Report from Dagstuhl Seminar 13312

"My Life, Shared" – Trust and Privacy in the Age of Ubiquitous Experience Sharing

Edited by

Alessandro Acquisti[1], Ioannis Krontiris[2], Marc Langheinrich[3], and Martina Angela Sasse[4]

1   **Carnegie Mellon University – Pittsburgh, US**, `acquisti@andrew.cmu.edu`
2   **Goethe-Universität Frankfurt am Main, DE**, `ioannis.krontiris@m-chair.net`
3   **Università della Svizzera italiana, CH**, `marc.langheinrich@usi.ch`
4   **University College London, GB**, `a.sasse@cs.ucl.ac.uk`

─── **Abstract** ───────────────────

Many researchers have already begun using personal mobile devices as personal "sensing instruments" and designed tools that reposition individuals as producers, consumers, and remixers of a vast openly shared public data set. By empowering people to easily measure, report, and compare their own personal environment, such tools transform everyday citizens into "reporting agents" who uncover and visualize unseen elements of their own everyday experiences. This represents an important new shift in mobile device usage – from a communication tool to a "ubiquitous experience sharing instrument". This report documents the program and the outcomes of Dagstuhl Seminar 13312 *"My Life, Shared" – Trust and Privacy in the Age of Ubiquitous Experience Sharing*, which brought together 33 researchers and practitioners from multiple disciplines – including economics, psychology, sociology, as well as various fields within the discipline of computer science dealing with cryptographic feasibility, scalability and usability/acceptability – to discuss opportunities and challenges of sharing information from the pervasive environment.

1   Executive Summary

Alessandro Acquisti
Ioannis Krontiris
Marc Langheinrich
Martina Angela Sasse

Advancements in smart phones and sensing technology have bolstered the creation and exchange of user generated content, resulting in new information flows and data-sharing applications. Through such applications, personal mobile devices are used to uncover and share previously private elements of people's own everyday experiences. Examples include using smartphones or wearable sensors to collect and share context information (e.g., activities, social context, sports performance, dietary or health concerns). These flows of personal

information have two distinct characteristics: they happen seamlessly (in real time, without necessarily the conscious participation of the user), and they are shared with a user's family, social circles, or even publicly.

This new paradigm repositions individuals as producers, consumers, and remixers of a vast set of data with potential many economic and societal benefits. However, as sharing practices become more fluid than in desktop-based online environments, control over personal information flows becomes harder to maintain.

The goal of Dagstuhl Seminar 13312 *"My Life, Shared" – Trust and Privacy in the Age of Ubiquitous Experience Sharing* was to advance a research agenda in trust and privacy that addresses not only the evolution of the pervasive technologies underlying these trends (e.g., smartphones, wearable sensors), but also the surrounding societal and economic context, and to identify the resulting qualitative changes to the privacy landscape.

With that in mind, the seminar created an interdisciplinary discussion forum and a set of organised presentations around four broad areas: 1) tools and protocols, 2) usability and control tools, 3) behavioural decisions, and 4) social implications. Each area saw a selected set of participants present their work and views in the context of a short presentation, followed by an in-depth discussion session. From these discussions the organizers collected the main challenges and opportunities, and grouped them around four major themes: "Personal Data Services", "Social Justice", "Tool Clinics", and "Consequence-based Privacy Decision-making". Each theme was subsequently discussed during one and a half days in four individual working groups, which presented their findings at the end of the seminar.

This report not only contains the abstracts of the initial presentations (section 3) but also the findings of the four thematic working groups. Below we summarize the main findings from these working groups – a more analytical description can be found in section 4.

Theme 1: Personal Data Service (PDS)

A "Personal Data Service (PDS)" represents a trusted container for aggregating, storing, processing and exporting personal data. In principle, all data regarding the user (either user-generated or obtained from other sources, e.g. service providers) should be accessible to this container, including data about the user collected and published by others. Users are in control of all data stored in the PDS, which includes the option to share or sell parts of this data. In addition to storing data, the PDS can execute code to process this data locally.

By considering both a household- and a health-related scenario, the working group identified some of its properties and functionalities and sketched a possible system architecture that would include such a container. In a detailed discussion of benefits and risks, the working group concluded that there were still several issues to be investigated and real challenges that needed to be addressed before a PDS framework could be implemented and deployed, such as:

- Creating *incentives* to initial data providers to engage and open up the personal data APIs that are needed to fuel the PDS and associated applications.
- Creating *utility* from stored data: data fusion, sense making, and visualization that will lead to meaningful and actionable and sustainable engagement of the end user with their data.
- Addressing *privacy*: even though the PDS can increase transparency, awareness and engagement of users with their data, it is neither obvious nor guaranteed that PDS will resolve user privacy problems and several of them remain open.

Theme 2: Social Justice

Privacy issues in participatory sensing are symptoms of broader concerns about the impact of sensing on social justice. Framing a social justice research agenda for participatory sensing requires the operationalization of concepts like fairness, human flourishing, structural change, and balances of power for system design, use, and regulation. The working group discussed how one might begin to operationalize these concepts for the design of data collection features, processing, sharing, and user interfaces. The group developed an analysis tool – a social justice impact assessment – to help system designers consider the social justice implications of their work during the design phase. The participants identified and presented several open questions that could spark future research, such as:

- If one assumes that participatory sensing will lead to *greater transparency*, will such transparency equally impact individuals, powerful people, and institutions?
- Do the powerful always end up *subverting transparency* schemes? Or can sensing change that tendency, for example by making facts visible to consumers and citizens, enabling organized responses (unionization)?
- What are the forums for *encouraging collective action* in participatory sensing? Can one encourage system designers to consider social justice during design by framing design as a collective action problem? Can participatory sensing open new avenues for consumers and citizens to organize collective action?

Theme 3: Tool Clinics

Privacy researchers and practitioners are working largely in isolation, concentrating on people's use of different user interfaces for privacy control, largely ignoring existing cross-disciplinary collaboration techniques. A "tool clinic" could encourage a collaborative (re)consideration of a technological solution, research technique or other artefact, in order to critically assess its design, development and deployment from multiple perspectives. A tool clinic can be used to provide a setting for those who are developing the solutions to rethink the framing and presentation of their solutions. The objective is to reflect from different perspectives on practices around the development, encoding, use, domestication, decoding and sustainability of a tool to gain quasi-ecological validation. The working group recommended to develop a tool clinic as a new event format for a scientific conference, ideally at a renowned computer-science conference. This would combine the tool-centric nature of a demo session, the protected space of work-in-progress afforded by a workshop, and the mentoring spirit of a doctoral workshop. The format of a tool clinic session could typically consist of three steps:

1. Identifying particular affordances of the technological solution, research technique or other artefact and possible (unintended) consequences for people and society;
2. Gathering perspectives and practices of different experts, disciplines and/or stakeholders (e.g. users, policy makers, industry, etc.) linked with the development, deployment and sustainable evolution of a particular tool, solution, technique or artefact;
3. Informing and advising on technological design of the tool or solution, in order to avoid negative consequence and to further positive outcome.

Theme 4: Consequence-based Privacy Decisions

Recent research shows that people not only want to control their privacy but are actually trying to do so. An appropriate privacy-respectful user interface should thus show users the consequences of making different privacy choices, rather than framing the choices only in technical terms regarding system parameters, which users often do not understand and do

not care about. Providing tools to increase user comprehension of potential consequences is one of the next big challenges to be addressed in the field of privacy respectful user interfaces. In addition to helping users make better choices in terms of privacy protection, this will also allow them to make better informed decisions and hence, implement the notion of informed consent. The attempt to develop user interaction in this direction requires research on a number of issues that have so far received relatively little attention and concern, such as:

- *Expression of potential consequences:* The consequences should be expressed in a way that is comprehensible by different user categories from novices to experts.
- *Decision support:* Users could be further helped in their privacy decisions by external information sources. Studies to determine the responses to different kinds of information sources, different formats, and information from different groups of users will be necessary.
- *Minimal effort:* Introducing additional tools to help users make informed decisions may add significant overhead to the interaction. While this overhead may be the price to pay for better privacy protection, it should be limited to the minimum.

2   Table of Contents

Working Groups

3   Overview of Talks

The first three days of the seminar saw a range of short presentations from a subset of seminar participants, grouped into four areas: 1) tools and protocols, 2) usability and control tools, 3) behavioural decisions, and 4) social implications. The goal was to create a common understanding between the diverse set of participants and to stimulate discussions about opportunities and challenges in the space. Below, we provide the abstracts of the 15 talks that were given by individual seminar participants, in alphabetical order.

3.1   Gone in 15 Seconds: The Limits of Privacy Transparency and Control

Alessandro Acquisti (Carnegie Mellon University, US)

I will present some results from privacy experiments inspired by behavioral economics and decision research. The experiments investigate the role of control and transparency in privacy decision making. The results suggest that even simpler or more usable privacy notices and controls might not improve users' decision-making regarding sharing of personal information: Control might paradoxically increase riskier disclosure by soothing privacy concerns; transparency might be easily muted, and its effect even arbitrarily manipulated, through simple framing or misdirections.

