs how many colleagues a working individual interacts with during a working day. To “interact with” here means that there is an opportunity for infection, given that either the individual or the colleague is ill. Even though micro data are available for each workplace—including the number of employees at each company—it is defensible not to use these data in full, since large workplaces have so many employees that it makes no sense to assume that the individual interacts with them all. In reality, the individual might not even see more than a fraction of the total number of colleagues on a given day. The workplace size is therefore set to a precise value, meant to capture an average number of colleagues, which is kept constant throughout a set of runs.

By definition, macro models do not represent local interaction. However, in any dynamic model utilizing micro data, including SIR-inspired individual-based models (Roche, Guegan, and Bousquet 2008), local interaction will affect the output. If there appear discernible patterns in the output that are not explicitly stated by the model descrip-

tion at the outset, they are referred to as emergent patterns. In the described model, all output logs are mapped onto a real population. This means that every discernible pattern has an interpretation that can be understood in the epidemiological context, using terms such as “spread” and “giant component”, and also in the societal context, using terms like “number of infected” and “absenteeism”. Hence, patterns discernible at the macro level resulting from local interactions at the micro level are easily made understandable to the PM.

The meso layer (Liljenström and Svedin 2005) includes everything that is more general than the properties of single individuals but less general than the properties of the whole society. In the model at hand, this is most visible in neighborhoods, defined by the geographical proximity of different households. Adding the meso layer to an epidemiological model enables researchers to represent a crucial part of human interaction: social contacts outside the home or workplace. This includes encountering others while shopping, and social gatherings of neighbors.

Variables in the executable model represent properties of the real population, but many of them cannot be observed directly. Therefore, the argument goes, a suitable value for the executable must be determined by experimenting with the simulator. In the implementation phase, the CE strives to get a handle on the parameter space, i.e. the value space for all parameters that can be subject to variation. To illustrate this, a sample case is now considered.

To find a suitable value for the parameter *context contacts*, representing the average number of contacts outside the home or the workplace, the behavior of the simulated outbreak is observed using the total number of infected individuals per week for a large interval of context contact values. The interval is set to start from zero, where the model behavior is undefined, to where the parameter no longer has an observable impact, i.e. when it is high enough to exhaust the population regardless of all other parameters. Within the [8,20] interval, changing the *context contacts* parameter had, in this example, a significant effect on the behavior of the model. Repeating the same series of experiments with a smaller step size within the [8,20] interval, a smaller region of interest was obtained within the [14,16] interval. Finally, the analysis was repeated one last time for the [14,16] interval with a smaller step size. Figure 6.2 shows the number of infections per week for five runs where all parameters except *context contacts* were




Figure 6.2: Number of infections per week for five runs where all parameters except context contacts were kept constant.

kept constant. Further simulations were run to observe the effects of variation due to random seeds when contacts was set to 15. Figure 6.3 shows the number of infections per week for three runs with different random seeds where all other parameters were kept constant.

Other variables in the executable model that should be decided using a similar process include (but are not limited to): number of initially infected, office size, place choice based on disease level, place choice based on age, length of a work day, and the probability of receiving a symptomatic disease profile.

Stochasticity

An outbreak of pandemic influenza is a rare event. To trigger such an outbreak, either the simulations must be run repeatedly for a long period until an outbreak occurs, or the model must be configured in such a way that outbreaks will occur with higher frequency than in the real world. The former is not practical since it might take millions of runs before anything happens, and the latter comes with the risk




Figure 6.3: Number of infections per week for three runs with different random seeds where all other parameters were kept constant. Each random seed is a vector of numbers generated by a pseudo-random number generator.

of compromising the validity of output by introducing exogeneous variables that change the effects of the simulated outbreak.

All random events in the model use a series of numbers that are generated at run-time using the initial seeds provided by the user. Therefore, the outcome of every “random” event in a simulation run depends only on the initial seeds. By using the same seeds, identical results can be obtained using different computers, operating systems, or compilers.

In the present model, one highly influential parameter is the number of initially infected. When 50 randomly selected individuals are infected, an outbreak is triggered in nearly every run. If only three individuals are selected instead, the outbreaks become much more rare. This is due to the heterogeneity of the population: individuals with more contacts are more likely to initiate outbreaks if infected, and it is more likely that a highly connected individual would be infected if

50 rather than three are infected initially.

It is often assumed in executable models that in a few generations, a simulation with three infected would reach the stage with 50 infected, and that the difference between them would be negligible. Certainly every simulation with three initially infected would reach a stage with 50 infected, given that an outbreak occurs during the run. Therefore simulations can be started from the stage where 50 individuals are infected since that is the minimum number at which the simulation platform produces outbreaks in the majority of runs. This assumption is far from ideal. The simplest observable effect is that no runs will have less than 50 infected. This is acceptable because the object of study is nation-wide outbreaks. However, the difference between the two approaches is not negligible because 50 randomly selected individuals will not have the same geographical distribution as 50 individuals whose infections originate from three individuals. The 50-from-three group will most likely have overlapping social networks because they were all infected by three individuals, as opposed to being randomly selected from a population of nine million. As the outbreak grows to one thousand or one hundred thousand infected, the difference may lose its significance, but quantifying that significance remains challenging for all executable models that use heterogeneous populations. Hence, this is a good example of a simulation in which the CE makes an assumption about things beyond the PM's control, or even grasp. Good software development requires that such assumptions be made explicit and communicated to the PM.

6.3 Conclusion

The lessons learned from the software development project described above can be summarized in the form of a check list. Even if the list is not exhaustive, developers of computational epidemiology models could check off the items on the list, as applicable to their project. The presented workflow and checklist do not include surveillance in computational epidemiology and instead focus on modeling and simulation. A more comprehensive workflow for computational epidemiology would have to incorporate computer-assisted infectious disease surveillance, often performed using complex software platforms tailored to the task (Espino et al. 2004; Crubezy et al. 2005; Abramson

et al. 2010; Cakici et al. 2010), and the interaction of its users with the actors already identified in the preceding sections.

Computational epidemiology is a new area, and many of the methods and theories employed have yet to benefit from thorough scientific investigation. Even if important steps towards amalgamating models and performing alignment experiments have been taken (see, e.g., Halloran et al. 2008), the area is in need of extensive methodological advancement. The following checklist is intended to be a contribution to such development. Not every item in the check list introduces new issues for policy makers or computational epidemiologists, but, depending on the reader's area of expertise, one or two are highly likely to be more significant than the others. Much of it is part of the folklore of the area, and could be classified as procedural and pragmatic know-how. More specifically, the contribution is to have these items made explicit as one concise list, and tied to working procedures as demonstrated by our workflow description (Figure 6.1).

1. All population data sets are regional

To have access to data on the entire population on the planet is not a realistic goal. Hence, most studies are limited to one geographic region, such as a city, a state, or a country (Chao et al. 2010). This means that the universe of discourse includes not only the individuals in this geographic region, but also that a certain proportion of the individuals must be allowed to leave the region. Moreover, visitors and immigrants from other regions should be included in the population data. Some computational epidemiology projects employing micro data use census data, others extrapolate from samples, and yet others use synthetic data. In the rare cases where registry data is available for a large population—as is the case for the Swedish population—hard methodological questions must still be answered regarding the generalizability of results: which parts of a scenario execution in Sweden are likely to be analogous to ones in Norway, Iceland, or the state of Oregon?

2. Population data are sensitive

Even after extensive post-processing, any data set with real population data is subject to privacy and integrity concerns. In almost all coun-

tries, this means that running a simulator with the data set is subject to applying to an ethics board. If approved, data must be kept safe and experiments may be run in designated facilities only. This makes replication studies difficult, and it also restricts alignment studies to less interesting data sets.

3. Verifying the simulator is a serious engineering challenge

To formally verify that the simulator produces adequate results, is free from programming bugs, and can handle the computational complexity of modeling large outbreaks is, in general, not possible. The software is too large, as is the variation of possible input values and the spectrum of sensitivity analyses. Extensive testing—varying the hardware environment and the parameter values, including the random seeds for stochastic processes—yields evidence for adequacy, but no guarantees. This does not entail that the simulator is without use, or not to be trusted, but merely that its construction and maintenance is an engineering challenge.

4. Validating the simulator output is hard

Pandemics have been few and far between. Modeling a future scenario on a real outbreak of the past has been done with some success in the area of epidemiology. The structural properties of current and future societies may vary greatly from those studied in the past, however. Air travel, hygiene, and working conditions are three out of many factors that affect the spread of communicable disease and that vary greatly in the historical perspective. The low probability of catastrophic events such as a pandemic makes it very hard to validate any simulation experiment against real-world events.

5. Assumptions and hypotheses should be stated and controlled by the policy maker

Placing assumptions on top of assumptions will only create a gap between the policy maker and the computational epidemiologist. As illustrated by the example of selecting different initially infected individuals, the description of a single assumption can be interpreted in multiple ways, and the implementation of different interpretations

can diverge significantly from the respective intention. The complexity of communicating all assumptions implied by the decisions of the policy maker arises from the tremendous difficulty in identifying implicit assumptions at every step of development. Because every addition to the model carries the risk of modifying the interaction of existing parameters, ensuring that all assumptions have been made by the policy maker becomes a formidable challenge.

6. Triggering outbreaks in the simulator is nontrivial

To implement a simulator that always produces outbreaks is easy. Increasing the infectiousness of a disease (as done, e.g., Roche, Guegan, and Bousquet 2008) or the number of initially infected, quickly yields a disease pattern affecting the entire giant component, i.e. every individual connected to other individuals through the social network or by geographical proximity (cf. Youssef, Kooij, and Scoglio 2011), forming the largest connected subgraph of the population graph (cf. Newman 2003). If such settings are inconsistent with empirical data, or with assumptions and hypotheses declared, however, then the adequacy of the model should be questioned. There is evidence for the fact that the initial stages of a pandemic require a different kind of modeling than the later stages (Bonabeau, Toubiana, and Flahault 1998). It would therefore be naïve to think that increasing the number of initially infected—in order to trigger outbreaks in a larger proportion of runs—would not affect the model of the entire pandemic.

7. Hybrid models need constant refinement

A model in which the micro, meso, and macro properties are integrated has the potential to mirror reality in a relatively accurate way. Under the proviso that model adequacy yields better prediction, one could discard the simplest models in favour of such hybrid models. The level of ambition, however, comes at the price of the model never being finished, and model-dependent artifacts becoming more difficult to identify. Since the world to be modeled is a moving target, and since macro data can often be replaced by micro data as it becomes available, there are always refinements to be made. The devil is in the details.

Acknowledgements

The authors would like to thank the current leader of the MicroSim project at the Swedish Institute for Communicable Disease Control, Lisa Brouwers. The authors also thank Olof Görnerup, Eric-Oluf Svec, the editor, and the anonymous reviewers for their constructive comments.

References

- Abramson, David, Blair Bethwaite, Colin Enticott, Slavisa Garic, Tom Peachey, Anushka Michailova, and Saleh Amirriazi. 2010. "Embedding Optimization in Computational Science Workflows." *Journal of Computational Science* 1 (1): 41–47.
- Axtell, Robert, Robert Axelrod, Joshua M. Epstein, and Michael D. Cohen. 1996. "Aligning Simulation Models: A Case Study and Results." *Computational & Mathematical Organization Theory* 1, no. 2 (February): 123–141. doi:10.1007/BF01299065.
- Balcan, Duygu, Bruno Goncalves, Hao Hu, Jose J. Ramasco, Vittoria Colizza, and Alessandro Vespignani. 2010. "Modeling the Spatial Spread of Infectious Diseases: The Global Epidemic and Mobility Computational Model." *Journal of Computational Science* 1 (3): 132–145.
- Boman, M., and E. Holm. 2004. "Multi-Agent Systems, Time Geography, and Microsimulations." Chap. 4 in *Systems Approaches and their Application*, edited by M-O. Olsson and G. Sjöstedt, 95–118. Springer.
- Bonabeau, Eric, Laurent Toubiana, and Antoine Flahault. 1998. "The Geographical Spread of Influenza." *Proceedings of the Royal Society B* 265:2421–2425.
- Brouwers, L., M. Boman, M. Camitz, K. Mäkilä, and A. Tegnell. 2010. "Micro-simulation of a Smallpox Outbreak Using Official Register Data." *Eurosurveillance* 15 (35).
- Cakici, B., K. Hebing, Grünewald M., P. Saretok, and A. Hulth. 2010. "CASE: A Framework for Computer Supported Outbreak Detection." *BMC Med Inform Decis Mak* 10 (14).
- Chao, Dennis L., M. Elizabeth Halloran, Valerie J. Obenchain, and Ira M. Longini Jr. 2010. "FluTE, a Publicly Available Stochastic Influenza Epidemic Simulation Model." *PLoS Comput Biol* 6, no. 1 (January): e1000656.
- Coelho, Flavio, Oswaldo Cruz, and Claudia Codeco. 2008. "Epigrass: A Tool to Study Disease Spread in Complex Networks." *Source Code for Biology and Medicine* 3 (1).

- Crubezy, M., M. O'Connor, Z. Pincus, M.A. Musen, and D.L. Buckridge. 2005. "Ontology-Centered Syndromic Surveillance for Bioterrorism." *Intelligent Systems, IEEE* 20 (5): 26–35.
- Epstein, Joshua M. 1999. "Agent-Based Computational Models and Generative Social Science." *Complexity* 4 (5): 41–60.
- Espino, J.U., M. Wagner, C. Szczepaniak, FC Tsui, H. Su, R. Olaszewski, Z. Liu, et al. 2004. "Removing a Barrier to Computer-Based Outbreak and Disease Surveillance – The RODS Open Source Project." *MMWR Morb Mortal Wkly Rep.* 53 Supplement (September): 32–39.
- Eubank, S., H. Guclu, V.S.A. Kumar, M.V. Marathe, A. Srinivasan, and N. Toroczkai Zand Wang. 2004. "Modelling Disease Outbreaks in Realistic Urban Social Networks." *Nature* 429:180–184.
- Ferguson, Neil M, Derek A T Cummings, Simon Cauchemez, Christophe Fraser, Steven Riley, Aronrag Meeyai, Sophon Iamsirithaworn, and Donald S Burke. 2005. "Strategies for Containing an Emerging Influenza Pandemic in Southeast Asia." *Nature* 437, no. 7056 (September): 209–214. ISSN: 1476-4687.
- Halloran, M. E., N. M. Ferguson, S. Eubank, I. M. Longini, D. A. T. Cummings, B. Lewis, S. Xu, et al. 2008. "Modeling Targeted Layered Containment of an Influenza Pandemic in the United States." *PNAS* 105 (12): 4639–4644.
- Kermack, W. O., and A. G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 115, no. 772 (August): 700–721.
- Lakatos, Imre. 1977. "Science and Pseudoscience." In *Philosophical Papers Vol. 1*, 1–7. Cambridge University Press.
- Liljenström, H., and U. Svedin, eds. 2005. *Micro, Meso, Macro: Addressing Complex Systems Couplings*. World Scientific.

