D4.6 & D4.8 Design and evaluation of context privacy controls

Project	SWELL
Project leader	Wessel Kraaij (TNO)
Work package	WP4
Deliverable number	D4.6 & D4.8
Authors	Bob Hulsebosch (Novay), Milena Kooij (TNO), Ruud Kosman (Novay), Joris Janssen (Sense-OS)
Reviewers	Saskia van Dantzig (Philips) & Jan Geert van Hall (Almende)
Date	December 2013
Version	0.4
Access Rights	Public
Status	Final

SWELL Partners:

Noldus, Novay, Philips, TNO, Radboud Universiteit Nijmegen, Roessingh Research and Development, Almende, Sense Observation Systems, Universiteit Twente.

Summary

Applications for well-being and well-working typically consume all kinds of context or sensor information such as location, calendar info, and heart rate. Due to the privacy-sensitive nature of this information, end-users must be given control over which applications and users get access to (which part of) this information. Current solutions do not offer this or only provide some form of (policy-based) pre-provisioning of privacy policies. These solutions are too static, lack transparency and accountability, and do not empower the user.

To meet these requirements for implementing user controlled privacy in the CommonSense platform, several mock-ups were designed and evaluated by a small group of users to obtain feedback on the designs. Though the users had a lot of suggestions for improvement of the designs, they expressed their confidence in the taken approach of creating transparency, control and accountability. Additional enhancements are required for making data sharing with other users and applications easier and more trustworthy.

Sharing should be made easier in terms of reduced intrusiveness. Future work in SWELL will therefore focus on the design and implementation of an intelligent privacy reasoner component that uses context information to make the control features more adaptive and less intrusive.

Furthermore, sharing should be made more trustworthy in terms of being confident about the real identity of the user/application the data is shared with: “Am I sharing my data with my doctor or with someone else?”. Somehow the user should be able to authenticate the user or application he/she is willing to share data with.

These two functionalities and the other outcomes of the user evaluation will be taken into account in a next iteration of new wireframe designs. These will be used as starting points for the actual implementation of privacy control in the CommonSense platform.

Contents

1 Introduction

In the SWELL project, a range of well-being and well-working applications are envisioned that involve sharing of personal information in different contexts, including medical, working and home contexts, and involving different social contexts. In order to preserve the privacy of the individuals using the SWELL applications, the policies that govern what information may be shared must be maintained in a manner sensitive to the context in which the application is used. In the previous deliverables of the privacy work package of the SWELL project (Bokhove et al., 2011; Hulsebosch et al., 2012a; Hulsebosch et al., 2012a), we discussed general requirements for following the Context Aware Adaptive Privacy approach (D4.1), a broad inventory and discussion of privacy control mechanisms (D4.2), technological issues and challenges when implementing this approach (D4.3), and guidelines for (context-based) informed consent in wellbeing applications.

The present deliverable combines this knowledge base for an initial design and implementation of privacy control features for well-being and well-working applications that consume sensor data. Privacy *control* over the way sensor data is handled is a major requirement from the perspectives of user acceptance and legal compliance. But it is not the only one. Two other major requirements must be taken into account as well:

- *transparency* on how data is handled,
- and *accountability* of companies.[1]

This report focuses on the design and evaluation of these requirements in the SWELL sensor data platform that provides the user's well-being and well-working applications with data.

1.1 Distributed vs centralised privacy control

There is a vast number of different mobile apps. An average smartphone user is reported to have installed about 37 apps. These apps are able to collect large quantities of data from the device (e.g. data stored on the device by the user and data from different sensors including location) and process these in order to provide new and innovative services to the end user. Moreover, data may not only be limited to that on the smartphone. Numerous other data sources from external sensors or online process are available and useful for consumption as well.

Being able to be in control over what kind of data is shared with which applications, when and for what purposes is not trivial. The mobile smartphone offers only limited functionality in this area. Moreover besides lack of control due to lack of transparency there are other weaknesses of current mobile app and sensor data approach. One of them is the lack of free and informed consent. Once the app is downloaded, consent is often reduced to a tick box enabling the end user to accept a more or less all or nothing terms and conditions concerning data consumption without having a more granular choice. Another weakness is the poor security measures of the app that may lead to

[1] Dutch Ministry of Economic Affairs, Brief Kabinetsvisie op e-privacy: op weg naar gerechtvaardigd vertrouwen, 24 May 2013. Accessible online (in Dutch): http://www.rijksoverheid.nl/bestanden/documenten-en-publicaties/kamerstukken/2013/05/24/kamerbrief-met-kabinetsvisie-op-e-privacy/brief-kabinetsvisie-op-e-privacy-op-weg-naar-gerechtvaardigd-vertrouwen.pdf.

unauthorised processing of (sensitive) personal data. Poor security measures, an apparent trend towards data maximisation and the elasticity of purposes for which personal data are being collected further contribute to the data protection risks found within the current app environment. Furthermore, there is often little awareness of the types of processing an app may undertake, combined with a lack of meaningful consent from end users before processing takes place. Once downloaded, mobile apps are able to collect large quantities of personal data from the users' device, for example by having access to the photo album or using location data.

Based on these observations it does make sense to look for more centralised privacy control solutions that are designed to give a web user a unified control point for authorizing who and what can get access to their online personal data (such as identity attributes), content (such as photos), and services (such as viewing and creating status updates), no matter where all those things live on the web. Further, such a central solution should allow a user to make demands of the requesting side in order to test their suitability for receiving authorization. These demands can include requests for information (such as "Who are you?" or "Are you over 18?") and promises (such as "Do you agree to these non-disclosure terms?").

In SWELL therefore a centralised approach to sensor and other data aggregation and sharing is adopted.

1.2 SWELL sensor platform

In SWELL sensor data is aggregated by the CommonSense platform. CommonSense is a platform that helps the user to keep track of all his sensor data, store it in a central location, and play with it. CommonSense also processes raw sensor data into meaningful things like sleep, exercise, or top locations. This results in new information feeds that applications can use for their benefits. Via the CommonSense Dashboard the user gains insights into his raw sensor data and derived streams (see Figure 1).

With the CommonSense Tracker mobile app the user can turn his mobile phone into an advanced tracking device. Via this app data of sensors that are available in the mobile phone, like GPS-location, accelerometer, magnetometer, gyroscope, etc. are uploaded to the CommonSense platform. Furthermore, other sensor data sources like Fitbit and Twitter can be added as well. The data streams of these sensors are processed to get steps, sleep, exercise, locations, and sociality scores.