3.2   Privacy and Location Sharing: Challenging Technical Problems in Search for Conscious Users

Claudio Bettini (University of Milan, IT)

Among the many types of data that are shared as part of social online interactions, we focus on spatio-temporal data and on the privacy issues involved in sharing users' location and movements. We illustrate some of the technical challenges in the design of location privacy protection and location sharing monitoring tools, considering in particular the threats involved in posting geo-tagged resources in OSN, and the ones determined by location sharing in proximity services. We also report on our recent experience with the design, implementation and offering on the global market of PCube, a friend-finder mobile app that totally hides location and proximity information to the service provider. We will discuss the feedback obtained by users, the actual challenges we faced, and, in general, the lessons learned.

3.3 Collective Exposure: Peer Effects in Voluntary Disclosure of Personal Data

Rainer Böhme (Universität Münster, DE)

I will report empirical evidence for peer effects in privacy behavior using field data from Germany's largest online social lending platform. The study applies content analysis to measure personal data disclosure on and identifiability of borrower profiles with tailored scales. A logistic regression analysis suggest that individuals tend to copy observable behavior of others in their decisions on

- how much to write about oneself,
- whether to share custom pictures,
- what personal data to disclose, and
- how identifiable to present oneself.

I will frame this finding in the theory of descriptive social norms and explore moderating effects, such as similarity of context, social proximity, and mimicry of success factors. Peer effects in disclosure behavior seem to be an important factor to explain the formation and change of apparent social norms and attitudes towards information privacy. [1]

References
1   Rainer Böhme and Stefanie Pötzsch. Collective exposure: Peer effects in voluntary disclosure of personal data. In George Danezis, editor, *Financial Cryptography and Data Security*, volume 7035 of *Lecture Notes in Computer Science*, pages 1–15, Berlin Heidelberg, 2011. Springer-Verlag.

3.4 Transparency, Control, Minimisation, Law – How to Build Societal Trust in Ubiquitous Sharing?

Ian Brown (University of Oxford, GB)

Edward Snowden's dramatic revelations about the broad surveillance activities of the US National Security Agency have been called the most important leaks in American history. It seems the UK's equivalent intelligence agency, GCHQ, is equally busy wiretapping the world's Internet traffic.

What implications do these surveillance systems have for ubiquitous experience sharing systems? How can data privacy be protected in a world with such voracious intelligence agencies? And what responsibilities does this place on life logging system designers?

3.5 Tools of the Trade: Enhancing Privacy in Participatory Sensing

Tassos Dimitriou (Athens Information Technology, GR)

In this presentation we review some of the private-preserving mechanisms applied in participatory sensing applications. We start with a definition of privacy for participatory sensing and the threats associated with uncontrolled disclosure of sensitive information. We then consider typical countermeasures used by the various architectural elements of sensing applications. Finally, we highlight and discuss some interesting research directions that must be addressed to enhance user privacy and encourage user participation.

3.6 Sharing Private Location Information via Non-trusted Servers

Frank Dürr (Universität Stuttgart, DE)

Private location information is essential for many modern location-based services like geo-social networks. Often, such services are hosted on third- party server infrastructures ("in the cloud"). We argue that it is practically impossible to guarantee that private location information stored on such infrastructures is perfectly protected from unauthorized access. Many incidents in the past have shown that private data managed on "trusted" and protected servers was stolen or "leaked". Consequently, we think there is no such thing as a trusted server, and we have to think about technical concepts to protect private information shared through non- trusted server infrastructures.

Besides highlighting the problem, we are going to present technical concepts for the management of private location information on non-trusted servers. Our concepts include some very interesting features such as no single point of failure, graceful degradation of privacy with the number of compromised servers, and the possibility to define different levels of privacy for different location-based applications to defined a trade-off of quality of service and privacy. Moreover, we briefly present a second concept for protecting information derived from movement trajectories, namely speed information, by limiting the accuracy of spatial-temporal information.

3.7 Location Privacy – Myth or Reality?

Raghu K. Ganti (IBM TJ Watson Research Center – Yorktown Heights, US)

As spatiotemporal data generated by mobile devices become readily available and increase in volume, a critical and important question that arises is that of location privacy. On one hand, location and time data enable novel applications and services to the common man (e.g.,

smarter cities, smarter telcos). On the other hand, such information can result in serious privacy breaches (e.g., pleaserobme.com). In this talk, I will examine the implications of availability of large volumes of spatiotemporal data on privacy. I will present observations from our past work on the privacy analysis of spatiotemporal data, examining various techniques to obfuscate such data to achieve privacy at a community-wide level and also examining techniques that break the privacy based on different channels of information (e.g., text, maps). Finally, I will also discuss a system that we are currently building that can ingest large volumes of data and analyze it in real-time to infer motion behavioral patterns.

3.8   Toying with Facebook: An Online Experiment of Privacy Authorization Dialogues for Social Applications

Jens Grossklags (Pennsylvania State University, US)

Several studies have documented the constantly evolving privacy practices of social networking sites and users' misunderstandings about them. Researchers have criticized the interfaces to "configure" privacy preferences as opaque, uninformative, and ineffective. The same problems have also plagued the constant growth of third-party applications and their troubling privacy authorization dialogues. In this talk, I report the results of an experimental study examining the limitations of current privacy authorization dialogues on Facebook as well as four new designs which were developed based on the Fair Information Practice Principles (FIPPs). Through an online experiment with 250 users, the effectiveness of installation-time configuration and awareness-enhancing interface changes are studied. The experimental results are complemented with data from a measurement study on Facebook third-party applications.

3.9   Two Tales of Privacy in Online Social Networks

Seda F. Gürses (KU Leuven, BE)

Privacy is one of the friction points that emerges when communications get mediated in Online Social Networks (OSNs). Different communities of computer science researchers have framed the 'OSN privacy problem' as one of surveillance, institutional or social privacy. In tackling these problems they have also treated them as if they were independent. We argue that the different privacy problems are entangled and that research on privacy in OSNs would benefit from a more holistic approach. During my talk, I will first provide an introduction to the surveillance and social privacy perspectives in computer science emphasizing the narratives

that inform them, as well as their assumptions, goals and methods. I will then juxtapose the differences between these two approaches in order to understand their complementarity, and to identify potential integration challenges as well as research questions that so far have been left unanswered.

3.10   Planning your Digital Afterlife – User Needs and Control Tools

Thomas Heimann (Google – München, DE)

While we have established procedures for inheriting tangible items, our practices for inheriting digital assets are far less developed. Giving users control over their "digital estate" is becoming more and more important as the volume and importance of digital artifacts grows. Like physical possessions, digital artifacts may carry significant sentimental value for bereaved family members, and tools for managing the digital estate may help to remember, commemorate, and reminisce about the deceased and find closure.

Google's Inactive Account Manager provides Google users with the opportunity to manage if and how their data is made available to specified trustees in the case of death or temporary unavailability. In this presentation, I present some of the research, design and implementation challenges that accompanied the development of this feature.

3.11   Move Over, Westin

Anthony Morton (University College London, GB)

Researchers and practitioners have continued to use Westin's categorisation of people into privacy fundamentalists, privacy pragmatists and privacy unconcerned, even though it has not yet been shown to be a reliable predictor of people's privacy behaviour. A simple, three-level categorisation is flawed for two reasons: 1) people's privacy concern is dynamic – not static – and is influenced – amongst other things – by the perceived sensitivity of the information collected or requested, the purpose of collection, and the party requesting it; and 2) with the increasing power and ubiquity of technology to collect, process, store and disseminate information, a more comprehensive representation of people's'privacy concern is required. To address these weaknesses, we propose a richer model of privacy concern, encompassing: a) dispositional privacy concern; b) privacy concern due to environmental factors (e.g. the experiences of friends and family); and c) privacy concern specific to the technology- mediated interaction with the other party. In addition, as privacy concern is widely recognised as being subjective, we also describe an innovative use of Q methodology –

a research method which combines qualitative and quantitative research methods to identify people's subjective viewpoints – to determine if it is feasible to group people into segments representing different privacy attitudes, and hence specific configurations of our proposed richer model of privacy concern. Finally, we consider what is a 'reasonable' level of privacy concern? When does a person's desire for privacy, or concern about information collection, become paranoia?

References

1 Adams, Anne, and Martina Angela Sasse. *"Privacy in multimedia communications: Protecting users, not just data." People and Computers XV – Interaction without Frontiers.* Springer London, 2001. 49- 64.

3.12 Theatre, Mirror, or Laboratory: Cross-Purposes in Platforms of Intimate Knowledge

David Phillips (University of Toronto, CA)

In this presentation I briefly suggest several reasons why individuals create, collect, share, and analyze intimate data about themselves. I then discuss the tools with which they pursue these interests, and the infrastructures by which they access these tools. I review the economic and social models which support those infrastructures, and finally, probe some of the tensions among these various interests.