- Longini, Ira M, Azhar Nizam, Shufu Xu, Kumnuan Ungchusak, Wanna Hanshaoworakul, Derek A T Cummings, and M Elizabeth Haloran. 2005. "Containing Pandemic Influenza at the Source." *Science (New York, N.Y.)* 309, no. 5737 (August): 1083–1087. ISSN: 1095-9203. doi:10.1126/science.1115717. <http://www.ncbi.nlm.nih.gov/pubmed/16079251>.
- Newman, Mark E. J. 2003. "The Structure and Function of Complex Networks." *SIAM Review* 45:167–256.
- Popper, K.R. 1957. "Philosophy of Science: A Personal Report." *British philosophy in mid-century*:182–83.
- Roche, Benjamin, Jean-Francois Guegan, and Francois Bousquet. 2008. "Multi-Agent Systems in Epidemiology: A First Step for Computational Biology in the Study of Vector-Borne Disease Transmission." *BMC Bioinformatics* 9 (1).
- Ross, Ronald. 1916. "An Application of the Theory of Probabilities to the Study of a priori Pathometry. Part I." *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 92, no. 638 (February): 204–230.
- SCB. 2010. *Statistics Sweden*, October. <http://www.scb.se>.
- SMI. 2010. *MicroSim Source Code*, October. <https://smisvn.smi.se/sim/>.
- Thom, R. 1993. *Structural Stability and Morphogenesis: An Outline of a General Theory of Models*. Addison-Wesley.
- Youssef, Mina, Robert Kooij, and Caterina Scoglio. 2011. "Viral Conductance: Quantifying the Robustness of Networks With Respect to Spread of Epidemics." *Journal of Computational Science* In Press, Accepted Manuscript.

Chapter 7

Sustainability Through Surveillance: ICT Discourses in Design Documents*

Baki Cakici[†]

Abstract

In this paper, I examine design documents from three different ICT design and development projects. I argue that they present intersecting visions of sustainability entailing the widespread use of ICT, describe the properties of users compatible with such ICT, and provide ways of judging the users. In the design documents, the inhabitants are made individually responsible for living sustainably, and surveillance is positioned as integral to this future with the help of ICT. Underlying the visions, I identify a translation process that captures the traces of the inhabitants' lives, classifies them according to different criteria of sustainable living, and returns them to the tapestry of everyday life to convince the users to behave differently. In the discourses of these documents, surveillance translates the traces, and the translations exert new pressures on existing power relations.

*Cakici, Baki. 2013. "Sustainability Through Surveillance: ICT Discourses in Design Documents". *Surveillance & Society* 11(1/2): 177–189.

[†]Stockholm University, DSV, 164 40 Kista, Sweden

7.1 Introduction

Just before they get home her youngest daughter asks her mother to stop in front of the public installation. They all stop and one of the daughters drags her card on the installation and sees their own household blinking. “Mum our symbol is much bigger this week than last week. And it seems that our block is doing much better than my friend Milla’s block” she says. “That’s good news” her mum answers. (SRS 2011, appx. III, 51)

The fictional conversation between Anne and her daughter Maria appears in a report that details how information and communication technology (ICT) can be developed to encourage sustainable living within a residential city district. As an example, the report describes a brief interaction between the inhabitants and a kiosk that visualises the energy consumption within the district. While using the kiosk, the characters reflect on their behaviour, guided by the signals that indicate the value placed on their energy consumption by the system. The fiction of Anne, her daughter, and the kiosk provides a window into a broad discourse on using technology to contribute to sustainable futures.

In current European Union initiatives for sustainable development, ICT plays a major role. The European Commission’s recommendation report on the issue states that “ICTs can enable energy efficiency improvements”, and “provide the quantitative basis on which energy-efficient strategies can be devised, implemented, and evaluated” (European Commission 2009, 3). In the report, the umbrella term ICT covers a range of technologies designed to monitor, sort, classify, visualise, and optimise. These technologies are defined as tools that can aid in overcoming environmental, social, and economic challenges. To promote them, the recommendation presents three measures: inviting the ICT sector to reach a common measurement methodology, to identify where and how ICTs can play a role in reducing emissions, and a call for the Member States to “enable roll-out of ICT likely to trigger a shift in the behaviour of consumers, businesses and communities” (5). Partly due to vast financial resources that the framework programmes provide for scientific research, the motivations of a wide

variety of projects combining ICT and sustainability are both echoed in and shaped by these measures.

In this paper, I examine three different ICT design and development projects with similar motivations. I argue that they present intersecting visions of sustainability entailing the wide-spread use of ICT, describe the properties of users compatible with such ICT, and provide ways of judging the users. First, the three projects share a vision where sustainability can be achieved through the employment of ICT. In the vision, the primary agents responsible for change are identified as the users of the technologies, and they are made individually responsible for creating sustainable lifestyles with the help of ICT. Second, the users are viewed as rational, technologically competent decision makers who are cooperative with the vision of ICT as an aid for behavioural change. Third, the users are judged according to different criteria that are established by the systems. In the reports, achieving sustainability through the use of ICT is expressed as the right way to live, and the users are provided with social and financial incentives to live in this way. Connecting these three issues, I identify a process of translation where values are inscribed into the systems, and the surveillance of energy use is positioned as integral to sustainability with the help of ICT.

Surveilling the Smart Home

A highly vibrant intersection of sustainability and ICT is called the smart home. While it seemingly refers to an isolated housing unit, it is physically connected to infrastructures as much as any other urban structure, and its information networks reach even further into central servers that collect data at the street, district, city, and nation level. In the smart home, a wide variety of surveillance systems gather sensor and usage data from the surroundings. The gathered data are used by system designers, managers, other technological systems as well as the inhabitants themselves to interpret the activities within the home and to manipulate the home environment. In this description, the home is smart because it is populated by sensor systems. However, it is also a home, because it is populated by the inhabitants. While they go about their daily lives, the sensor systems continually record the interior temperature, water and electricity consumption, and the movements within the home. These measurements are gath-

ered to create representations of the inhabitants' behaviour over time using central databases and statistical methods for analysis. The data, collected from many smart homes simultaneously, are used to reconfigure the image of a standard home: how much energy is consumed, how much movement occurs in it, how many humans live together, etc. The data collection is made invisible, but certain results are communicated back to the inhabitants in the form of logs, graphs, and bills.

In the smart home and its surroundings, sustainability is linked to ICT through surveillance. Following David Lyon's definition, surveillance takes the form of routine attention to personal details with the intention to sort and classify (Lyon 2007, 14). Personal data are collected not only to make systems more efficient, but also to provide ways of creating categories, comparing different individuals, and sorting individuals into groups. The consequence of bringing ICT into the smart home resembles that of any other surveillance system anywhere else; it generates and expresses power (23). I refer here to the Foucauldian notion of power: it is immanent in all relations, inherently productive, and possible to resist, or rather, constitutes resistance (Foucault 1978, 92–96). The power expressed by surveillance links sustainability and ICT by making possible the formulation of knowledge about a population of smart home inhabitants. In the smart home domain, it appears as categories, groupings, and classifications. These can collectively be seen as normalisations:

In a sense, the power of normalization imposes homogeneity; but it individualizes by making it possible to measure gaps, to determine levels, to fix specialities and to render the differences useful by fitting them one to another. (Foucault 1977, 184)

As Foucault describes the process, the creation of averages also provides a way of combining differences. The “inhabitant” category is constructed by measuring the differences in consumption and linking them using statistical methods to produce a whole. In the case of energy consumption, those that consume less and those that consume more (themselves categories constructed through surveillance), can be connected to form a single category under the label “smart home inhabitants”. With such a construction, the differences of the

category are made useful in the quest for sustainability. For example, by regularly monitoring the events in and around the home, it becomes possible to translate traces of everyday life into values such as “avoiding excessive consumption” that the inhabitants are encouraged to recognise and support in their lives.

Categories and classifications are imbued with values, because they make some things visible while concealing others (Bowker and Star 2000). Donaldson and Wood (2004) have emphasised the importance of categories by defining surveillance itself as a process of translating worldviews, denoting systems of categorisation, into materialities. Categories embedded into ICT have also been understood as attempts to control and discipline those who use the systems (Suchman 1994a), although as with any other expression of power, they can also be contested and resisted. In the case of the smart home, surveillance systems categorise and classify traces of consumption behaviour. However, the categories themselves have to be created somewhere. In this paper, I examine acts of category creation in design documents.

Designing ICT

In the projects I analyse, the design process is characterised by the occurrence of many meetings, the production of a large number of reports, and development of deliverables including the material form of the technologies. I use the term design documents to refer to all the texts written during this lengthy process. These design documents express a collective position derived from texts written by a large group of experts including system designers, project managers, developers, funding agencies, etc., who may have different or even conflicting goals individually. For the types of ICT under consideration, the majority of the design processes take place after the funding for the project has been secured, but before the technology has been developed. The things worthy of monitoring and analysing in the smart home are discussed, selected, and written down in this phase. The participants, methods, and the length of the process vary greatly depending on the requirements of each project and the configurations of their actors both geographically and organisationally.

Although the design process aims at prescribing the uses of technologies, the eventual uses cannot be inferred from it. Scholars in the field of Science and Technology Studies (STS) have demonstrated

that technologies are often used in ways that are not intended or foreseen by the designers, and that no single, essential use exists for any technology (Albrechtslund and Glud 2010; Bijker and Law 1992; Oudshoorn and Pinch 2005). Furthermore, the designers constitute only one of many groups of actors that play a role in creating the technologies, and even the act of categorising designers or users as distinct groups tends to obscure the continually shifting boundaries of design and use (Suchman 1994b, 2002; Woolgar 1991).

In my analysis, I trace the discursive construction of the surveillance technologies in the design documents which precede their material form, with the awareness that once these systems are operational they are bound to be shaped and configured by a multitude of actors in many different ways that are impossible to predict. Although users actively co-construct the surveillance systems they encounter by appropriating, modifying, and resisting (Dubbeld 2006), the systems themselves are much more open to change, although by a smaller number of actors, before they fully materialise in their software and hardware incarnations. Starting from classifications, Bowker and Star highlight the negotiations that lead from decisions to technologies:

Someone, somewhere, must decide and argue over the minutiae of classifying and standardizing. . . Once a system is in place, the practical politics of these decisions are often forgotten, literally buried in archives (when records are kept at all) or built into software or the sizes and compositions of things. (Bowker and Star 2000, 44–45)

The extended process of arguing over minutiae binds project managers, system designers, funding agencies, programmers, as well as software standards, keyboards, chairs, and compatible word processor extensions. Somewhere between stand the design documents, describing the decisions for those who will construct the systems. From a collection of discussions, decisions, and limitations both material and temporal, the designers craft texts in the form of project proposals, specification documents, and pre-study reports. In the construction of ICT, the design process is a bridge from discourse to materiality. Often, material is constructed discursively in use cases where fictional scenarios depict future users and usage. However, sometimes the bridge allows movement in the opposite direction; discourse is

materially grounded in non-functional prototypes, objects that stand for the idea of a thing that is yet to be imbued with function.

The documents produced during the design process aim to describe the systems fully, to ensure that they can be developed in the future. Many choices that become invisible later when the system is operational are clearly described in the design documents. The ambition to produce a complete textual description aids the application of discourse analysis. Their discourses describe what is worthy of observation and contain valuable pointers when trying to critically analyse the consequences of living with the systems.

7.2 Material

I draw my material from three different ICT development projects with similar goals regarding sustainability. All three projects are contemporary collaborations between 15 or more partners, and the majority of these are European companies developing ICTs. Other members include universities, research institutes, and electric utility companies. These three projects were chosen primarily based on their goals to develop information infrastructures, particularly for residential areas, with the participation of a large number of organisations.

The first, Home Gateway Initiative (HGI), is a consortium of broadband service providers, and manufacturers of digital home devices, chips, and software (HGI 2012). HGI publishes guidelines and requirements on the digital home infrastructure. The report I analyse, “Home Gateway Initiative—Use cases and architecture for a home energy management service” (HGI 2011), includes eight use cases that aim to reduce the energy consumption in homes.

The second, Future Internet for Smart Energy (FINSENY), is part of the Future Internet Public-Private Partnership Programme (FI-PPP) launched by the European Commission (FINSENY 2012). In the projects, partners from ICT and energy sectors aim to identify the requirements of smart energy systems. The report I use in this paper, “Future Internet for Smart Energy—Smart buildings: Use cases specification” (FINSENY 2011), includes forty use cases that list the requirements for a smart energy infrastructure in five different contexts: homes, residential buildings, office buildings, data centres, and hotels.

The third, Smart ICT, is part of the Stockholm Royal Seaport (SRS) project. SRS is one of Europe's largest urban development projects, and its goal is to develop a new city district in Stockholm, Sweden (SRS 2012). Smart ICT, a much smaller sub-project within SRS, aims to detail a generic ICT infrastructure for SRS. A report from the Smart ICT pre-study, "Stockholm Royal Seaport—Smart Communication" (SRS 2011), describes possible ICT applications that can be deployed within the district. It includes eighteen use cases, called demonstrators, that describe fictional events in which future inhabitants interact with the proposed technologies.

The HGI and FINSENY reports describe technologies that can be deployed, in theory, at many different sites regardless of context. The Smart Communication report, on the other hand, describes applications specifically designed for the Stockholm Royal Seaport district, and provides much finer detail about user interaction and the imagined futures of the proposed technologies. To avoid repetition, I use the terms "HGI report", "FINSENY report", and "SRS report" when referring to the three documents.

7.3 Analysis

The analysis is divided into three parts: first, I describe the intersecting visions of sustainability in the three projects that entail the widespread use of ICT. I then list the properties of the users who are tasked with using the ICTs proposed by the projects within and around the smart homes. Finally, I illustrate the mechanisms that the reports propose to judge these users according to different criteria. Additionally, I identify a process of translation where values are inscribed into the systems.

Visions of Sustainability

The three projects share the common goal to design and develop ICT to accomplish sustainability goals. These goals are defined under four headings in the SRS report: climate change, ecological sustainability, economic sustainability, and social sustainability (5). The HGI report only refers to the concept of sustainability in the abstract, favouring the terms "energy efficiency" and "reduction in energy consumption" instead (HGI 2011, 9). The FINSENY report motivates its fo-

cus through its understanding of the users, where the “generic home dwellers” are willing to accept optimising technologies if the services are kept at the same level, and a few who are “energy conscious” and thus more likely to be proactive (FINSNEY 2011, 20). I return to the definition of the user in the next section.

In all three reports, similar future urban environments are described as being rendered more environmentally sustainable by introducing ICT. Many of the technologies rely on the wealth of sensors proposed for inclusion in smart homes. The properties of the inhabitants of these homes are made visible by the wide-spread use of surveillance systems in the form of sensor networks. Additionally, in the SRS report, the phrase “the ease of doing the right thing” appears in several sections, denoting a specific right thing, a way of behaving sustainably with the help of ICT.

Regarding the evaluation of a population and the ordering of individuals through ICT, the HGI report cites the European Union directive 2006/32/EC in its introduction, which states that member states should ensure that energy distributors make available “comparisons with an average normalised benchmarked user of energy in the same user category” (European Parliament 2006, 72). In accordance with the directive, the use case “Visualization of historical data” proposes to allow the customers to “compare their own energy consumption with other similar customer/communities types” (HGI 2011, 27). These comparisons are motivated in the “business rationale” section of the HGI report:

Environmental degradation and global warming are among the major challenges facing society. . . The most pressing challenge is to reduce the rate of increase of greenhouse gases in the atmosphere and ultimately to decrease the absolute level of these gases. . . ICT technologies can help reduce energy consumption and manage scarce resources, improve efficiency and contribute to cutting carbon emissions. . . Smart Metering, Smart Buildings and Smart Grids, are among the most important ICT-enabled solutions with the highest potential to reduce CO2 emissions. (15)

The HGI report constructs a particular society in which all in-

habitants of the smart homes are compelled to act to counter environmental degradation through the use of technological solutions. By framing the reduction in the rate of increase of greenhouse gases as the most pressing challenge, and proposing management and efficiency as potential solutions, the report links ICT to environmental sustainability.

