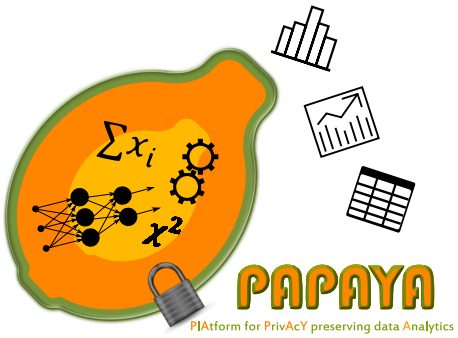

D3.4 - Transparent Privacy Preserving Data Analytics

Work Package	WP3, Privacy Enhancing Technologies for Data Analytics
Lead Author	Simone Fischer-Hübner (KAU)
Contributing Author(s)	Matthias Beckerle (KAU), Simone Fischer-Hübner (KAU), Nuria Ituarte (ATOS), Jonathan Magnusson (KAU), Patrick Murmann (KAU), Angel Palomares Perez (ATOS), Tobias Pulls (KAU), John Sören Pettersson (KAU), Patrick Murmann (KAU), Jonas Frei (KAU), Christian Weis (KAU)
Reviewers	Dominique Le Hello (ORA), Orhan Ermis (Eurecom)
Due Date	30.04.2020
Delivery	30.04.2020
Version	0.14
Dissemination Level	Public

The research leading to these results has received funding from the European Union's Horizon 2020 Research and Innovation Programme, through the PAPAYA project, under Grant Agreement No. 786767. The content and results of this deliverable reflect the view of the consortium only. The Research Executive Agency is not responsible for any use that may be made of the information it contains.

Revision History

Revision	Date	Author	Description
0.1	11.02.2020	Simone Fischer-Hübner (KAU)	Template with ToC created.
0.2	09.03.2020	Jonathan Magnusson (KAU)	Added 1st version of chapter 2.
0.3	10.03.2020	Patrick Murmann (KAU)	Added first version of sections 3.1 and 3.2.
0.4	13.03.2020	Angel Palomares Perez (ATOS)	Added chapter 4.
0.5	18.03.2020	Matthias Beckerle (KAU)	Added chapter 3.3 and 3.4.
0.6	18.03.2020	Simone Fischer-Hübner (KAU)	Added chapter 1.
0.7	27.03.2020	Tobias Pulls (KAU)	Added to chapter 2.
0.8	30.03.2020	Simone Fischer-Hübner (KAU)	Did editorial review for chapters 2 and 3. Added content to chapter 2 and sections 3.1 and 3.2.
0.9	02.04.2020	John Sören Pettersson, Patrick Murmann, Matthias Beckerle (KAU)	Added and changed content in chapter 3.
0.10	02.04.2020	Simone Fischer-Hübner (KAU)	Added chapter 5 and editorial changes.
0.11	05.04.2020	Simone Fischer-Hübner (KAU)	Added glossary and executive summary. Smaller changes to chapters 1 and 5.
0.12	17.04.2020	Matthias Beckerle (KAU)	Added to chapter 3. Harmonised headlines and captions. Some smaller changes in the document.
0.13	1.04.2020	All authors (KAU)	Review comments fixed
0.14	27.04.2020	All authors (KAU)	Review comments fixed
0.15	28.04.2020	Orhan Ermis (EURC)	Quality check completed

Executive Summary

This report on Transparent Privacy Preserving Data Analytics (PAPAYA deliverable 3.4) presents the results of our development of user interfaces (UIs) for making privacy-preserving data analytics, which we develop in the PAPAYA project, transparent to data subjects and other stakeholders and for increasing means of control for data subjects.

Our work implements end user and legal requirements that we previously elicited in PAPAYA deliverable D2.2 and builds upon the state of the art that we presented in deliverable D3.2, in which we elaborated what to explain and how to explain privacy-preserving data analytics.

For the development of most of the user interfaces, we followed a human-centred design approach, consisting of iterations of prototyping and evaluations involving user tests and expert walkthroughs. The UIs were finally implemented in React Native for Android mobile phones.

The user interfaces presented in this deliverable will be utilised in PAPAYA's data subject tools and dashboard and include:

- **User interfaces for presenting Risk Management artefacts for assessing the impact of privacy-preserving data analytics on privacy risks.** For creating risk management artefacts, the Privacy Impact Assessment (PIA) Tool by the French Data Protection Authority CNIL was enhanced in three ways: Firstly, we added support for an initial threshold analysis for a PIA, which in turn resulted in UIs for explaining why a PIA needed to be conducted. Secondly, the privacy protection goals *transparency, unlinkability* and *intervenability* were added to the security goals *confidentiality, integrity* and *availability* resulting in more comprehensive PIA results to be shown. Thirdly, the output of the risk matrix shown was improved for making the information more easy to view for data subjects, especially if displayed on mobile phone screens.

- **User interfaces for explaining and exemplifying how privacy-preserving data analytics work.** A focus was on UIs for explaining privacy-preserving neural networks for classification and collaborative training. They present explanations in multiple layers of details, including explanations for the privacy-preserving building blocks of homomorphic encryption and differential privacy. Furthermore, UIs for explaining multi-party computation and functional encryption are provided as further advanced cryptographic techniques used by the PAPAYA platform.

- **User interfaces for enhancing end user control via the Privacy Preference Manager and Data Subject Rights Manager of PAPAYA's Policy Engine.**

The presented user interfaces for enhancing transparency should be integrated in multi-layered policy notices, following a recommendation by the Art. 29 Data Protection Working Party for enhancing the usability of privacy notices. Information about the conducted PIA, residual privacy risks and explanations of how the utilised privacy preserving data analytics work could be provided on lower layers with a clickable link to information on the PIA and Privacy by Design approach taken on the top layer of the privacy notice.

Contents

List of Figures

Glossary of Terms

CNIL Commission nationale de l'informatique et des libertés

DPIA Data Protection Impact Assessment

DSRM Data Subject Rights Manager

GDPR General Data Protection Regulation

HCI Human Computer Interaction

MPC Multi Party Computation

PAPAYA Platform for Privacy-Preserving Data Analytics

PE Privacy Engine

PET Privacy Enhancing Technology

PIA Privacy Impact Assessment

PPM Privacy Preference Manager

UC Use Case

UI User Interface

UX User Experience

1 Introduction

1.1 Aims and Scope

The PAPAYA project is designing and developing dedicated modules for privacy preserving data analytics, which can be performed by untrusted third-party data processors while still adequately addressing privacy concerns and guaranteeing accuracy of analysis results. For achieving this, data analysis may be performed obliviously on protected data (e.g., encrypted data), or other types of privacy-enhancing technologies (PETs), such as PETs based on differential privacy, are utilised.

For addressing end user requirements and legal privacy requirements for transparency and intervenability derived from the EU Data Protection Regulation (GDPR) [3], the project also develops data subject tools with user interfaces (UIs) for transparent privacy data analytics. These UIs are explaining to data subjects how the privacy preserving data analytics work and what impact they have on (residual) privacy risks.

Moreover, the project develops UIs for the Privacy Engine (PE) that enables data subjects to define and manage their privacy preferences and allow them to exercise their rights pursuant to the GDPR.

This deliverable reports about the development work of these user interface prototypes that were designed and implemented in the second project year, and which will be part of data subject tools (e.g., intervenability tools) developed within the scope of the project.

1.2 Our Approach and Relation to other Deliverables

For following a user-centered design approach, we investigate "*what to explain*" and "*how to explain*" [8].

The question of "what to explain" was earlier addressed by the deliverables D2.2 [7] and D3.2 [8]. D2.2 analysed the users' mental models and elicited end user and legal privacy requirements, whereas deliverable D3.2 identified risk management artefacts that can be shared with data subjects to inform them about the risks associated with having their personal data processed by the PAPAYA platform. D3.2 also provided a state-of-the-art analysis on how to convey trade-offs between privacy and utility in privacy-preserving data analytics and made a first suggestion of "how to explain". Moreover, it presented the Privacy Engine design.

Based on the elicited requirements presented in D2.2 and results from D3.2, this deliverable further elaborates how to explain privacy preserving data analytics and related risks to data subjects. For this, we have developed user interface prototypes, which we have partly tested with end user lab tests and/or evaluated via expert walkthroughs and implemented as independent React Native[1] components. The results are reported in this deliverable and will provide input for the development of the data subject tools and data subject dashboard to be developed within Work Package 4 (Platform Design and Development).

[1] `https://reactnative.dev/`

1.3 Outline

The remainder of this deliverable is structured as follows:

- In Chapter 2, we present our work on improving the Privacy Impact Assessment (PIA) tool by the French data protection authority CNIL for creating risk management artefacts, and show how we use these artefacts for developing user interface components explaining risks of privacy-preserving data analytics.
- Chapter 3 is then presenting our development of user interfaces for explaining privacy-preserving data analytics to data subjects. In particular, we present user interfaces explaining privacy-preserving neural networks for data classification based on homomorphically encrypted data to be utilised in PAPAYA's healthcare use case on Arrhythmia detection (UC1). These include UIs for explaining homomorphic encryption, which can also be used for PAPAYA's Mobility analytics use case (UC3) for privacy-preserving counting. Moreover, we present user interfaces explaining privacy-preserving neural networks for collaborative training with differential privacy, which we utilise in PAPAYA's healthcare use case UC2 on Stress detection. Finally, we present user interfaces for explaining functional encryption as building blocks to be utilised in PAPAYA's use case on Mobile usage analytics use case (UC4) and for explaining multi-party computation (MPC) that can be used for the Mobile usage analytics use case (UC3) and, in the form of Two-Party Computation, for the Arrhythmia detection case.
- Then, in Chapter 4, we will present user interfaces implemented for the Privacy Engine, including user interfaces allowing data subjects to manage their privacy preferences and user interfaces allowing data subjects to exercise their data subject rights pursuant to the GDPR.
- Chapter 5 is finally summarising the contributions and provides the main conclusions from this deliverable.

2 Explaining Risks of Privacy-Preserving Data Analytics

The Privacy Impact Assessment (PIA) tool from CNIL is a tool that guides an organisation in the process of conducing a PIA of a Data Protection Impact Assessment (DPIA). Pursuant to Art. 35, a DPIA needs to be performed for a type of personal data processing that is likely to result in a high privacy risk for individuals. The goal of the tool is to understand and mitigate the privacy risks to individuals whose personal data are processed by the organisation, so called data subjects. A privacy risk is estimated in terms of its *severity*—the potential impact on data subjects—and its *likelihood*, i.e., the probability of the risk occurring. Towards the aforementioned goal of the CNIL PIA tool, a user of the tool evaluates the causes and consequences of privacy risks and the measures taken to decrease the likelihood and/or the severity of those risks. For example, measures can serve to decrease the likelihood of a risk through protecting personal data using encryption, or to reduce the severity of a risk by minimising the collection of personal data.

In PAPAYA, we develop novel technical privacy protections in the form of privacy-preserving data analytics. In general, these technologies should reduce the privacy risks for data subjects by keeping personal data confidential towards the PAPAYA platform taking part in the analytics. However, such technologies may also increase some risks, such as reducing transparency or intervenability controls of data subjects over their personal data, because the data are now encrypted or obfuscated with the result that data subject right request may not be answered directly. The ultimate goal of our work here by extending the features of the CNIL PIA tool (Section 2.1) and developing user interface components (Section 2.2) is to communicate these risks to data subjects.

2.1 Improving the PIA Tool by the CNIL

Our earlier work, as described in D3.2 [8], identified the following possible modifications to the CNIL PIA tool:

1. Add support for performing an initial threshold analysis
2. Add transparency, unlinkability and intervenability as protection goals
3. Improve the output of the risk matrix

The initial threshold analysis is a checklist which recommends whether a PIA has or does not have to be performed. Transparency, unlinkability and intervenability are three privacy protection goals which can be used to extend the existing security protection goals in the PIA to encompass more risks. The risk matrix in the PIA was only visualising the current risk likelihood and severity. By adding two more questions for each risk, a visualisation of the likelihood and severity *before current controls were implemented* can be displayed. We describe the implementation of the threshold analysis in Section 2.1.1, additional risks in Section 2.1.2, an improved risk matrix in Section 2.1.3, and other improvements in Section 2.1.4.

2.1.1 Threshold Analysis

The threshold analysis was added in the first section of the PIA. This section is responsible for getting an overview of the system under assessment and listing data items, processes and supporting assets. This lays the foundation needed to proceed with a threshold analysis since it is first after understanding the context that a threshold analysis should take place. Thus the threshold analysis was inserted at the end of the first stage of the PIA (see Figure 1). The threshold analysis was created from the recommendation in Article 29 Data Protection Working Party [12] shown in Figure 2. The criteria from Article 29 were allocated a check-box each in the threshold analysis. Updating a check-box by ticking or unticking the criteria will affect the recommendation at the bottom of the threshold analysis. It is changed from *"a PIA might be unnecessary"* to *"a PIA should be considered"* when two or more criteria are ticked.

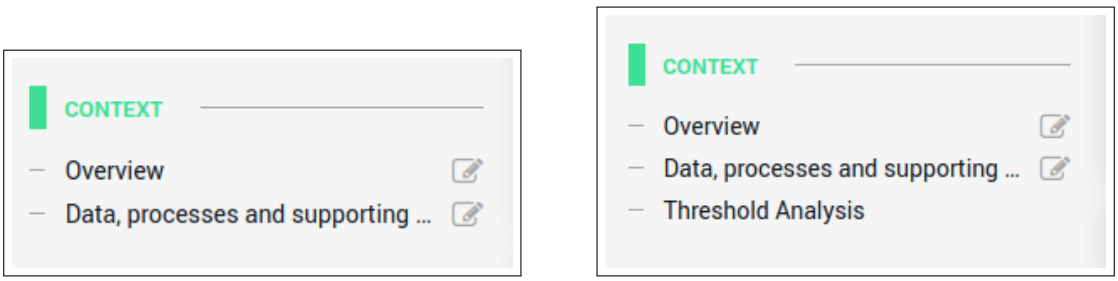

(a) Before (b) After

Figure 1: Before and after adding the threshold analysis.