3.13 The Right Privacy Controls for Social Search

Sören Preibusch (Microsoft Research – Cambridge, GB)

We share our lives, as the motto of our seminar evokes, but Web search, the most ubiquitous of Web activities, is still very much a solitary pursuit. This is particularly surprising, as our lives can be told in queries: interests and preferences, goals and unsatisfied needs. As we prepare for social search, how should we engineer a shared search experience? Social search promises to improve the relevance of results by taking into account the social connections of the searcher. One aspect of social search is the collaborative search experience, where one can inspect and learn from friends' queries, results, and clicked links. Results are discovered and ranked by borrowing from similar queries that our friends have issued in the past, and the Websites they subsequently visited. However, a sizeable proportion of everyday queries are better kept private, such as health issues, surprise gifts or porn. Acceptance of social search will hinge on users keeping control over which queries to share and which queries to keep private. I will report on a recent large-scale laboratory experiment into ways user manage their privacy when using a Web search engine: privacy features are universally appreciated and usage is high. A sizable proportion of users is willing to pay for added privacy. Although usage decisions were found not to be systematically dependent on

consumers' privacy preferences, the sensitivity of the topic has a significant impact on the demand for privacy.

3.14   User-Controllable Privacy: An Oxymoron?

Norman Sadeh (Carnegie Mellon University, US)

Increasingly users are expected to evaluate and configure a variety of privacy policies (e.g. browser settings, mobile app permissions, or social networking accounts). In practice, research shows that users often have great difficulty evaluating and configuring such policies. As part of this presentation, I will provide an overview of research aimed at empowering users to better control their privacy in the context of a family of location sharing applications we have deployed over the years. This includes technologies to analyze people's privacy preferences and help design interfaces that are capable of effectively capturing their desired policies. This research helps explain why, with the possible exception of Foursquare, applications in this space have failed to gain traction and what it will likely take to go beyond the mundane scenarios captured by Foursquare. Part of this talk will be devoted to user-oriented machine learning techniques intended to reduce user-burden and help users converge towards policies they feel more comfortable with. Beyond location sharing, this talk will also discuss our longer-term goal of developing personalized privacy assistants (or "agents") capable of engaging in dialogues with users to help them semi-automatically evaluate privacy policies and configure privacy settings.

3.15   Experience Sharing & Social Justice

Katie Shilton (University of Maryland – College Park, US)

I will argue that privacy issues in participatory sensing are symptoms of broader concerns about the impact of sensing on social justice. Framing a social justice research agenda for participatory sensing challenges us to operationalize concepts like fairness, human flourishing, structural change, and balances of power for system design, use, and regulation. I will discuss how we might begin to operationalize these concepts for the design of data collection features, processing, sharing, and user interfaces. And I will explores how we can encourage participatory sensing designers to consider these challenges as collective action problems.

4   Working Groups

After three days of presentations and discussions, the participants identified four major themes that are of relevance in pervasive experience sharing: "Personal Data Services", "Social Justice", "Tool Clinics", and "Consequence-based Privacy Decision-making". Each theme was subsequently discussed during one and a half days in four individual working

groups, which presented their findings at the end of the seminar. The four sections below constitute the output of each of these working groups.

4.1 Personal Data Service: Accessing and Aggregating Personal Data

Alessandro Acquisti (Carnegie Mellon University, US)
Claudio Bettini (University of Milan, IT)
Rainer Böhme (Universität Münster, DE)
Claude Castelluccia (INRIA Rhône-Alpes, FR)
Tassos Dimitriou (Athens Information Technology, GR)
Frank Dürr (Universität Stuttgart, DE)
Deborah Estrin (Cornell Tech NYC, US)
Michael Friedewald (Fraunhofer ISI – Karlsruhe, DE)
Raghu K. Ganti (IBM TJ Watson Research Center – Yorktown Heights, US)
Jens Grossklags (Pennsylvania State University, US)
Renè Mayrhofer (University of Applied Sciences Upper Austria, AT)
David Phillips (University of Toronto, CA)
Kai Rannenberg (Goethe-Universität Frankfurt am Main, DE)
Norman Sadeh (Carnegie Mellon University, US)
Marcello Scipioni (University of Lugano, CH)

4.1.1 Motivation

The last decade has seen a major trend towards personal data analysis: in addition to purely digital services such as social networks, we have observed an increasing integration of real-world sensing from various sources, including digital traces from online search, shopping, social, entertainment and financial services; fitness applications; medical sensing; as well as environmental, vehicular, and household data. So far, these data sources are mostly handled as separate streams: they are collected, stored, and analyzed within their own commercial niches – typically out of view of the individual whose life they describe. One foreseeable trend is the aggregation of different data streams about the individual, in order to facilitate better data utilization by correlating and generally combining across data streams, and bringing the individual in control over the aggregated use of these data.

Current digital service providers like Google or Facebook may be more than likely to provide such an aggregation, analysis, and reporting service for their users free of charge – in return for getting an even more detailed picture of the individual, but thereby raising privacy concerns. Personal Data Services (PDS) are an alternative aggregating platform under control of the end user.

We use two common scenarios (a household shared by a family, and health data about a single individual) to define PDS, discuss their potential benefits and risks, and propose potential architectures for implementation.

4.1.2 Household Scenario

Consider the following data streams that are relevant to a typical household (be that an individual or a family unit):

- Transaction streams concerning payments, potentially spanning multiple vendors, different people, and different payment cards. As a key example, consider the possibility of a family wanting to reduce the amount of processed foods containing high salt and sugar content, or to increase the regularity of family meals. To track progress or to create social games that incentivize families or individuals within families, these transaction traces could be used as the measurement feedback. For independent living seniors, these transaction traces can also contribute to a wellbeing pulse shared with close family or friends, to help them stay aware of important but subtle changes that might benefit from intervention.

- Data on transportation, including public transportation, car usage, gas purchases, parking and toll fees, etc. A family wishing to reduce costs or increase sustainability could similarly use apps that analyze digital traces to make recommendations and incentive the desired behaviour. Communities and employers could sponsor "drives" or challenges in which progress as a community gets matched by corporate dollars to develop community resources, such as community gardens or green spaces.

- Household usage data, including utilities such as electricity, water, gas, etc. as typically measured by smart meters. These data can be aggregated into transportation based applications described above to create a bigger picture sustainability app. They can also be combined with transaction data to contribute to the overall wellbeing pulse because they capture what might be significant changes in behaviour such as decline home preparation of food, or change in diurnal patterns (when the coffee pot is used or increased night time activity).

- Time-activity-location data, including data from cellular network providers, location based services (foursquare), activity monitoring devices (Fitbit, Nike plus, Jawbone Up), and mobile apps. These data provide a baseline of individual and family patterns that can fuel family management, wellness, health, financial and other applications in combination with the other data streams described.

4.1.3 Health scenario

In this scenario we consider the streams of data that may be useful to collect and analyse to enable useful services related to an individual's health and well-being. We identify the following relevant sources of data:

- Mobile phone. A lot of potentially useful data comes from the digital traces we leave in our interaction with smartphones including temporal patterns of interactions with the apps, the user location and movements, user activities and more.

- Home appliances and sensors. Home automation is making available data that provide information about activities of individuals within their home (e.g., TV usage patterns, use of appliances for cooking, lights, heating, room occupation, etc.)

- Medical and fitness-specialized sensors. Wearable sensors like Nike FuelBand, Jawbone Up, Fitbit Flex, or Misfit Shine are becoming common and can reveal levels of physical activity, type of activity, sleep patterns and more. Other useful data can be morning and evening blood pressure or morning weight. All of this data can be integrated into the PDS, perhaps through commercial third parties (e.g., Qualcomm Life 2net).

- External data sources from service providers (e.g., data about purchased food products obtained from merchants, data about media consumption from a cable company, data about physical training from the gym, data from EHR, etc.)
- Public and environmental data sources (e.g. pollution maps, weather reports, etc.)

4.1.4   Personal Data Service

We define a Personal Data Service (PDS) as a user trusted container for aggregating, storing, processing and exporting Personal Data. In principle, all data regarding the user (either user-generated or obtained from other sources, e.g. service providers) should be accessible to this container, including data about the user collected and published by others. Users are in control of all data stored in the PDS, which includes the option to share or sell parts of this data. In addition to storing data, the PDS can execute code to process this data locally, and will in turn store the processed data alongside the raw source data.

4.1.5   PDS functionality and architecture

By considering the two scenarios illustrated above we expect a PDS to perform the following functions:

- To access, protect, and analyze the incoming data streams for different sources for overall fusion of the data and local storage of resulting derived data.
- To provide an open third-party marketplace of apps with local processing, filtering, and auditing of user privacy preferences and actions of the applications.
- To support exporting locally derived data (e.g. about car usage in the household scenario) to third-party application providers with filtering and transformation for privacy reasons, and strong auditing.

In Figure 1 we illustrate a possible architecture for the interaction between a PDS and other software components and entities. The architecture in Figure 1 refers in particular to the household scenario.

In the health related scenario, we can foresee a similar architecture, with the PDS playing a major role for the following:

- For medical use of the data by patients as well as by clinicians through dedicated software, exporting data either filtered (for privacy control) or raw.
- For individual summaries/reports with aggregated data or to generate specific alerts.
- To export to existing healthcare systems (e.g., Ginger.io).

4.1.6   Privacy Impact of PDS

We explicitly note that the main goal of a PDS is utility as emphasized in the description of its functionalities with respect to our scenarios.