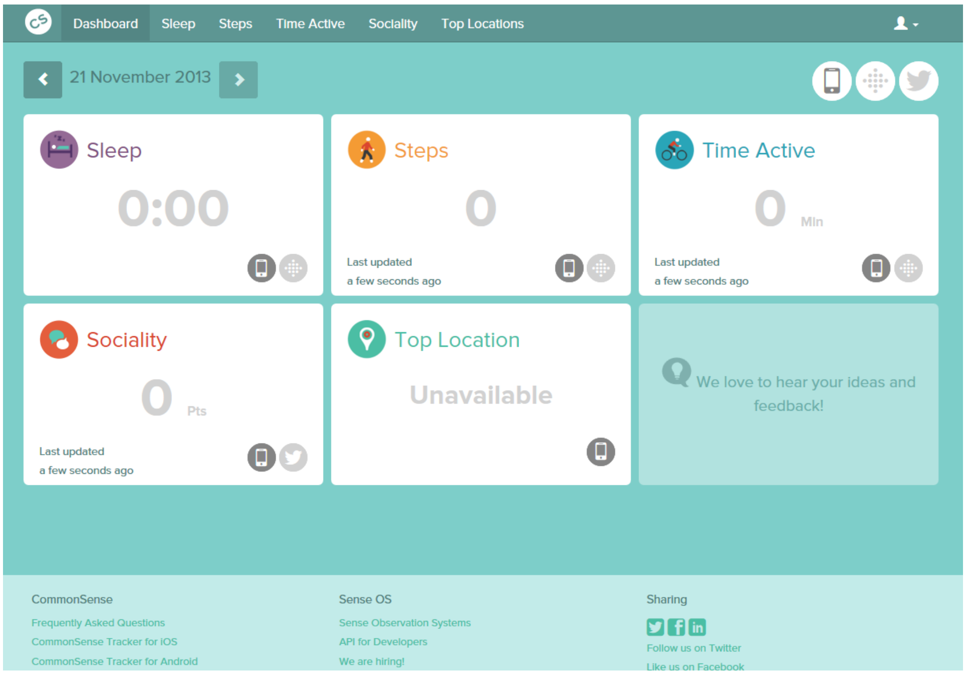

Sleep
0:00

Steps
0
Last updated
a few seconds ago

Time Active
0 Min
Last updated
a few seconds ago

Sociality
0 Pts
Last updated
a few seconds ago

Top Location
Unavailable

We love to hear your ideas and feedback!

Figure 1: CommonSense Dashboard.

1.3 Privacy in CommonSense

The current publicly available version of the CommonSense platform barely provides privacy features. The only feature is the ability to disable the connection of sensors from the platform and to remove an account including its sensor data. This does not satisfies the requirements for the privacy management within SWELL. Hence, new functionality has to be added to the CommonSense platform.

Expected future extensions of the CommonSense platform functionality that include sharing of sensor data with apps and other users will require additional privacy functionality, particularly in the areas of privacy control, transparency and accountability. This deliverable provides some initial designs for privacy features in CommonSense that meet these requirements. An additional challenge for these privacy features is that they must be user friendly and intuitive, otherwise the user will lose confidence in the platform or will even stop using it.

2 Design of privacy in CommonSense

Key privacy features are control, overview and accountability. For each of these features several design solutions will be presented.

2.1 Control features

2.1.1 Consent

It is important for users of SWELL applications to understand why, where and how their data is being processed. They also must have the possibility to give or withhold consent to such processing. Otherwise the SWELL users don't have meaningful control over their personal data, which is essential in the sensitive context of health-related personal data. This combination of transparency and control is usually described as 'informed consent'. Consent is meaningless if the user is unable to know or understand what he or she gives consent to.

In D4.4 of SWELL a number of components for informed consent were assessed and guidelines were provided. The components are:

- Disclosure
- Comprehension
- Voluntariness
- Competence
- Agreement
- Minimal distraction[2]

The first two components (disclosure and comprehension) refer to *informed* consent, while the following three components (voluntariness, competence and agreement) refer to *consent*. The final component recognizes the importance of the tasks that users want to perform in an application or using a service. These six components of informed consent describe in essence a best practice for informed consent mechanisms: a good mechanism incorporates all the components.

The main challenge in designing a consent mechanism, viewed in the context of the assessment framework used here, is in providing full disclosure to the user, offering a meaningful choice, while distracting the user as little as possible from the task at hand. While there is a tension between these two aims (involving the user in making an informed choice and not interrupting the user), some of the examined mechanisms appear to overcome this tension.

Some key aspects of good consent mechanisms are:

1. using intuitive, visual controls;
2. offering a number of layered choices;
3. providing a brief and easy to comprehend overview of the consequences of each choice.

Each of the consent mechanisms is based on setting a *policy* that a system or platform should follow, rather than providing consent at each and every instance in which personal information is shared. One example of this is the authorization for third party access used by Twitter, and some

[2] Friedman, Batya, Peyina Lin, and Jessica K Miller, "Informed Consent by Design," No. 2005, 2005, pp. 503–530.

comparable services. By authorizing the third party's access, the third party may, under certain conditions, use the personal information disclosed by Twitter. This consent, and hence the policy of sharing personal data with the third party, may be revoked at any time.

Based on these consent components and the guidelines for implementing them several wireframes for consent have been designed. These wireframes are initially based for desktop screens as they are primarily meant to get a good insight in the required and desired privacy control functionalities. Later on they will have to be tailored to smartphone screens as well. The limited size of these screens will pose additional requirements to the designs in order to keep them user-friendly and practical.

In the wireframe below, a pending data sharing request is shown. First, it identifies the requesting party by e-mail address. Next, for each different measured state (e.g., steps, sleep, locations, activity) the sharing policy is specified. The sharing can be limited based on the context, both in space and in time. In other words, one can share on certain days or times, and/or when one is at a certain location. Subsequently, the timeframe parameter specifies which period of time from the data can be accessed (e.g., only future data, only historical data, or any data). These features provide the user with a lot of control over when exactly what data should be shared.