In the three reports I have analysed, energy is conceptualised primarily as a commodity to be bought and sold. The HGI report states that “instead of measuring energy use at the end of each billing period, smart meters provide this information at much shorter intervals” and “[e]nergy companies will also be able to innovate and offer their customers new types of tariffs that will allow customers to take advantage of cheaper deals at off-peak times” (HGI 2011, 15). In the FINSENEY report, the consumer is defined as having “signed a contract with the electricity provider to access electricity” (FINSENEY 2011, 23), and that “for many customers, monitoring energy consumption is in fact monitoring the bill” (21).

In these proposals, sustainability is interpreted as something that can be achieved in the future and only through change. Since the ability to sell energy, and hence the structure of the participating organisations, is conserved, the partner that is designated as being compatible with change is users. After locating the potential of change in the users, ICT solutions are proposed to utilise that potential and to effect change. Viewed from this perspective, the status quo is preserved for the organisations that provide the energy, and the home dwellers become individually responsible for creating sustainable lifestyles with the help of ICT.

Describing the User

Alongside the descriptions of the proposed systems, the three reports also describe the users of these technologies. These descriptions necessarily refer to a group that does not exist at the time the documents are written, since the technologies themselves have not yet materialised. In other words, the texts discursively construct their users, and assign them certain properties that are essential for operating the proposed technologies.

The necessity of gathering and processing of personal data is a widely held assumption (Murakami Wood 2006). Building on the as-

sumed necessity, the three reports describe a type of user whose data are always available for gathering. The following paragraph from the FINSENY report, describing the differences between the home domain and the residential building domain, illustrate some properties of this user:

In the home domain use cases there is a single user or a number of users with commonly aligned interests (being members of the same household) so they can be assumed to be capable to make timely and coordinated economic decisions. . . . In contrast, in the building use cases there is not a single agent who makes cost-optimising decisions so many such scenarios [*sic*] and use cases are not applicable. In other words, there exists a multitude of different economic agents with access to common resources, giving rise to situations where these individuals, rationally consulting their self-interest, might engage in what will ultimately be wasteful behaviour. Individualised metering can be used to maintain economies of scale while providing incentives to avoid waste. (FINSENY 2011, 38)

The user is viewed as being rational, willing to operate the interfaces provided by the projects, and ultimately, cooperative with the systems. In the “home domain use cases”, the user exists as part of a group of agents that coordinate and decide rationally. When the users’ goals are identified as not necessarily overlapping with those of the systems, as in “building use cases”, the aim of the systems is stated as a way of steering the inhabitants’ rationality towards financial decisions that align with the vision of the project. The user is characterised as *homo economicus*, a self-interested agent making decisions individually based on a personal utility function, and users compatible with the systems are constructed based on the assumption that a shared motivation mechanism is inherent to all individuals. While the reports do not describe how the proposed systems would interact with non-compatible users, scholars of surveillance have thoroughly documented systems that explicitly assume the existence of the non-cooperative user, and the complications that await such “neoliberal deviants” (Gilliom 2001; Maki 2011).

The assumption that financial gain and self-interest guide the inhabitants plays a key role in the argument that ICTs for scheduling electricity usage to take advantage of lower electricity prices at night can contribute to sustainability goals. The rational cooperative user assumption is compatible with the existing commercial activities of the project partners because it does not challenge the practice of buying and selling electricity or other forms of energy. The SRS report extends the commercial frame even further with the following statement: “Research is to use money to create knowledge. Innovation is to use knowledge to create money” (SRS 2011, 23). In the three reports, commercial interests of the project partners form the central pillar around which other concerns such as sustainability and comfortable living are structured.

The necessity of discursively constructing compatible users becomes visible primarily in scenarios that describe the interactions between the users and the systems. The HGI report includes an example:

The system will provide the end user with an easy way to access the appliance configuration web page. When the end user opens the remote page general options to control an appliance are available (stop and start), and other more specific ones for various appliance (e.g. for a washing machine skip the spin cycle or use a lower temperature). The user will be informed of the current consumption and the impact of any configuration change. (HGI 2011, 32)

The user referred to in this scenario is both a technically competent individual that can operate the system, and one that wishes to operate the system. The proper functioning of the system depends on both: if the user cannot understand how the system works, or does not wish to invest time in its functioning, then none of the options provided in the application configuration web page accomplish any positive outcome (Darby 2010). The scenarios detailed in the reports emphasise individual choice, but the systems constrain the consequences of that choice severely. In most cases, the benefit to the customer is a lower energy bill, and possibly lower energy consumption. However, to accomplish the outcome, the inhabitants are expected to modify their behaviour to suit the systems, that is,

to understand and configure them, while the practices of the energy providers remain largely unchanged. Even for the inhabitants who fit the described user profile and choose to cooperate with the systems, there remains the possibility that the financial gain of using the systems would not be worth the time they need to invest in it. Additionally, the emphasis on choice has a very significant consequence that remains undiscussed in the three reports: inhabitants might make the “wrong” choice. Depending on the nature of the choice, this failure may lead to undesired, or unintended consequences for all actors involved in the system. The possibility of choice does not solve problems by itself, and even more critically, it places the responsibility for the wrong choices on the inhabitants using a narrative of consumer empowerment (Ottinger 2010).

Within the three reports, the users are also described as gendered. Anne-Jorunn Berg has criticised a much earlier incarnation of the smart home as “unlikely to initiate any developments that would substitute or save time in housework” (Berg 1994, 312). While acknowledging the impossibility of determining the use of technology solely from design, she observes the gendering of design in the smart home. She argues that the smart home is “gendered in what it leaves out—its lack of support for changes in the domestic sexual division of labour” (312). This argument finds support in the three reports.

The scenarios in the reports discuss running dishwashers and washing machines as examples of housework. The proposed projects make these tasks more energy efficient by scheduling them at different hours. The HGI report does not use any pronouns or names, the activities it describes are performed by “the customer” or “the user”. The majority of the FINSENY report follows the same pattern, but the scenarios that do include a user always use the masculine pronoun, and women are entirely absent. The pronoun “he” is used in multiple descriptions to refer to the “home dweller” category (FINSENY 2011, 21, 27, 39). The SRS report differs from the previous two by providing names for the users that appear in their demonstrators. However, the division of labour remains clear as the report describes women engaged in housework and care work much more often than men. For example, Beda and Bertil’s daughter follows the updates on their health generated by the smart home (SRS 2011, appx. III, 15). Jimmy’s daughter helps him adapt to changes in the smart home systems after he suffers a stroke (appx. III, 19). Anna, in her role as Bertil’s “contact

person”, warns him to take a shower and hurry to the day centre using a speaker placed in the home (SRS 2011, appx. III, 19). Julia ensures that the stove is off and the washing machine is running using her mobile phone, and goes shopping for food while both Anna and Anne pick up their or somebody else’s kids from school (appx. III, 20, 51). In other examples, Charlotte and Anna use the car pool while Lars remains unconvinced about its benefits and drives to work (appx. III, 32, 41), and Emil, aged 11, becomes an “agent against power waste” in an activity at his school (appx. III, 52). These constitute the majority of the characters that appear in the SRS demonstrators. As noted earlier, while the use of the technologies cannot be determined from their designs or descriptions, both still exert influence over how these technologies should be perceived by others, and what is expected of them as users.

The systems that make the home smart also describe the properties of the inhabitants that matter. The SRS report includes a detailed description of the site itself and the available tools for the surveillance of the inhabitants: sensors for monitoring water and electricity consumption, sensors for monitoring temperature in all rooms, near-field communication access control system in the outer door, light control system to monitor power consumption of all outlets (appx. III, 23). Within the home, the everyday behaviour of the inhabitants is monitored continually, and the traces of their behaviour, relieved from their context by the sensors, are transported to databases for storage and further analysis. By generating measures such as the mean, the median, and the standard deviation, the similarity and the difference of the recorded instances are quantified (e.g., the average smart household consumes 21 litres of water per day). The statistical measures enable both the evaluation of the performance of a population collectively, and the ordering of individuals separately. They are used to create the smart home inhabitant category, where differences are fitted to one another to form a whole.

In describing their users, these project proposals also illustrate how the systems begin to translate the traces of energy consumption into a form of knowledge that is used to judge the inhabitants. The translation proceeds according to values inscribed onto the systems, emphasising technological expertise, rational choice, and individual responsibility, which are used to position the ICT-based surveillance of energy use as integral to sustainability in these proposals.

Judging the User

The translation process continues with the communication of the quantified differences back to the inhabitants in the form of usage logs, recommendations, and visualisations. This communication allows the judging of the users, both by the users themselves and by others, on the basis of different criteria.

In the HGI report, the mechanisms for judging appear only when the user engages in an activity that involves energy consumption. For example, the use case “Alarms” describes a way of informing the user in case of “abnormal appliance consumption (e.g.: possibly indicating a fridge with a door left open)” (HGI 2011, 28), and the use case “Visualization of current energy and power data” describes how to “[i]nform the consumer of the energy implications of selecting different operational modes, in particular different washing machines cycles, before initiating the activity” (26). In both cases, the activities of the user are observed by the system, categorised according to the processes described earlier, and the final evaluation of the system is communicated back in the form of a judgement about whether the activity is abnormal, or with a list of consequences (“energy implications”) of consumption. The latter stops short of delivering a final judgement, but in the report it is followed by the statement “[s]tudies have shown that an energy saving of 10-15% could be achieved as a result”, implying that some users will choose to behave differently if the system judges the future activity as involving high energy consumption.

The SRS report broadens the scope of both the activities that can be judged and the judgements that can be delivered by stating that “[t]he overarching sustainability goals of SRS stipulate that it should be easy to do the right thing” (SRS 2011, 12). In the report, the right thing is assumed to be the same for all inhabitants, and the role of ICT in the district is interpreted as encouraging the doing of the right thing. One of the proposed projects, titled “A sustainable community”, illustrates the interpretation:

We propose a set of mechanisms for increasing the social belongingness of a sustainable lifestyle. By applying a point based system on the use of resources (such as water and electricity), bonus points of active choices of trans-

portation, and otherwise effective uses of green alternatives where others are available, this information can be used and presented in different ways. . . Furthermore publicly available top lists can be made available on a community web site, and also on public displays. (SRS 2011, appx. III, 47)

The proposal combines the data collected from the smart home with a categorisation of the activities by assigning scores and declaring winners. In the description, desirable behaviour is defined by the system designers, and the inhabitants are encouraged to participate in the vision with a promise of the public acknowledgement of their activities. Although this proposal only mentions “top lists”, the public acknowledgement of scores can also serve as a threat for those who do not score high. More broadly, “increasing the social belongingness of a sustainable lifestyle” implies that the proposed systems are charged with making some lifestyles belong in the district. Conversely, the lifestyles that do not rank high in the categorisation are made to belong less. The passage continues by motivating the need for such systems: “It is not enough that ‘it’s easy to do the right thing’—it must also be shown what that this [*sic*], how one can do it, and why (incentives)” (appx. III, 47). In the SRS report, the method that shows the inhabitants what the right thing is and how it should be done often takes the form of a visualisation. It is formed by processing surveillance data using statistical methods to categorise the inhabitants (“your energy consumption is below average”), and it provides a value judgement for the observers (“your energy consumption meets the sustainability goals of the city”). Starting from the traces of energy consumption, surveillance aids the creation of a social marker that can be reflected on, discussed, and connected to value judgements. A device used for these purposes appears in the demonstrator titled “Participatory Installation”:

Input for the installation are peoples’ environmentally related decisions, concerning registered [*sic*] on “smart cards”. . . Output of the installation will give some kind of indication of the environmental effects of people’s actions. This output will serve as feedback to participating people and visitors. Although individuals will be able to

identify their own contribution in the installation, contributions of others will not be recognizable. (appx. III, 50)

The demonstrator describes devices that can be placed in public spaces within the district to reveal the traces of consumption behaviour and the associated value judgements. In this case, the visual representation of the collected data is taken not only from the smart home, but also from the surveillance of other activities with the help of “smart cards”. The results are displayed to the inhabitants to invoke reflections, as exemplified by the youngest daughter in the opening quote of this paper, taken from the same demonstrator: “[I]t seems that our block is doing much better than my friend Milla’s block” (appx. III, 51). The mother verbalises the value judgement provided by the system with her reply: “That’s good news”. In the fictional scenario, the inhabitants reflect on the graphical representation of the collective consumption of the neighbourhood with the help of surveillance technologies. The exchange between the mother and the daughter illustrates the final step of the translation where the inscribed values are inserted into the social fabric of everyday life in the form of value judgements.

The system proposed in the “Participatory Installation” demonstrator sets up a way for inhabitants to monitor one another to encourage energy-saving behaviour. This type of activity where individuals are provided with surveillance tools to keep track of one another has been called lateral surveillance (Andrejevic 2005). In lateral surveillance, the populace is made responsible for monitoring itself, and everyone is “invited to become spies” (494) for their own good. Another example of lateral surveillance appears in the demonstrator titled “Educational Game” (SRS 2011, appx. III, 52). In the game, participants play the role of “special agents” who are responsible for reducing energy consumption in their homes, which is measured by connecting their mobile phones to the electricity meters in their homes. They are organised into teams, and the example from the report describes one team as fifth grade students in a school in the district:

The team competes with another team of agents located in another town. A successful player persuades everyone in the household to conserve as much electricity as

they can during the mission, which on most days takes place between 17.00 and 22.00. Throughout the mission, the game monitors electricity consumption in the participants' homes. The winning team is the one who made the combined largest relative decrease in energy consumption. (SRS 2011, appx. III, 52)

In this demonstrator, surveillance is used as the basis of a competition, and the participants are encouraged to align their behaviour with the goals of the system to overcome their competitors. During the "mission" hours, the players become responsible for the traces of consumption that originate from their smart homes. Once again, the consumption act is freed from its original context through surveillance, and translated into an orderable quantity that is associated with a player contributing to the performance of the team as a whole. The players in this educational game are placed under immense social pressure to either change the behaviour of other members of the household, or risk contributing negatively to the performance of the team.

The FINSENY report also includes a system that utilises social pressure to change the behaviour of the inhabitants. The project titled "Support Online Community" describes a system where all the tenants in the building are invited to participate in the judging of consumption behaviour:

The basic idea behind this use case is that energy conservation and a greener life-style can also be encouraged by facilitating the formation of an online community of like-minded people who can use the online platform to share ideas, experiences, know-how, or even participate in competitions. . . . [T]he building can comprise ICT infrastructure that collects information / statistics of the various apartments, possibly also including the bills, and exports this information to a web front-end where it can be charted and visualized. This can allow the tenants to monitor the historical trend of their apartment in terms of energy and bills paid and also compare it against apartments of a similar profile. (FINSENY 2011, 42)

Although this system is meant to be accessed from private devices such as smartphones and personal computers, it shares the same goal

as the public and semi-public systems described in the SRS report. It provides the inhabitants with methods of comparing themselves to others on the basis of the categories defined by the system designers, and to form value judgements about consumption behaviour. The quantified differences in the monitored traces are provided to the inhabitants who are expected to identify other apartments with “similar profiles” against which to compare themselves. They are also encouraged to communicate with each other about the similarities and differences in the traces of consumption behaviour. The “support online community” provides support for the construction of the smart home inhabitant category by the inhabitants themselves and allows the fitting of the differences within the group to one another.