Many of state-of-the-art privacy preserving methods, based for example on obfuscation or generalization of sensitive personal information, can be applied in the PDS as part of the filtering functionality for exporting personal data. They can also be coupled with advanced access control methods. However, despite filtering, a PDS is NOT by itself a privacy preserving technology. Indeed, relevant personal data are also independently acquired and held by third parties.

The PDS is a tactical move to increase transparency, awareness, and engagement of users with their data. This by itself is a significant step along the lines of the recent EU proposal for a new data protection regulation that emphasizes the need to increase transparency, user
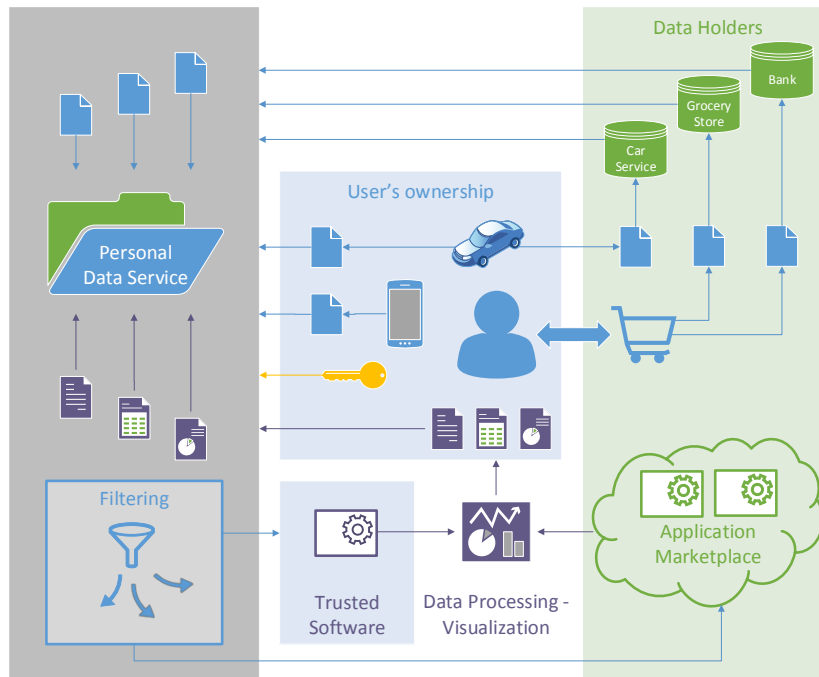
Figure 1 The PDS architecture in the household scenario.

awareness and control together with the principles of privacy by default and the right to be forgotten.

The PDS may also contribute to shift the market to a different behaviour through standardization of personal data formats and APIs to store and retrieve personal data. However, ultimate impact on privacy and alignment with diverse national policy frameworks will still depend on regulation of data service providers to support the particular country's sovereignty to enforce local policy.

4.1.7   Benefits of a PDS framework

We can identify the following benefits:

- Transparency to users of, and engagement of users with, data collected about them, e.g., increase privacy when it comes to processing across diverse data streams including counter-surveillance and the opportunity for the individual to validate the integrity of the data.
- A context in which to develop and support domain specific identity management so that individuals can configure appropriate roles and associated data sharing and not have one size fits all identity and data sharing with third party services.
- Catalyzing a new ecosystem of innovations in personal services from commuting to consumption, entertainment and health/wellness. This will have the dual benefit of utility to the individual and economic growth of this new sector.
- The potential to support consumer groups and movements by giving them more collective power through their access and use of the individual and community level data. Interest groups will be able to more effectively combine their data in ways to make their community's case.

4.1.8   Risks and Challenges

The biggest challenges are:

- First and foremost, creating the incentives to initial data providers to engage and open up the personal data APIs that are needed to fuel the PDS and associated applications. This is a standard first mover problem: until there are such data, the app developers cannot engage; and until there are engaged app developers, the utility of the data is minimal. Moreover, some business models may feel challenged by the accessing of data by their subscribers either because they feel it discloses performance issues in their services, or they are concerned about other third parties from getting access to monetizable data, or they are concerned about subscriber perceptions once they see and have to carefully handle these data, or they are concerned about liability issues if data is inaccurate or processed without legitimate basis, or they are concerned that the data may allow third parties to reverse engineer algorithms that are perceived as trade secrets.

- A second challenge, both in conceptual and perception terms, arises from the confusion between data access and data ownership. That is: is it my data, or data about me? PDS may have to align with national policies and regulations.

- A third, more technical challenge, relates to creating utility from the data: data fusion, sense making, and visualization that will lead to meaningful and actionable and sustainable engagement of the end user with their data. Moreover, this utility must compete in the market with closed platform commercial alternatives that provide similar function to the end user without control over their data.

- A fourth challenge arises from the potential for PDS (and the apps that will be created to work on PDS data) to subtly influence user behaviour. Apps developed for PDS will reflect the coders' value systems (for instance, what data the coder thinks is most important to consider or use, or even what user behaviour should be encouraged or discouraged). Paternalistic or soft paternalistic interventions may reduce user autonomy. This scenario is particularly troublesome, considering that data streams and the algorithms apps will apply onto the data may not be visible to, or understandable by, the end user.

- A fifth challenge arises from the observation that it is neither obvious nor guaranteed that PDS will resolve user privacy problem. PDS, for instance, may provide users with local control over the data, but may not be able (both for technical and economic reasons) to stop third parties from collecting and exploiting user data. In a worst case scenario, PDS may even end up facilitating third parties' collection of user data, by creating a central repository for the user's diverse data streams.

- A sixth challenge arises from the fact that data aggregated in the PDS may be accessed and used against the user's interest (and potentially without her knowledge) for the purpose of law enforcement or other governmental use. It may also be subject to discovery in civil lawsuits. Even without legitimate reason, individuals may be more likely to fall for trickery or coercion than third parties who routinely handle requests and know the conditions under which they have to comply.

- The PDS architecture itself represents an important challenge in terms of storage, processing and security:
  - Storage of data must be flexible to accommodate local, hosted and cloud based alternatives.
  - Processing: providing the flexibility of local processing, simple APIs for third party application developers, and flexibility as to where accessed and derived data are actually stored (locally, hosted, cloud). Some processing services will be real-time and some more offline and retrospective; both modalities need to be supported.

- Security of the PDS itself, in terms of key management, audit mechanisms, and robustness (avoiding single point of failure and vulnerability to mobile device loss/theft/damage).
- Security of third party data extraction APIs to grant access only to the PDS of authorized users. This is particularly challenging if the user is only weakly identified or data refers to more than one user.
- Security architecture of the PDS app platform with regard to granular and usable access permissions and the suppressions or detection of covert channels.

4.1.9 Conclusions

We discussed the notion of a Personal Data Service (PDS) as a trusted container for aggregating, storing, processing and exporting personal data. By considering a household and a health related scenarios, we identified some of its properties and functionalities and sketched a possible architecture including such a container. Our discussion of benefits and risks shows that there are still several issues to be investigated including privacy and security aspects, and real challenges need to be addressed before a PDS framework can be implemented and deployed.

4.2 Social Justice

Mads S. Andersen (Aarhus University, DK)
Ian Brown (University of Oxford, GB)
Ioannis Krontiris (Goethe-Universität Frankfurt am Main, DE)
Sören Preibusch (Microsoft Research – Cambridge, GB)
Martina Angela Sasse (University College London, GB)
Katie Shilton (University of Maryland – College Park, US)
Sarah Spiekermann (Universität Wien, AT)

This working group explored the idea that privacy issues in participatory sensing are symptoms of broader concerns about the impact of sensing on social justice. Framing a social justice research agenda for participatory sensing challenges us to operationalize concepts like fairness, human flourishing, structural change, and balances of power for system design, use, and regulation. We discussed how we might begin to operationalize these concepts for the design of data collection features, processing, sharing, and user interfaces. We developed an analysis tool – a social justice impact assessment – to help system designers consider the social justice implications of their work during the design phase.

4.2.1 Definitions of Social Justice

We first began with a discussion of definitions of social justice. We drew primarily on the work of John Rawls [3] and Amartya Sen [4]. Rawls wrote that we could best achieve a just society by setting rules that anyone would agree would be fair from behind a "veil of ignorance," in which no one knows where in that society they will be placed. Sen wrote that justice requires that individuals have the basic capabilities (education, health) they need to flourish and make the most of their lives.

4.2.2 Domains Impacted By Ubiquitous Experience Sharing and Participatory Sensing

We next discussed a variety of social domains in which we believe justice will be impacted – positively or negatively – by ubiquitous information sharing.

On the positive side, we believe that ubiquitous sharing platforms hold great promise for leveling the playing field in a variety of social domains. It seems clear that participatory sensing will have a positive impact on health, as patients are empowered with respect to institutions and new forms of data enable new kinds of diagnosis, monitoring and treatment. Similarly, care and independent living support for older citizens may be positively impacted by increased experience sharing. Sensing might even be useful for improving psychological health and mindfulness. We discussed whether context-aware applications might even be helpful for pursuing personal virtues, such as honesty or generosity.