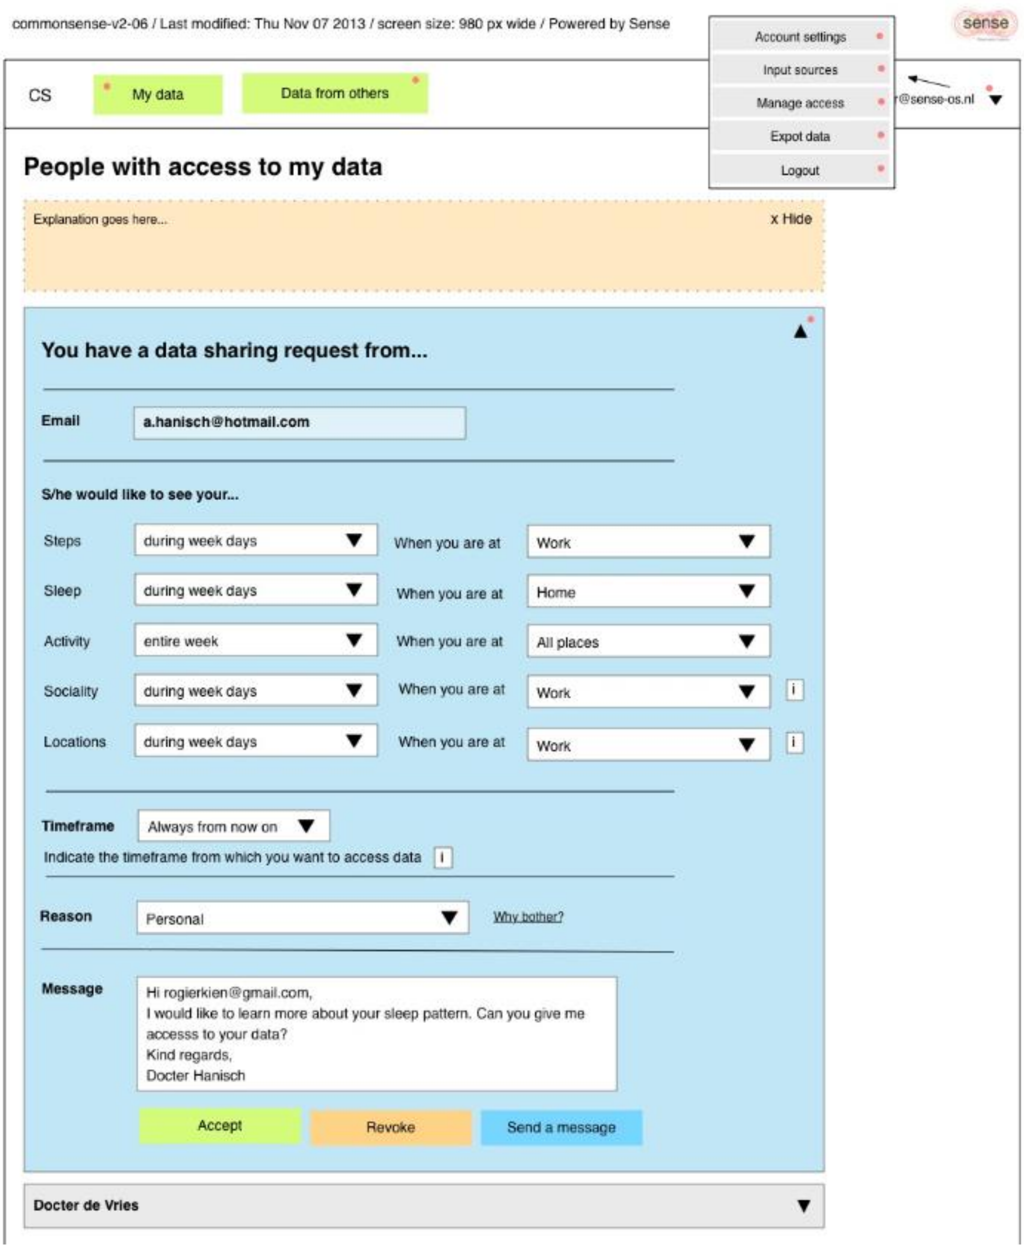

commonsense-v2-06 / Last modified: Thu Nov 07 2013 / screen size: 980 px wide / Powered by Sense

People with access to my data

Explanation goes here...

You have a data sharing request from...

Email: a.hanisch@hotmail.com

S/he would like to see your...

Steps	during week days	When you are at	Work
Sleep	during week days	When you are at	Home
Activity	entire week	When you are at	All places
Sociality	during week days	When you are at	Work
Locations	during week days	When you are at	Work

Timeframe: Always from now on

Indicate the timeframe from which you want to access data

Reason: Personal

Message:
Hi rogierkien@gmail.com,
I would like to learn more about your sleep pattern. Can you give me accesss to your data?
Kind regards,
Docter Hanisch

Each request requires a specified reason. This can range from medical reasons all the way to marketing. Application developers can use these reasons in their terms of service to limit the use of the data to a certain specified reason. There can be a lot of reasons for sharing. Unfortunately there isn't an agreed or generic semantic framework that can be used to express 'reasons' in a standardised manner. For the moment it is open to individual interpretation as to what a reason entails. Finally, there is room for a message and the accept of ignore buttons. Given the complexity of the conditions, contradictions may easily be introduced and must be resolved in the CommonSense platform.

In the wireframe the user requesting access to the data is identified via an email address. For the owner of the data this may not be sufficient to be sure that it really is his doctor. A number of obvious use cases exist that require a high level of security or carry high liability consequences. In these cases the designed access control mechanisms in which a user types email addresses into an access control list and that is used by the app to email a "secret link" or "private URL" to those people, does not suffice for controlling access to healthcare records, tax data, or perhaps even university transcripts. A practical system is needed for authenticating the source of third-party-asserted claims, evaluating their quality, ensuring that they apply to the requesting party in question, and transferring claims in a privacy-sensitive way only for policy decision-making purposes – all in a way that reaches web-wide scalability. Moreover, delegating access to third party users not always suffices; in most cases it is an application of a certain user that wants access: How can Alice type “bob@gmail.com” into an access control list to govern access to her travel calendar, such that whatever requester app Bob is using can lead him down a path that ends up with e.g. Google generating a trustworthy claim that he’s the guy who matches the policy? Similarly, when the survey company is the requesting party, how can Alice ensure that the party agreeing to her privacy and data usage policies is really representing that company, so she can successfully sue the company for breach of contract if they sell her data? A solution direction that can be thought of are the use of third party authentication service providers. These providers can authenticate a user or app and can make identity assertions towards the CommonSense platform. The solution direction is described in more detail by the Kantara initiative’s User Managed Access work group[3]. The possibilities of this and potential other solution directions need to be analysed in future work of SWELL.

Following the symmetry privacy principle (see D4.2), the user may also want to request access to data of others. This is shown in the wireframe below. Also shown is a pending request for such data access; jeroenhoorn@gmail.com still has to confirm a request. Here another weakness of consent emerges: lack of responsiveness may delay service effectiveness. Particularly in cases were data sharing need to be arranged on the fly it is not desired to wait too long for consent confirmation. In such cases, communication channels other than the CommonSense platform should be exploited to make the user aware of the pending request. In the CommonSense implementation, we do this through e-mail reminders, and by giving the requesting party the option to send a reminder with one simple click.

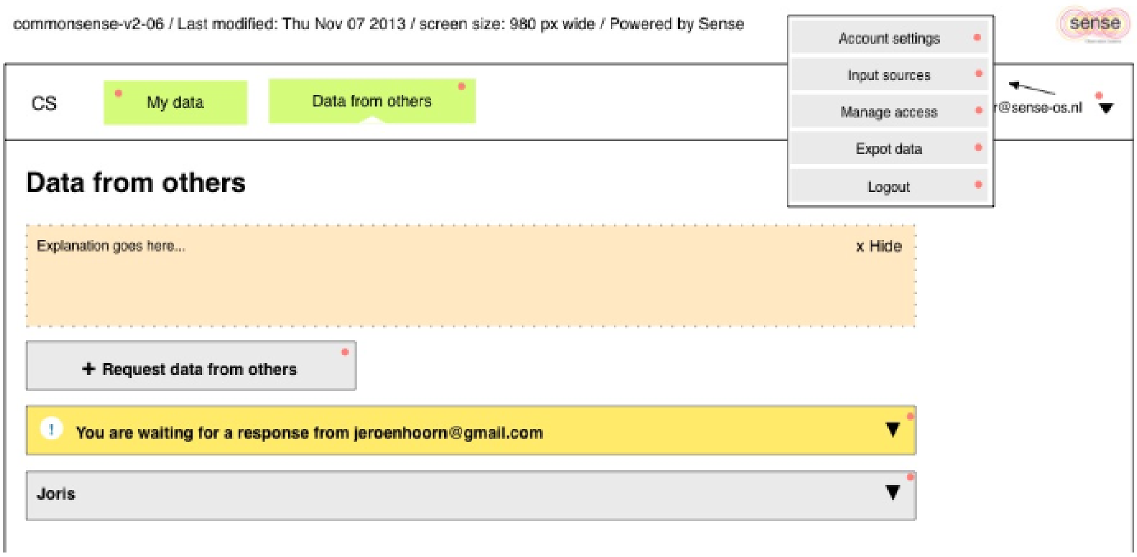

commonsense-v2-06 / Last modified: Thu Nov 07 2013 / screen size: 980 px wide / Powered by Sense

Data from others

Explanation goes here...