7.4 Conclusion

In this paper, I examined a shared vision of a future society from three different ICT design and development projects where inhabitants are made individually responsible for living sustainably, and I argued that the surveillance of energy use is positioned as integral to this future with the help of ICT. Underlying the vision, I identified a translation process that captures the traces of the inhabitants’ lives, classifies them according to different criteria of sustainable living, and returns them to the tapestry of everyday life to convince the users to behave differently. The systems described in the design documents make the differences in populations commensurable with each other, and aid the creation of new categories such as the smart home inhabitant and the sustainable city district. In the discourse of these documents, surveillance translates the traces, and the translations exert new pressures on existing power relations.

The examples of ICT-driven surveillance I analysed present an ordering of society where surveillance and sustainability are linked to one another. Moreover, echoing the findings of Graham and Wood (2003), they support processes of individualisation and emphasise consumption. A significant shortcoming of this vision of a future society is the invisibility of the institutional and systemic causes of the problems addressed under the banner of sustainability. Even if the necessary change is defined merely as reduced energy consumption and carbon emissions, holding the inhabitants individually responsible for

effecting this change is not only unlikely to succeed, but also risks obscuring a deeper understanding of the problems.

As the design documents illustrate, ICTs do not materialise from thin air: considerable resources are expended to propose, design, and develop them. If they become operational, they require constant maintenance, and continue to consume energy. Any proposal to use ICTs to create sustainable or “green” futures must weigh such costs against the theoretical gains. Additionally, the categories and the classifications that the proposed systems use to describe all users privilege a particular type of user that is already privileged within the context of ICT use. As Oudshoorn, Rommes, and Stienstra (2004) note, design practices to create ICTs for everybody remains “an inadequate strategy to account for the diversity of users”(54).

If the challenge to be tackled is reducing energy consumption, or creating more sustainable ways of living, the answer does not necessitate the development of new technologies. ICT does not need to be everywhere, and it does not need to be involved in solving every problem. Sometimes ICT might be the wrong answer. Even in cases where ICT development simply has to be involved in sustainability initiatives, it can be used for purposes other than the surveillance of inhabitants. For example, ICT can be used to understand how other technologies in residential spaces can be constructed differently to last longer, or to waste less energy, without falling back on the common solution of monitoring the users. Finally, if ICT simply has to be used for surveillance, that surveillance can be aimed at larger institutions rather than the individual inhabitants. It can provide ways for the inhabitants to hold accountable companies that develop wasteful technologies, or energy suppliers that attempt to classify and sort their customers using smart metering schemes.

Given the role of ICT design documents as plans for constructing technological systems, their descriptions of the envisioned users serve as a method of discursively constructing new categories such as the smart home inhabitant. In the texts, these categories serve to demonstrate to the reader how values related to technological expertise and rational choice fit into the visions of sustainability. Clearly, the consequences of living with these systems cannot be determined by only examining the design documents: once the systems are operational, they are likely to be used, resisted, and configured in ways never intended by their designers. However, their influence on the agencies of

the inhabitants, and how they encourage certain uses and discourage others, remains worthy of attention, because the development processes of these technological systems are much more open to change before they assume their material forms.

References

- Albrechtslund, Anders, and Louise Nørgaard Glud. 2010. "Empowering Residents: A Theoretical Framework for Negotiating Surveillance Technologies." *Surveillance & Society* 8 (2): 235–250.
- Andrejevic, Mark. 2005. "The Work of Watching One Another: Lateral Surveillance, Risk, and Governance." *Surveillance & Society* 2 (4).
- Berg, Anne-Jorunn. 1994. "A Gendered Socio-technical Construction: The Smart House." In *The Social Shaping of Technology*, Second Edition, 1999, edited by Judy Wajcman and Donald MacKenzie, 301–313. Open University Press.
- Bijker, Wiebe E., and John Law, eds. 1992. *Shaping Technology / Building Society: Studies in Sociotechnical Change*. Cambridge, MA: The MIT Press. ISBN: 0262521946.
- Bowker, Geoffrey C., and Susan Leigh Star. 2000. *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: The MIT Press. ISBN: 0262522950.
- Darby, Sarah. 2010. "Smart Metering: What Potential for Householder Engagement?" *Building Research & Information* 38 (5): 442–457.
- Donaldson, A., and D. Wood. 2004. "Surveilling Strange Materialities: Categorisation in the Evolving Geographies of FMD Biosecurity." *Environment and Planning D* 22 (3): 373–392.
- Dubbeld, L. 2006. "Telemonitoring of Cardiac Patients: User-Centred Research as Input for Surveillance Theories." In *Theorizing Surveillance: The Panopticon and Beyond*, edited by D. Lyon, 182–205. New York: Routledge.
- European Commission. 2009. "Communication from the Commission on Mobilising Information and Communication Technologies to Facilitate the Transition to an Energy-efficient, Low-carbon Economy." COM (2009) 111 (12.03). Accessed January 17, 2013. http://ec.europa.eu/prelex/detail_dossier_real.cfm?CL=en&DosId=198049.

- European Parliament. 2006. "Directive 2006/32/EC of the European Parliament and of the Council of April 2006 on Energy End-use Efficiency and Energy Services and Repealing Council Directive 93/76/EEC." *Official Journal L* 114 (66). Accessed January 17, 2013. <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:32006L0032:en:NOT>.
- FINSENY. 2012. *Future Internet for Smart Energy*. Accessed January 17, 2013. <http://www.fi-ppp-finseny.eu>.
- . 2011. *Future Internet for Smart Energy – Smart Buildings: Use Cases Specification*. Technical report. FLICT-2011-285135 D4.1 v1.2. FINSENY.
- Foucault, Michel. 1977. *Discipline & Punish: The Birth of the Prison*. Second edition, 1995. Translated by Alan Sheridan. Vintage. ISBN: 9780679752554.
- . 1978. *The Will to Knowledge: The History of Sexuality Volume 1*. New Ed, 1998. Translated by Robert Hurley. Penguin. ISBN: 0140268685.
- Gilliom, John. 2001. *Overseers of the Poor: Surveillance, Resistance, and the Limits of Privacy*. Chicago, IL: University Of Chicago Press. ISBN: 0226293610.
- Graham, Stephen, and David Wood. 2003. "Digitizing Surveillance: Categorization, Space, Inequality." *Critical Social Policy* 23 (2): 227–248.
- HGI. 2012. *HGI Homepage*. Accessed January 17, 2013. <http://www.home-gatewayinitiative.org>.
- . 2011. *Home Gateway Initiative – Use Cases and Architecture for a Home Energy Management Service*. Technical report. HGI-GD017-R3. HGI.
- Lyon, David. 2007. *Surveillance Studies: An Overview*. Cambridge: Polity. ISBN: 074563592X.
- Maki, Krystle. 2011. "Neoliberal Deviants and Surveillance: Welfare Recipients Under the Watchful Eye of Ontario Works." *Surveillance & Society* 9 (1/2): 47–63.

- Murakami Wood, David, ed. 2006. *A Report on the Surveillance Society*. For the Information Commissioner by the Surveillance Studies Network. Accessed January 17, 2013. http://www.ico.gov.uk/upload/documents/library/data_protection/practical_application/surveillance_society_full_report_2006.pdf.
- Ottinger, Gwen. 2010. "Constructing Empowerment Through Interpretations of Environmental Surveillance Data." *Surveillance & Society* 8 (2): 221–234.
- Oudshoorn, Nelly, and Trevor Pinch, eds. 2005. *How Users Matter: The Co-Construction of Users and Technology*. Cambridge, MA: The MIT Press. ISBN: 9780262651097.
- Oudshoorn, N., E. Rommes, and M. Stienstra. 2004. "Configuring the User as Everybody: Gender and Design Cultures in Information and Communication Technologies." *Science, Technology & Human Values* 29 (1): 30–63.
- SRS. 2012. *Stockholm Royal Seaport – A World-Class Environmental Urban District*. Accessed January 17, 2013. <http://stockholm-royalseaport.com>.
- . 2011. *Stockholm Royal Seaport – Smart Communication: Final Report*. Technical report. Smart ICT. Accessed January 17, 2013. <http://smartict.swedish-ict.se/files/2012/06/SRS-SC-pre-study-final-report-2011-05-25-1.00.pdf>.
- Suchman, Lucy. 1994a. "Do Categories Have Politics?" *Computer Supported Cooperative Work (CSCW)* 2 (3): 177–190.
- . 1994b. "Working Relations of Technology Production and Use." *Computer Supported Cooperative Work (CSCW)* 2 (1-2): 21–39.
- . 2002. "Practice-based Design of Information Systems: Notes from the Hyperdeveloped World." *The Information Society* 18 (2): 139–144.
- Woolgar, Steve. 1991. "Configuring the User: The Case of Usability Trials." In *A Sociology of Monsters: Essays on Power, Technology and Domination*, edited by John Law, 58–99. New York: Routledge.

Chapter 8

Changing Behaviour to Save Energy: ICT-Based Surveillance for a Low-Carbon Economy in the Seventh Framework Programme*

Baki Cakici^{†,‡}, Markus Bylund[‡]

Abstract

In research and development of information and communication technologies for sustainability, there is a strong belief that human behaviour can be monitored at the individual level to generate different signals, and that these signals can be used to influence individuals to behave differently. We analyse Seventh Framework Programme policy documents published by the European Commission, and descriptions of research projects granted funding from it, to highlight the uncritical development and application of surveillance technologies to change human behaviour. We argue that EU-financed

*Unpublished manuscript.

[†]Stockholm University, DSV, 164 40 Kista, Sweden

[‡]Swedish Institute of Computer Science (SICS), 164 29 Kista, Sweden

projects dealing with sustainability and information and communication technology use models of social change that have been widely criticised as unlikely to lead to substantial changes in resource consumption. Additionally, we show that these texts discuss only the potential positive effects of technological surveillance, but neither acknowledge nor require the handling of the potential negative effects of surveillance.

8.1 Introduction

The assumption that energy consumption behaviours can be altered by prompts from computer-supported systems is common in research and development of information and communication technologies (ICTs) for sustainability. Most often, it is assumed that behaviour can be monitored at the individual level to generate different signals, and that these signals can be used to influence individuals to behave differently.

Shove (2010) identifies this assumption as belonging to a dominant paradigm of understanding social change as a combination of attitudes, behaviour, and choices within environmental policy. In earlier work, she states that “policies designed to promote sustainable consumption are generally founded upon an extraordinarily narrow understanding of human behaviour” (Shove 2005, 111) and argues that this model overemphasises the influence of individual behaviour and individual choice while ignoring many other factors involved in resource consumption. In a similar direction, Dourish (2010) details the limitations of systems that frame problems of sustainability as issues of personal choice, and argues for a broadening of the theoretical framework to allow systems to be developed for scales other than the individual.

Focusing on systems that aim to persuade their users to behave differently, Brynjarsdóttir et al. (2012) state that such systems are based on a narrow vision of sustainability as resource optimisation, and argue that many such systems tend to provide technological solutions to social problems. These systems reflect specific definitions of sustainability chosen by the designers, and “their framing of sustainability as optimization of a simple metric places technologies incorrectly as objective arbiters over more complex issues of sustainability” (947).

Other studies within human-computer interaction highlight different targets for intervention aside from individual choices. For example, researchers emphasise the importance of studying everyday interactions (Pierce, Schiano, and Paulos 2010; Strengers 2011), and recommend designing systems to actively promote trust among users (Rodden et al. 2013). Additionally, Dillahunt et al. (2009) show that research on sustainability technologies tends to focus on affluent households which make up a narrow demographic, and neglects marginalised groups such as low-income households. The authors also demonstrate that these households engage in energy saving behaviours even when there are no financial incentives.

To these critiques we add an analysis from the perspective of surveillance studies. As DiSalvo, Sengers, and Brynjarsdóttir (2010) point out, it is important to engage with questions about “who gets to decide what change should happen and how, whose needs are met and whose values matter in the end” (1981) when designing ICTs for sustainability. We argue that insights from the field of surveillance studies are well-suited to considering such questions. Ball, Lyon, and Haggerty (2012) describe the field as follows:

The contribution of surveillance studies is to foreground empirically, theoretically and ethically the nature, impact, and effects of a fundamental social-ordering process. This process comprises the collection, usually (but not always) followed by analysis and application of information within a given domain of social, environmental, economic, or political governance. (1)

Surveillance systems are used to categorise and classify, and questions about who is included in which categories, and who remains excluded follow from their use. Lyon (2007) states that surveillance depends on “modes of categorizing populations and treating people differently according to socio-economic status, ethnicity, gender, region, age, and so on” (177). These categorisations may benefit their subjects unequally, and as Gandy (2006) argues, the differential treatment may become discriminatory.

Surveillance systems and the classifications that they rely on are not neutral. The key point is that there are moral and political consequences of classification (Bowker and Star 2000, 324). Classifications

that divide information into different categories are affected by the contexts in which these systems are designed in, and they reflect the worldviews of their designers. Cakici (2013) discusses issues of classification, and the judgements that result from their use, in the design of ICT-based surveillance systems for smart homes.

In this paper, we are primarily concerned with the uncritical development and application of surveillance technologies. We describe how reliance on the attitude-behaviour-choice model and the use of ICT-based surveillance appear in the Seventh Framework Programme (FP7) policy documents published by the European Commission. With a total budget of 32.4 billion Euros (EC 2013a), FP7 is highly influential in steering research within European research institutions, and its prominence makes it a good starting point for a critique of ICT-based surveillance. We focus on Challenge 6, *ICT for a low carbon economy*, within the ICT theme of the Co-operation sub-programme. Based on an empirical investigation of FP7 policy documents and project descriptions that have been financed within the challenge, we argue that EU-financed projects dealing with sustainability and ICT use models of social change that have been widely criticised as unlikely to lead to substantial changes in resource consumption. Additionally, we show that these texts discuss only the potential positive effects of technological surveillance, but neither acknowledge nor require the handling of the potential negative effects of surveillance within everyday interactions.

8.2 Methodology

FP7 collects all research-related European Union (EU) initiatives (EC 2013a). It is composed of four sub-programmes: Co-operation, People, Capacities and Ideas. Co-operation is defined as the core programme of FP7. It aims to support collaborative research within and beyond the EU, and it is divided into ten themes: health; food, agriculture and fisheries and biotechnology; information and communication technologies; nanosciences, nanotechnologies, materials and new production technologies; energy; environment; transport; socio-economic sciences and humanities; space; and security.

Our analysis covers the ICT theme of the Co-operation sub-programme within FP7. Each call within Co-operation is associated with

a set of documents describing both the content of the call and the application procedures. The primary policy reference in our analysis is the *Cooperation Work Programme: Theme 3 – ICT* document, which has been revised six times over the course of FP7, initially appearing in 2007, and with the most recent revision in 2012. Three of these revisions (in 2007, 2009, and 2011) introduce major changes to the challenges and objectives listed in the document, such as renamed or merged items, whereas the intermediate years (2008, 2010, and 2012) bring only minor modifications. The ICT calls are broken into smaller categories called *Challenges*. Each challenge is further divided into sub-categories called *Objectives*, and these use a consistent numbering scheme throughout FP7. For example, the label *ICT-2009.1.4* refers to the fourth objective within the first challenge of the 2009 version of the policy document.

In our analysis, we used the three major versions of the policy document, and identified the objectives that are of interest for each challenge. To choose the relevant objectives, we searched primarily for descriptions of surveillance targeting humans, that is, systems which monitor, sort, and classify properties of humans and human behaviour directly or indirectly. We conducted the analysis by marking all references to the hypothesis that individual behaviour can be shaped through a combination of surveillance and feedback mechanisms in the collected documents. From this material, we compiled a collection of expected societal impacts of these types of technologies, in particular those that referenced optimisation and rationalisation. We also compiled references to different surveillance activities and expected behavioural changes. From the three versions of the policy document, we selected 15 objectives for further investigation using the same criteria. These objectives were distributed over four challenges which are modified and moved during the three revisions (see Appendix for a detailed list). Using the search engine of the Community Research and Development Information Service (CORDIS) (EC 2013b) we identified 173 projects which were financed under the objectives we selected.