2.1.2 Additional Privacy Triad

The widely accepted security triad is *confidentiality, integrity* and *availability*. In the original layout of the PIA this triad is the foundation of the associated risks to personal data. The three security protection goals have been inverted in order for the PIA emphasise on the risk aspect: *Illegitimate access to data, Unwanted modification of data*, and *Data disappearance*. A privacy complement to this established security triad has been proposed in [10], further refined in [4, 9] and summarised in [2]. This additional privacy triad consists of *transparency, unlinkability* and *intervenability*. Figure 3 illustrates how the six protection goals are established by combining the two triads of security and privacy.

Transparency makes sure that the processing of privacy sensitive data is explained and understood for the target audience. The audience could be the user, the data controller or the supervisory authority, to name a few (even though strictly speaking, the transparency principle and rights pursuant to the GDPR have the objective to protect data subjects). The information should be tailored to the target audience and thus not give too much or too little details. The processing information should be available both before future processing (ex-ante transparency) as well as for the past processing already done (ex-post transparency).

Unlinkability addresses the problem of correlating privacy sensitive data across multiple do-

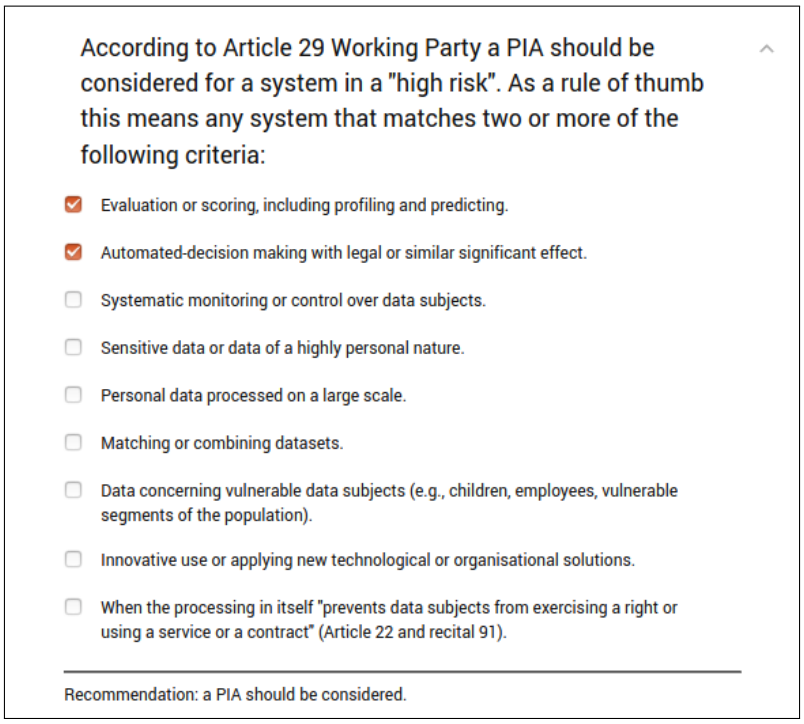

Figure 2: Threshold analysis.

mains. Linking privacy-relevant data between domains increases the risk of extracting more information from the data beyond purpose. If age is collected for one purpose, and income is collected for another purpose, these should be unlinkable. Otherwise it could be possible to extract additional information lacking a stated purpose. Data avoidance, anonymisation, pseudonymisation and separation of contexts are a few different ways to support unlinkability.

Intervenability is particularly important for the data subjects in order to intervene in ongoing or future planned processing of their privacy sensitive data. But even the data controllers should have the possibility to intervene if they deem it necessary. In the case of the data subject, intervenability ensures for instance the individual's rights to withdraw consent and erase or rectify data or to object.

Similarly to the security triad, the three new risks in the PIA are added by inverting the three privacy protection goals: *Opaque data processing, Linkable data processing*, and *Lack of control*. The collection of the original risks in the PIA, corresponding to the security triad, can be found in Figure 4a. The result of adding the privacy triad to the PIA can be seen in Figure 4b.

2.1.3 Risk Matrix

The dots in the risk matrix are representing the likelihood (x-axis) and the severity (y-axis) of a risk (shown later). These coordinate values are set by pair-wise sliders on each of the risks in the PIA. An additional pair of sliders have been added to each risk representing the risk

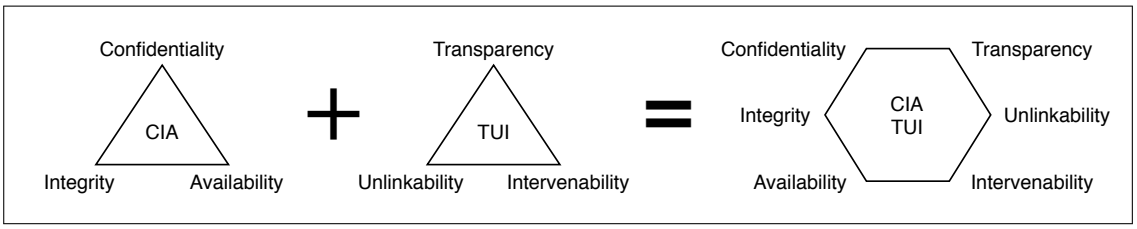

Figure 3: The two triads of security and privacy combined.

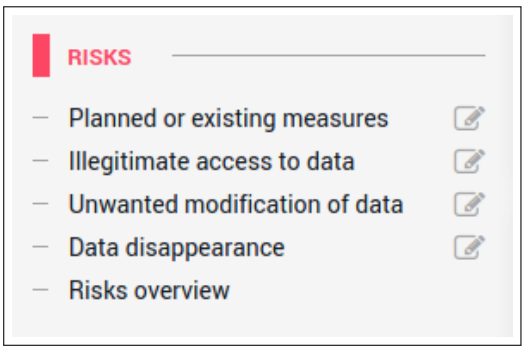

(a) Before

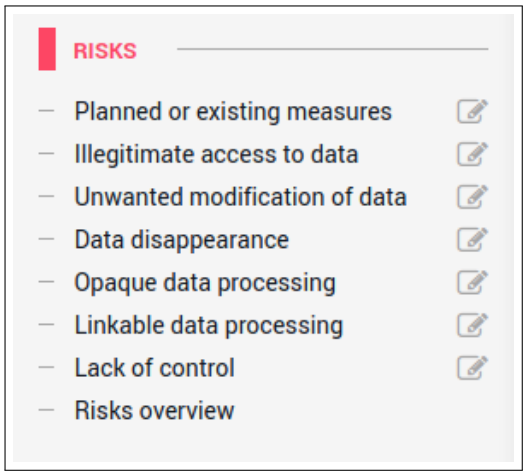

(b) After

Figure 4: Before and after adding privacy protection goals.

severity and risk likelihood if *no planned controls* were implemented (worst case scenario). This enables us to later highlight the effect of controls. Note that these effects can be negative, i.e., increase risk. The new slider for likelihood with no planned controls can be seen in Figure 5 and for severity with no planned controls in Figure 6.

The matrix has also been improved, shown in Figure 7. The previously mentioned new sliders are represented by the red dot, from which an arrow points towards the risk severity and risk likelihood with planned or existing measures. A dynamic legend was implemented to only display the necessary information needed to understand the matrix. If there are no blue evaluation dots, then there is no legend for blue dots. If there are no red dots—representing the risk without existing measures—then there is no legend explaining the red dot. This is to keep the matrix as simple and clear as possible. For the risks legend, if only two risks have been assessed in previous sections of the PIA then only those two risks show up in the matrix legend on the left.

Figure 5: Risk likelihood slider for no planned controls.

Figure 6: Risk severity slider for no planned controls.

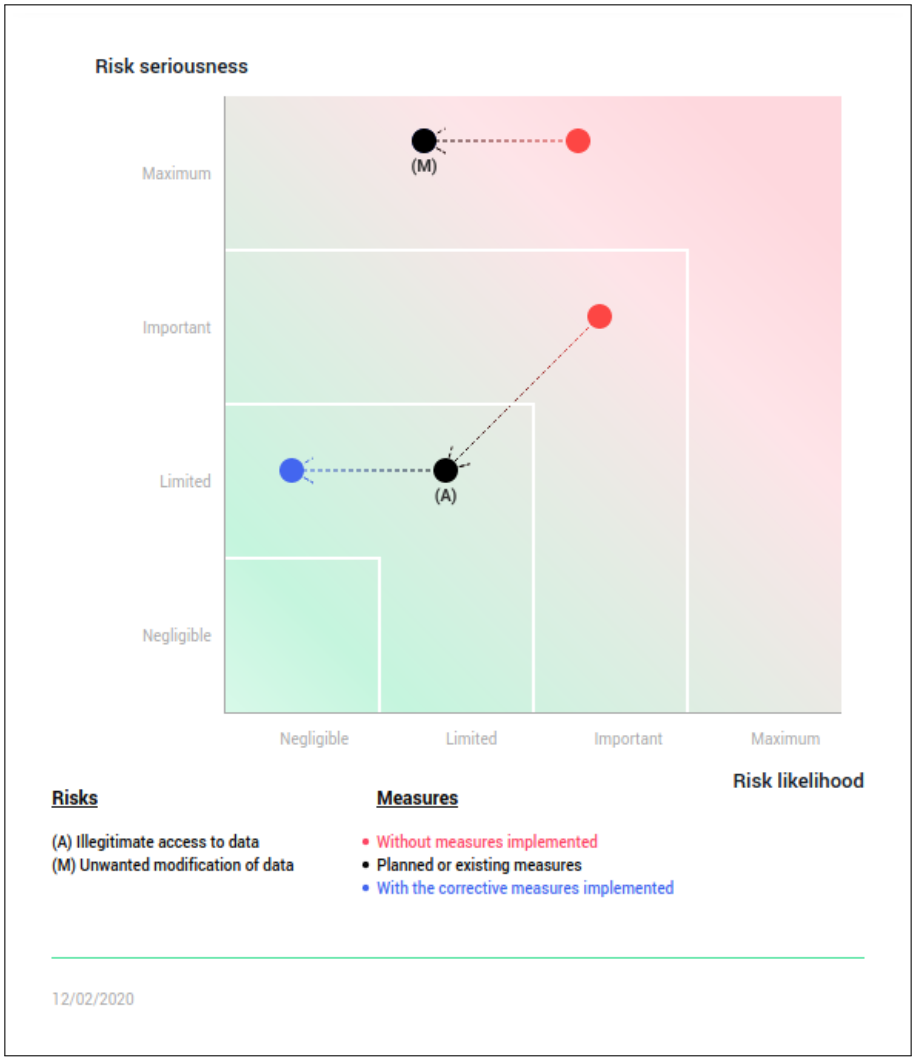

Figure 7: Improved risk matrix.

2.1.4 Other Improvements

The code for generating the matrix (drawing the dots and the arrows) has been refactored to reduce redundancy. A JSON-like structure was already used internally by the PIA in order to store the answers throughout the process, so a button has been added in order to download the data structure containing the answers as a JSON file.

2.2 UI Components for Communicating Risk

Based upon the improvements to the CNIL PIA tool and its JSON output, the next step is to take that output and use it as input for mobile UI components that communicate risk to data subjects. The UI components are implemented as independent React Native[2] components. React Native is an open-source mobile development framework that can be used to create native apps for Android and iOS. Components in React Native can be integrated into existing mobile applications with relative ease[3]. UI components for communicating the threshold analysis is presented in Section 2.2.1 and for the risk matrix in Section 2.2.2.

2.2.1 Threshold Analysis

The threshold analysis criteria in the JSON are used in the application to display *re-formulated criteria*, coinciding with the purpose of transparency in Section 2.1.2. The information presented in the application is tailored to the data subject by using a simpler language in order to understand why a PIA was either needed or deemed unnecessary. To keep the UI simple and avoid cluttering, only the criteria of the ticked check-boxes in the PIA are shown in the application. Figure 8a shows a threshold analysis which had all its check-boxes ticked. Figure 8b shows another threshold analysis where only some of the check-boxes had been ticked. No usability tests has yet been done on the user interface of the application at stage. As part of integration during the last year of PAPAYA, we expect to tailor and integrate this UI component as part of larger component, see Section 5.1 for an example.