You are waiting for a response from jeroenhoorn@gmail.com

[3] UMA work group charter and specifications, see https://kantarainitiative.org/confluence/display/uma/Home.

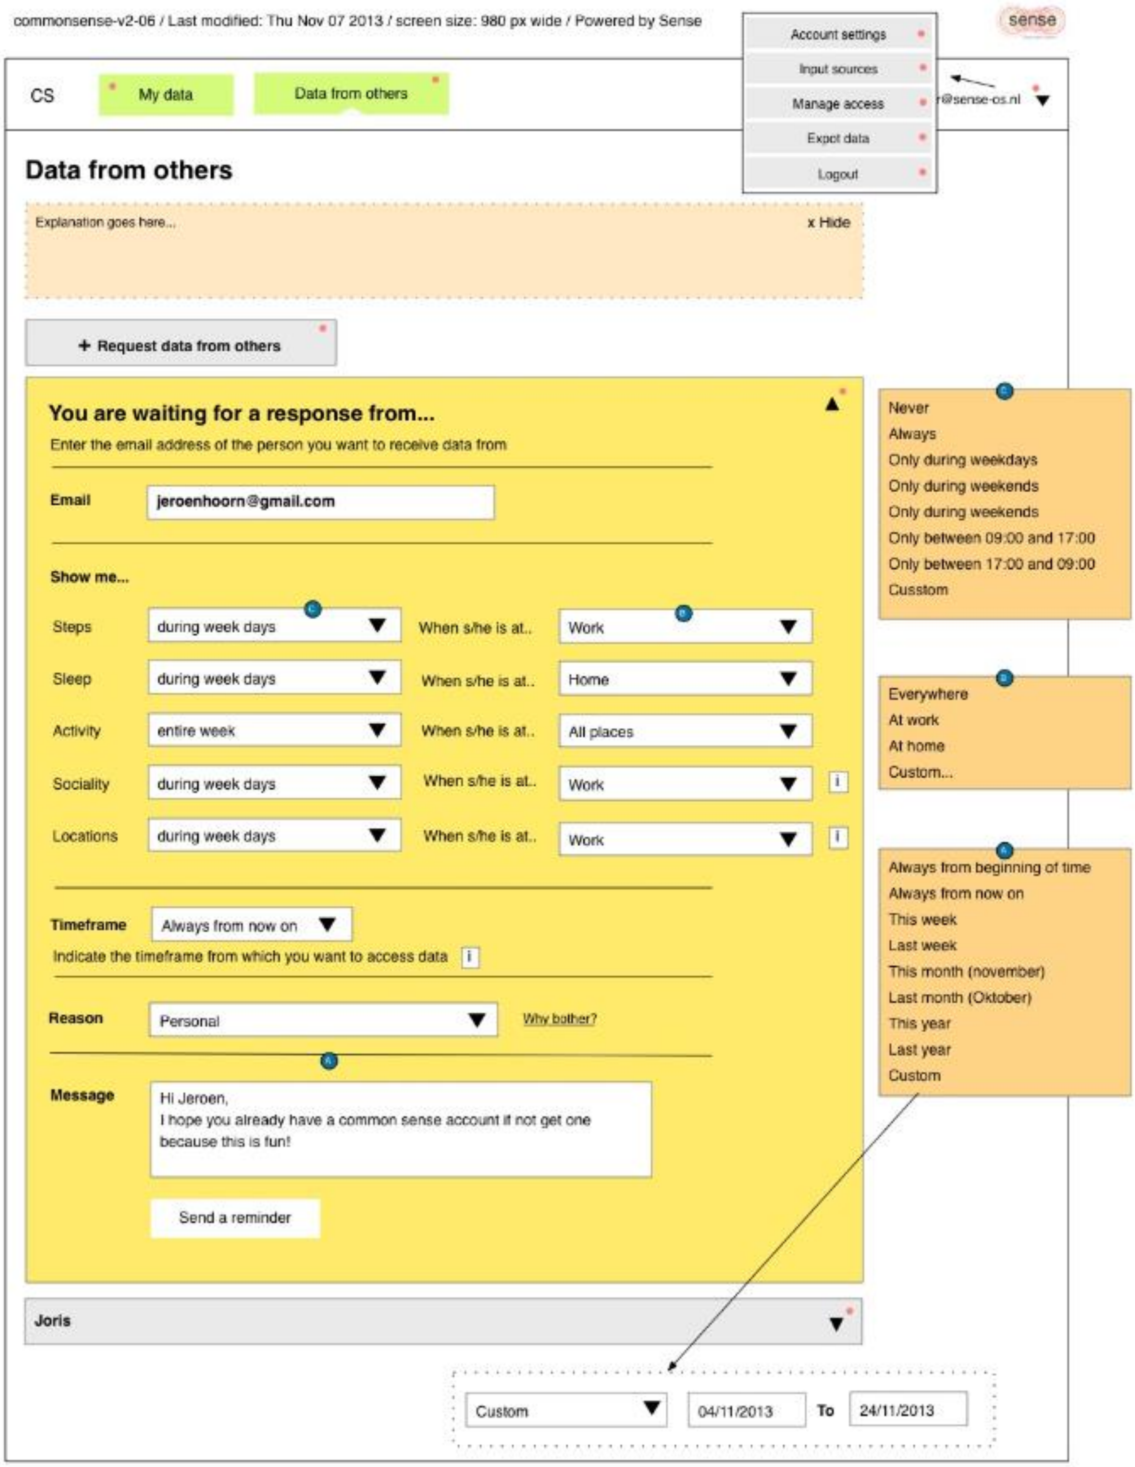

commonsense-v2-06 / Last modified: Thu Nov 07 2013 / screen size: 980 px wide / Powered by Sense

Data from others

Explanation goes here...

You are waiting for a response from...

Enter the email address of the person you want to receive data from

Email: jeroenhoorn@gmail.com

Show me...

Steps	during week days	When s/he is at..	Work
Sleep	during week days	When s/he is at..	Home
Activity	entire week	When s/he is at..	All places
Sociality	during week days	When s/he is at..	Work
Locations	during week days	When s/he is at..	Work

Timeframe: Always from now on

Indicate the timeframe from which you want to access data

Reason: Personal

Message: Hi Jeroen,
I hope you already have a common sense account if not get one because this is fun!

Custom 04/11/2013 To 24/11/2013

2.1.2 Policy management

In SWELL policy should be user driven: An individual should be able to configure their own policies required for the access relationship service to make access control decisions. As such, a user should be able to apply the same policy across distributed web resources, using a consistent interaction experience.