To limit our material for this paper, we chose to focus on a single challenge originally titled *Challenge 6: ICT for Mobility, Environmental Sustainability and Energy Efficiency*, and renamed in 2011 to *Challenge 6: ICT for a low carbon economy*. Out of the 15 objectives we had originally picked, five fell under this challenge, and they contained a

total of 55 projects. In terms of the number of financed projects, *Challenge 5: ICT for Health, Ageing Well, Inclusion and Governance* included a comparable amount of projects. Given the two options, we chose to proceed with Challenge 6 in our analysis primarily due to our greater familiarity with ICT within the domain of sustainability.

We retrieved the project descriptions for the 55 projects from COR-DIS. Our analysis is based on these project descriptions as well as the policy documents described earlier. In some cases where we required more detail or found the descriptions insufficient, we also consulted the web pages and the deliverables produced by the projects, but due to the high number of projects under consideration, this has been the exception rather than the rule.

8.3 Analysis

ICT for Surveillance

Within the policy documents, ICT is described as a tool to monitor different phenomena, to classify them, and to act upon those classifications. For example, the 2009 work programme states that “[t]he power grid needs new ICT-based monitoring and control systems to take on its growing complexity and distribution” (EC 2008, 62). In this statement, the value of ICT is in its capability to monitor complex exchanges as they happen. The same idea also appears in a different context in one of the objectives within the 2011 work programme:

Home energy controlling hubs that will collect real-time or near real-time data on energy consumption data from smart household appliances and enable intelligent automation. (EC 2010, 75)

The role of the technologies remain similar in the home context. Both examples begin from traces of energy consumption. In the first case, ICTs control power grids by monitoring activities of different sources that produce or require energy, and intervene upon that process based on the categorisations they produce. In the second case, ICTs are used to monitor the energy consumption of inhabitants as they use different appliances, and to categorise the activities to coincide with different automation schemes.

The projects financed by the work programmes within Challenge 6 retain the same role for ICT. For example, the SmartHG project gathers “real-time data about energy usage from residential homes and exploit[s] such data for intelligent automation pursuing two main goals: minimizing energy usage and cost for each home” (SmartHG 2012). Similarly, the FIEMSER project aims to build “a monitoring and control system to optimize in near-real time the local generation-consumption matching” (FIEMSER 2013), and the Adapt4EE project “treat[s] occupants as the central reference point” and “analys[es] occupancy behaviour (presence and movement)” (Adapt4EE 2011).

The term *optimisation* appears regularly in the energy domain as a challenge that can be solved using ICT. In optimisation problems, an ongoing process considered sub-optimal is improved by monitoring and categorising. The policy documents request the optimisation of different activities, such as “generation–consumption matching” (EC 2008, 66), the “dynamics of energy supply and demand” (EC 2010, 74), and the “use of energy” (79). Although the term is commonly used, a clear definition of what would be considered optimal in these contexts is not provided. The variables that should be considered for the cases of optimisation, and the evaluation criteria by which the optimisations would be measured are also not described. The texts focus on demand and consumption, hinting that the variables to be optimised are found at the consumer level. Statements such as “behavioural changes in the society at large” (EC 2008, 67) point in the same direction as well. This statement on behavioural change is noteworthy, and we explore this issue further in the next section.

Moving from policy documents to the project descriptions, the term optimisation continues to occur frequently. Although in some cases the term remains ambiguous, for example in the case of the project description for IDEAS (“a neighbourhood energy management tool to optimise energy production and consumption” (IDEAS 2012)), in other cases the projects provide different targets for optimisation which entail different forms of evaluation. For example, the SmartHG description states that it aims to “minimise the home energy bill and usage (local optimisation) with respect to a given price policy computed to attain global (grid level) optimisation” (SmartHG 2012).

In more general terms, the policy documents suggest several optimisation processes which aim to shape consumers’ energy consump-

tion using ICT-based surveillance. This basic model is then reflected by numerous projects, all of which propose to implement some variants depending on what type of surveillance is used and which mechanisms are considered to shape energy consumption.

Changing User Behaviour

In their roles as surveillance tools, ICTs are used in attempts to change user behaviour. In the policy documents, this change is described as raising user awareness, and aiding users in decision making. Under the *ICT for Energy Efficiency* objective from the 2009 work programme, the aim of the projects are described as producing “[i]ntuitive user interfaces that help end-users save energy while maintaining the desired comfort levels” (EC 2008, 66). The intuitiveness of interfaces are emphasised as important properties for saving energy, and discomfort is positioned as the opposing force that must be avoided. The possibility of creating discomfort by introducing new ICTs into the lives of the inhabitants is preempted by this statement. The same work programme provides a little more detail on what users might require in a different objective, *Environmental Services and Climate Change*: “Projects should be driven by the possibility for a range of users, including non ICT-skilled users, to plug-in their own use cases and get access to customised information and decision support.” (67) These users, now positioned as information seekers, provide their needs to the ICTs and receive advice from them.

Although the above statements hint at an aim to change behaviour, the policy documents themselves rarely mention behaviours explicitly, except in the case of driving, where projects are expected to develop “[n]ew tools, systems and services supporting energy-efficient driving and driver behaviour adaptation” (EC 2010, 80). The intent to alter behaviour using ICT is much more apparent in the project descriptions. The ENERSip project provides a typical example: “The outcome of the adoption of ENERSip will allow setting new behavioural patterns in the society and reduce overall intense economic dependence on energy” (ENERSip 2013). Similarly, the FIEMSER project proposes “solutions based on a rational consumption of energy, local generation and an increase in the consciousness of the building owners towards their energy consumption habits” (FIEMSER 2013). Both projects also echo the aim set by the work programmes to build in-

tuitive interfaces while maintaining the desired comfort levels using the same words as the programme. Some projects propose different mechanisms for promoting behavioural change such as financial incentives (E-Price 2013), mixed reality technologies (IDEAS 2012), as well as gaming and social networking (EEPOS 2012).

The project descriptions construct a chain of causality beginning from ICT-based surveillance, continuing with behavioural change, and concluding with energy savings or reduced CO₂ emissions. The tension in this chain is most visible when we consider the set of decisions that must be made to construct these technologies. For example, what is required to set new behavioural patterns? Which patterns would be preferred, and how would those who do not behave in accordance with the patterns react to it? Considerations along these lines are not present in the texts we have examined.

Consequences

The primary outcome that Challenge 6 aims at is using ICT to “assist in reshaping the demand side of our energy-dependent society, reducing energy consumption, and subsequently CO₂ emissions” (EC 2010, 74). In addition to environmental goals, the work programmes mention different bureaucratic requirements in reference to international commitments, and economic goals such as maintaining competitiveness, reinforcing industrial and technological positions, and establishing an open energy market (EC 2008).

More specifically, the programmes refer to the need to aid communication and the transfer for information as goals that should be met by the projects. For example, the 2011 document describes part of its focus as the “[f]uture electricity distribution grids applying seamless communications systems to increase the connectivity, management, automation and coordination between suppliers (including renewable sources), consumers and networks” (EC 2010, 74). In this statement, ICT makes different contexts more amenable to management by connecting them. Consumers and suppliers of energy are linked by the flow of information, and the connection reshapes the *demand side* of society.

The projects financed under different objectives within the challenge begin with these goals and provide more specific solutions. The GreenCom project describes a smart grid that provides “[a] value

based demand control based on individual consumer contracts with attractive tariffs, reward/penalty clauses and other elements will allow intelligent energy demand management and control” (GreenCom 2012). In line with the goals of the work programme, in this project ICT allows the consumers to enter into contracts where they can be rewarded or penalised, aiding the management of their demands on the supplier side.

Surveillance technologies allow the projects to begin from residential contexts and link them to energy goals as well. The SmartHG project states its expected outcome as: “gathering real-time data about energy usage from residential homes and exploiting such data for intelligent automation pursuing two main goals: minimising energy usage and cost for each home” (SmartHG 2012). In contrast to the diverse ICTs proposed for saving energy, the home itself is represented as a homogeneous space for energy consumption surveillance. The PEBBLE project extends the homogeneity further:

Through user-interfaces humans act as sensors communicating their thermal comfort preferences to the PEBBLE system, and in return the PEBBLE system returns information with the goal of enhancing energy-awareness of the users. The generality of the proposed methodology affords a universality that transcends regional, behavioral, environmental or other variations. (PEBBLE 2012)

In this statement, ICT is used to abstract away the irrelevant properties of the inhabitants, and to construct their involvement as a vehicle for voicing input. By discarding local contexts, it becomes possible for the system to transcend variations and to offer a universal solution.

Not all projects describe their users primarily as sensors. For example, the INERTIA project offers “fine grained control [...] while also protecting privacy and autonomy on the local level, fully respecting prosumers preferences and needs.” (INERTIA 2012). The inhabitants are described using the term prosumers, which this project defines as consumers who produce energy using the proposed technologies, although there is also a recognition that the systems come into contact with users who have different needs and different preferences. These users are made into a coherent group in this brief statement by referring to how they are expected to interact with the system. At the

same time, the designers recognise the potentially different needs and preferences of the users, although the specifics of that difference are not defined in the descriptions.

In general, the work programmes and the project summaries reveal very little concern about the heterogeneity of the stakeholders involved in their visions. Goals pertaining to aiding trust, safety, and inclusion, or preventing privacy breaches, marginalisation, and exclusion are not voiced by these texts. The underlying assumption regarding outcomes is that energy savings and the reduction of emissions are bound to occur if the users cooperate with ICTs and individually make decisions compatible with the continued functioning of these systems.

8.4 Discussion

In the texts we have analysed, surveillance systems generate information by monitoring individuals and recording traces of their behaviour. These traces are then communicated back to the individuals in an attempt to change their behaviour. Depending on the goals of the project, the individuals are either provided financial incentives to convince them to consume energy differently, or they are taught to behave differently using feedback generated by surveillance systems. As we have stated previously, the assumption that behaviour can be altered by providing information presumes a theoretical stance that identifies the individual's outcome of deliberation as the cause, and the behaviour as the effect. The theory postulates that the individual can be informed using ICT, and the individual, who is presumed to be acting rationally by applying a cost-benefit analysis to decisions, is expected to change their behaviour after receiving information from the surveillance systems. We argue that there are two problems with assuming this theoretical stance when designing and developing ICTs for sustainability. First, the assumption that behaviour can be manipulated using information about energy consumption is questionable. Second, it neglects to address the negative societal effects of employing surveillance to influence behaviour while arguing for increased surveillance.

The assumption about behaviour change through consumption feedback is questionable primarily because it overemphasises individ-

ual choice while neglecting a whole range of other factors involved in shaping human behaviour. Information may be important in influencing how people behave, but there are a series of events between rendering information and behavioural change. First, the information must reach the user. This is a matter of both timing and presentation. The user must have the time required to perceive and interpret the information, and it must be presented in a way that allows the user to do so while in the midst of other ongoing events. In a home setting this accounts for high variation of activities, not all of which come with a large degree of spare time. Second, the user must possess the experience as well as the analytical expertise to interpret the information and to figure out how to act upon it. When targeting virtually everyone, ranging from children to elders, from formally uneducated to professionals and academics, from people that have spent their whole life in their current setting to people who just arrived from a life spent on the other side of the planet, the user population displays an extreme variation with regard to ways of interpreting and acting upon information. In addition to taking the time and possessing the expertise to make sense of the information, the user must also have an interest in doing so. Both the policy texts and many of the project descriptions from our analysis seem to assume that people have a general interest in changing their behaviour to save energy. Alternatively, the texts assume that financial or altruistic incentives can motivate people to take an interest. It is hard to find empirical evidence in support of this assumption. Even if people are informed and motivated to adopt new behaviour, a number of other factors may work against change. Established practices related to behaviour do not only constrain what type of behaviour is possible, they also help shape motivations to behave in certain ways. Social norms play an important role in affecting behaviour, as do physical constraints. One popular example in the energy saving domain is the possibility of automatically delaying the start of the laundry machine until early mornings when there is a surplus of electricity in the network, when it is cheap, or when the CO₂ emissions from electricity production are low. However, that might imply having to hang the wet laundry to dry in the morning instead of at night, which would conflict well-established practices of morning routines filled with activities like having breakfast, sending children to school, and commuting to work. In addition, the noise from the washing machine centrifuging during early morning hours is likely to

interfere with the common practice of sleeping at night. Therefore, the assumption that users will change their behaviour as a result of information pertaining to energy consumption being made available through an “intuitive” user interface, without careful consideration of other factors that affect this process, can in general be considered flawed.

The definition of the term *optimal* in the materials we have examined reflects the viewpoints of their designers and stakeholders. To identify optimal resource consumption, surveillance systems are proposed for monitoring ongoing events, and for classifying the recorded traces. However, categorisations and classifications are not neutral; they encode the preferences of their designers, show what their designers have deemed should be visible, and hide what they have deemed not worth seeing. For example, a shortcoming of the financial incentives is made apparent through the emphasis on the proposed exchange of lowering energy consumption to obtain financial benefits. The exchange contains an internal contradiction: the saving of energy motivated by financial gain can only be an optimal behaviour if it generates more income per unit time than other methods of wealth generation such as salaries. For those with higher incomes, the time and attention spent on acquiring the best deal from the system is less likely to be higher than their current income per unit time. On the other hand, those with lower incomes, those who would benefit more financially from the financial incentives, are only able to participate if they invest time and acquire the technological competence required to operate the systems. For those who are able to learn to operate the system, interpret its results, and make the necessary changes, the system grants certain benefits such as lower energy costs. Marginalisation becomes visible at this level, where those who are not able to interpret the system become unable to enjoy its benefits, primarily expressed as a financial gain in the documents (e.g., they pay more because they do not know when it is cheapest to run the washing machine). Thus, the technologically and financially privileged can afford to ignore the system and disregard the disadvantages of lost profit while the underprivileged are marginalised further. By targeting the overlap between technological competence and positive financial gain, the idea to provide financial incentives to motivate behavioural change minimises the number of users likely to benefit from the system.

The focus on individual behaviour and choice by different EU

projects is questionable on its own. Given the limitations of the model, the possibility of change using such frameworks seems small. However, the problem is not limited to model choice. After all, different disciplines use different methods and theories to understand the world, and when those methods do not yield results, they may be revised, or other methods can be used in their place. The problem with using these particular theories, of providing information to change behaviour, is that it is done with very little attention to possible negative social consequences. Large sums of European research funding in ICT design and development, generated largely through taxation (EC 2013c), are expended on technologies that may not contribute with any positive effect for society, but also induce negative social effects, partially due to their uncritical use of surveillance. We argue that the lack of focus on possible negative social effects like marginalisation, exclusion, privacy breaches, and reduced trust raised by these quantifying methods, both at the policy level and at the project level, points to a general lack of engagement with these issues among ICT projects financed under FP7.

By opening the judgements of the designers to critical debate, concerns can be identified, voiced, discussed, and ultimately resolved through negotiation, preferably as an integral part of the design process. In this view, conflict is assumed to be a necessary component of social negotiation, but debates about assumptions can aid in minimising structural conflict that deeply advantages some subjects over others. If the assumptions are not challenged early in the design process, it becomes possible for the failure of the system to be blamed on the non-cooperative behaviour of the inhabitants. The system goes unused, and the end result does not contribute towards sustainability goals.