On a social level, we think ubiquitous sharing can benefit community integration through applications like neighbor-to-neighbor sharing of goods, services, and experiences. We discussed scenarios in which ubiquitous experience sharing could benefit food systems and reduce waste by connecting producers more directly to consumers. Sensing can enable mapping of citizen's relationship with cities, as projects on walkability, bikeability, and even the emotions evoked by places demonstrate.

Sensing might also help us improve environmental justice, for example by helping communities make the case about unequal air pollution levels in underserved communities. Citizens armed with sensors might be better equipped to surface inconvenient truths about quality of life in their neighborhoods, and provide the impetus for like-minded individuals to start talking, and proceed to other organized responses to achieve change ("unionization" of citizens and consumers). Finally, we discussed a number of ways that experience sharing and participatory sensing (or "sousveillance") can increase transparency and accountability of powerful organizations to the citizen. Poll watching (such as that performed using mobile phones by Ushahidi) and cop watching [1] are two practices in which phones are used to keep governments accountable.

More complicated scenarios included sensing's impact on community rules and enforcement. We discussed digital vigilantism, and whether it could empower local communities or lead to increasing conformity and stigmatization. It is not yet clear under which circumstances "peer to peer" infrastructures for policing social norms increase or reduce justice. Education is another social context in which the impact of sensing is unclear. Sensing programs such as Mobilize[1], which use participatory sensing to teach data literacy in underserved communities, hold promise for improving justice in the educational systems of data-intensive societies. Sensing could conceivably increase accessibility for children with some types of learning challenges. But will these measures benefit underprivileged communities, or already well-resourced schools? And will participatory sensing also allow for increased monitoring and measurement, further quantifying student learning outcomes and potentially enabling stigmatization and pressure for conformity?

Criminal justice is another area in which ubiquitous information sharing is more likely to produce greater inequalities than less. We've already seen the tracking of sex offenders after serving out a prison sentence made easier by technology, and it's not hard to imagine apps that would help citizens avoid all contact with former inmates. What does this do to the concept of rehabilitation and a "second chance"?

---

[1] http://www.exploringcs.org/about/related-grants/mobilize

Participatory sensing also raises the specter of increasing inequality in a variety of social sectors that involve profiling and demographic sorting. Previously, indicators such as skin color, gender, etc. were used to sort people into categories. With increased sensing capabilities, will we see the emergence of new marginalizations, new visibilities, and new indicators? With sensing, you can "see" so much more. Will those factors be used to divide and discriminate? For example, we discussed possible negative impacts on the insurance industry. The current trend in insurance is away from spreading the risk among a population, and towards profiling to quantify individual risks. Participatory sensing data ranging from driving habits to location-based indexing of environmental data could all increase the granularity of personal profiles. Similar sorting could impact the financial industry, risk management, and price discrimination (charging different prices based on the ability to pay).

In cases of social sorting (such as price discrimination), research has shown that people tend to find such sorting fair provided they can understand the system behind the sorting. With ubiquitous sensing, we accumulate piles of big data for mining. New characteristics emerge, which are used as proxies for willingness to pay, health risk, etc. These new categories may be perceived as unfair if they are difficult to understand. Indeed, the complexity of the algorithms used may be beyond explanation to non-statisticians and machine learning specialists.

4.2.3 Contribution: A Social Justice Impact Assessment

After this high-level discussion, we decided to construct a method for drilling down on specific applications to evaluate their potential impact on social justice. We discussed the factors we would need to make these classifications, including such questions as:

- Who is the target of the data collection? (Individuals, groups, things?)
- Who collects the data? Who analyzes it?
- What is the intended goal of the application?
- What forms of feedback are given (motivating vs. punishment)?
- Is the data aggregated?
- How distributed is control over the data?
- Are incentives given? Financial? Is the data collected with or without knowledge or consent?
- Who might be caught up or implicated without knowledge or recourse? Are there negative externalities to data collection?

Using these questions as a loose guide, we built on earlier assessment techniques suggested by Oetzel and Spiekermann [2]. The scenario of neighborhood sensing was chosen to exemplify a social justice impact assessment: users would contribute air quality data and self-reports of issues like allergies and asthma to challenge a city's existing air quality models.

The first step was to break the concept of "social justice" into smaller component parts. The group listed:

- Fairness
- Flourishing
- New opportunities
- Structural change
- Power Dynamics
- Plurality / diversity

We next chose "fairness" to break into even further sub-components (in a comprehensive assessment tool, a similar exercise would be carried out for each of the other five components). We defined these as:

- Equality
- "Just desserts" (meritocracy)
- Distributive justice
- Chance to reply
- Procedural justice
- Transparent processes

We then set out to see how each of these smaller concepts might impact the case of a neighborhood sensing application which included a location tracker to index a person's personal environmental impact, crowd-sourced measures of air pollution, crowd-sourced data about the state of roads, and self-reports of asthma rates. Again focusing on the first sub-component, we asked: how could the neighbor sensing app threaten equality? We identified threats (T's) including:

- T1: Distortion of facts leading to unequal funding
  - T1.1: Analysis done in a biased way
  - T1.2: Creation of biased samples
- T2: Indirect negative externalities to individuals
- T3: Direct negative externalities

We then identified control or mitigation strategies (C's) that could be built into the application.

- C1: Make raw data available (T1.1)
- C2: Statistician ensures the representative data sampling (T1.1)
- C3: Analysis algorithms should be published for scrutiny (T1.1)
- C4: Collect data that allows for meaningful transparency (T1.2)
- C5: Privacy controls (T2)
  - C5.1: Anonymization (and aggregation?) of individual data (T3)
  - C5.2: Giving the individual the choice to participate (T3. Influences T1.2, so a tradeoff exists)

4.2.4   Open Questions

During the course of the work, we identified several open questions that could spark future research. These include:

- If we assume that participatory sensing will lead to greater transparency, will such transparency equally impact individuals, powerful people, and institutions? For instance, should powerful officials or celebrities be subject to the same transparency needs as institutions?
- Do the powerful always end up subverting transparency schemes? Or can sensing change that tendency, for example by making facts visible to consumers and citizens, enabling organized responses (unionization)?
- When thinking about individual liberties vs. social action, does sensing technology push in one direction or the other? As someone nicely put it: "Ask not what sensing can do for you, but what sensing can do for your country."

- What are the forums for encouraging collective action in participatory sensing? Can we encourage system designers to consider social justice during design by framing design as a collective action problem? Can participatory sensing open new avenues for consumers and citizens to organize collective action?
- Could sensing data help us "diagnose" people's moral predispositions? (And therefore political behavior?)
- What factors in sorting and categorization processes make people feel that resulting algorithmic treatment is fair or unfair?

References

1 Laura Hueya, Kevin Walby, and Aaron Doyle. *Surveillance and security: technological politics and power in everyday life*, chapter Cop watching in the downtown eastside: exploring the use of (counter) surveillance as a tool of resistance, pages 149–165. Routledge, 2006.
2 Marie C. Oetzel and Sarah Spiekermann. Systematic methodology for privacy impact assessments. *European Journal of Information Systems*, July 2013.
3 John Rawls. *A theory of justice.* Cambridge, MA: Harvard University Press, 1999.
4 Amartya Sen. *The idea of justice.* Cambridge, MA: Harvard University Press, 2009.

4.3 "Tool Clinics" – Embracing Multiple Perspectives in Privacy Research and Privacy-Sensitive Design

Anthony Morton (University College London, GB)
Bettina Berendt (KU Leuven, BE)
Seda Gürses (KU Leuven, BE)
Jo Pierson (Free University of Brussels, BE)

4.3.1 Focalism – The Challenge

Computer scientists or engineers are continually asked to "solve problems" or "improve" existing situations, by selecting from available design features to produce the "best" technical solution. For example, a software developer faced with the problem of securing data must choose between different encryption algorithms – each with different characteristics. Factors such as strength of encryption, speed of encryption, usability, key management and hardware requirements must all be considered. Other requirements such as the sensitivity and amount of data to be protected, the estimated resources of potential attackers, the operational context of the required solution, etc. must also be taken into account. It is impossible for any solution to be 100% perfect, e.g. encrypting data with no detectable delay using an algorithm which cannot be broken. Trade-offs during the design and development process are therefore inevitable as requirements are balanced, e.g. speed vs. strength of encryption. These trade-offs are dilemmas faced by the specialist in arriving at the final design. However, what is the "best solution", and who decides what "best" means, requires more involved discussion and reflection. The engineer, with their narrow focus on solving the technical problem, might not be best equipped to solely decide what the optimum solution is, particularly if there are likely to be unintended consequences when the solution is deployed, or the proposed technology is decoded differently by users, those directly or indirectly affected, and other stakeholders.