2.2.2 Risk Matrix

The risks are split up into their own matrices for two reasons. The first reason is to make the information as easy as possible to digest for the intended audience: the data subjects. The second reason is to be able to select a subset of risks that have high likelihood and severity, or risks with the most significant improvement/deterioration. Two separate matrices for two arbitrary risks can be seen in Figure 9. Each matrix has a bar on the x-axis representing the likelihood of the risk, and a bar on the y-axis representing the severity of the risk. The bar either contains an arrow or a dot. The dot represents that no change has happened between *no planned measures* and *implemented measures*. The arrow shows if the implemented measures are improving or worsening the state before the measures where implemented. The matrix between these two bars is showing the combined trend of likelihood and severity. The colours

[2] `https://reactnative.dev/`
[3] `https://reactnative.dev/docs/integration-with-existing-apps`

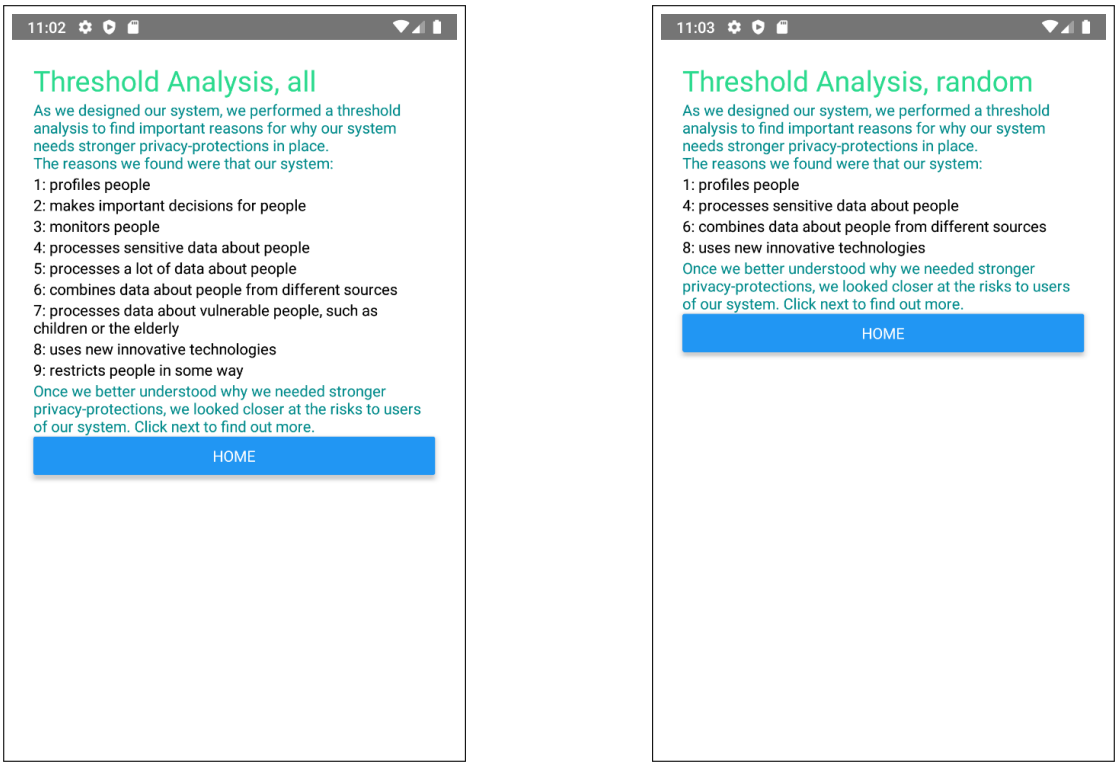

(a) A full threshold analysis

(b) A partially filled threshold analysis

Figure 8: Threshold analysis on the application.

used in these new matrices are changed from continuous (fading) to discreet (distinct squares). This change aims to improve the ease of understanding for the data subject. A colour scale between red and green indicates the risk of the system, which is desired to have low likelihood and low severity (bottom left). There are a couple of main differences between the risk matrix in the PIA by the CNIL and the new risk matrix in the application. First the separation of risks into their own matrices. Second, the use of only the red and the black dot from Figure 7. The blue evaluation dot is less interesting for the data subject and might be implemented and represented as the black dot in the future. The component is kept minimal, but easy to extend for later integration.

2.3 Summary

We improved the CNIL PIA Tool by adding support for a threshold analysis, extending the risks to also include the three privacy protection goals, and adding an extra stage to the risk matrix representing the risks *before* current controls were implemented. In addition to these improvements, a refactoring of the code responsible for drawing the risk matrix was done, and we added support for exporting the PIA into JSON. Based on the exported JSON output, we created two UI components in React Native to communicate risks to data subjects. One

(a) Likelihood and severity decrease

(b) Likelihood decrease only

Figure 9: Risk matrices intended for data subjects.

component is based on the threshold analysis, and the other on communicating risk-specific changes with and without controls.

The modifications to the CNIL PIA tool has been open sourced and can be found on the PAPAYA-H2020 github[4]. We also sent two pull requests to the main CNIL PIA repository[5,6] with our changes.

[4] https://github.com/papaya-h2020/pia
[5] https://github.com/LINCnil/pia/pull/438
[6] https://github.com/LINCnil/pia/pull/439

3 Explaining how Privacy Preserving Data Analytics work

This section presents the development of user interfaces (UIs) for explaining privacy-preserving data analytics to data subjects for the different categories of PETs (privacy enhancing technologies) that PAPAYA is developing. In particular, we present UIs for privacy-preserving neural network for classification based on homomorphic encryption and for collaborative training utilising differential privacy. In addition to explaining homomorphic encryption, UIs for multi-party computation and functional encryption are presented as well as basic building blocks for PAPAYA's PET solutions for privacy-preserving clustering and counting.

While the UIs can be utilised in the different use cases of PAPAYA, which were presented in D2.1, we designed and tested the UIs use-case independent. This enables us to use the UIs in different types of applications, beyond our five use cases, which are based on PAPAYA's PETs.

As the UIs of the different PETs developed for PAPAYA will later all appear in different applications, we did not strive for a consistent design for the UIs of the different categories of PETs.

An important criterion for all prototypes presented in this section was to have them implemented in the form of a mobile app, as mobile phones represented the platform that would most likely reflect a future user's terminal device. This meant that the form factor and interaction paradigms that had to be considered while conceptualising user experience were constituted by the general ergonomics of contemporary mobile phones. It meant, e. g., that the layout of information presented to users had to be optimised for portrait format, and that the amount and complexity of information presented on screen was comparatively small.

Consequently, all information pertaining to a particular narrative that needed to be conveyed to facilitate understanding of the underlying phenomenon had to be segmented into units that were manageable and intelligible by a casual user without domain knowledge. Each segment was modelled as a dedicated interaction phase presented on a dedicated screen.

3.1 UIs for Privacy preserving Neural Networks – Explaining Classification based on Homomorphic Encryption

3.1.1 Background and Motivation

This section presents UIs for explaining privacy-preserving neural networks for data classification. This category of PETs and the UIs can for instance be used for PAPAYA's use case on Arrhythmia detection (UC 1 under the health care umbrella). However, as mentioned above, we chose to have use case-independent UI designs that can be more generally be used for different types of use case deploying this category of PETs.

The UIs presented in the next sub sections include UIs for explaining privacy-preserving neural network classification for data classification. In particular, we tried to convey to users that data are sent by a client to the analytics platform in encrypted (i.e. protected form), are analysed in encrypted form and also the analysis results are encrypted and can only be decrypted by the client. While this functionality can be achieved with different underlying cryptographic schemes, we have for our UIs chosen a (fully) homomorphic encryption scheme as an example. However,

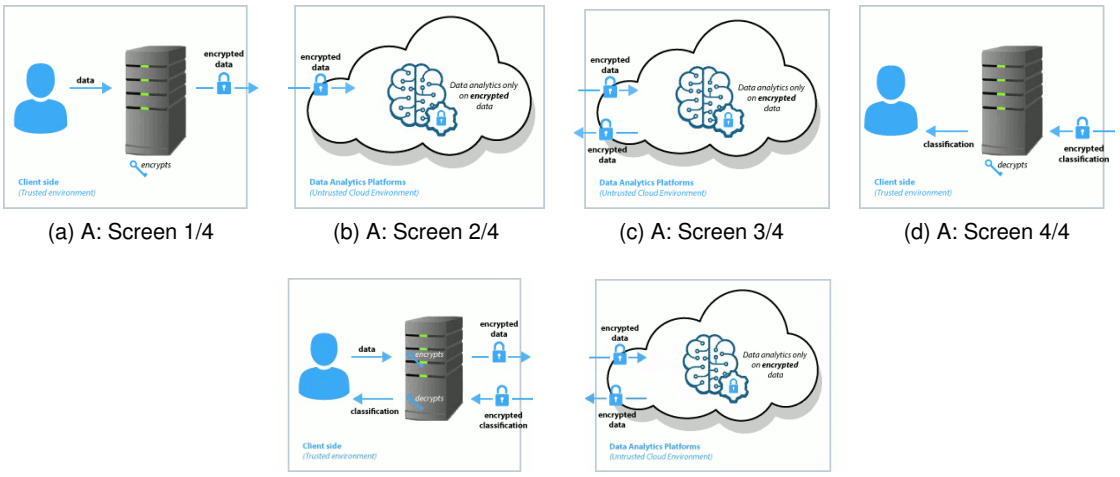

(a) A: Screen 1/4 (b) A: Screen 2/4 (c) A: Screen 3/4 (d) A: Screen 4/4

(e) B: Screen 1/2 (f) B: Screen 2/2

Figure 10: Mockups to explain classification based on homomorphic encryption: (a–d) variant A, (e,f) variant B.

the UIs for explaining homomorphic encryption could be replaced by other UIs for explaining alternative schemes, such as two-party computation.

3.1.2 Mockups

The mockups for this PET were implemented in two variants: Variant A (Figs. 10 a–d) consisted of four screens, each of which represented a segment of a linear narrative (Fig. 11a). The navigational flow of variant A was strictly rightward directed, i. e. users moved rightward to proceed the narrative, and leftward to track back. The narrative spanned across four operational phases, each of which was depicted on a dedicated screen. Despite the navigational and temporal flow facing rightward, all data sent from the analytics platform back to the client side were facing leftward. Conversely, variant B consisted of two screens representing the client side and server side, respectively, and users could navigate back and forth between the phases. In variant B, the direction of the navigational flow and data flow are identical, which was one of the reasons why this variant was preferred over variant A by the majority of the test subjects during the first iteration of the user tests (Section 3.1.3).

The illustrations of both variants were accompanied by an itemised list of four short paragraphs (where the word ”data” was treated as a singular to conform with a widespread colloquial language):

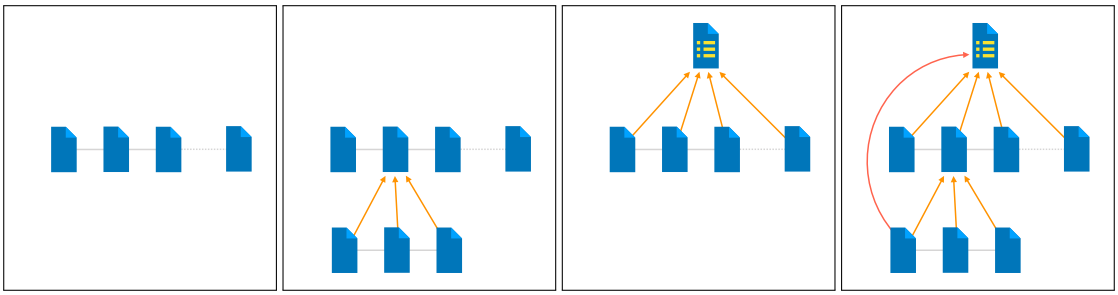

(a) Linear/sequential (b) Nested/layered (c) Initial root node (d) Root plus layering

Figure 11: Topologies of navigation. Arrows indicate navigational pointers from subordinate nodes (i. e. individual interaction phases) to superordinate nodes.

- Data is encrypted at the client side and sent encrypted to the analytics platform. (What is encrypted?)
- Data is analysed in encrypted form at the analytics platform (the platform cannot decrypt the data – at no time is the data available in decrypted form/clear text there) (How does this work?)
- Also the resulting classification is in encrypted form.
- Only the client can decrypt the classification.

For variant A, each item referred to and described the contents of one of the four illustrations displayed above the text. The semantic connection between text and graphics was emphasised by rendering the current paragraph in boldface, while a regular font weight was used for the three other paragraphs that referred to the phases currently not displayed. Due to the fact that variant B consisted of two phases only, a semantic connection to one of the items could not be made unequivocally. Hence, the text did not change when users navigated back and forth between the two phases of the narrative.

The text contained two hyperlinks that led to another screen, which provided further details on the subject matter at hand. The first link led to a screen that described how encryption works in general. The second link led to a screen that provided information on how homomorphic encryption works, as homomorphic encryption served as the cryptographic basis to facilitate the analysis, i. e. computation, of encrypted data. The example used for explaining how homomorphic encryption works was adapted from [5]. The latter explanation included two additional hyperlinks that lead to sample calculations whose purpose was to illustrate the basic functionality of homomorphic encryption. The screens are shown in Fig. 32 in Appendix A.1, and represent the layout used for iterations 1–3 of the user tests. Conceptually, the navigational structure of nested screens used for the design constituted multilayered information (Fig. 11b).

3.1.3 Testing

The mockups (Fig. 10) were implemented as click-through wireframes running on a portable computer. The purpose of the prototype was to inform potential users of a system that operates on the basis of the aforementioned analysis platform, and allow them to make informed

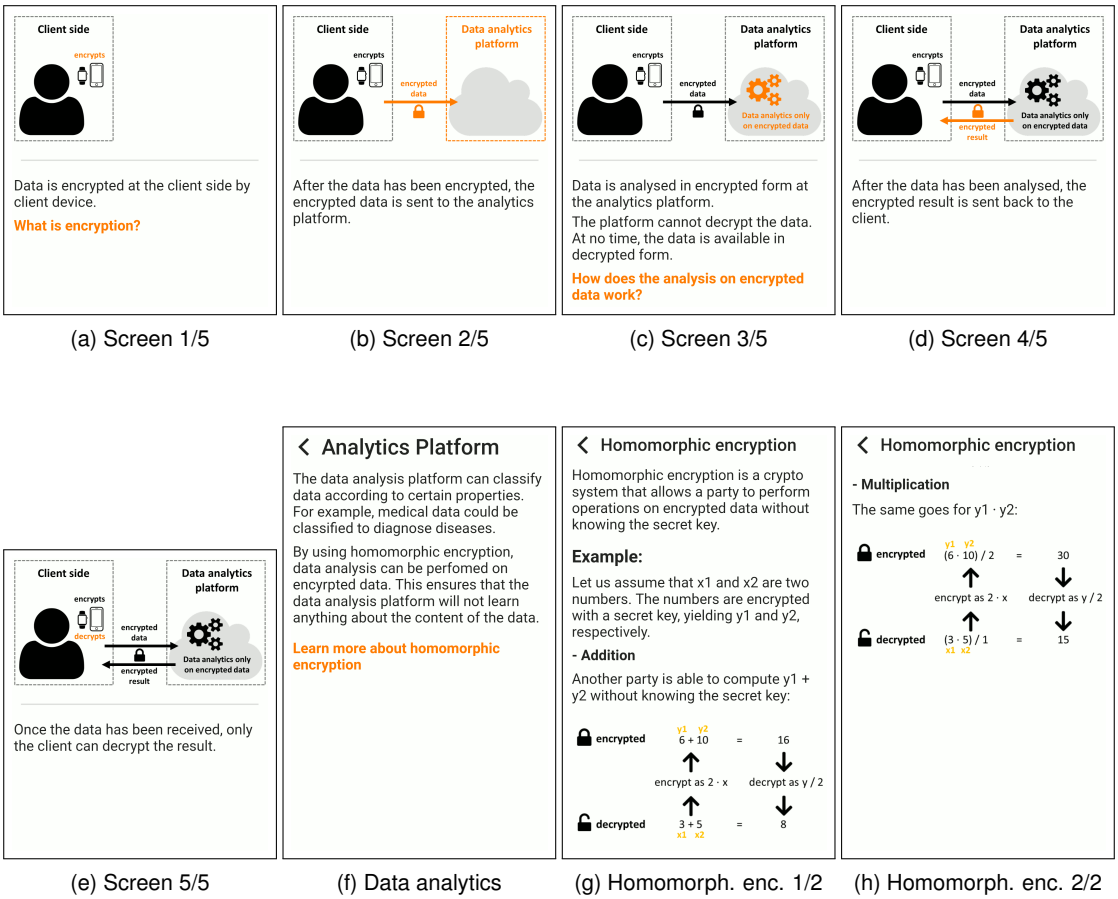

(a) Screen 1/5 (b) Screen 2/5 (c) Screen 3/5 (d) Screen 4/5

(e) Screen 5/5 (f) Data analytics (g) Homomorph. enc. 1/2 (h) Homomorph. enc. 2/2

Figure 12: Use case 1/iteration 4: (a–e) analytics of encrypted data; secondary information on (f) the nature of the analytics platform, and on (g,h) homomorphic encryption (adapted from [5]).

decisions about whether they wanted to use the service in question.