8.5 Conclusion

In summary, our analysis shows that the belief that human behaviour can be monitored at the individual level to generate different signals, and that these signals can be used to influence individuals to behave differently is well represented in both FP7 policy documents and descriptions of funded FP7 projects related to sustainability and energy efficiency. However, this is promoted and proposed almost entirely

without references to the great number of other factors that determines human behaviour. Failing to do so, the texts provide little evidence to the fact that the potential behavioural change caused by this chain of events holds any promise for positively affecting energy consumption.

Meanwhile, both FP7 policy documents and descriptions of funded projects present a view of human behaviour and society seen primarily from a quantifying, managerial position. This position privileges counting, tracking, and efficiency, but it fails to address potential negative effects introduced by the surveillance technology, such as marginalisation, exclusion, and discrimination. In the policy documents and the project descriptions, surveillance systems are used to assist in reaching sustainability goals, and the potential positive effects are highlighted without considering the undesirable consequences of surveillance for those that are to be monitored.

This is not to suggest that there is no potential in applying surveillance technologies for the purpose of achieving sustainability goals with regard to environmental and energy efficiency. This may still very well be the case. However, in order to realise that, the design space must be expanded to cover more factors involved in shaping human behaviour, and the judgements of the designers and policy makers must be voiced in a critical debate where concerns for social effects can be identified, analysed, and discussed.

Appendix

Projects	Objective	Name
12	ICT-2007.4.1	Digital libraries and technology-enhanced learning
9	ICT-2007.5.1	Personal health systems for monitoring and point-of-care diagnostics
13	ICT-2009.4.2	Technology-enhanced learning
10	ICT-2009.5.1	Personal health systems
10	ICT-2009.6.1	ICT for safety and energy efficiency in mobility
13	ICT-2009.6.3	ICT for energy efficiency
8	ICT-2011.1.3	Internet-connected objects
16	ICT-2011.5.1	Personal health systems (PHS)
8	ICT-2011.5.3	Patient guidance services (PGS), safety and healthcare record information reuse
12	ICT-2011.5.4	ICT for ageing and wellbeing
13	ICT-2011.5.5	ICT for smart and personalised inclusion
9	ICT-2011.6.1	Smart Energy Grids
14	ICT-2011.6.2	ICT systems for energy efficiency
9	EEB-ICT-2011.6.5	ICT for energy-positive neighbourhoods
17	ICT-2011.8.1	Technology-enhanced learning

The distribution of the 173 projects over the 15 selected objectives.

References

- Adapt4EE. 2011. *Projects : Search*. CORDIS. Accessed May 8, 2013. http://cordis.europa.eu/projects/rcn/100798_en.html.
- Ball, Kirstie, David Lyon, and Kevin Haggerty, eds. 2012. *Routledge Handbook of Surveillance Studies*. Routledge.
- Bowker, Geoffrey C., and Susan Leigh Star. 2000. *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: The MIT Press. ISBN: 0262522950.
- Brynjarsdóttir, Hrönn, Maria Håkansson, James Pierce, Eric Baumer, Carl DiSalvo, and Phoebe Sengers. 2012. "Sustainably Unpersuaded: How Persuasion Narrows Our Vision of Sustainability." In *Proceedings of the 2012 ACM Annual Conference on Human Factors in Computing Systems*, 947–956.
- Cakici, Baki. 2013. "Sustainability Through Surveillance: ICT Discourses in Design Documents." *Surveillance & Society* 11 (1/2): 177–189. ISSN: 1477-7487. <http://library.queensu.ca/ojs/index.php/surveillance-and-society/article/view/sustainability>.
- Dillahunt, Tawanna, Jennifer Mankoff, Eric Paulos, and Susan Fussell. 2009. "It's Not All About Green: Energy Use in Low-Income Communities." In *Proceedings of the 11th International Conference on Ubiquitous Computing*, 255–264.
- DiSalvo, Carl, Phoebe Sengers, and Hrönn Brynjarsdóttir. 2010. "Mapping the Landscape of Sustainable HCI." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1975–1984.
- Dourish, Paul. 2010. "HCI and Environmental Sustainability: The Politics of Design and the Design of Politics." In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, 1–10.
- EC. 2008. *European Commission. Work Programme 2009, Cooperation Theme 3, ICT – Information and Communication Technologies*. C(2008)6827, November 17.

- EC. 2010. *European Commission. Work Programme 2009, Cooperation Theme 3, ICT – Information and Communication Technologies*. C(2010)4900, July 19.
- . 2013a. *European Commission. CORDIS: FP7*. Accessed May 7, 2013. <http://cordis.europa.eu/fp7/cooperation>.
- . 2013b. *European Commission. CORDIS: Homepage*. Accessed May 7, 2013. http://cordis.europa.eu/home_en.html.
- . 2013c. *The Budget Explained – Where Does the Money Come From?* Accessed August 10, 2013. http://ec.europa.eu/budget/explained/budg_system/financing/fin_en.cfm.
- EEPOS. 2012. *Projects : Search. CORDIS*. Accessed May 8, 2013. http://cordis.europa.eu/projects/rcn/105854_en.html.
- ENERsip. 2013. *Projects : Search. CORDIS*. Accessed May 8, 2013. http://cordis.europa.eu/projects/rcn/93727_en.html.
- E-Price. 2013. *Projects : Search. CORDIS*. Accessed May 8, 2013. http://cordis.europa.eu/projects/rcn/93745_en.html.
- FIEMSER. 2013. *Projects : Search. CORDIS*. Accessed May 8, 2013. http://cordis.europa.eu/projects/rcn/93803_en.html.
- Gandy, Oscar H. Jr. 2006. “Quixotics Unite! Engaging the Pragmatists on Rational Discrimination.” In *Theorizing Surveillance: The Panopticon and Beyond*, edited by D. Lyon, 182–205. New York: Routledge.
- GreenCom. 2012. *Projects : Search. CORDIS*. Accessed May 9, 2013. http://cordis.europa.eu/projects/rcn/105697_en.html.
- IDEAS. 2012. *Projects : Search. CORDIS*. Accessed May 8, 2013. http://cordis.europa.eu/projects/rcn/105738_en.html.
- INERTIA. 2012. *Projects : Search. CORDIS*. Accessed May 9, 2013. http://cordis.europa.eu/projects/rcn/105543_en.html.

- Lyon, David. 2007. *Surveillance Studies: An Overview*. Cambridge: Polity. ISBN: 074563592X.
- PEBBLE. 2012. *Projects : Search*. CORDIS. Accessed May 9, 2013. http://cordis.europa.eu/projects/rcn/93732_en.html.
- Pierce, James, Diane J. Schiano, and Eric Paulos. 2010. "Home, Habits, and Energy: Examining Domestic Interactions and Energy Consumption." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1985–1994*.
- Rodden, Tom A., Joel E. Fischer, Nadia Pantidi, Khaled Bachour, and Stuart Moran. 2013. "At Home with Agents: Exploring Attitudes Towards Future Smart Energy Infrastructures." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 1173–1182*.
- Shove, Elizabeth. 2005. "Changing Human Behaviour and Lifestyle: A Challenge for Sustainable Consumption?" In *Consumption - Perspectives from Ecological Economics*, edited by I. Røpke and L. Reisch, 111–132. Elgar.
- . 2010. "Beyond the ABC: Climate Change Policy and Theories of Social Change." *Environment and Planning, A* 42 (6): 1273–1285.
- SmartHG. 2012. *Projects : Search*. CORDIS. Accessed May 8, 2013. http://cordis.europa.eu/projects/rcn/105541_en.html.
- Strengers, Yolande AA. 2011. "Designing Eco-Feedback Systems for Everyday Life." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, 2135–2144*.

Chapter 9

Detecting the Visible: The Discursive Construction of Health Threats in Syndromic Surveillance System Design*

Baki Cakici^{†,‡}, Pedro Sanches^{‡,§}

Abstract

Information and communication technologies are not value-neutral tools that reflect reality; they privilege some forms of action, and they limit others. We analyse reports describing the design, development, testing, and evaluation of a European Commission co-funded syndromic surveillance project called SIDARTHa. We show that the reports construct the concept of a health threat as a sudden, unexpected event with the potential to cause severe harm, and one that requires a public health response aided by surveillance. Based on our analysis, we state that when creating surveillance technologies, design

*Unpublished manuscript.

[†]Stockholm University, DSV, 164 40 Kista, Sweden

[‡]Swedish Institute of Computer Science (SICS), 164 29 Kista, Sweden

[§]Royal Institute of Technology (KTH/ICT/SCS), 164 40 Kista, Sweden

choices have consequences for what can be seen, and for what remains invisible. Finally, we argue that syndromic surveillance discourse privileges expertise in developing, maintaining, and using software within public health practice, and it prioritises standardised and transportable knowledge over local and context-dependent knowledge. We conclude that syndromic surveillance contributes to a shift in broader public health practice, with consequences for fairness if design choices and prioritisations remain invisible and unchallenged.

9.1 Introduction

The term syndromic surveillance describes a collection of methods within the field of public health surveillance. These are characterised by their use of secondary sources, referring to data collected for a variety of purposes only indirectly related to population health. For example, over-the-counter medicine sales, or records of ambulance dispatches from a hospital, or records of emergency room visits can be considered secondary sources because these data are not collected with the primary intent to perform public health surveillance. Using this formulation, syndromic surveillance is often positioned as an efficient use of already collected data (Mandl et al. 2004).

Within the discourse of syndromic surveillance, numerical representations of populations and the application of statistical analyses are essential to the task of constructing meaning from diverse sources of surveillance data. The use of these types of surveillance methods follow a long tradition of health discourses where individuals and populations are constructed as knowable and governable by establishing comparative systems to analyse their similarities and differences (Foucault 1977, 190). Equally important is the idea of disease as a collection of symptoms and signs observed by a *medical gaze* (Foucault 1973, 88-106). These discourses inform and guide the practice of syndromic surveillance.

We analyse reports describing the design, development, testing, and evaluation of a European Commission co-funded syndromic surveillance project called SIDARTHa (System for Information on Detection and Analysis of Risks and Threats to Health). We start with the assumption that information and communication technologies (ICTs) are never neutral. They necessarily reflect the viewpoints of their de-

signers, the influence of others involved in their construction, and consequences of the material limitations encountered during their development. With our analysis, we aim to identify how a syndromic surveillance discourse privileges or excludes specific types of knowledge and practices. We show that the discourse of syndromic surveillance constructs the *health threat* concept as a sudden, unexpected event with the potential to cause severe harm, and one that requires a public health response aided by surveillance. Based on our analysis, we state that when creating surveillance technologies, design choices have consequences for what can be seen, and for what remains invisible. Finally, we argue that syndromic surveillance discourse privileges expertise in developing, maintaining, and using software within public health practice, and it prioritises standardised and transportable knowledge over local and context-dependent knowledge. We conclude that syndromic surveillance contributes to a shift in broader public health practice, with consequences for fairness if design choices and prioritisations remain invisible and unchallenged.

9.2 Background

Partially due to its origins in communicable disease surveillance, most syndromic surveillance systems are engaged primarily in the detection of unexpected events. This type of detection is typically performed by first calculating the proportion of the number of individuals in the population suffering from a health condition to the total size of the population, called prevalence, which is then compared at different time points. Many statistical methods exist to perform much more complicated versions of this calculation, taking into account other factors such as season, geography, migration, etc., but almost all retain the notion of prevalence comparison.

The earlier applications of syndromic surveillance were characterised by the use of health-related data that precede diagnosis, and the continual monitoring of disease indicators to detect outbreaks of communicable diseases earlier than traditional methods (Mostashari and Hartman 2003; Buehler et al. 2004; Henning 2004). More recently, the definition of syndromic surveillance has broadened to include the monitoring of non-communicable diseases and other health conditions such as heat-related illnesses, injuries caused by tornadoes,

or respiratory illness after wildfires (Buehler et al. 2009).

Traditionally, these activities are performed by a central health authority that conducts surveillance for a particular geographical region, and reports to another institution that is tasked with monitoring a wider region, terminating at the health authority with the widest geographical coverage. For example, when a local clinic encounters a patient with a reportable communicable disease, it reports the case to the regional centre that collects the cases. The regional centres report to the national centre, and the national centre might report to an international health body, such as the European Centre for Disease Prevention and Control (ECDC).

A well-known example of syndromic surveillance is Google Flu Trends (Ginsberg et al. 2009; Google 2013), which uses search queries entered by users into Google as indicators of influenza outbreaks. The developers of the system have claimed that they can “accurately estimate the current level of weekly influenza activity in each region of the United States” (Ginsberg et al. 2009, 1012), although the estimates have not always been accurate (Cook et al. 2011; Butler 2013). While the system itself is well-known, it is highly unusual as it is run by a private company. The majority of syndromic surveillance systems are used by public institutions such as regional and national health authorities.

Many different types of ICTs are used for communicable disease surveillance. Bravata et al. (2004) provide a broad review of 115 systems and identifies those that provide syndromic signals. Other reviews focus on syndromic surveillance, and include more recent systems (Chen, Zeng, and Yan 2009; Cakici 2011). Additionally, the Triple-S project provides an extensive list of syndromic surveillance systems used within Europe (Triple-S 2013).

Challenges

Syndromic surveillance has been challenged both by public health practitioners as well as scholars from various disciplines studying public health and related domains. One of the earliest challenges was by Reingold (2003) who asked: “If syndromic surveillance is the answer, what is the question?” Reingold points to the scarce resources within the field of public health, and challenges those who call for the further development of syndromic surveillance to detail not how detec-

tion could be improved, but if public health response would be improved in any way if these systems are implemented. He emphasises the importance of building capacity within the public health departments instead of collecting and analysing data externally, an activity often performed by for-profit corporations in departments without academic partners. In a discussion on the difficulties of creating a nation-wide syndromic surveillance Mostashari voices concerns about the relationship between detection and response capacity:

We have 80 per cent of the nation covered but we really have nothing covered, because signals come and go, and an e-mail maybe is sent out, and there's no local capacity ... If you're doing analysis for a thousand different towns, villages, cities, whatever, every day you're going to find alarms. (USMI 2003; cited in Fearnley 2008b)

Similarly, Heffernan et al. (2004) state that although their syndromic surveillance systems have been beneficial in detection, they are “essentially ‘smoke detectors’ and call for prompt investigation and response if they are to provide early warning of outbreaks” (Heffernan et al. 2004). They continue to recommend that efforts for bioterrorism preparation should focus on hiring well-trained public health professionals with responsibilities that are broader than just bioterrorism.

Fearnley (2008b) analyses the debates surrounding the development and use several nation-wide syndromic surveillance systems in the U.S., and identifies a central conflict: the early syndromic surveillance systems were designed primarily for national bioterrorism preparedness, but the data collection and reporting were assigned to public health institutions tasked with maximising the health of human populations. These systems collected data from many different, non-traditional sources, and transferred them to health departments. However, the collected data were not immediately helpful, because “[m]ore information means more interpretive work, without certain benefits; and more detected events requires more epidemiological responses, without (at this point) the necessary epidemiological resources to undertake them” (Fearnley 2008a, 84). Fearnley points out that without the resources to interpret the incoming data, the newly developed systems were not helpful to epidemiologists.