How this could be implemented is shown in Figure 2.

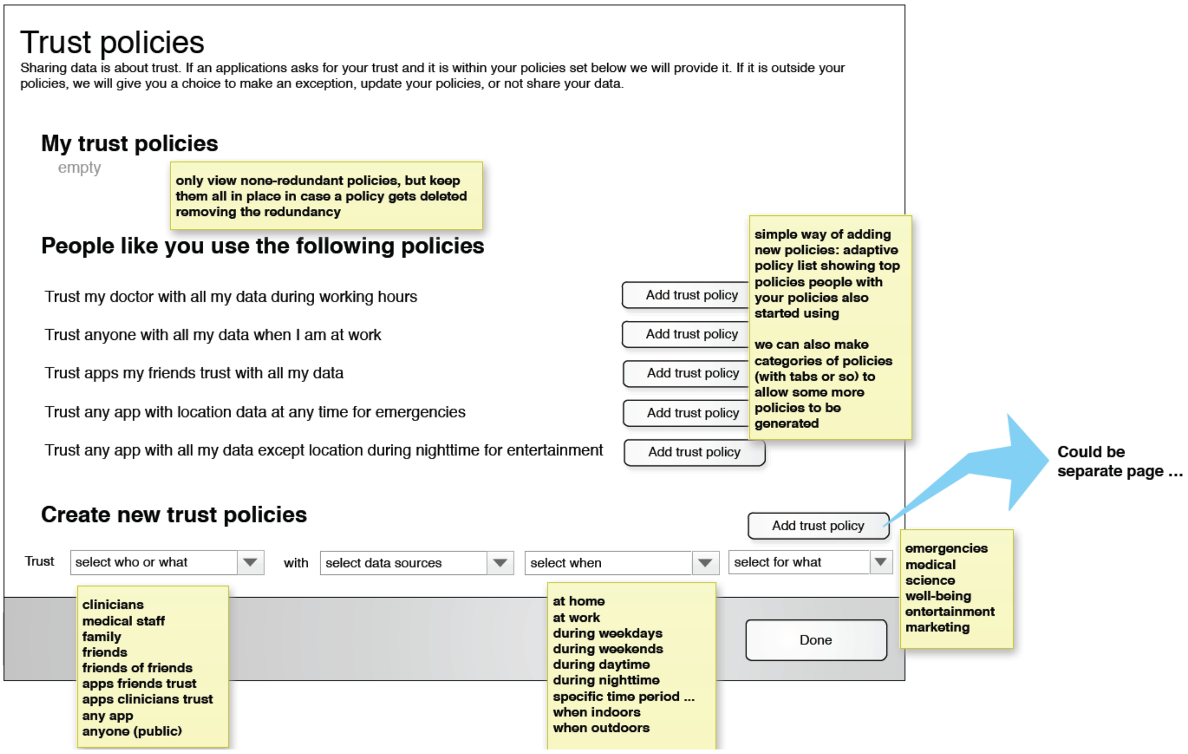

Trust policies

Sharing data is about trust. If an applications asks for your trust and it is within your policies set below we will provide it. If it is outside your policies, we will give you a choice to make an exception, update your policies, or not share your data.

My trust policies

empty

only view none-redundant policies, but keep them all in place in case a policy gets deleted removing the redundancy

People like you use the following policies

Trust my doctor with all my data during working hours

Trust anyone with all my data when I am at work

Trust apps my friends trust with all my data

Trust any app with location data at any time for emergencies

Trust any app with all my data except location during nighttime for entertainment

simple way of adding new policies: adaptive policy list showing top policies people with your policies also started using

we can also make categories of policies (with tabs or so) to allow some more policies to be generated

Could be separate page ...

Create new trust policies

Trust select who or what with select data sources select when select for what

clinicians
medical staff
family
friends
friends of friends
apps friends trust
apps clinicians trust
any app
anyone (public)

at home
at work
during weekdays
during weekends
during daytime
during nighttime
specific time period ...
when indoors
when outdoors

emergencies
medical
science
well-being
entertainment
marketing

Figure 2: policy management wireframe.

In order to fit well into the highly dynamic and open web environment, access control should not be based solely on preliminary identification and authentication of requesters; servers may not know possible identities of clients accessing data in advance. Policies should allow a user to define properties that clients must possess before authorization can be granted. Moreover, a user should be able to impose contract terms that govern access rights, as well as data storage, further usage, and further sharing on the part of requesting services.

The wireframe in Figure 2 describes an example of a trust policy management interface. It enables users to manage access control not only for each individual or app that wants to have access but also through more general policies that or might not be applicable to a certain case. Trust policies specify for a certain type of relation (to the user) which data, at what moment, and for what reason will be shared. In the example above, users can create their own trust policies, or they can use commonly used trust policies to make them easier to manage. Like with the conditions for sharing data, a conflicts between polices may easily occur. Policy conflict resolution functionality will be required in the CommonSense platform to detect conflicts and notify the user about them.

2.2 Transparency

In SWELL the user can have various overviews of what data is shared with whom and which apps. Based on these overviews the user should be able to modify the conditions of access and terminate relationships easily. It should also be possible to audit and monitor various aspects of such relationships at any time. If they wish, the user should not have to be directly involved in interactions between services accessing data and services hosting data. Rather, it might be possible to guide such interactions by applying the user's defined policies alone.

The wireframe below shows a sensor-centric view. It allows the user to select a particular sensor, to see its output, the applications that are consuming its data, and the other users the data is shared with.