To ascertain to what extent the prototype fulfilled that role, and what factors had to be improved to fulfill it more adequately, user tests were conducted in the form of lab studies. During the study, lay persons were first asked to try out and interact with the prototype for themselves. After having experienced the prototype first-handedly, the participants were interviewed about what they had experienced and learned.

The user tests followed the following conceptual structure:

1. Brief introduction to the topic (data analysis on the Internet).
2. Interaction with the prototype.
3. Questions about the knowledge obtained.

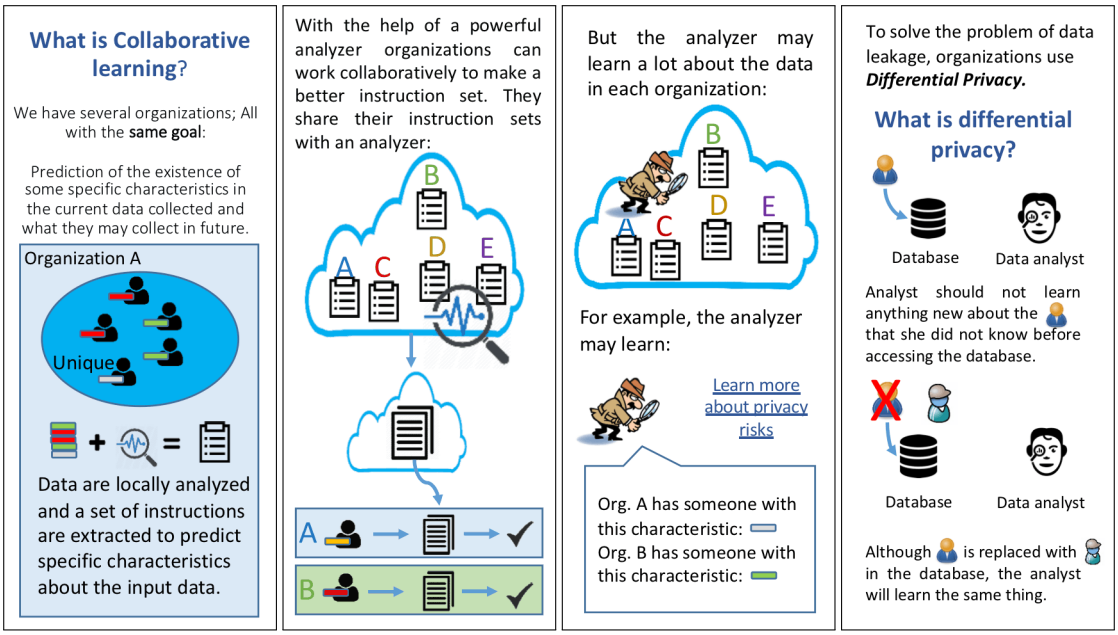

(a) Introduction (b) Analysis (c) Risks (d) Diff. privacy

Figure 13: Second iteration of mockups dealing with differential privacy applied to collaborative learning.

4. Discussion of the prototype together with the moderator(s).

The questions asked to the participants (see Appendix B) covered the following topics, all of which were related to the analysis platform described previously in the prototype:

Comprehension. To what extent test subject had comprehended the individual components that constituted the analysis platform?

Security. To what extent data that were disclosed to the analysis platform was secure?

Trust. To what extent the test subject would trust her data to the analysis platform?

The last part of the study, the joint discussion of the prototype between participant and moderator, served the purpose of resolving recognitive gaps that might have prevailed during the interaction and subsequent answering of the questions. It helped the moderator understand in which cognitive facets the test subject's mental model deviated from the designer's, which provided valuable insight in terms of possible improvements for future versions of the prototype.

The user tests were conducted in three successive iterations. The first iteration was based on an implementation of variant A of the mockups (Figs. 10 a–d). As part of the study, participants were asked in the end, whether they would have preferred variant B instead of the one they had been using. As the majority of participants favoured variant B, iteration 2 (Fig. 33 in the

appendix) was conducted based on the implementation of B instead of A. Iteration 3 (Fig. 34 in the appendix) was based on a revised design that accommodated the issues detected and points raised during the previous tests.

Each iteration collected the input from six participants, all of which were international guest or exchange students currently staying at Karlstad University. They hailed from a variety of countries, pursued different subjects, and were aged 20–26. The purpose and process of the user study were registered at the research ethics advisors at Karlstad University, who confirmed that no application for ethics approval needs to be submitted for this kind of study.

There were a number of ways participant could (and did) misunderstand the concept. Even if the UI is rather generic for homomorphic encryption, it helps to have a concrete framing story for most test participants. The UI was gradually improved and the user evaluation of iteration 3 shows an improvement in comprehension compared to previous iterations. The difference between App, Encryption & Analyses was understood by 5 of the 6 participants. The Client-Server model (sending and receiving of data) was understood by four participants. Still there are problems to fully understand the analysis platform even at the superficial level given by the UI. Only 2 of the 6 participants in the last round can be said to comprehend. Nevertheless, 4 participants would not mind their data being treated in this way, and only 1 was against it. This person would like to have some assurance that ”they used only algorithms not employees”.

While the samples were small and therefore also for that reason biased, the user tests were conducted for the purpose of iterative user interface prototyping. According to Nielsen [6], 5 test users for each iteration are considered sufficient for an iterative design process. Due to the Corona pandemic, user tests of the fourth iteration could not be conducted before completing this deliverable. Nonetheless, further tests with at lest 15 persons is planned to be conducted for the final user interfaces in the last project year.

3.1.4 User Interfaces

Figure 12 shows the fourth and final iteration of the prototype, which is implemented in React Native for Android mobile phones.

3.2 UIs for Privacy preserving Neural Networks – Explaining Collaborative Training with Differential Privacy

3.2.1 Background and Motivation

This section presents User Interfaces for explaining privacy-preserving neural networks for collaborative training utilising differential privacy. This category of PETs and the UIs can for instance be used for PAPAYA’s use case on stress detection (UC 2 under the health care umbrella).

These UIs were more complex and challenging than the other UIs that we developed for explaining PETs, because different aspects needed to be explained, including: What is collaborative learning and what are privacy risks of collaborative learning? What is differential privacy, how can it minimised privacy risks and how can it be applied to collaborative learning with what kind of trade-offs?

For conveying these different aspects, a multi-layered approach was chosen, which we will present below.

3.2.2 Mockups

The criteria for explaining collaborative training with differential privacy to users of respective data services were as follows:

Utility. The utility of data analytics shall be explained. It shall be conveyed that using the service will bring added value for users.

Risk. The privacy risks that will or might arise due to data analytics shall be explained. It shall be clarified that using the service might expose users to certain risks.

Mitigation. It shall be explained how the aforementioned risks can be mitigated while making acceptable utility trade-offs. Knowing both utility and risk shall enable users to make informed decisions as to whether they want to use the service.

The navigation of the prototype was conceptualised as a layered approach (Fig. 11b), consisting of an overarching narrative that aimed at clarifying the aforementioned interplay of utility, risks, and risk mitigation. Multiple subordinated screens were dedicated to clarifying individual facets pertaining to each of these topics. The purpose of conceptualising secondary information as nested pages that were accessible upon requests was to facilitate a layered approach of conveying information [1]. This not only reduced the cognitive load imposed on users while reading and processing the information on any given screen, but also to enable them to access the kind of information they were actually interested in.

Multiple iterations of mockups were designed by individual researchers, discussed in a group in cognitive walkthroughs, and then refined based on the decisions made. Fig. 13 (the full series is provided in Figs. 35–37 in Appendix A.2) shows the second iteration of mockups created to explain the concepts of collaborative learning, differential privacy, and the application of the latter to mitigate the risks that may arise in the former. These mockups served as the basis for the refined design presented in Section 3.2.3.

3.2.3 User Interfaces

Once the contents and scope of one of the drafts had been agreed upon by all members of the design team (Section 3.2.2), the illustrations were refined such that they conveyed a consistent theme and style throughout the entire narrative (Figs. 14, 15, 16).

The topology of the navigation was extended from a layered approach (Fig. 11b) to a narrative that commenced on a superordinate root node (Fig. 11c). This start page served the purpose of a ‘Questions & Answers’ page that was supposed to help address specific questions users might have about the scenario at hand (Fig. 14a). Individual topics were reachable via hyperlinks that referred the reader to subordinated nodes.

The design featured a fixed navigation bar at the bottom of the screen that implemented the following four controls, as shown at the bottom of Fig. 17:

Home. Navigate back to the root node.

Up. Navigate back to whatever node served as a referrer that linked to the current screen.

Backward. Navigate to the predecessor of the current screen in the current linear narrative (Fig. 11a).

Forward. Navigate to the next screen in the current linear narrative (Fig. 11a).

Having these options available at all times ensured that users would be able to navigate quickly between the topics. Consequently, the topology represented in the final prototype reflected layers of linear narratives (Fig. 11a), secondary information in the form of multilayered tiers (Fig. 11b), a start page in the form of a root node (Fig. 11c), and quick navigation between individual screens of the prototype (Fig. 11d).

Figure 17 shows the fourth and final iteration of the prototype, which is currently implemented as an React Native app running on Android phones.

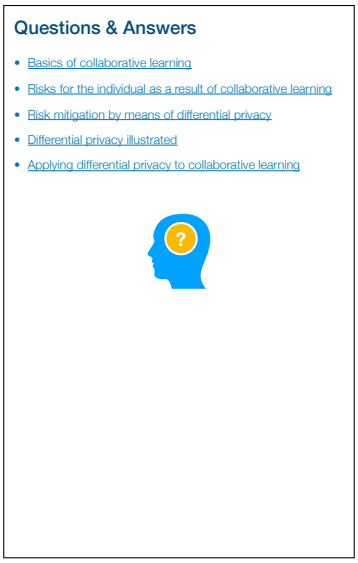

(a) Q & A (root node)

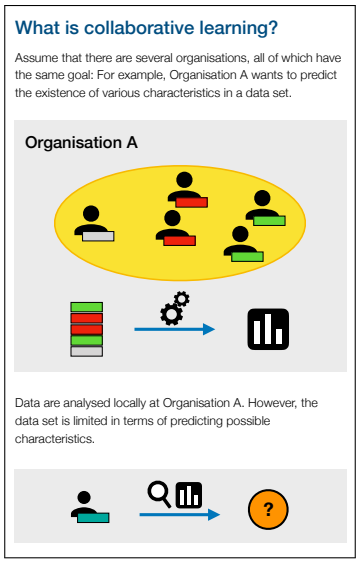

(b) Collaborative learning 1/3

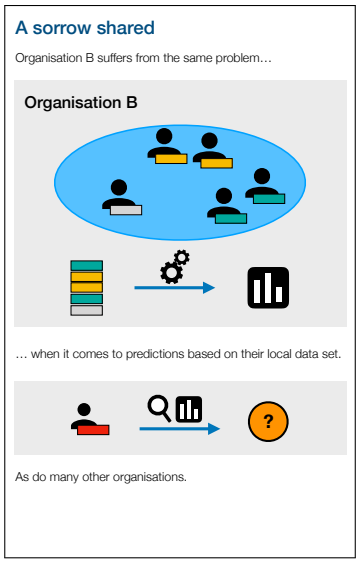

(c) Collaborative learning 2/3

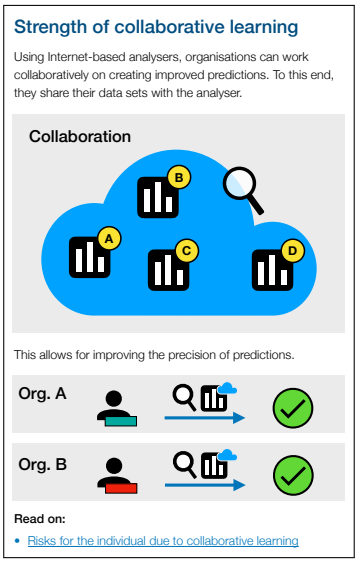

(d) Collaborative learning 3/3

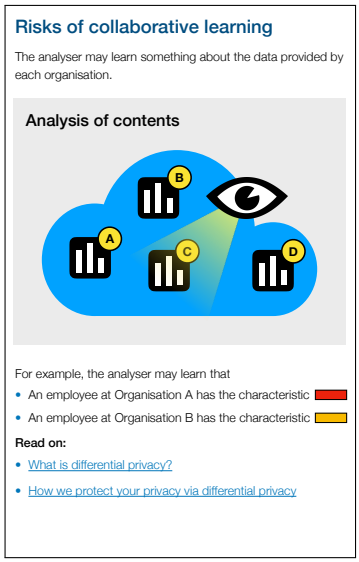

(e) Risks of collaborative learning

Figure 14: Utility and risk of collaborative learning: (a) Q&A: access to various sub topics, (b–d) basics of collaborative learning, (e) risks of collaborative learning.