In syndromic surveillance systems, collecting data is the key to tackling public health problems. All syndromic surveillance systems

are constructed with the assumption that gathering more data helps public health authorities make decisions. Analysing the practices that accompany ICT-based surveillance systems in public health, French (2009) identifies the assumption that more data are helpful as a consequence of *an immaterial conception of information*:

[... A]n immaterial conception of information, whether implicit or explicit, assumes that information signifies in the same way regardless of time or place. As a consequence of this assumption the significance of information-processing practices is minimized; the material diversity of such practices, and the effort required to unify them, is underestimated. (111).

As French argues, data collection also comes at a cost. It requires significant amounts of work to standardise data, to connect systems, and to ensure that data are transferred correctly from one place to another. Ignoring the cost of collection, or more broadly, the material dimensions of information, causes problems for ICTs. The very immediate consequence is that it requires more work. This is the work of maintaining and interpreting.

In their extensive study of classification, Bowker and Star point out that data entry is never a trivial task: it requires trained staff to perform, it is prone to mistakes with respect to the classifications being used, and there are always cultural variations in what is interpreted as worthy of recording and what is omitted (Bowker and Star 2000, 107). Additionally, they state that there is always a tension between the standardisation of lists centrally, and their use locally (139). Finally, they emphasise that “all category systems are moral and political entities” (324). As syndromic surveillance systems rely heavily on rigid classification schemes, they exhibit these characteristics.

Based on a study of medical records, Berg and Bowker also argue that data collection and recording methods influence how work is organised, and which practices are considered to be part of that work (Berg and Bowker 1997, 532). Similarly, emphasising data collection and information generation in public health surveillance modifies the definition of activities that belong to public health, and some types of expertise are prioritised over others. For example, when syndromic

surveillance is used, familiarity with developing and maintaining software becomes an important requirement for public health practice.

Given the costs and risks of immaterial conceptions of information, and the work standardisation, the question of benefits, or what there is to gain remains. The concept of immutable mobiles, introduced by Latour (1988, 227), provides a way to engage the question. Latour uses the term to refer to objects that are archivable and comparable regardless of their age, place of origin, and context of use. For example, an entry in a table listing the admissions to the emergency room can be considered an immutable mobile because, unlike the locally bound emergency room, it can be transported easily, compared to other descriptions, or even be organised in charts and statistics to construct an aggregated view of events that occur in emergency rooms. The complexity of the world is overcome by translating it into intelligible and stable objects, and for this reason, immutable mobiles are capable of exerting power over great distances. As their power is tied to their simplicity, the work of keeping the object stable is often invisible. However, the structures that maintain it, such as the layout of the rows and columns in a table, or the procedure for filling out the details of each entry must also remain constant across places and contexts for the object to remain immutable. Therefore, its immutability can only be accomplished by mobilising people and material resources, and by keeping them together. By categorising syndromes, collecting data, and creating graphs, the output of the peripheral surveillance systems are brought closer to each other, but constant work is required to uphold the categorisations.

9.3 Analysis

SIDARTHa is designed to be installed at health institutions to monitor sources of emergency data, including records of emergency dispatches and reports from emergency practitioners. The system has previously been tested in Göppingen, Germany and in Leuven, Belgium. According to the list of European syndromic surveillance systems provided by the Triple-S project, SIDARTHa is currently being used in the Autonomous Region of Cantabria in Spain and in the state of Tyrol in Austria (Triple-S 2013).

We chose to analyse the SIDARTHa project as it is one of the

largest syndromic surveillance systems developed in Europe, and it involves multiple partners from different member states. Similar projects which focus on single health institutions separately instead of one centralised system have also been developed in the U.S., such as RODS (Espino et al. 2004), but the organisation of public health authorities in U.S., divided between the federal and the state levels differs from the European context where the division is between member states and EU-wide public health efforts. Although the majority of syndromic surveillance systems in the last decade have been developed and used in the U.S., we consider ourselves more qualified to analyse a European project due to our previous experience in public health surveillance in an EU member state (Cakici et al. 2010).

Our analysis is based primarily on the seven project reports published by the SIDARTHa project between the years 2009 and 2010. They describe the process of design, development, testing, and evaluation of the SIDARTHa syndromic surveillance system. These reports form a suitable corpus for discourse analysis because the authors describe different stages of the project in detail, and they state their reasons for their decisions. The earliest reports focus on literature reviews and pre-studies, later reports describe the design and development of the system, and the most recent ones discuss testing and evaluation. This progression allows us to observe changes and shifts in the project and see how different ideas evolve throughout the design and development of the system.

We began our analysis by reading the seven project reports, and marking all sentences that make knowledge or truth claims related to the overall argument, state hypotheses, or present research questions. We continued by marking sentences which inform the reader about the existence of entities outside of the text. In this category we included any statements that refer to uncertainty, because such statements also provide insight into what is considered to exist only partially outside the text. We also marked all sentences that refer to a group of human subjects, because these statements show how the text constitutes subject positions that are relevant to its own argument. After this process was completed, we read only the marked sentences for each report, and tried to group them according to common themes based on either the topics they discuss, or the frequency of the words they use. We examined these themes in relation to knowledge, truth, and subject positions. We repeatedly asked the question “what is con-

sidered true in this statement?”, and wrote down our interpretations while keeping in mind the identified themes for each report, the collection of marked sentences, as well as the unmarked sentences referenced by the marked ones.

What is a Health Threat?

In the SIDARTHa reports, the aim of the project is stated as detecting *public health threats* (Ziemann et al. 2009, 1). The reports further describe their goal as helping public health authorities to become aware of health threats earlier, or even to prevent them. The scope of the term *health threat* within the project reports is very broad: it refers to diseases, both communicable and non-communicable (1), but it also refers to other events such as floods, heat waves, and even volcanic eruptions (Rosenkötter, Ziemann, et al. 2010). Additionally, it can refer to acts of *bioterrorism*, which most commonly refers to the intentional release of *biological agents* such as anthrax into the air in large quantities with the intent to harm others (Ziemann et al. 2009, 7).

From the perspective of syndromic surveillance systems, health threats first become visible when other surveillance systems capture their traces. For example, several people living in the same district might write about suffering from shortness of breath and digestive problems on social media platforms, or a person experiencing high fever and a sore throat may visit the emergency room where her visit is recorded in the admission logs. In both cases, the experiences of these individuals are recorded, and those records are then accessed by syndromic surveillance systems. For these systems, the health threat is best defined in terms of data and signals: “Early detection of public health threats in general relies on the components: timely and reliable data; the sensitivity, specificity; and timeliness of signals detected.” (Rosenkötter, Kahl, et al. 2010, 3)

Syndromic surveillance relies on the presence and the functioning of other surveillance systems. For example, the primary motivation for tracking ambulance dispatches is not to generate signals for syndromic surveillance. Ambulances are dispatched to retrieve patients, and the signals are used for a variety of purposes. Ambulance drivers may need to track their working hours, or hospital administrators may need to know how many ambulances are available to respond in

case of emergencies, etc. Every departure and arrival of the ambulance is logged for a multitude of reasons, and syndromic surveillance systems depend on these logs, but the logs themselves are not produced specifically, or at least not primarily, to support the practice of syndromic surveillance.

The SIDARTHa reports describe a system that can be used to collect emergency data in different health institutions in multiple European countries independent of one another, each with their own rules and regulations for gathering data. To form a more homogeneous set of records, the designers propose a coding standard to convert the *local* data to the SIDARTHa-compatible version. Organisationally, the *raw data* remains locally stored, as one of the aims of the system is to ensure that the collected data does not leave the collecting institution or regional authority.

The standardised data format is a string of numbers with different fields. The designers state that “[t]he minimum data set for syndromic surveillance must contain enough information to produce the number of cases per day for temporal syndromic surveillance” (Garcia-Castrillo Riesgo et al. 2009, 6). It includes seven variables: the anonymous case identifying number, date, geographic reference, syndrome, age, gender, and severity. The final three variables are called *modifiers*, because these are not essential to constructing the number of cases, but the information contained within them may be relevant for certain syndromes such as gastroenteritis (or stomach flu) in children or heat-related illness in the elderly. These particular variables are chosen based on a survey of availability within the countries that participate in the SIDARTHa project, and it is the minimal set of properties that all the participants are able to provide.

Using the standard, it is possible to represent each field using numbers, even when receiving data from different sources. For example, the case identification can be filled using the call identification code if the data are provided by an emergency medical dispatch centre, or it can be filled using the patient identification code if the provider is an emergency department. The date is converted to a series of unambiguous numbers by specifying the order that day, month, and year appear in the string. Geographic reference is more heterogeneous, the numbers contained in this field can refer to X and Y coordinates generated using the global positioning system (GPS), health zone codes, post codes, or community codes. Attempting to fit gender into the

form brings the ubiquitous problem of coding gender as a binary, and the transition from sex to gender in the system vocabulary. The original specification of the standard provides “0” for male and “1” for female, and the final implementation of the system changes the order and adds a third category, “unknown” (Garcia-Castrillo Riesgo et al. 2010, 33).

All of the variables we have described above pertain to the identification of persons and their locations. The actual work of describing syndromes occurs in the remaining two variables: syndrome and severity. The syndrome category requires the largest amount of work to convert into a number. Syndromes are divided into six categories: influenza-like, gastrointestinal, respiratory, intoxication, environment-related (heat-related), and unspecific. Since the system is designed to be compatible with many different sites, a multitude of conversion tables are provided to make possible the homogeneous coding of cases for any SIDARTHa system. The coding manual provides a series of tables to allow this conversion, where the codes of other, more established standards such as the International Classification of Diseases (ICD) can be converted to the SIDARTHa standard (Garcia-Castrillo Riesgo et al. 2009).

At the end of this conversion, the *health threat* is narrowed from a wide variety of uncontrollable events to a series of numbers that can be transmitted without change, and one that refers to the same thing regardless of context. Unlike the data collected in each institution, which carry marks of local practices particular to each case, the resulting string is constructed to function independently from its context. It is easily comparable and combinable across different databases. In Latour’s terms, this standardisation creates an immutable mobile that can be used for statistical analysis in any institution. Freed from their earlier contexts, these strings of numbers start expanding their contexts in a different direction to affect the world. They are collected by the SIDARTHa software, and classified using different detection algorithms.

The classification may be only of change over time, or it may also include a spatial component that connects the data points to particular places using GPS coordinates or postal codes. Although the algorithms vary depending on the input and the intent, the primary purpose of this analysis is to divide the data points into two groups: expected and unexpected. The data points that end up in the *expected*

category are not of high importance to the system designers. These may indicate low-risk diseases in low volumes, or a seasonal variation in an illness (implying that the change occurs every year). The designers of SIDARTHa are much more interested in the *unexpected* category. This category includes all the data points that have been marked by the system as deviating from the norm. These unexpected events are all potential health threats within the system.

The designers state that one of SIDARTHa's strengths is its flexibility: "[T]he SIDARTHa system can easily be adjusted to cover additional health threats, in this case the volcanic ash cloud with new syndromes such as traffic accidents and cardio vascular syndrome." (Pineiro et al. 2010, 10) This flexibility is not only due to the way the system has been constructed, but also because of how the discourse of syndromic surveillance describes health threats. In another report, they point out that "[o]ne important feature of syndromic surveillance is flexibility, which allows the generation and monitoring of syndromes according to suddenly emerging, potential health threats" (Rosenkötter, Kauh, et al. 2010, 3). In this definition, a health threat is a sudden event which has the potential to affect population health. The words "suddenly" and "potential" highlight the issue of the absent subject in the definition: For whom is the event sudden, or unexpected, and what is meant by the potential to affect? The concern about sudden events can be viewed from the perspective of public health authorities who would be expected to respond to an event. Suddenness can then be stated relative to how long it would take for the authorities to act (by contacting the patients, by visiting the clinics, by issuing public notices, etc.) when the health threat appears. Defining the limits of the word "potential" are much more difficult. For example, an influenza outbreak can sometimes spread rapidly in a population but result in only minor suffering as the infected experience the discomfort of a sore throat. However, it can also cause major illness in those who are infected, and even become life-threatening for those who already suffer from other conditions. Complicating the decision even further, the difference between the two is often not clear until its effects are experienced by individuals. The public health response is required to negotiate this conflict, and come to a decision about the limits of the potential of a disease to affect the population.

In the SIDARTHa reports, a particularly unusual event serves to highlight the vital need for the systems to be able to intervene in re-

sponse to health threats, even in cases of non-detection. Following the eruption of the Eyjafjallajökull volcano in Iceland in 2010, the project consortium developed a method to evaluate the potential public health impact of the ensuing ash cloud. The event is described as offering a suitable scenario for testing the system, although the authors also point out:

This report uses the term ‘volcanic ash cloud’ without determining if the ash cloud was a cloud or rather a contamination. Therefore it should be understood that the term ‘ash cloud’ used throughout this report is not to be understood as a scientific term. It should be further noted that the authors do not intend to give any prejudice on the question if there was any risk to health at all caused by the ash cloud as such. The intention of this rapid assessment was to test the capability of the SIDARTHa concept and pilot syndromic surveillance system to be timely adjusted for monitoring a suddenly occurring event potentially affecting health [emphasis in the original]. (Rosenkötter, Ziemann, et al. 2010, 3)

A sudden event, and its potential to affect health both appear in this justification for performing the ash cloud assessment using the SIDARTHa system. Establishing an event as a health threat occurs before the investigation begins, but the results of the investigation do not fully resolve the status of the event either. In the four SIDARTHa implementation sites in Austria, Denmark, Germany, and Spain, the system identifies no significant correlation between the volcanic ash cloud and the unusual signals of respiratory conditions reported during the same period. The authors state that “further in-depth analysis of case characteristics is necessary” (16), but also mention that their investigation, and the fact that it was possible to perform it at all, demonstrates the “flexibility of syndromic surveillance systems to be used for ad-hoc surveillance after suddenly occurring events” (16). In the booklet *Generic Public Health Preparedness in Europe*, a brief description of the SIDARTHa project is accompanied by several quotes from Dr Thomas Krafft, the scientific-technical coordinator of the project. Commenting on the volcano investigation, he states:

The European Centre for Disease Prevention and Control asked us to test out SIDARTHa during this time to see if there were any health impacts from the volcanic ash plume We found no increased demand for emergency care services. It is important to be able to distinguish between ‘real threats’ and ‘perceived threats’. (EU Health Programme 2011)

This division between real and perceived threats positions SIDARTHa, and by extension syndromic surveillance, as the arbiter of truth for public health practice in the implementation sites. The system does not establish a correlation, but it does not establish a definitive lack of correlation either. The operations of SIDARTHa are also perceptions, although they are perceptions of expertise, supported by advanced ICTs and public health authorities. The *real* in the quote can be interpreted as describing events that public health authorities should act upon to improve the health of a population; the *perceived* can also be interpreted as events that they should act upon to reassure the population that there is no health effect. In either case, the detection possibilities offered by SIDARTHa are geared towards shaping the type of public health response following an event.

In the beginning, the traces of different individuals are collected in various public health surveillance systems, and converted into the SIDARTHa standard to create a more uniform unit that is compatible with the different methods of statistical analysis. Up until this point, the traces of individuals shed their context until only the bare minimum remains. The contents of that bare minimum are defined by the SIDARTHa standard. After reaching that point, the work of inscribing a new context into the numbers begins. The first step is the statistical analysis, which constructs the objects that public health institutions can act on, followed by the visualisation of the results which show the traces in the unexpected category. The process that begins with the set of recorded traces and eventually becomes the evidence for the existence or the non-existence of the health threat depends on a long chain of translations between different systems and different mediums. When the authors assert their claim as “[e]mergency care demand shows a pattern allowing for detection of unusual aberrations from the expected demand” (Baer et al. 2009, 20), they refer to the

large web of surveillance systems that individually construct and link the patterns, detections, aberrations and expectations.