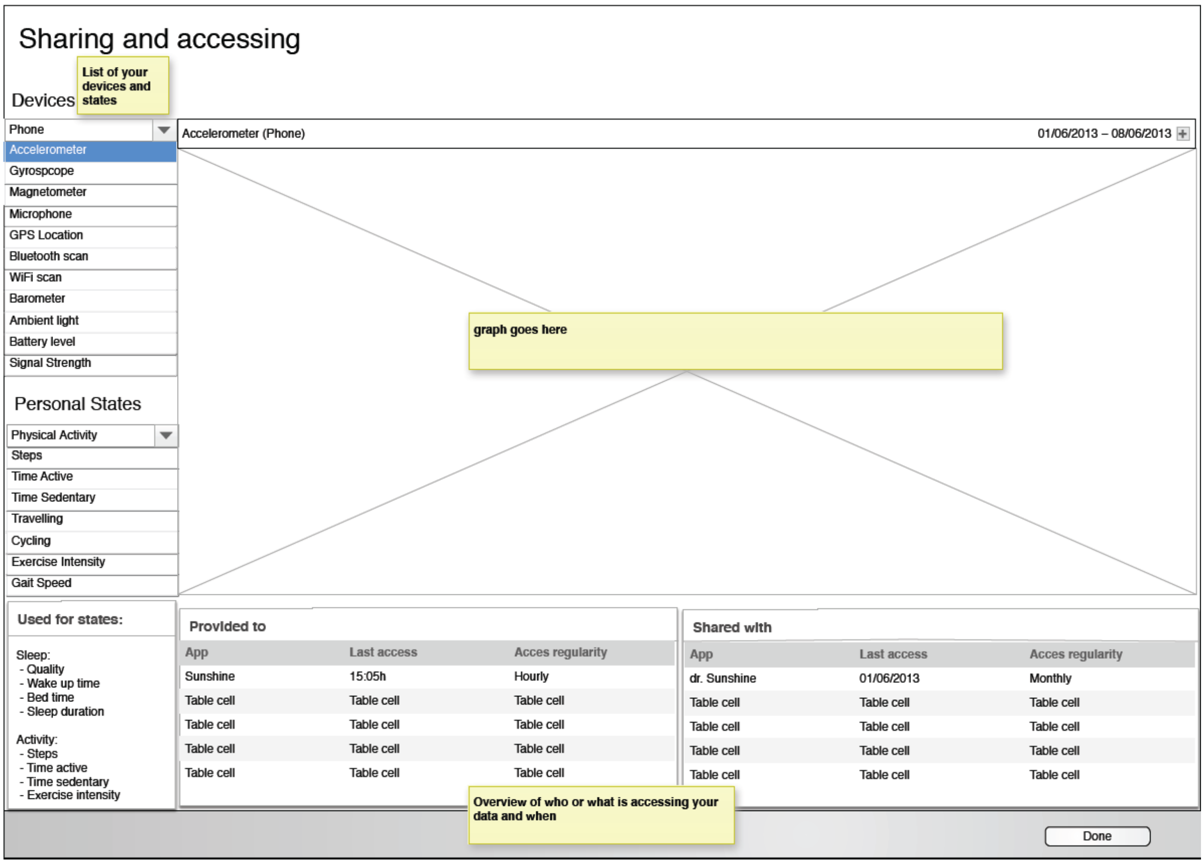

List of your devices and states

graph goes here

Used for states:

Sleep:
- Quality
- Wake up time
- Bed time
- Sleep duration

Activity:
- Steps
- Time active
- Time sedentary
- Exercise intensity

Provided to

App	Last access	Acces regularity
Sunshine	15:05h	Hourly
Table cell	Table cell	Table cell
Table cell	Table cell	Table cell
Table cell	Table cell	Table cell
Table cell	Table cell	Table cell

Shared with

App	Last access	Acces regularity
dr. Sunshine	01/06/2013	Monthly
Table cell	Table cell	Table cell
Table cell	Table cell	Table cell
Table cell	Table cell	Table cell
Table cell	Table cell	Table cell

Overview of who or what is accessing your data and when

Another approach takes a more application-centric view. The corresponding wireframe is shown below. It shows the applications and the sensor data they are consuming. Furthermore, it shows with whom data is shared and how frequently these users are consulting the data.

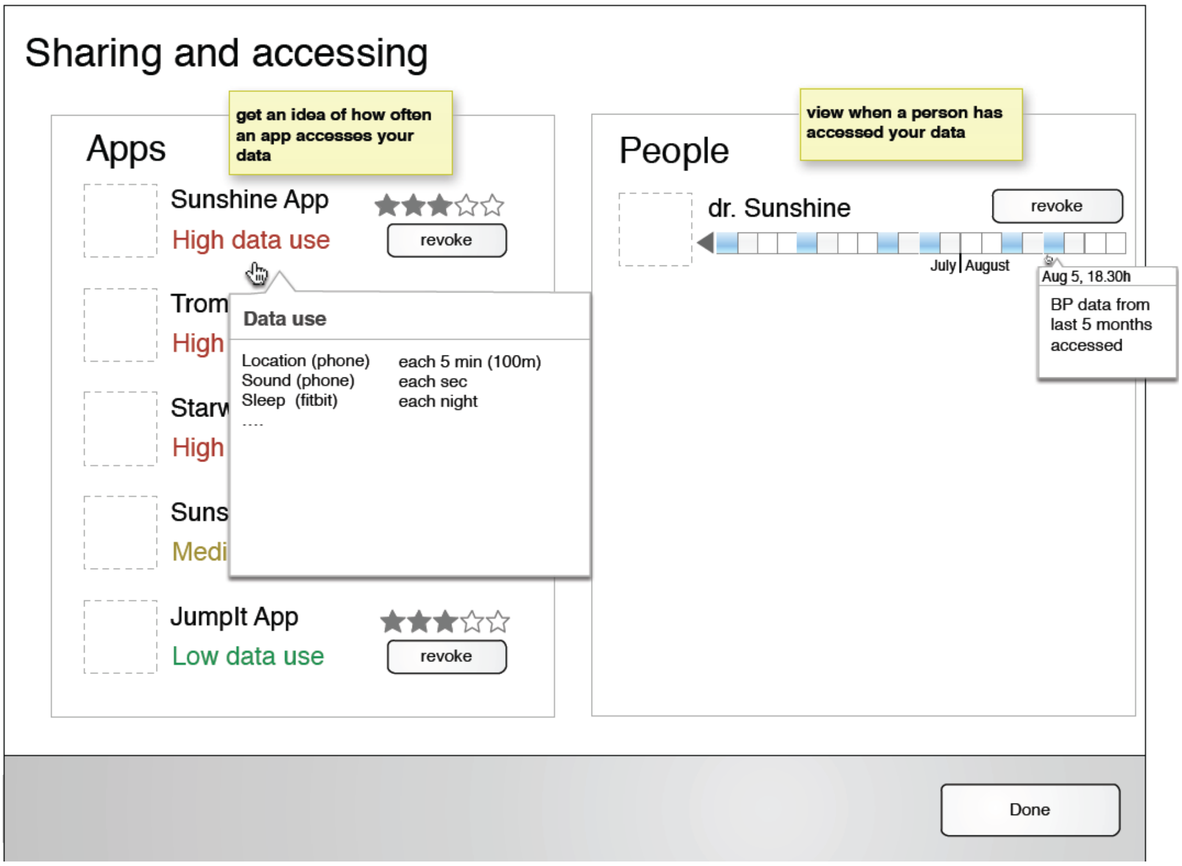

get an idea of how often an app accesses your data

view when a person has accessed your data

Data use

Location (phone)	each 5 min (100m)
Sound (phone)	each sec
Sleep (fitbit)	each night
....	

Aug 5, 18.30h

BP data from last 5 months accessed

2.3 Accountability

Accountability is difficult to implement. One way to force apps and clients that consume data to behave appropriately is to monitor their behaviour. In case of unexpected behaviour they can be asked for justification or their access to data could be blocked. The wireframe below indicates when a person has viewed your data and how often that data was viewed. This creates a certain level of accountability. The user has control about who has consumed his/her data and the consumer knows it is being watched by the user and cannot allow inappropriate data consumption behaviour.