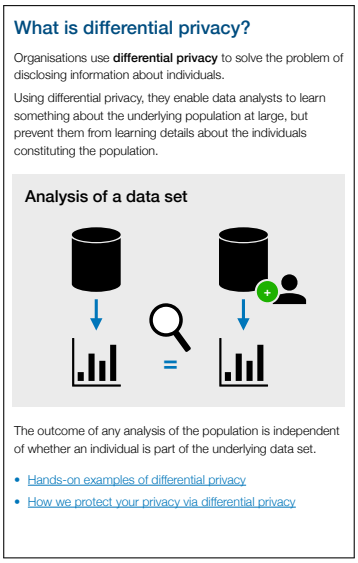

(a) Overview

(b) Hands on examples 1/3

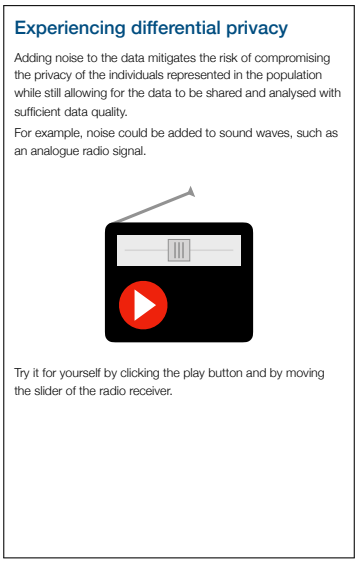

(c) Hands on examples 2/3

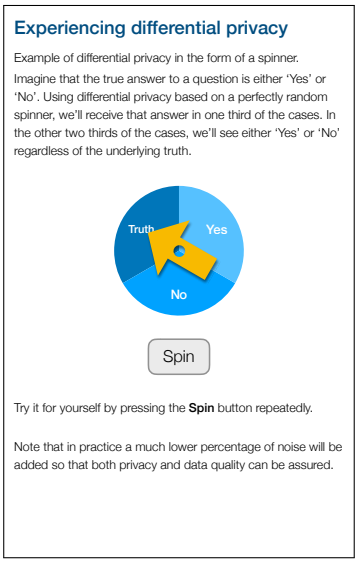

(d) Hands on examples 3/3

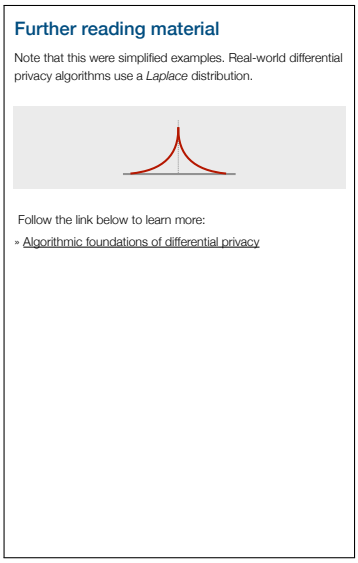

(e) Access to further material

Figure 15: Explaining differential privacy by means of multi-modal and interactive examples.

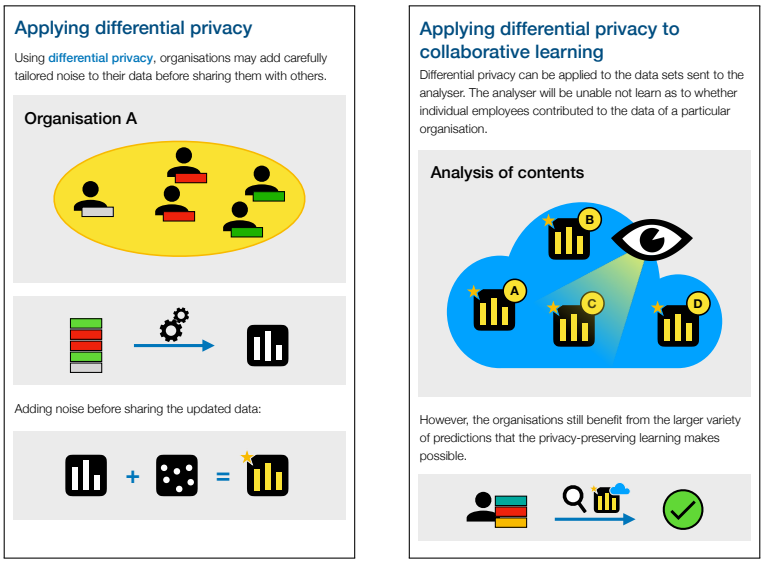

(a) Applied to collab. learning 1/2 (b) Applied to collab. learning 2/2

Figure 16: Risk mitigation via differential privacy.

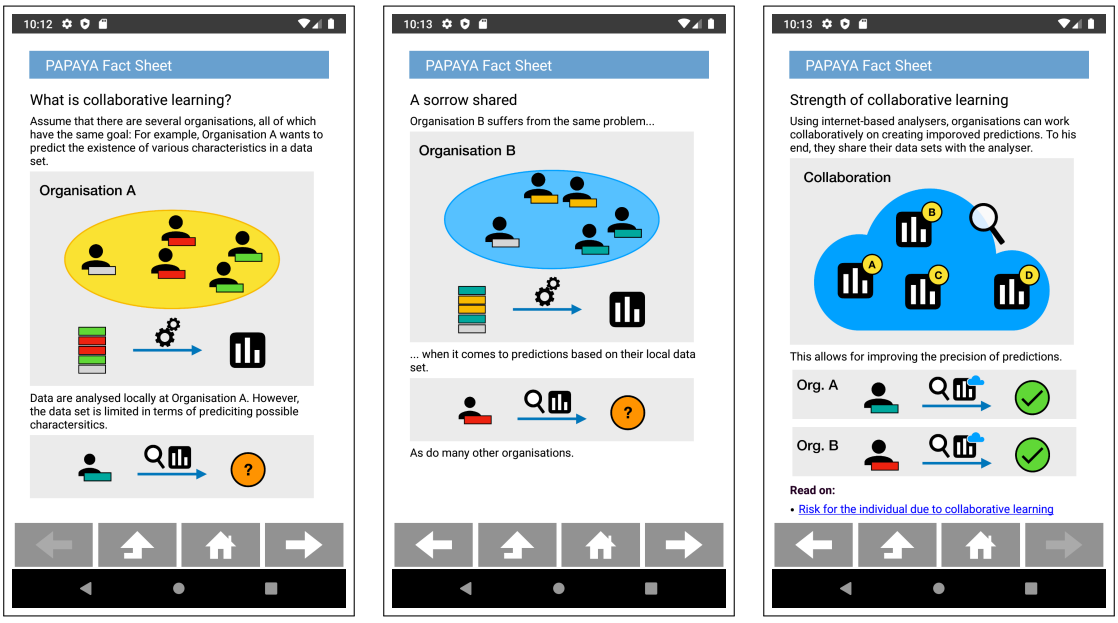

(a) Collaborative learning 1/3 (b) Collaborative learning 2/3 (c) Collaborative learning 3/3

Figure 17: Collaborative learning: screen shot of the native prototype running on Android.

3.3 UIs for Explaining Multi Party Computation

Multi Party Computation (MPC) is one of the underling cryptographic primitives that is used for Privacy Preserving Neural Networks, Privacy Preserving Clustering, and Privacy Preserving Counting . Since MPC is the mechanism that ensures privacy, the focus will be on MPC in this section.

3.3.1 Background and Motivation

MPC is used in the PAPAYA's Mobile usage analytics use case (UC3) and, in the form of Two-Party Computation, also used in the Arrhythmia detection case (UC1). However, the explanation here aims not at a particular use case but tries to explain the underlying mechanism of MPC in general as easy as possible.

The mockups were generated with two main goals in mind:

- First, the descriptions should be understandable without deeper technical knowledge while being still correct and close enough to the technical implementation.
- Second, the descriptions should be short to not overstrain the attention span of readers and provide visual aids to make it easy to follow.

3.3.2 Mockups

The mockups were generated with rapid prototyping and evaluated with impression gorilla testing. Gorilla testing is a very informal but useful type of user experience (UX) testing that focuses on qualitative feedback by interviewing or observing people in a casual setup, often performed in the office corridor, a coffee place, or over video chat. Impression gorilla testing is mostly meant to catch logical errors and ask people about their first impression and general understanding. Some non-research test users with different non-technical background have been included in the evaluation process. In total six people participated. One of the six participants is an UX-designer that provided valuable expert feedback.

It became obvious how different the preferences of people are. Some would have preferred more technical details but even then, the focus of different people with that preference was not uniform. The good news is that everyone, in this limited test set, understood the mockups and could correctly explain afterwards, what the described mechanisms do.

In this limited test setup, users liked having an example more than a general description, which is the reason why an example driven approach was chosen for explaining MPC.

The resulting mockups can be seen in Figure 18. If different preferences were expressed by the testers, the most non-technical approach was chosen that still enabled testers to correctly understand the mechanism. The resulting example explains how the average of two numbers can be calculated without exposing these two numbers. The Idea is to combine textual description next to a visualisation. Colours and symbols are used to emphasise correlations between text, visuals and different steps in the example. However, additional testing would be required to see if that approach sufficiently works with a broader audience.

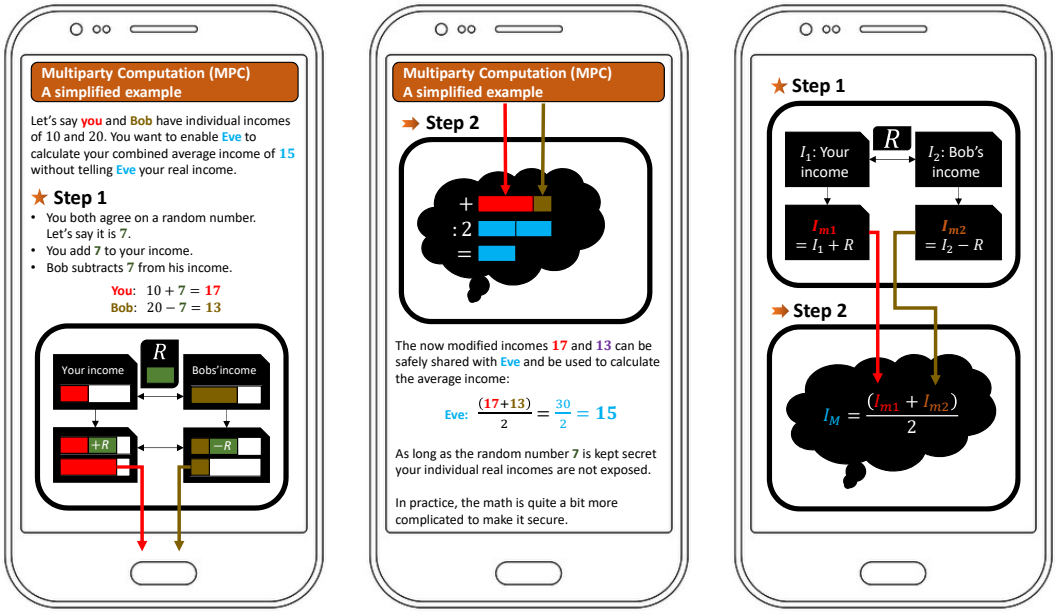

Figure 18: Explaining Multi Party Computation mockup.

3.4 UIs for Explaining Functional Encryption

This Section describes how the advanced cryptographic primitive Functional Encryption can be explained to non-expert users. Functional Encryption is a powerful tool that is used for Privacy Preserving Clustering and Privacy Preserving Counting.

3.4.1 Background and Motivation

The motivation, approach, and testing is mostly the same as described in Section 3.3.1 and Section 3.3.2. In addition, Function Encryption is to some degree more challenging to explain without using technical terms or advanced mathematical concepts since it has the indirectness of the functional part. However, avoiding the technical terms and using the frequently used key and lock metaphor to describe encryption worked here well enough in the gorilla testing.

3.4.2 Mockups

The goal was to avoid technical or mathematical terms like ”function” or ”encryption” since initial tests showed that some people do not understand these terms. Future research should investigate if that approach was too restrictive and a more technical explanation, that explains the functional part more detailed, would be an improvement. The following narrative to explain Functional Encryption to a non-technical users was chosen:

Functional Encryption explained

Let us say your user profile contains a lot of private information. Since you do not want all that information available to everyone, you lock the information away, so only you can read it. Now only you, who has the master key, can read your user profile.

However, sometimes it is convenient or simply nice to share some of this information:

- *If you have a medical emergency, you most likely are fine with sharing your medical records with the doctors.*
- *To find the best restaurant in the area, you might be okay to share your location area one time.*

Functional Encryption allows you to do that by only sharing what you need to share and not more. For that, Functional Encryption allows you to make specialised subkeys from your master key. If you give such a specialised subkey to someone else, together with your locked away user profile, this other person can only get the specific information that the key was made for. Only for that, this specific subkey works. If this person would try to get other information about you with this subkey, it would simply not work.

Thanks to computer technology, most of that happens automatically as long as you agree with what information is shared.

For Functional Encryption, an narrative driven approach was chosen. The slightly more technical graphical mockups can be seen in Figure 19. They are meant to transform the narrative driven learned knowledge into structured knowledge and serve as a repetition at the same time.

In addition, a first example driven explanation was generated which can be seen in Figure 20. However, this approach is not yet evaluated and more testing is needed.