Design Choices and Expertise

The SIDARTHa system uses emergency care data to monitor health threats. The project researchers divide it into four data sources: emergency medical service (EMS), emergency medical dispatch (EMD), emergency physician service (EP), and emergency department (ED) data. Each source has different properties that make it more or less suitable for inclusion in a syndromic surveillance system. For example, the authors note that ED data are often collected electronically, while EMS data are filled in paper forms that are then transferred to the computer, making the latter more difficult to connect to an ICT-based surveillance system (Ziemann et al. 2009, 20). The authors also mention that emergency data covers only severe cases because patients with mild symptoms are not very likely to call the emergency medical service or to visit an emergency department (22). At the same time, emergency departments receive patients and respond to calls outside the working hours of other health services (22).

The decision to use emergency data for syndromic surveillance has consequences for the types of events that can be detected, and for those that are likely to remain invisible. In this case, one consequence is being able to detect severe cases, and not being able to detect mild ones. For example, the authors state that “since gastrointestinal problems do not mainly lead to the need of emergency medical care treatments these outbreaks can only be identified by a syndromic surveillance system if the outbreak occurs under special circumstances (i.e., symptoms in a group from abroad)” (Pinheiro et al. 2010, 15). The system itself makes events that manifest with acute symptoms more visible. This is in line with the vision of syndromic surveillance as a tool of preparedness for sudden, unexpected, and highly dangerous events. By beginning with the aim to detect health threats, and then setting up a system that is suited to the detection of severe illness rather than mild illness, the designers contribute to the definition of a health threat. In the previous examples, the health threat was a sudden event with the potential to affect health, whereas now a health threat is specified further as an event that can cause severe illness with sharp and sudden symptoms.

Keeping the agency of users in mind, once the system is operational it could also be used to detect mild cases depending on how the users choose to use the system. However, the design choices do constrain the use in this case, since the system itself is designed to run on its own automatically, and changing it requires expertise of the kind that is much more commonly associated with a software developer than an epidemiologist. To modify the system to detect mild cases, the system would need to be connected to different data sources which would require knowledge about how and where the data are stored, how they can be imported into the SIDARTHa system, as well as an understanding of the statistical methods to make them compatible with the newly connected data source. Although such a scenario is certainly possible, it is also clear that a different skill set is privileged in public health practice when syndromic surveillance is involved.

During the development of the SIDARTHa system, the ICT company BeValley programmed the system to match the specifications of the designers. In the evaluation report, the authors state that “Be-Valley agreed to adjust and update the system in the future but the question remains how this can be sustained also with additional funding” (Pinheiro et al. 2010, 10). They also ask: “If the regional system cannot easily be repaired by the future users how does that affect the usefulness and acceptance of the system?” (11). The authors anticipate that future users in emergency care institutions will have ICT staff who can install the software and program the data transfer (11), but they do not elaborate on neither whether emergency care institutions do in fact have such capacity, nor if that capacity should be found within emergency care. The expertise required to perform public health surveillance shifts if SIDARTHa or a similar syndromic surveillance system is introduced, and that new practice privileges ICT expertise.

Finally, the work of classifying and sorting does not end with the programming and installation of the system. During the data analysis study, Rosenkötter, Kauh, et al. (2010) report numerous mismatches that need to be resolved locally by the practitioners attempting to perform analyses based on emergency data. For example, the authors find that it is not possible to analyse unspecific syndromes using the Austrian emergency physician service data due to repeated entries for the same incident, while using the emergency department data the only option is to analyse unspecific syndromes because the source does not

list the reasons for care. The Belgium emergency physician service data arrive with delay, causing difficulties in developing a syndrome-specific coding, and the emergency department data is not possible to analyse fully because the systems recording the data have changed recently. The authors also describe the necessity of taking into account ICD shortlists used by physicians in different sites. These lists lump the detailed categories of the ICD into larger groups, allowing the physicians to quickly assign codes without referring to the extensive ICD documentation every time. However, the mappings of these shortlists differ from site to site, and they do not combine easily when centralising the data. It is not possible to ignore the shortlists either, because similar cases may be assigned to different syndromes if the mappings in the shortlists are not taken into consideration. Although these events may sound as exceptions to the rule of smooth classification, as Bowker and Star point out, the work of classification always includes these complexities and tensions.

9.4 Conclusion

We have demonstrated how a discourse of syndromic surveillance constructed the *health threat* concept as a sudden, unexpected event with the potential to cause severe harm, and one that requires a public health response aided by surveillance. Based on our analysis of reports from the SIDARTHa project, we stated that when creating surveillance technologies, design choices have consequences for what can be seen, and for what remains invisible.

We argued further that syndromic surveillance discourse values standardised and transportable knowledge more than local, context-dependent knowledge, and it privileges expertise in developing, maintaining, and using software within public health practice. We argued further that syndromic surveillance discourse privileges expertise in developing, maintaining, and using software within public health practice, and it prioritises standardised and transportable knowledge over local and context-dependent knowledge.

ICTs are not neutral tools that reflect reality. They privilege some forms of action, and they limit others. They are imbued with values, and different people benefit or suffer from their use differently. Fearnley states that “[s]yndromic surveillance itself, with its orientation

towards unexpected events and nonspecific objects, inevitably moves epidemiology in new directions.” (Fearnley 2008a, 84). French’s analysis of ICT use within public health defines some of these new directions:

[...A]n over-arching immaterial conception of information imbues some kinds of information with more import, for surveillance, than other kinds of information. Specifically, this conception encourages the collection of abstract, digitized signifiers while simultaneously marginalizing other kinds of embodied, contextual information. Indeed, the pursuit of immaterial information for public health surveillance produces a dominant but superficial epidemiology at the expense of other potentially more effective epidemiologies. (French 2009, 6)

Syndromic surveillance systems are not costless solutions because gathering information and sustaining surveillance requires work. They also require a different kind of expertise that is not always found in public health institutions. Moreover, syndromic surveillance relies on a particular definition of health threat which focuses on catastrophes and rare events. To perform syndromic surveillance within a public health institute is to modify public health as a practice. This is not problematic in itself since practices can always change, but the question that we need to ask remains: What does syndromic surveillance mean for fairness? If it orients public health towards the detection of catastrophes, or sudden events with severe consequences, what happens to other types of events that do not produce such signals? For example, to what extent can syndromic surveillance deal with chronic illness, or aid those who suffer from health issues due to poverty? Public health practice is a large field, and there may be a place for different surveillance systems that monitor different types of illness. However, discussions of syndromic surveillance should not stop at sensitivity, specificity, and timeliness, but also address questions of fairness.

References

- Baer, M., A. Ozguler, L. Garcia-Castrillo Riesgo, T. Krafft, M. Fischer, A. Krämer, P. Pinheiro, F. Lippert, G. Vergeiner, and H. for the SIDARTHa project group Brand, eds. 2009. *Using Emergency Care Data for Syndromic Surveillance in Europe. Public Health Threats Detectable by and Availability of Routinely Collected Emergency Care Data in Europe. Results from the SIDARTHa project.* Bad Honnef.
- Berg, Marc, and Geoffrey Bowker. 1997. "The Multiple Bodies of the Medical Record: Toward a Sociology of an Artifact." *The Sociological Quarterly* 38 (3): 513–537.
- Bowker, Geoffrey C., and Susan Leigh Star. 2000. *Sorting Things Out: Classification and Its Consequences.* Cambridge, MA: The MIT Press. ISBN: 0262522950.
- Bravata, Dena M., Kathryn M. McDonald, Wendy M. Smith, Chara Rydzak, Herbert Szeto, David L. Buckeridge, Corinna Haberland, and Douglas K. Owens. 2004. "Systematic Review: Surveillance Systems for Early Detection of Bioterrorism-Related Diseases." *Annals of Internal Medicine* 140, no. 11 (June): 910–922.
- Buehler, J. W., R. S. Hopkins, J. M. Overhage, D. M. Sosin, and V. Tong. 2004. "Framework for Evaluating Public Health Surveillance Systems for Early Detection of Outbreaks." *MMWR Morbidity and Mortality Weekly Report* 53:1–11.
- Buehler, J. W., E. A. Whitney, D. Smith, M. J. Prietula, S. H. Stanton, and A. P. Isakov. 2009. "Situational Uses of Syndromic Surveillance." *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science* 7 (2): 165–177.
- Butler, Declan. 2013. "When Google Got Flu Wrong." *Nature* 494, no. 7436 (February): 155–156.
- Cakici, Baki. 2011. *Disease Surveillance Systems.* Licentiate thesis. Royal Institute of Technology (KTH). <http://kth.diva-portal.org/smash/record.jsf?pid=diva2:416962>.

- Cakici, Baki, Kenneth Hebing, Maria Grunewald, Paul Saretok, and Anette Hulth. 2010. "CASE: A Framework for Computer Supported Outbreak Detection." *BMC Medical Informatics and Decision Making* 10 (1): 14.
- Chen, Hsinchun, Daniel Zeng, and Ping Yan. 2009. *Infectious Disease Informatics: Syndromic Surveillance for Public Health and Bio-Defense*. Springer. ISBN: 1441912770.
- Cook, Samantha, Corrie Conrad, Ashley L. Fowlkes, and Matthew H. Mohebbi. 2011. "Assessing Google Flu Trends Performance in the United States during the 2009 Influenza Virus A(H1N1) Pandemic." *PLoS ONE* 6, no. 8 (August).
- Espino, Jeremy U., M. Wagner, C. Szczepaniak, F. C. Tsui, H. Su, R. Olszewski, Z. Liu, et al. 2004. "Removing a Barrier to Computer-based Outbreak and Disease Surveillance: The RODS Open Source Project." *MMWR Morbidity and Mortality Weekly Report* 53 Suppl (September): 32–39.
- EU Health Programme. 2011. *Generic Public Health Preparedness in Europe*. European Union. doi:10.2772/59355. Brochure.
- Fearnley, Lyle. 2008a. "Redesigning Syndromic Surveillance for Biosecurity." In *Biosecurity Interventions: Global Health and Security in Question*, edited by Andrew Lakoff and Stephen J. Collier, 61–88.
- . 2008b. "Signals Come and Go: Syndromic Surveillance and Styles of Biosecurity." *Environment and Planning A* 40 (7): 1615–1632. doi:10.1068/a4060.
- Foucault, Michel. 1973. *The Birth of the Clinic: An Archaeology of Medical Perception*. Vintage, 1994. Translated by Alan Sheridan. 1994. New York: Vintage Books. ISBN: 0679753346 9780679753346.
- . 1977. *Discipline & Punish: The Birth of the Prison*. Second Edition, 1995. Translated by Alan Sheridan. Vintage. ISBN: 9780679-752554.
- French, Martin. 2009. "Picturing Public Health Surveillance: Tracing the Material Dimensions of Information in Ontario's Public Health System." PhD diss., Queen's University. <http://hdl.handle.net/1974/1689>.

- Garcia-Castrillo Riesgo, L., N. Rosenkötter, A. Ziemann, T. Krafft, M. Fischer, F. Lippert, G. Vergeiner, A. Krämer, and H. for the SIDARTHa project group Brand, eds. 2009. *The SIDARTHa Coding Manual – How to Generate Syndromes Based on Routinely Collected Emergency Care Data for the European Syndromic Surveillance System SIDARTHa*. Bad Honnef.
- Garcia-Castrillo Riesgo, L., A. Ziemann, T. Krafft, R. Conde, M. Fischer, G. Vergeiner, F. Lippert, A. Krämer, P. Pinheiro, and H. for the SIDARTHa project group Brand, eds. 2010. *SIDARTHa Technical Guidelines. How to Implement the European Emergency Data-based Syndromic Surveillance System SIDARTHa*. Bad Honnef.
- Ginsberg, Jeremy, Matthew H. Mohebbi, Rajan S. Patel, Lynnette Brammer, Mark S. Smolinski, and Larry Brilliant. 2009. “Detecting Influenza Epidemics Using Search Engine Query Data.” *Nature* 457, no. 7232 (February): 1012–1014.
- Google. 2013. *Google Flu Trends*. Accessed July 10, 2013. <http://www.google.org/flutrends/>.
- Heffernan, R., F. Mostashari, D. Das, A. Karpati, M. Kulldorff, and D. Weiss. 2004. “Syndromic Surveillance in Public Health Practice, New York City.” *Emerging Infectious Diseases* 10 (5): 858–64.
- Henning, Kelly J. 2004. “What is Syndromic Surveillance?” *MMWR Morbidity and Mortality Weekly Report* 53 Suppl:7–11.
- Latour, Bruno. 1988. *Science in Action: How to Follow Scientists and Engineers Through Society*. Harvard University Press. ISBN: 06747-92912.
- Mandl, Kenneth D., J. Marc Overhage, Michael M. Wagner, William B. Lober, Paola Sebastiani, Farzad Mostashari, Julie A. Pavlin, et al. 2004. “Implementing Syndromic Surveillance: A Practical Guide Informed by the Early Experience.” *Journal of the American Medical Informatics Association* 11 (2): 141–150.
- Mostashari, Farzad, and Jessica Hartman. 2003. “Syndromic Surveillance: A Local Perspective.” *Journal of Urban Health* 80 (1): i1–i7. doi:10.1093/jurban/jtg042.

- Pinheiro, P., Ziemann A., N. Rosenkötter, A. Krämer, L. Garcia-Castrillo Riesgo, T. Krafft, M. Fischer, G. Vergeiner, F. Lippert, and H. for the SIDARTHa project group Brand, eds. 2010. *SIDARTHa Syndromic Surveillance System: Test, Evaluation & Recommendations*. Bad Honnef.
- Reingold, A. 2003. "If Syndromic Surveillance is the Answer, What is the Question?" *Biosecurity and Bioterrorism: Biodefense Strategy, Practice, and Science* 1 (2): 77–81.
- Rosenkötter, N., B. Kauh, L. Garcia-Castrillo Riesgo, F. J. Llorca Diaz, J. Kraan, A. Ziemann, M. Schorbahn, T. Krafft, and H. for the SIDARTHa project group Brand, eds. 2010. *Retrospective Data Analysis and Simulation Study as Basis for an Automated Syndromic Surveillance System. Results from the SIDARTHa project*. Bad Honnef.
- Rosenkötter, N., A. Ziemann, L. Garcia-Castrillo Riesgo, G. Vergeiner, M. Fischer, T. Krafft, H. Brand, F. Lippert, A. Krämer, and P. on behalf of the SIDARTHa project group Pinheiro, eds. 2010. *SIDARTHa Volcanic Ash Cloud Rapid Public Health Impact Assessment. Regional Public Health Impact of Volcanic Ash Cloud Covering Europe after Eruption of Eyjafjallajökull, Iceland Starting April 14th, 2010. Results as of May 15th, 2010*. Bad Honnef.
- Triple-S. 2013. *Human Health Projects and Systems Identified Through the Triple-S Inventory*. Accessed July 9, 2013. <http://syndromic-surveillance.eu/systems-in-europe/humansys>.
- USMI. 2003. *Weaving a National Surveillance System – The Role of Federal Healthcare*. Technical report. Washington, DC: US Medicine Institute for Health Studies.
- Ziemann, A., T. Krafft, L. Garcia-Castrillo Riesgo, M. Fischer, A. Krämer, F. Lippert, and G. for the SIDARTHa project group Vergeiner, eds. 2009. *Early Detection of Health Threats: International State-of-the-Art & European Health Surveillance Landscape. Results from the SIDARTHa project*. Bad Honnef.