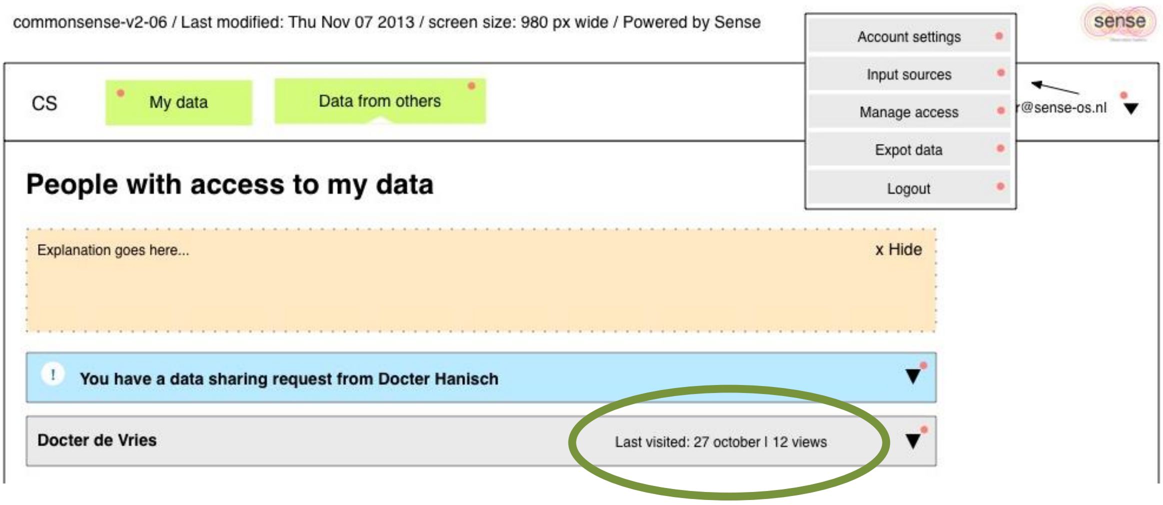

commonsense-v2-06 / Last modified: Thu Nov 07 2013 / screen size: 980 px wide / Powered by Sense

People with access to my data

Explanation goes here...

You have a data sharing request from Docter Hanisch

Docter de Vries

Last visited: 27 october l 12 views

3 Evaluation

In this chapter we describe the approach and outcomes of a small scale user evaluation experiment. With 10 test participants the designed wireframes were discussed.

3.1 Preparation

3.1.1 Participants

For this evaluation we carried out an interview with 10 people. These were employees of TNO and Novay. To get multiple perspectives from different people, there is chosen to select a divided group: from highly technology and privacy minded people to people who have less technology and privacy experience. Moreover the test group consisted of 2 females and 8 males, varied in age between 25 – 55, and all members were not familiar with the CommonSense platform.

3.1.2 Approach

One of the basic lessons learned in the area of HCI is that usability evaluation should start early in the design process, optimally in the stages of early prototyping. The earlier critical design flaws are detected, the greater the chance that they can and will be corrected.

For this usability evaluation, there is chosen to use the ‘walkthrough’ method (in combination with ‘Talk Out Loud’ method). This method is developed to give design teams a chance to evaluate early mock-ups of designs quickly, and therefore to identify some of the problems that might arise in interactions with the system.

The ‘Walkthrough’ method is a practical evaluation technique that describes human- computer interaction in terms of four steps:

1. The user sets/gets a goal to be accomplished with the system.
2. The user searches the interface for currently available actions.
3. The user selects the action that seems likely to make progress toward the goal.
4. The user performs the selected action and evaluates the system's feedback for evidence that progress is being made toward the current goal.

The following goals are given to the user which they had to accomplish:

1. Start the Sense-OS platform: http://prototypes.dev.sense-os.nl/v2-06 and click on the green button: “Enter prototype”
2. Look at your ‘sleep’ data from this week, only from your mobile phone. Share this data.
3. You have a data sharing request from someone. Look at the request and accept it.
4. Look at the data Joris is sharing with you.
5. Share your data with a person with whom you don’t share your data yet.
6. Look at the status from your request to see Pim’s data.

Another version (http://prototypes.dev.sense-os.nl/v2-04) is also used for showing alternative options of operating.

3.2 Analysis: main findings

General Impression:

- Overall all participants agree that the system provides transparency and that they have the feeling that they are in control about what data will be shared with whom, provided the data is indeed only shared with users they themselves authorized (and not with, e.g. ”trusted parties”).
- Especially people with a high affinity for privacy and technology wanted to see more information about the person who sends a request for getting access to his/her data. An email address is not satisfying enough (“how do I know that doctor Hanisch is really doctor Hanisch?”).
- Because of the very personal data that will be stored, multiple people would like to have more explanation about where the data will be stored (“I would like to see more information about where the data is stored. Is it safe, can I trust the system?”)

User experience:

- Especially for the less technology experienced, but also for some of the higher experienced people, the system was not always clear to use. This is mainly caused by the inconsistency of the UI symbols and the navigation structure.
- Multiple people discussed the unclear menu structure. There are 3 main parts: ‘my data’, ‘people with access to my data’ and ‘data from others’. This is what should appear in the main menu.

My data:

- Majority of the participants found the dashboard to be clear. However, a number of them indicated that they missed the definition of the data presented. Some examples:
    - What is a definition of “sociality”? What exactly does it represent, how is it computed? Is it how often I ‘whatsapped’ with my friends? Or with how many of them? Of how long my phone conversations were?
    - What exactly is the definition of “locations”. Places I visited today?
    - What exactly is “time active”? How is it computed?
    - Number of steps is presented, but how about if a user is cycling?
- Most participants expressed that they would like to see on the dashboard itself an overview (a list) of people with access to their data, to which data exactly they have access to, and if/how often they indeed accessed this data!
- Some participants indicated they would like to see a message in the dashboard page that they received a sharing request they should attend to.
- Some participants indicated they would like to see some other data, if they would use this platform to share with e.g. their health care specialists, fitness trainers or sleep therapists: blood pressure level, glucose level were some parameters mentioned.
- A participant with pollen allergy indicate he would like to know if some of the locations he visited is particularly affected by pollen.
- Some participants indicated they missed heart rate statistics in the dashboard. (e.g. a week average of 24 minutes in “fatburn” zone daily, etc.).
- The source tails (e.g. sleep) are clickable, most people did not understand why there is also a (grey) menu bar above.

- Most people expect to have an option to see the same views as another person would see when they share their data with them:
    - "I would expect to see my data with the same filters as those I can control when I 'share data with others'…what does doctor Hanisch see when I share my data?"
    - "I want to see visuals per person which data he/she can see about me"
    - "They can combine data and insights in matters I can't see"
    - "Only share aggregated data"
- "There are different levels of abstraction about having control. This is not always clear."
- "What does happen when I suspend one of the data sources…how does it affect the 'people with access to my data'?"

Data from others:

- Most, but not all, people see this as a nice-to-have option.
- Those that do state they would like to see data from friends, or from family members. E.g. one can challenge each other to make the most steps per week. Some people envisage contacting their friends to check up on them if, for example, they notice their friends lack sleep lately.