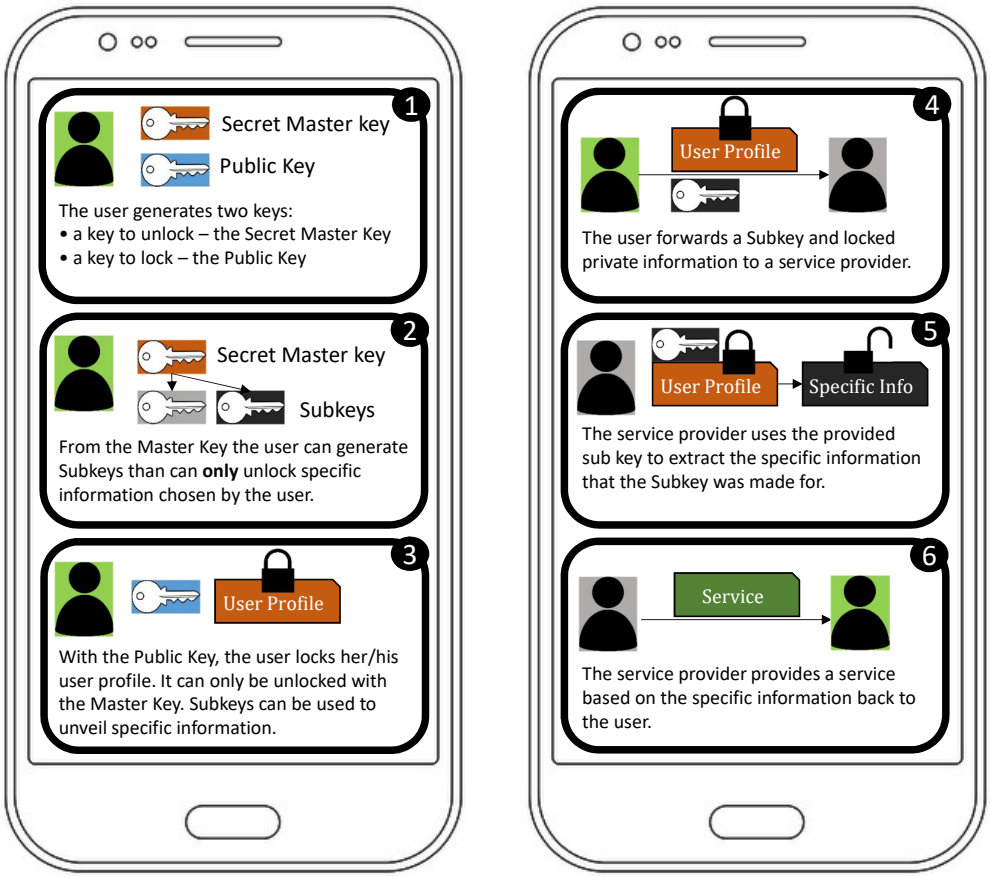

Figure 19: Explaining Functional Encryption mockup.

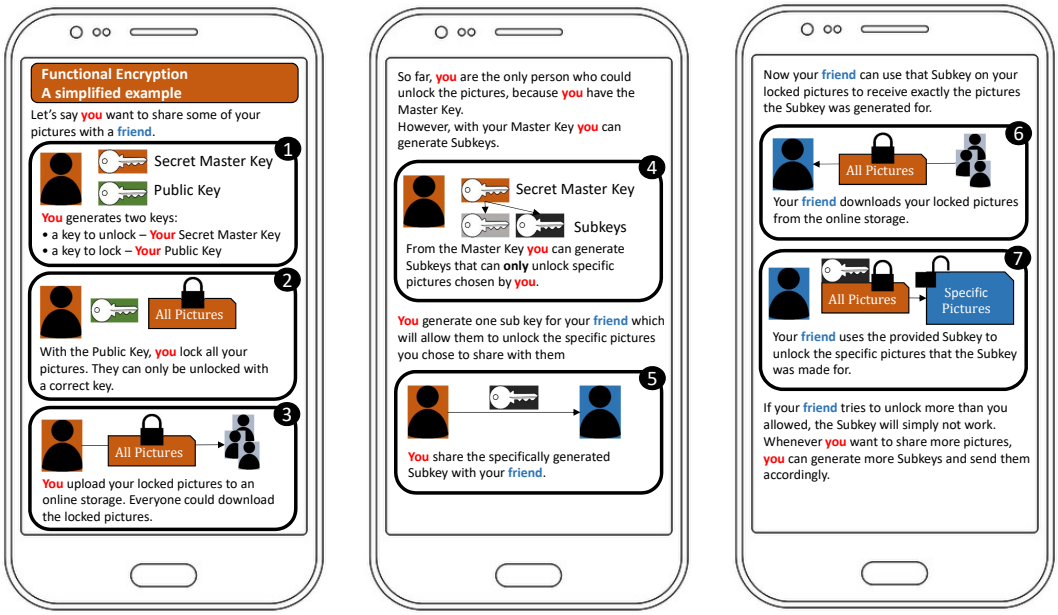

Figure 20: Functional Encryption example mockup.

4 User Interfaces for the Privacy Engine

The Privacy Engine has been designed and developed in order to provide a set of tools to the end user (or Data Subject) to help in maintaining his/her privacy preferences and in providing an easy interface to exercise the data subjects' rights defined in current legislation (e.g. GDPR). Furthermore, the Privacy Engine also provides interfaces for Data Controllers and Data Privacy Experts to define and configure its functionalities. This chapter describes the User Interfaces (UIs) that we implemented for those actors who will use the services of the Privacy Engine. In order to obtain technical details on the analysis and the design of the different components of the Privacy Engine, the reader can refer to the deliverable [8]. The functionalities of the Privacy Engine can be divided into two main blocks as follows:

- **Privacy Preference Manager:** where the data subjects can define their privacy preferences on the management of their personal data.
- **Data Subject Rights Manager:** to facilitate the data subjects to exercise their rights expressed by the GDPR.

The following subsections detail the User Interfaces for both of these functionality blocks.

4.1 Privacy Preferences Manager User Interfaces

The Privacy Preferences Manager (PPM) has been designed and developed to allow the end users or Data Subjects to express their preferences when sharing their personal or sensitive data with the Data Controller. For that purpose, the PPM provides two different interfaces devoted to two different actors:

- **Privacy Expert:** the Privacy Expert will be designed by the Data Controller and he/she will be in charge of the definition of the different questions to be answered by Data Subjects.
- **Data Subject:** using a mobile interface, the end user will be able to answer the questions defined previously by the Privacy Expert.

The following subsections will define both interfaces in detail.

4.1.1 PPM - Privacy Expert Interface

As a first step, the Data Controller has to nominate a Privacy Expert to be in charge of the development of the question definitions. These questions must be defined taking into consideration that a regular Data Subject (not a technical or legal expert) should be able to understand the content of the question and, at the same time, will be easily answered using the dichotomic option: to grant access or to deny access. In addition, as part of the definition process, the Privacy Expert must classify different aspects associated with each question, such as the type of data, type of data processing, etc. This characteristic classification is based on the main results obtained by the project Special[7]

[7] https://www.specialprivacy.eu

Figure 21 shows the interface for the Privacy Expert for defining the question. On this screen, the Privacy Expert can define the title of the question, the content of the question and the target URL where the data will be sent when the Data Subject will answer the question. As mentioned before, the Privacy Expert must take into consideration the acknowledgement of the Data Subject on the matter, and, therefore use simple language.

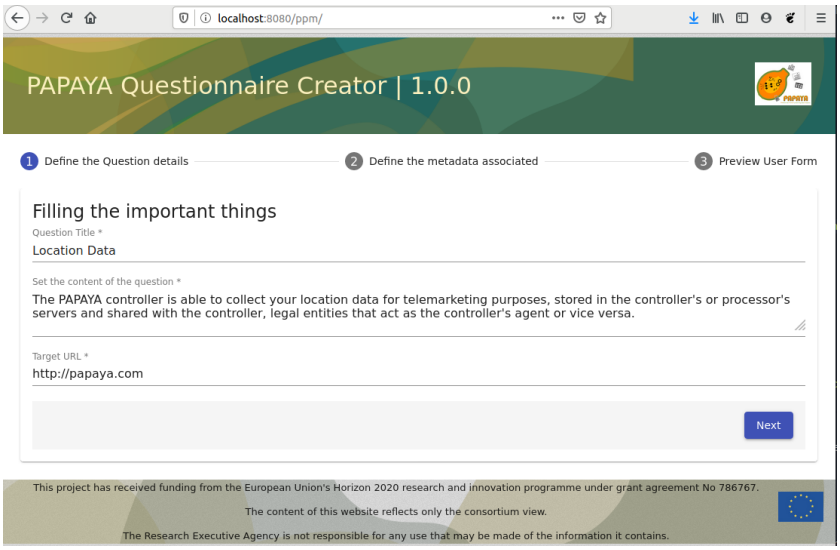

Figure 21: Question definition.

Once the body of the question has been defined, the Privacy Expert, by clicking on the next button can start with the classification of the characteristic associated with the question. Using this interface the Privacy Expert is able to configure the following characteristics:

- **Data**: the data processed by the operation
- **Processing**: a description of the operation itself (e.g. “query”, “classification”, “disclosure”, etc.)
- **Purpose**: the purpose of the operation
- **Recipient**: the entities that can access the result of the operation (recipients)
- **Storage Location and Storage Duration**: a description of where the result is stored and for how long

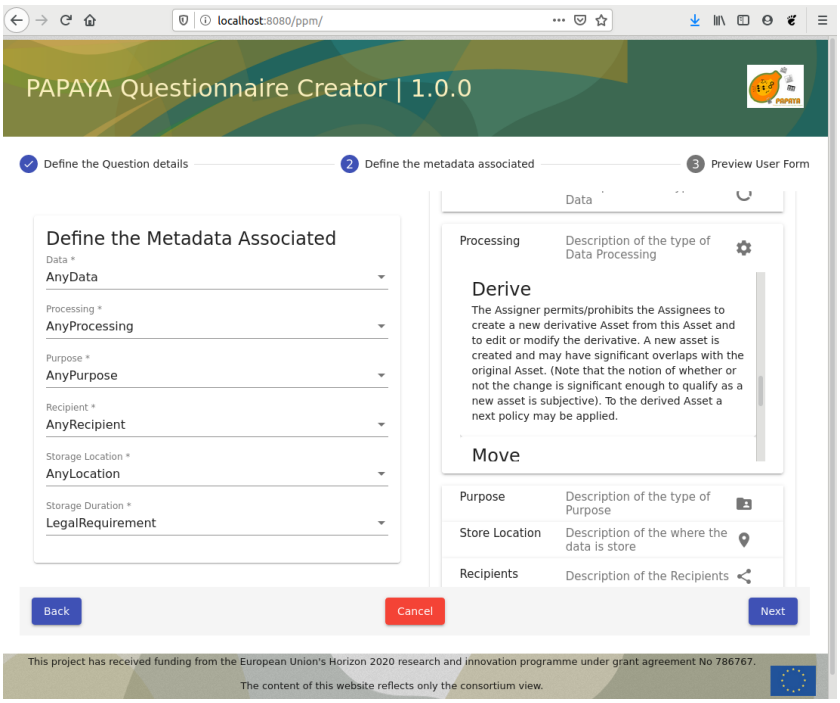

Figure 22: Configuration of the metadata associated to the question.

This classification has been obtained from the definition of the usage policy established in the [11] deliverable of the Special project. Furthermore the Special project defines within the cited deliverable a whole privacy-related vocabulary for specifying the possible attributes of the characteristics associated with each question. Figure 22 shows the interface for the Privacy Experts and how they easily select for each characteristic the most suitable attribute for the question. In addition, on the left side of the screen, the Privacy Expert can consult the attribute definitions in order to help him/her to select the proper one.

After configuring the proper metadata to the questions, the Privacy Expert can click on the next button in order to carry out the following step. At this step the PPM transforms the information filled in by the Privacy Expert to a HTML form file (including the metadata). In addition, on the right-hand side of the screen, the Privacy Expert can easily view a preview of how the HTML form will look like in the end user's browser. The Consortium chose to export the question information to a HTML form because the broad use of this protocol will be easily stored and integrated on any of the Data Controller's final Service Provider or mobile application.

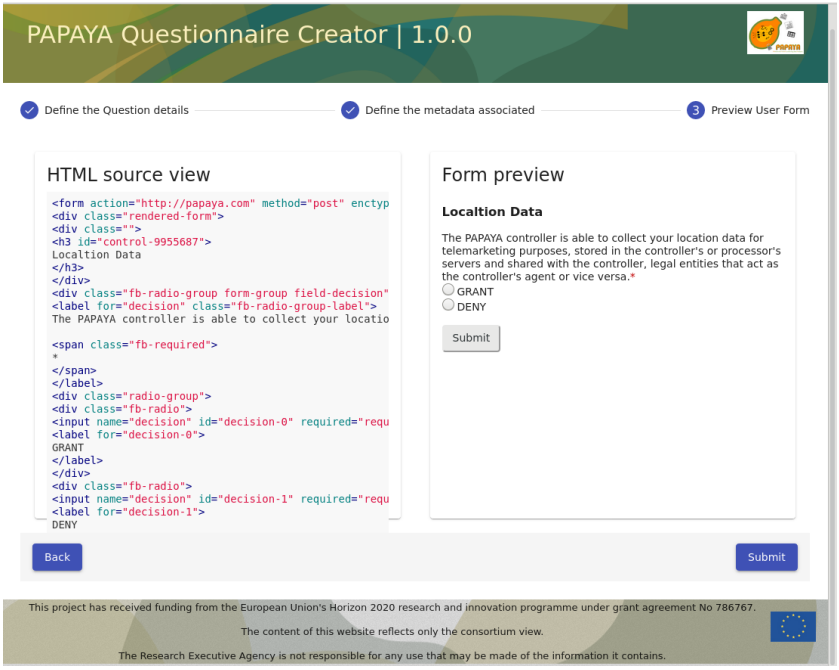

Figure 23: Question in HTML format and form preview.

If the Privacy Expert validates the final format of the question, by clicking on the Submit button, he/she can store the question into the server, then it will be available for the Data Subjects. Figure 23 and Figure 24 displays the end user interface provided by the PPM. Otherwise, the Privacy Expert can always click on the Back button on any of the previous steps to make all the changes necessary.