The desire of specialists – particularly those in the fields of science or technology – to frame complex and messy situations as a single problem to be solved by technology – for which only they have the answer – often leads to overconfidence in the envisaged solution, an overemphasis on intended consequences, and a tendency to focus narrowly on one or a few aspects of the problem. This is typically identified as a form of "technological determinism", a perspective which consists of two parts: (1) the belief that technological developments take place outside society, independently of social, cultural, economic and political forces; and (2) the assumption that technological change causes or determines social change [31]. This kind of technologically deterministic approach can result in bigger problems than the one originally being solved because the understanding of the original problem situation was incomplete or wrong; Tenner [26] calls these the "unintended consequences" of technological innovation, e.g. the increasing resistance of certain strains of bacteria to antibiotics. However, unintended consequences are not restricted to technological innovation, but occur in political science, organisations, medicine and public health, ecology and social systems [11, 26]. Ehrlinger and Eibach [11] observe:

> "[F]ocalism, or a tendency to focus narrowly on one or a few variables, [...] with respect to the intended consequence can result in a neglect of important information regarding alternative, unintended consequences – including information that is knowable and plainly relevant to predictions" (p. 60)

Using a computer simulation, Ehrlinger and Eibach [11] showed that participants who were "defocused" by being encouraged to consider a wider system of variables, tended to make more accurate predictions and were less optimistic about the proposed solution. This suggests that viewing problems more holistically – particularly from multiple perspectives – can improve decision-making and increase the chances of successful technology development. Focalism – probably first suggested by Wilson et al. [30] – is essentially the same as "focusing illusion" proposed by Schkade and Kahneman [21] and Loewenstein and Schkade [14]. They found that when people are asked to predict their emotive reaction to a major event (e.g. the loss of employment), they typically concentrate on their likely responses to the focal event, to the exclusion of possible effects of other non-focal events (e.g. new opportunities to start a business or retrain). A practical example of people's tendency to ignore other events when their attention is focused elsewhere – inattentional blindness – is described the study by Simons and Chabris [23] in which most people missed a gorilla appearing during a video, when asked to concentrate on the number of times the ball was passed between particular basketball players.

We propose that the notion of focalism is equally applicable to scientists and technologists, who are often reluctant to challenge assumptions surrounding a problem, and principally concentrate on finding a solution to the problem as they perceive it, without adequate consideration of: (1) what it is that actually needs to be achieved – not from only one viewpoint; (2) any foreseeable consequences of the proposed solution; (3) and the viewpoints of other affected and/or interested actors who may have different priorities. We suggest this can be viewed as "solution focalism", and we propose that de-focusing may best be achieved by making other viewpoints salient As Genus observes, "*the employment of participatory approaches has been proposed to accommodate the interests of a wide range of actors holding different value positions, while minimising the potential risks associated with technology development.*" [12]

The problems of focalism are not restricted to technology development. It also reduces the efficacy of privacy research and privacy-sensitive design. For example, Privacy Enhancing

Technologies (PETs), such as Privacy Bird and Privacy Finder[2], appeared *prima facie* at the time to offer useful technical solutions to the problem of managing people's privacy. Both PETS use a protocol published in 2002 by the Platform for Privacy Preferences Project (P3P) [7] that enables web sites and applications to describe their privacy policy in XML. However, they have failed to become widely accepted and deployed. In 2003, the adoption rate of P3P was broadly flat at around 10% [10], partially due to the limited functionality of the first P3P user agents, and user interface problems [8]. Reay et al [18] observed that "*P3P adoption has stagnated in a niche position; it appears that browser implementers simply do not have enough market incentive to expend the resources needed to develop and integrate P3P 1.1 user agents*" (p. 162). Those browser implementers that did implement P3P made such fundamental technical mistakes that P3P was easily circumvented by publishing invalid policies [9]. Companies who chose not to use P3P suffered no consequences, which underlined the fact that P3P – albeit an elegant technical design – also required, as a minimum, enforcement external to itself, either through government regulation or industry self-regulation, both of which never materialised. The development of P3P may have benefited from collaborative design and development informed by a critical assessment of the perspectives of browser developers, the interests and technical capabilities of those who host and manage web sites, and the role of regulators. Certainly, there is much to be learned from the P3P experience that can be used to look at contemporary proposals for privacy-sensitive design. Focalism has also influenced the empirical aspects of privacy research. Many privacy studies have focused on the user experience with different interfaces and privacy controls, without thinking more holistically and considering the context in which the tool is used, the primary goals the user is trying to achieve, or the interaction of these goals with the interests of other affected stakeholders.

We propose a "tool clinic" to encourage a collaborative (re)consideration of a technological solution, research technique or other artefact, in order to critically assess its design, development and deployment from multiple perspectives. Another objective is to turn such solutions or artefacts into a tool for exploring the problem space. For example, what is the privacy problem when we look at it through a solution such as P3P? Finally, a tool clinic can be used to provide those who are developing the solutions with a setting to rethink the framing and presentation of their solutions. The term "tool clinic" emphasizes the motivation for embarking on this exercise. Athletes dedicated to improving some specific skill routinely go to a "rebound clinic" (in basketball) or a "dribbling clinic" (in football). The use of the word "clinic" does not indicate that a tool clinic provides a specific fix for problems, best practice guidelines, or solution templates – a typical panacea sought by those in the field of engineering. Rather, a tool clinic provides a framework and approach for multiple-perspective formative exploration and review of a technological solution, research technique or other artefact under development. The objective is to reflect from different perspectives on practices around the development, encoding, use, domestication, decoding and sustainability of a tool to gain quasi-ecological validation. In this sense, a tool clinic is more like a "law clinic", where law students study law and practice the adversarial legal process in context, or "design crits", during which designers learn to critique and receive critique of their work from others in the arts, academia or design practice.

---

[2] Privacy Bird was initially developed by AT&T. Privacy Bird and Privacy Finder are managed by Carnegie Mellon University's Usable Privacy and Security Laboratory.

4.3.2   Existing Uses of Multi-perspective Formative Exploration and Review

It is important to demonstrate that similar approaches to the suggested "tool clinic" are already used successfully in areas of industry and academia. This section describes some existing techniques that use a multi-perspective and collaborative approach.

In industry, disaster recovery practitioners often use corporate "war games" – a term originating from the military – to simulate a potential disaster situation (e.g. the loss of a data centre), and step through its disaster recovery plans to ensure they operate correctly. This avoids situations such as employees not being able to relocate to a cold-standby office building due to keys or swipe-cards not being readily available because the security department was excluded from disaster recovery planning. The use of disaster recovery simulations involving all affected areas of the business ensures disaster recovery plans are considered from multiple perspectives. A related technique to war games, the "Red Team"[3] review, also originated in the military as a means of assessing plans in an operational context from the perspectives of adversaries, affected areas of the military and their partners. Like war games, a Red Team review subjects a problem, plan, process, technique or artefact (e.g. tool, document, service, software product, etc.) to rigorous scrutiny by trained team members and experts. One of the authors of this report has been involved in Red Team reviews of complex commercial bid documents by the technical design and implementation, financial, service management and legal areas of a business organisation.

Gaining multiple perspectives is a technique also used by Soft Systems Methodology (SSM), which emerged in the 1980s from Checkland's work [5, 6]. SSM is a framework for organising the exploration of messy, complex problems as a learning *system*, and therefore failures in projects, processes etc. are viewed as a *systems failure*. Checkland [5] suggests that to fully understand a system it is necessary to consider its purpose from different viewpoints. This systemic pluralism represents one aspect of the "soft" systems approach, which aims to construct a rich picture of a problem, encompassing different viewpoints, rather than the reductionist focus of systems engineering. These different viewpoints, or *Weltanschauungen*, represent unquestioned models of the world that makes the system meaningful for study [5, 6]. It is important to stress that although SSM views problems as a *system*, it is not a representational model of reality; it is epistemological, not ontological; just because SSM views a situation *as if it were* a system, does not mean *it is* a system [6], e.g. a computer system.

To facilitate understanding of the reasons for failures, Checkland created the idea of a *formal system model* (FSM), which is a "*general model of any human activity system*" [5]. Comparison between the formal system model and the conceptual model of the problem situation under investigation is an intrinsic part of the SSM process, as it identifies flaws, weaknesses and omissions in the conceptual model, facilitating its improvement. The improved conceptual model can be compared with the real-world situation to determine which desirable or feasible changes are required [5, 6]. A project specific form of the FSM has been developed by Fortune et al [28] for use in analysing project failures, such as large-scale building projects [29].

The existing multi-perspective techniques described thus far, not only subject items to rigorous review, but encourage collaborative improvement and design. Soliciting the

---

[3] A "Red Team" is defined as "a team that is formed with the objective of subjecting an organisation's plans, programmes, ideas and assumptions to rigorous analysis and challenge. Red teaming is the work performed by the red team in identifying and assessing, inter alia, assumptions, alternative options, vulnerabilities, limitations and risks for that organisation." [1].

viewpoints of stakeholders, potential users of a technology or service, and those affected by it, can dramatically improve its quality. The notion of collaborative development and improvement to ensure effort is not expended on features or services that customers do not require, is key to the notion of "*the lean startup*" [19] used by many Internet companies. The lean startup philosophy suggests that companies release a "minimum viable product" – a "*version of a new product which allows a team to collect the maximum amount of validated learning about customers with the least effort*" [19] – to a subset of sympathetic customers, such as early adopters. The release of a minimum viable product is part of an iterative prototyping process, collecting suggestions for improvement, learning how customers use the product and what they want from it. The use of minimum viable products allows business to understand how customers actually decode the technology or service being provided; the product must be viable in that the customer must value what it provides. Use of minimum viable products should be an iterative learning process, generating ideas and collecting data about product use.