Sharing data:

- Most people indicated that at this point they feel reluctant about sharing this data with anyone. Most would use this platform for themselves only, without sharing data with anyone else. Only in case of a specific illness, and in the case particular data was relevant for effective treatment of that illness, they would consider sharing that specific data with their health care specialists.
- Some participants did however indicate to be willing to share their data with friends, in order to compete with and challenge each other.
- One participant who was reluctant to share his personal data would be motivated to use and share it anonymously with e.g. his department colleagues for the purpose of running some statistics on the whole group.
- Some participants indicated that they would be interested to see integrated data. E.g. if they are following behaviour of 20 other people, they would like to see some statistics available, or, e.g. which of their friends is sleeping worst, who walked most steps etc.
- Most participants were confused about what to do when asked to share the data: should they click the Facebook button? Or should they allow access via manage access tab? And in the latter case, how can they do that for only sleep data? Or via export data tab?
- Furthermore, most people recognized no added value in Facebook and Twitter button, as, according to their claims, they would not share this data with these platforms.
- There were many questions regarding the sharing request form (see paragraph below).

People with access to my data

- Participants had many remarks about the sharing form:
    - No indication whether your request to see data from other has been seen by the other (you can however send the reminder).

- Participants had questions on how system knows when or where the user works. Does a user have to set that himself? Is it associated with a place of work? And what if the user is working from home? And what if a user wants to share data with company doctor only on working days (only during weekdays 9 -17), but the user works part time?
- The current message in the request form is not accurate, because it is not only about making the account, but also about sharing the data.
- After a clicking the 'start sharing' button, it is expected to have an 'ok' button.
- There were questions about 'Reason for sharing' as well. Menu was not working, so participants could not identify different possible reasons. There were also questions about integrity: "what if a friend fills in 'for medical use' and he doesn't use it for that purpose? This is only useful when there is some sort of hierarchy that can be verified upon."
- Some participants expected a link to which sources are shared when a user shares his sleep data for example. E.g. "can I share my sleep data from my Fitbit but not from my phone?" (see also under Settings).
- There were also questions about possibility to push it to 'Share your data with': "the checkbox which says: 'I would like to share all possible data' is not clear. What if, over time, I add some extra sources. Can the other person also see that data?"
- Different participants posed questions about "revoke" field. E.g. "What does 'revoke' exactly mean?" Some users expected more explanation on this field: does it revoke from now on, or also the historical data, the latter being preferred. Also some users found "decline" to be more appropriate in the new request form.
- Different participants mentioned they would like the option where they can accept or decline request per data requested (e.g. "I allow him to see my sleep data, but he also wants to see my steps, and that I don't want to share"). However one participant found that: "the doctor determines what he thinks is needed to get the right insights. The user should not have the option to modify this. He can reject or accept the request only. Otherwise it would be useless."
- Some people (especially people with privacy and technology experience) like to have the option 'timeframe', for other people there are too many variables about time.
- Timeframe field provoked some more questions. E.g. is this a timeframe within which my data was collected, or within which the other user can access my data? Different users interpreted this field differently. Those that interpreted it in the latter way made the following comments:
  - "It is nice to have control over the time and locations when someone is allowed to see my data, but it might make it impossible to work with.
  - "To give people restrictions when they may see the data from me, does not match the trend of 'New Ways of Working'.
  - "I would like to have control about what data will be shared. It doesn't bother me when my doctor will look at it."
- Some participants found that sharing data should be in general (e.g. number of steps) not device specific. ("That's not relevant for others and could show more than I want to. Combining different sources on a detailed level, could provide insights I

    - don't know myself. So I don't want to give the possibility to others. However, it should be possible to share more detailed data, if people want that."
    - As already noted in previous paragraph, participants expressed it was important to them to see the 'number of views' and 'last view', but these should be clickable ("I want to see more information there.")

Settings:

- Some participants expressed that they would like to have control over which sensors can be used for which purposes. E.g. "I want to use my phone for sharing activity data, but not for sharing sleep data."

Aspects of trust in the system:

- Almost all participants observed the lack of https connection. The online tool that was used for showing wireframes was not secured during the evaluation. In the real Common Sense platform all communication is https secured. Nevertheless it is good to observe that users are aware of the connection and its security.
- Many participants had concerns about where a server is located and which party is maintaining it. Participants indicated they would want to have insight in the measures taken to secure data, such as encrypting the data etc. Some suggest visualization of this information (perhaps with a logo of ISO 27001 certificate?). Others had questions about privacy policy.
- Most participants were concerned about their data being shared with "trusted partners", and were posing questions about how could they be given guarantees that their data is not being shared with third parties?
- Almost all participants indicated that they may be more trusting if the platform was associated with a renowned company with a solid reputation, (e.g. Philips Medical).

3.3 Conclusions pre user study

10 individuals of different age, gender, and affinity for privacy issues have been interviewed. All participants acknowledged the added value this platform would provide. However, a significant number of participants was reluctant to share the information with others, and preferred to use the displayed information themselves to monitor the healthiness of their lifestyle. Those that were willing to share were only ready to do so if the information shared was relevant for a specific health treatment they were undergoing.

Participants had many suggestions for the improvement of the layout and navigation structure, as well as for the dashboard itself. Additional data in the dashboard (e.g. heart rate, distance cycled), the data definition in dashboard, a list of users authorized to view their information (including how often they access their information) are a few of the points mentioned. There were also many remarks regarding the sharing process and the sharing request form indicating that this is a critical functionality of the platform. Users want to be sure with whom they are sharing data with. Moreover, sharing comes with a lot of complexity (i.e. policy management) that easily tends to reduce the feeling of trust in the system if it is not intuitive or becomes too intrusive.

Related to the question of trust, participants expressed concerns about the location of hosting servers and the security measures that have been taken place. Most participants reported having a feeling of not being in control over, nor having any insight into, whether the information is nevertheless being shared with third parties, in spite of them not authorizing such actions, and indicated this as another major factor of trust. This might be partially alleviated, they added, by associating this platform with some well-renowned company with a reliable reputation (e.g. by displaying a trusted company logo).

Finally, all participants agreed that, if their remarks would be integrated in the platform, the platform would provide transparency in and control over which data is being shared and with whom (provided no data is being shared with the so-called “Trusted partners”).

4 Conclusions

Within the SWELL project, we are looking at privacy control solutions for sensor platforms. The aim is to provide more transparency and control to the end-users about what data will be shared with whom.

A number of mock-ups have been designed that constitute several components which provide more transparency and control to the end-user. The mock-ups were based on the results of earlier work done in SWELL (deliverables D4.1, D4.2, D4.3 and D4.4). The idea is to implement these mock-ups in the CommonSense platform that gathers data and gives the possibility to share this data with other persons.

In order to get a feeling if the designed mock-ups for user control and transparency were well-received by users, a small user experience evaluation was conducted. The evaluation showed that many improvements are necessary and that with these improvements users expressed they will be more confident in sharing data.

A particular challenge that need to be addressed is that of being able to authenticate a user or application that is making a request for data. Solutions are not trivial and will be researched in 2014 with the intention to integrate the best solution in future versions of CommonSense.

Another challenge that remains to be solved is the experienced intrusiveness of privacy control. Future work in SWELL will therefore focus on the design and implementation of an intelligent privacy reasoner component that uses context information to make the control features more adaptive and less intrusive.

Finally, a generic challenge will be to tailor the mock-ups to smartphone screens such that they remain practical to use.

The outcomes will be taken into account in a next iteration of new wireframe designs. These will be used as starting points for the actual implementation.