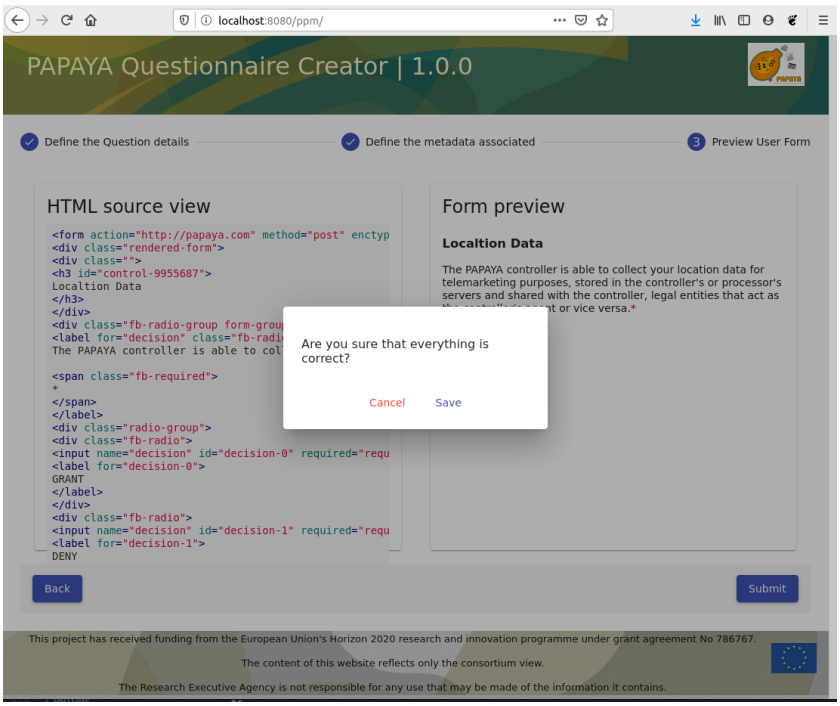

Figure 24: Operation confirmation.

4.1.2 PPM - Data Subject Interface

Once the different questions have been defined by the Privacy Expert, they are available for the Data Subjects. In this particular case, the Data Subjects have a mobile application interface. Figure 25.a) shows the initial screen welcoming the end user and informing him/her of the main functionalities available in this mobile interface. This welcome screen allows the Data Subject to consult the list of questions available or to go directly to answer the first one. In the case where the Data Subject will click on the list of answers, he/she can easily consult all the different questions defined by the Privacy Expert (as can be seen in Figure 25.b).

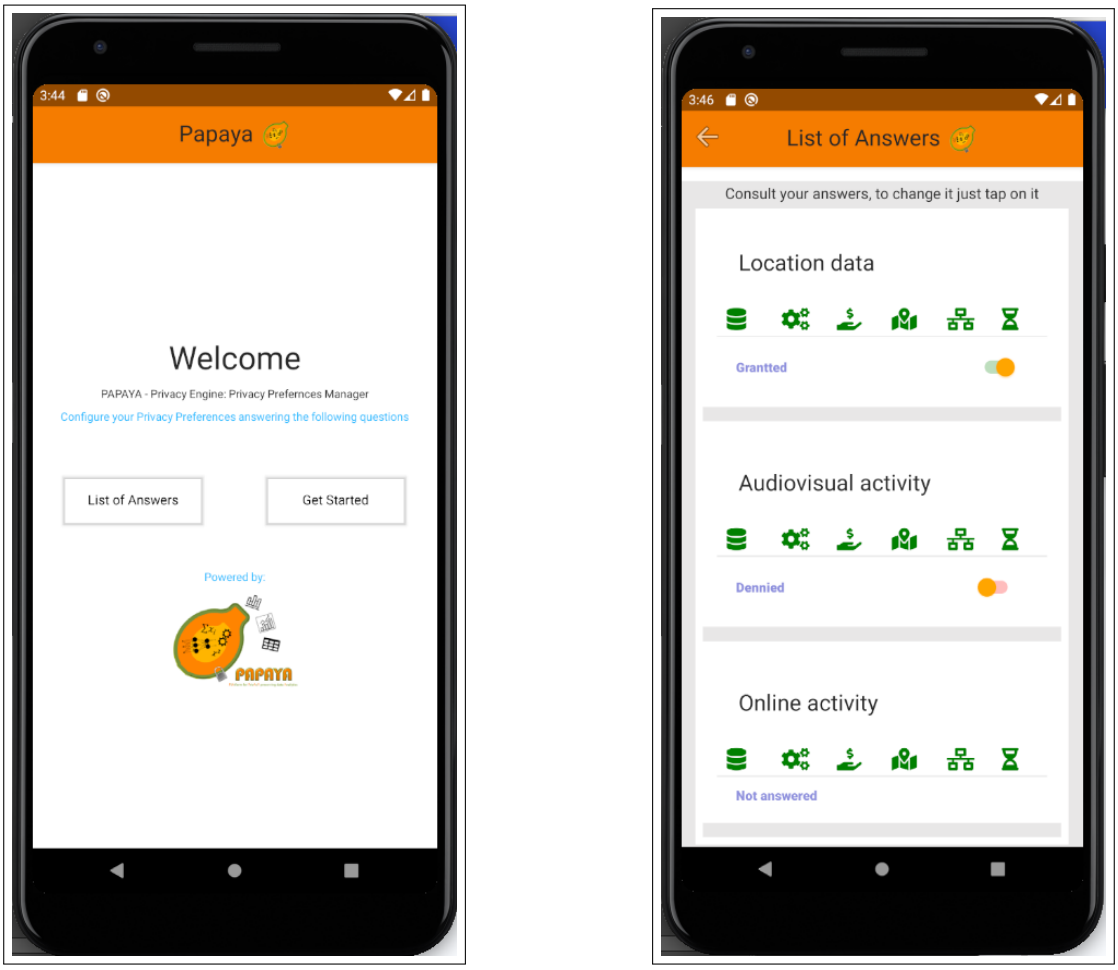

(a) Welcome Screen

(b) List of Answers

Figure 25: Privacy Preferences Manager mobile application I.

The list of answers shows the title of all the questions available, the metadata associated with each question and the decision taken by the Data Subject. It is worth mentioning that to follow the GDPR recommendations, the answers are not predefined by default and they will only show the decision of the Data Subject or ”Not answered” message in the case he/she has not selected any choice. This list view allows the end user to tap on each question, and then shift to the question details screen (as shown in Figure 26.a). On the Question detail interface, the end user can consult the number of this question and the number of questions available on the top part of the screen. In addition, the user can read the title and content of the question defined by the Privacy Expert. Based on the content of the question, the Data Subject can select between Grant or Deny access to his/her data, just tapping on the showed buttons. In this particular case, it is worth highlighting, to help the understanding of the end user, the colour of both buttons will change, depending on the selection the end user makes, making it

much easier to understand the selection taken by the user. On the bottom part of the Question details screen is available three different buttons to facilitate the end user to navigate through the mobile application. Therefore, the end user can always come back to the list of answers by tapping on the corresponding button, or going to the next question or either skip this question if he/she does not have a final answer to the question defined.

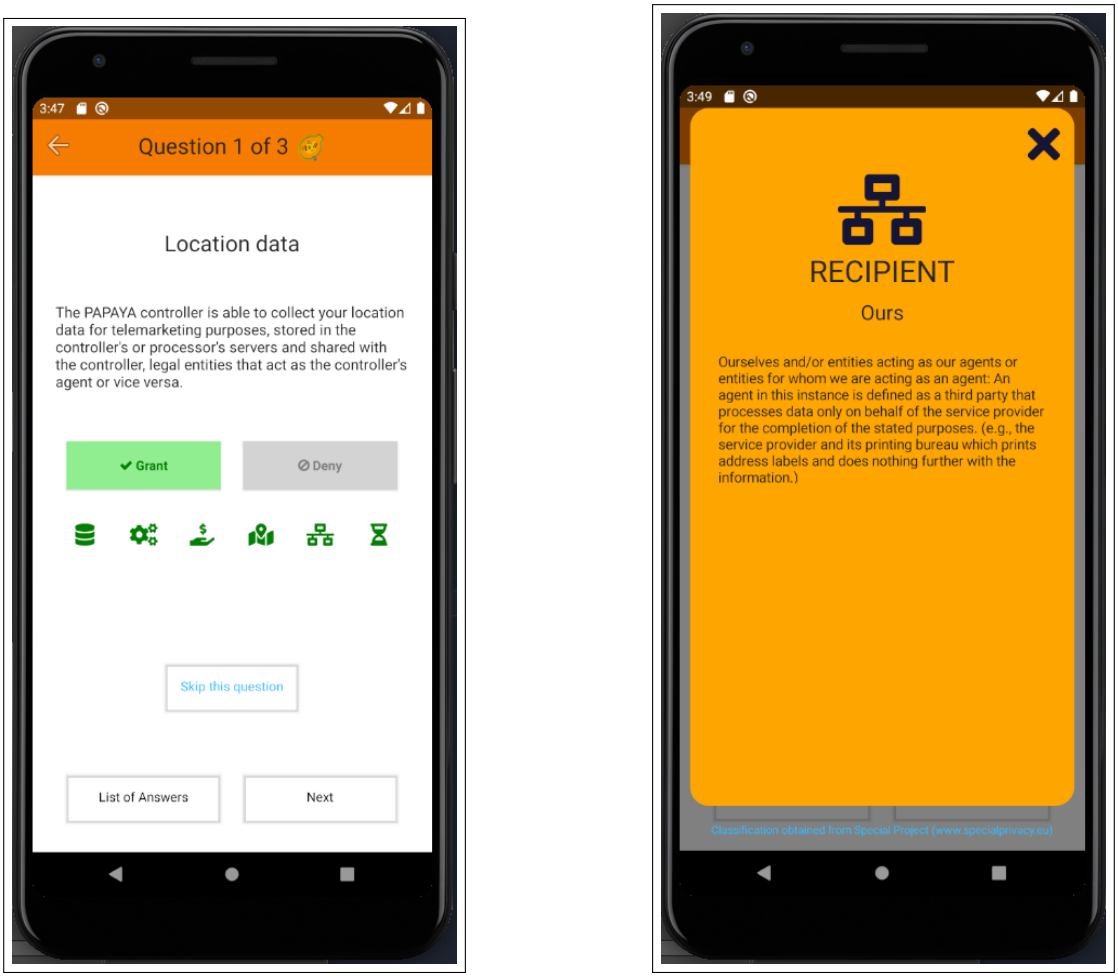

(a) Question details

(b) Metadata Details

Figure 26: Privacy Preferences Manager mobile application II.

Figure 26.a also shows a row of green icons, associated with the characteristics defined previously by the Privacy Expert, the Data Subject can tap on any of them, to open a modal screen displaying the specific detailed information of the specific attribute selected by the Privacy Expert, as shown in Figure 26.b. Each icon corresponds with a definition from the usage policy defined in the [11], a deliverable of the Special project. The icons (from left to right) stand for the following characteristics: Data, Processing, Purpose, Recipient, Storage Location

and Storage Duration.

4.2 Data Subject Rights Manager User Interfaces

The Data Subject Rights Manager (DSRM) is the second set of functionalities that the Privacy Engine provides. In this particular case, the DSRM is designed and implemented to help the Data Subject, but also the Controllers, to comply with the current legislation on privacy matters, where especially the GDPR defines some rights for the Data Subjects. These rights are defined in Chapter 3 in the [3] and they are as follows:

- The right to be informed (article 13[8] and 14[9])
- The right of access (article 15[10])
- The right to rectification (article 16[11])
- The right to erasure (article 17[12])
- The right to restrict processing (article 18[13])
- The right to data portability (article 20[14])
- The right to object (article 21[15])

Although the aim of this document does not go into detail with the legal aspects of these rights, it is worth highlighting that by law the Data Controllers must provide the services and means to guarantee the application of these Data Subjects' rights in all those operations using personal or sensitive data. The DSRM will help to the Data Controllers to provide these means as it decouples the application of the rights from the specific implementations of its systems. The DSRM allows the Data Controllers to select for each specific right how its system can react to it: sending an email, publishing a notification (following the the publish-subscribe pattern) or triggering an action on the Protection Orchestrator. This document is focused on the description of the user interfaces, nevertheless more technical details describing this decoupling mechanism can found in the deliverable [8]. Therefore, the DSRM has developed two different user interfaces for the two different actors involved:

- **Data Controller Administrator Interface:** where the administrator can select and define what type of reaction for each different right.
- **Data Subject Interface:** where the end user can be informed of the rights that they are in title and if they consider it necessary, they can be exercised

The following sections describe both interfaces in detail.

[8] https://gdpr-info.eu/art-13-gdpr/
[9] https://gdpr-info.eu/art-14-gdpr/
[10] https://gdpr-info.eu/art-15-gdpr/
[11] https://gdpr-info.eu/art-16-gdpr/
[12] https://gdpr-info.eu/art-17-gdpr/
[13] https://gdpr-info.eu/art-18-gdpr/
[14] https://gdpr-info.eu/art-20-gdpr/
[15] https://gdpr-info.eu/art-21-gdpr/

4.2.1 Data Controller Administrator Interface

This web interface is devoted to allow the Data Controller Administrator an easy way to configure for each different right the type of reaction that the system will do. For instance, the Administrator can configure that for the right to erasure and the system will send an email to him with all the information necessary and for the right to access the system publishing a notification that the corresponding department is subscribed to obtain the information to react to it. As Figure 27 shows, the Administrator can select on the top part of the screen, between the different tabs available, the right to configure. The same Figure also shows that in the central part of the screen the end user can click on three different options (in blue): Email configuration, Publish notification and Configuring the Protection Orchestrator. By clicking on each one a specific menu form will drop down with the specific attributes for either option.