One existing approach to answer the question posited earlier, "*Who decides what 'best/better' really means?*" is constructive technology assessment (CTA). The latter fits within the long-standing tradition of Science and Technology Studies (STS), which investigates how the things that it studies are being constructed. The STS domain has increased its scope over the years, starting with scientific knowledge and expanding to artefacts, methods, materials, observations, phenomena, classifications, institutions, interests, histories, and cultures [24]. One of the most prominent ways to apply the thinking in STS in the real world has been the CTA approach. The objective of CTA is to "*produce better technology in a better society*" [12] by taking a more social constructionist position, and moving "*beyond technological determinism towards an evolutionary view of technology development*" [12]. This is done by advising on interventions in early stages of technology development based on the assessment of possible problems and risks that these technologies could pose for society [25]. CTA emphasises the importance of including a wide range of actors to anticipate the potential impact of a technological development ("vermaatschappelijking" of technology [27]) and decide on improvements to it, thus facilitating social learning. It should be stressed that CTA is not a research method, but an overall approach into which participatory techniques may be placed. Genus [12] suggests moving away from the interventionist and prescriptive stance of existing CTA approaches towards a more discursive, democratic and reflective process because "*contention and openness to criticism are prerequisites for producing reflective socio-technical expertise*" [12]. This is also known as "participatory technology assessment" [13]. The use of a modified form of CTA to address the ethical problems caused by technology is proposed by Palm and Hansson [16] as part of a continuous dialogue between developers and affected actors. For emerging technologies, Merkerk and Smith [27] propose a three-step CTA approach, using permutated dialogue workshops attended by insiders and outsiders to the item under review to consider selected issues about the proposed technology and reflect on different technology scenarios.

In order to apply multi-perspective formative exploration and review of technological solutions or tools in early stages of development, different types of multi-method approaches have been developed. One of the most elaborate ones is the living laboratory approach. The 'living lab' is a specific type of test and experimentation platform (TEP), which refers to facilities and environments for (joint) innovation including testing, prototyping and confronting technology with usage situations [3]. Living labs are facilities for designing, developing, testing and evaluating communication technologies and services in early stages of the innovation process. They do so by involving (early) users, in line with the CTA

perspective. However they can also be configured as open and innovation-oriented platforms that involve various technology experts, disciplines and/or stakeholders in different stages of technology design, development and testing [17]. Thus, we discern three main ways to put living labs[4] into action as: (1) a platform for open innovation; (2) a user-driven research methodology; and (3) an experimental setting [20].

4.3.3   Perceived Research Gap in Privacy

Most privacy researchers agree that privacy is contextual and dependent upon information use, information sensitivity and the trust in the entity collecting, storing, processing and disseminating the information entrusted to it [2]. Furthermore, users engaged in technology mediated interactions with other parties will have expectations and assumptions about the technology, the providing organisation and other partners in communication [2]. If these assumptions and expectations are violated, the user is likely to have an emotional reaction and reject the technology and/or providing organisation [2]. A practical example of this was the launch of Google Buzz. Gmail users believed they were only signing onto Gmail as usual, when they were actually being enrolled in Google Buzz [22]. It would appear the developers of Buzz did not take into account: (1) that people's primary task was to access their e-mail and hence they would likely "swat away" any dialogue boxes without properly reading them; and (2) that people's mental model is that Gmail is a tool to access their e-mail and not a social networking service.

User studies may aid developers and designers in foreseeing likely troubles that users may have with a given design. However, the task of achieving an understanding of the complexity of the privacy problem, and translations of this problem into the technical solution space may benefit greatly from a multi-perspective approach. This is line with the notion of *contextual integrity* (CI) by Nissenbaum [15], which is used to answer whether a situation contained a privacy breach or not. CI is guided by norms of appropriateness (i.e. norms that govern what can be disclosed in a certain context or situation) and norms of distribution (i.e. norms which assess the transfer of personal information from one party or context to another context). This demonstrates how not all publicly revealed information or information collected in the public space, is meant for every form of public use. "*Just because something is publicly accessible does not mean that people want it to be publicized. Making something that is public more public is a violation of privacy.*" [4]

Addressing the privacy implications of increasingly complex, powerful and ubiquitous computing will be even more of a challenge than Buzz, as the potential for unintended consequences is even greater than before. However, privacy researchers and practitioners continue to work largely in isolation, concentrating on people's use of different user interfaces for privacy control, and have largely ignored existing cross-disciplinary collaboration techniques such as those described above.

4.3.4   Future Directions for Researchers and Practitioners

Tool clinics are essentially practices, and they need to be living practices – thus future directions are not only researching, but also must be *doing* tool clinics. We have performed a first *ad hoc* requirements analysis for tool clinics at the Dagstuhl Seminar itself (i.e. we "clinicked"

---

[4]   In Europe living labs are associated in the European Network of Living Labs (ENoLL) which was set up under the auspices of the Finnish EU presidency in 2006 and since the 6th wave of call for new members in March 2012 consists of over 300 accepted members.

the tool clinic idea) and have seen the challenges the concept poses. Most importantly, our clinic participants expressed concerns about exposing their methods, approaches and original ideas to a critical audience. Further issues were raised with respect to matters of intellectual property. Some of these problems are likely to stem from the employment requirements and the working conditions of senior and junior researchers. They also often associated the word "clinic" with doctoring their (software) artefacts with others, a goal that we only partially share.

Based on this experience, our next step will be to develop a tool clinic as a new event format for a scientific conference, ideally at a renowned computer-science conference. This will combine the tool-centric nature of a demo session, the protected space of work-in-progress afforded by a workshop, and the mentoring spirit of a doctoral workshop[5].

The format of a tool clinic session could typically consist of three steps (inspired by the CTA and Privacy by Design approach):

1. Identifying particular affordances of the technological solution, research technique or other artefact and possible (unintended) consequences for people and society;
2. Gathering perspectives and practices of different experts, disciplines and/or stakeholders (e.g. users, policy makers, industry, etc.) linked with the development, deployment and sustainable evolution of a particular tool, solution, technique or artefact;
3. Informing and advising on technological design of the tool or solution, in order to avoid negative consequences and to further positive outcomes.

We foresee three essentially needed incentives for participation: (1) enlisting big names in the field who can signal through their own example that "grown-ups too can learn"; (2) a broad-enough team of participants to represent a wide range of perspectives; and (3) a follow-up that makes it worthwhile to put oneself into the ring. For the first two, we can draw on our respective scientific networks. A special issue in a good journal is one option for creating the third incentive, and further developments of the tool clinic method described in the introductory article of this special issue are among the next intended research activities.

4.3.5   Acknowledgement

References
1   Red Teaming Guide. Technical report, UK Ministry of Defence (2nd Edition), 2013.
2   Anne Adams and Martina Angela Sasse. Privacy in multimedia communications: Protecting users, not just data. In *People and Computers XV – Interaction without Frontiers*, pages 49–64, 2001.
3   Pieter Ballon, Jo Pierson, and Simon Delaere. *Designing for Networked Communications: Strategies and Development*, chapter Fostering Innovation in Networked Communications: Test and Experimentation, pages 137–166. IGI Global, 2007.
4   Danah Boyd. Making sense of privacy and publicity. Technical Report MSR-TR-2010-25, Microsoft Research, 2010.

---

[5] In this way the tool clinic approach has some resemblance with a 'crit' as done in art schools. This is a critique session, in which a student's artwork is formally presented to and evaluated by a group of faculty and peers, responding with feedback: comments, questions, advice, cheers, jeers, and tears.
[6] www.spion.me

5   Peter Checkland. *Systems thinking, systems practice.* John Wiley & Sons Ltd., 1981.

6   Peter B. Checkland and Jim Scholes. *Soft Systems Methodology in Action.* John Wiley & Sons Ltd., 1999.

7   Lorrie Faith Cranor. *Web privacy with P3P - the platform for privacy preferences.* O'Reilly, 2002.

8   Lorrie Faith Cranor. P3P: Making Privacy Policies More Useful. *IEEE Security and Privacy*, 1(6):50–55, November 2003.

9   Lorrie Faith Cranor. Necessary but not sufficient: Standardized mechanisms for privacy notice and choice. *Journal of Telecommunications and High Technology Law*, 10(2), December 2012.

10   Lorrie Faith Cranor, Simon Byers, and David Kormann. An Analysis of P3P Deployment on Commercial, Government, and Children's Web Sites as of May 2003. Technical report, AT&T Labs-Research, 2003.

11   Joyce Ehrlingera and Richard P. Eibachb. Focalism and the failure to foresee unintended consequences. *Basic and Applied Social Psychology*, 33(1):59–68, 2011.

12   Audley Genus. Rethinking constructive technology assessment as democratic, reflective, discourse. *Technological Forecasting and Social Change*, 73(1):13–26, 2006.

13   Simon Joss and Sergio Bellucci, editors. *Participatory technology assessment: European perspectives.* Centre for the Study of Democracy, University of Westminster, 2002.

14   George Loewenstein and David Schkade. Wouldn't it be nice? predicting future feelings. *Well-being: The foundations of hedonic psychology*, pages 85–105, 1999.

15   H. Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life.* Stanford University Press, 2009.

16   Elin Palm and Sven Ove Hansson. The case for ethical technology assessment (eTA). *Technological Forecasting and Social Change*, 73(5):543–558, 2006.

17   Jo Pierson and Bram Lievens. Configuring living labs for a 'thick' understanding of innovation. *Ethnographic Praxis in Industry Conference Proceedings*, 2005(1):114–127, 2005.

18   Ian K. Reay, Patricia Beatty, Scott Dick, and James Miller. A Survey and Analysis of the P3P Protocol's Agents, Adoption, Maintenance, and Future. *IEEE Transactions on Dependable and Secure Computing*, 4(2):151–164, April 2007.

19   Eric Ries. *The Lean Startup: How Constant Innovation Creates Radically Successful Businesses.* Penguin Books Limited, 2011.

20   S.C. Sauer. *User innovativeness in living laboratories: everyday user improvisations with ICTs as a source of innovation.* PhD thesis, Universiteit Twente, Enschede, September 2013.

21   David Schkade and Daniel Kahneman. Does living in california make people happy? a focusing illusion in judgments of life satisfaction. *Psychological Science*, 9(5):340–346, September 1998.

22   Maggie Shiels. Google buzz 'breaks privacy laws' says watchdog. BBC News, February 2010.

23   Daniel J. Simons and Christopher F. Chabris. Gorillas in our midst: sustained inattentional blindness for dynamic events. *Perception*, 28(9):1059–1074, 1999.

24   Sergio Sismondo. *The handbook of science and technology studies*, chapter Science and Technology Studies and an Engaged Program. The MIT Press, 2008.

25   Wim A. Smit and Ellen C.J. van Oost. *De wederzijdse beïnvloeding van technologie en maatschappij: een Technology Assessment-benadering.* Bussum: Coutinho, 1999.

26   Edward Tenner. *Why things bite back: technology and the revenge of unintended consequences.* Vintage Books, 1997.

27   Rutger O. van Merkerk and Ruud E.H.M. Smits. Tailoring CTA for emerging technologies. *Technological Forecasting and Social Change*, 75(3):312–333, 2008.

28 Diana White and Joyce Fortune. The project-specific formal system model. *International Journal of Managing Projects in Business*, 2(1):36–52, 2009.

29 Diana White and Joyce Fortune. Using systems thinking to evaluate a major project: The case of the gateshead millennium bridge. *Engineering, Construction and Architectural Management*, 19(2):205–228, 2012.

30 Timothy D. Wilson, Thalia Wheatley, Jonathan M. Meyers, Daniel T. Gilbert, and Danny Axsom. Focalism: A source of durability bias in affective forecasting. *Journal of Personality and Social Psychology*, 78(5):821–836, 2000.

31 Sally Wyatt. Technological determinism is dead; long live technological determinism. In Edward J. Hackett, editor, *The handbook of science and technology studies*. MIT Press, 2008.

4.4 Consequence-based Privacy Decisions: a New Way to Better Privacy Management

Zinaida Benenson (Universität Erlangen-Nürnberg, DE)
Delphine Christin (TU Darmstadt, DE)
Alexander De Luca (LMU München, DE)
Simone Fischer-Hübner (Karlstad University, SE)
Thomas Heimann (Google - München, DE)
Joachim Meyer (Tel Aviv University, IL)

4.4.1 Introduction and Motivation

An increasing number of users contribute privacy-sensitive content, such as pictures, comments, or location information, to online services. In order to protect the privacy of the users or to comply with data protection regulations, most services enable the users to customize privacy and sharing preferences. For example, this includes determining who will be authorized to access or receive and process which content and for which purposes. However, management of privacy preferences is often a fairly complex procedure that even technically savvy users often fail to understand.

Recent research shows that people would like to control their privacy and actually do so. For example, the number of Facebook users with customized privacy settings has been growing in the last years. However, users are frequently unaware of consequences resulting from their selected configuration and cannot be sure that the changes will actually have the effects they are intended to have. In addition, many users do not set or adapt privacy settings as they cannot correctly grasp the consequences of their actions. For instance, tagging a person on a photo may cause this photo to appear in searches of this person, which may be at time unwanted. Although recently some tools for granular privacy management emerged, the problem of determining all the consequences at the system level and showing them to the users in an understandable and actionable way still remains largely unsolved.

We argue that an appropriate privacy-respectful user interface should show users the consequences of making different privacy choices, rather than framing the choices only in technical terms regarding system parameters which users often do not understand and do not care about.

We believe that providing tools to increase user comprehension of potential consequences is one of the next big challenges to be addressed in the field of privacy respectful user interfaces. In addition to helping users to make better choices in terms of privacy protection, this will also allow them to make better informed decisions and hence, implement the notion of informed consent (that is often required pursuant to Art. 7 European Data Protection Directive 95/46/EC) not only formally but also to live up to the spirit of this legal requirement.

4.4.2   Challenges and Research Directions

The attempt to develop user interaction regarding privacy in which the user is clearly aware of the consequences of actions requires research on a number of issues that have so far received relatively little attention:

Expression of potential consequences

Informing the users should take into account several parameters. Indeed, users may have different backgrounds and education levels. As a result, the consequences should be expressed in a way comprehensible by different user categories from novices to expert users. This may include translating potential consequences into different metrics. Such metrics do not exist at this time. For instance, if a user decides to share his/her location, the interface could display a list of people that will be able to see this location. This list could include close relatives as well as remote friends and unknown people. In the computation of relevant metrics, the context, e.g., the user's location, will need to be taken into consideration as the notion of privacy depends on the context. Finally, the consequences should be displayed in usable interfaces.

Decision support

In addition to displaying potential consequences, users could be further helped in their privacy decisions by external information sources. This could include showing the privacy decisions of their relatives, friends or other expert and non-expert users (e.g., via crowdsourcing). By doing so, users may have a social reference and make better informed decisions. On the other hand, this may influence them to disclose more sensitive data than they initially intended to disclose. As a result, studies to determine the responses to different kinds of information sources, different formats, and information from different groups of users will be necessary. For example, trade-offs between privacy and social compliance in case of crowdsourcing need to be investigated.

Minimal effort

Introducing additional tools to help users make informed decisions may add significant overhead to the interaction. While this overhead may be the price to pay for better privacy protection, it should be limited to the minimum. Otherwise users may be tempted to rush through the configuration and ignore this additional step. In this context, habituation effects are a serious problem that has to be taken into account. As a result, novel interaction schemes that are robust to such effects need to be investigated and developed to provide users with appropriate tools.

4.4.3   Activities of the Group

During the workshop the group identified the issue of consequence-based privacy decisions as a topic with great possible potential. After discussing the topic we developed a conceptual model and identified the major challenges that need to be addressed if one wants to implement consequence-based privacy decisions.

Furthermore we began to work on a joint conceptual paper, presenting the problem of users being unable to predict the implications of privacy decisions they make. The paper proposes some directions which may be taken to build a system that provides users with information about the consequences of their actions. We also discussed possible directions for future joint research resulting from the workshop. This includes organization of follow-up meetings of workshop participants and the search for possible funding sources for research, based on the ideas developed in the workshop.

Participants

- Alessandro Acquisti
  Carnegie Mellon University, US
- Mads Schaarup Andersen
  Aarhus University, DK
- Zinaida Benenson
  Univ. Erlangen-Nürnberg, DE
- Bettina Berendt
  KU Leuven, BE
- Claudio Bettini
  University of Milan, IT
- Rainer Böhme
  Universität Münster, DE
- Ian Brown
  University of Oxford, GB
- Claude Castelluccia
  INRIA Rhône-Alpes, FR
- Delphine Christin
  TU Darmstadt, DE
- Alexander De Luca
  LMU München, DE
- Tassos Dimitriou
  Athens Information Techn., GR
- Frank Dürr
  Universität Stuttgart, DE

- Deborah Estrin
  Cornell Tech NYC, US
- Simone Fischer-Hübner
  Karlstad University, SE
- Michael Friedewald
  Fraunhofer ISI – Karlsruhe, DE
- Raghu K. Ganti
  IBM TJ Watson Res. Center –
  Yorktown Heights, US
- Jens Grossklags
  Pennsylvania State Univ., US
- Seda F. Gürses
  KU Leuven, BE
- Thomas Heimann
  Google – München, DE
- Ioannis Krontiris
  Goethe-Universität Frankfurt am
  Main, DE
- Marc Langheinrich
  Univ. della Svizzera italiana, CH
- Renè Mayrhofer
  University of Applied Sciences
  Upper Austria, AT

- Joachim Meyer
  Tel Aviv University, IL
- Anthony Morton
  University College London, GB
- David Phillips
  University of Toronto, CA
- Jo Pierson
  Free University of Brussels, BE
- Sören Preibusch
  Microsoft Res. – Cambridge, GB
- Kai Rannenberg
  Goethe-Universität Frankfurt am
  Main, DE
- Norman Sadeh
  Carnegie Mellon University, US
- Martina Angela Sasse
  University College London, GB
- Marcello Paolo Scipioni
  University of Lugano, CH
- Katie Shilton
  University of Maryland –
  College Park, US
- Sarah Spiekermann
  Universität Wien, AT