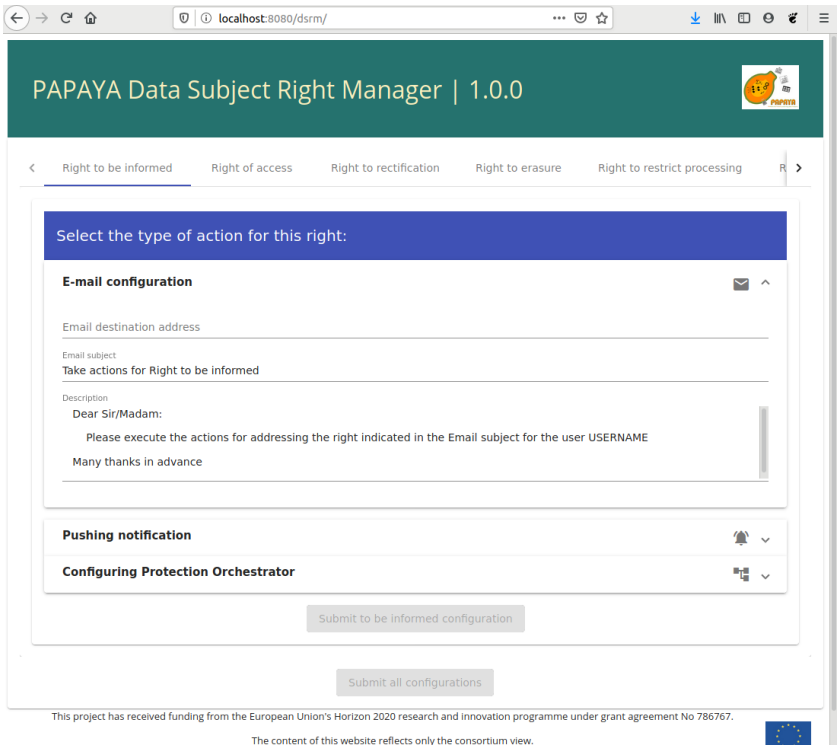

Figure 27: Recipient classification.

Going into detail, for the Email configuration (Figure 27), the Administrator can define the following attributes:

- **Email destination address:** the administrator can define to whom the email will be sent when the data subject wants to exercise this right.
- **Email subject**: the end user can define what will be the subject of the email sent

- **Email Description:** defining here the content or body of the email. Here it is worth highlighting that in the content of this email the reserved word USERNAME will be replaced with the real end user identifier when the event will be triggered.

With regards to the Notification configuration, Figure 28 shows how the administrator is able to configure the URL of the Kafka server where the notification will be published and the associated topic, or in other words, the channel where the subscriber will be waiting for the notification.

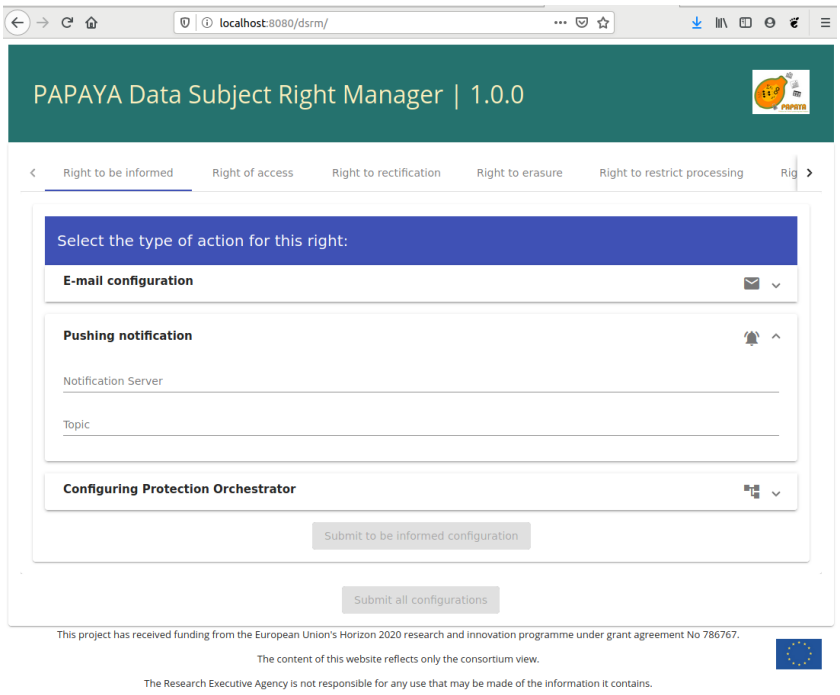

Figure 28: Recipient classification.

Last but not least, Figure 29 shows how the end user can provide the BPMN 2.0[16] file with the sequence of different steps to be executed by the Protection Orchestrator.

[16]https://www.omg.org/spec/BPMN/2.0/

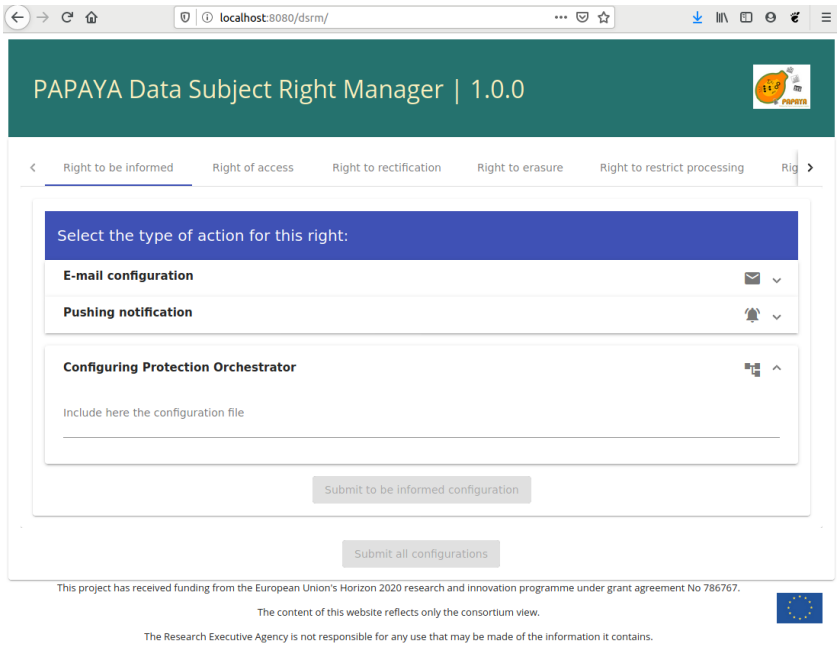

Figure 29: Recipient classification.

In addition, it is worth to mention that all the different forms and their fields have self-explanatory tool tips describing what the end user has to do in any situation. Once the Administrator configures the rights, clicking on the bottom buttons can either save the current configuration or save all of them at the same time.

4.2.2 Data Subject Interface

In order to allow to the end user or Data Subject to exercise his/her rights a mobile application has been designed and implemented. The mobile application, as it shown in Figure 30, allows the end user to list the rights that they are entitled to (Figure 30.a) and, just tapping on each one, the Data Subject can see the details associated with each right (Figure 30.b).

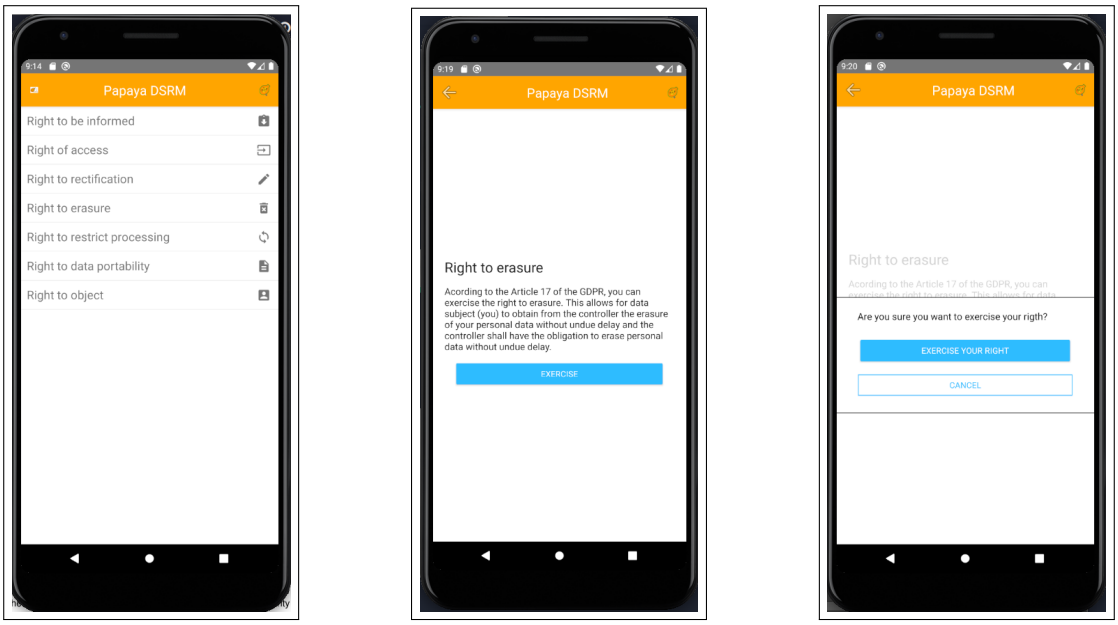

(a) Data Subject rights list (b) Right details (c) Confirmation Menus

Figure 30: DSRM mobile interface.

If the end users decide to execute any right, they can easily do it just by tapping on the bottom blue button. This action will then show a modal screen asking the end user for confirmation of the execution of the right selected (Figure 30.b). If the end user confirms it, the application will send a corresponding event to the server, which will trigger the corresponding action configured beforehand by the Data Controller Administrator.

5 Conclusions

This deliverable is presenting UIs for improving transparency and control for data subjects using PAPAYA or similar PET solutions.

Other types of stakeholder using such PETs will also profit from increased transparency. For instance for PAPAYA's use case on Arrhythmia detection (UC1), our previous user studies [7] revealed that the medical doctors receiving the data analytics results may appreciate more information about privacy protection and data utility impacts for establishing trust in PAPAYA and for being able to inform patients upon requests.

Moreover, the improved PIA tool presented in chapter 2 with the new UIs created by this enhanced CNIL tool can guide data controllers when conducting a PIA with a more comprehensive view on the assessment of privacy risks and impacts.

5.1 Putting UIs for Informing Users Together

Providing ex ante transparency to data subjects through privacy notices is especially needed when they are requested to give their consent to the processing of their data by a data analytics platform. Moreover, transparency via privacy notices presenting relevant details of the data controller's privacy policy is also essential for data subjects when taking decisions in regard to managing their privacy preference or exercising their data subject rights via the UIs of the Privacy Engine. In this subsection, we briefly outline how the UIs for explaining risks and explaining how privacy-preserving data analytics work can be put together and integrated into privacy notices.

Privacy policy information should be provided to the data subject in a "concise and transparent" (Art. 12 GDPR) manner for avoiding information fatigue. The Art. 29 Data Protection Working Party has for this reason suggested to use multi-layered privacy notices [1], which enable data subjects to easily navigate to the particular section of the privacy notice which they are interested to read.

In addition to policy information required by Art. 13 GDPR, the first (top) layer of such a layered notice should also contain information on the processing which has the most impact on the data subjects and that enables them to understand what the consequences of the processing in question will be for them and any information which could surprise them. The design and layout of the first layer should provide the data subjects with a clear overview of the information available to them on the processing of their data and where and/or how they can find that detailed policy information within the layers of the privacy notice [1].

While information on privacy risk assessment and privacy-enhancing measures implemented may have a high impact on the data subject, the top layer policy information should rather focus on informing about risks and consequences that the data subjects cannot foresee rather than on details of privacy protecting measures that the data subject may not expect. Nonetheless, the top layer should also briefly communicate how with PAPAYA (or with another similar PET) privacy risks are addressed and how any further information about the risk assessment and the PET can be accessed, for instance by displaying a clickable statement: "Information on our Privacy Impact Assessment and Privacy by Design approach". On a second layer, the following links could be provided for further information: "Why we conducted a Privacy Impact Assess-

ment?" (linking to UIs as displayed by Figure 8), "How are Privacy Risks impacted/reduced by our Privacy by Design solutions?" (linking to a page with UIs displayed by Figure 9, "How do our Privacy by Design solutions work and protect your privacy?" linking to a page displaying UIs presented in Chapter 3). The UIs that we presented in Chapter 3 are mostly also using a layered approach for informing data subjects with different layers of details corresponding to different levels of interests that different types of users may have (starting with information for lay user and leading to technical information for interested expert users on lower layers).

The approach of informing data subjects about privacy risks and privacy protection measures that are taken in different layers of details was also suggested by us in [7] based on user studies that we conducted in the first project year.

5.2 Future work

The UIs presented in this deliverable will be integrated into the data subject tools and data subject dashboard that the PAPAYA is currently developing.

As most of the UIs have not been thoroughly tested yet, further end user evaluations should still be conducted in the project's piloting phase.

References

[1] Art. 29 Data Protection Working Party. Guidelines on transparency under Regulation 2016/679, 2016.

[2] George Danezis, Josep Domingo-Ferrer, Marit Hansen, Jaap-Henk Hoepman, Daniel Le Metayer, Rodica Tirtea, and Stefan Schiffner. Privacy and data protection by design-from policy to engineering. *arXiv preprint arXiv:1501.03726*, 2015.

[3] European Commission. General Data Protection Regulation, `https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679`.

[4] Marit Hansen. Top 10 mistakes in system design from a privacy perspective and privacy protection goals. In *IFIP PrimeLife International Summer School on Privacy and Identity Management for Life*, pages 14–31. Springer, 2011.

[5] Namid Sinha. Homomorphic Encryption, `https://www.slideshare.net/iamrandomizer/homomorphic-encryption-53238006`.

[6] Jakob Nielsen. Why you only need to test with 5 users, 2000.

[7] Papaya Consortium. D2.2 Requirements specification, 2019.

[8] Papaya Consortium. D3.2 Risk Management Artefacts for Increased Transparency, 2019.

[9] Thomas Probst. Generische Schutzmaßnahmen für Datenschutz-Schutzziele. *Datenschutz und Datensicherheit-DuD*, 36(6):439–444, 2012.

[10] Martin Rost and Andreas Pfitzmann. Datenschutz-schutzziele—revisited. *Datenschutz und Datensicherheit-DuD*, 33(6):353–358, 2009.

[11] Special Consortium. D2.5 Policy Language V2, 2016.

[12] The European Commission. Guidelines on Data Protection Impact Assessment (DPIA) and determining whether processing is “likely to result in a high risk” for the purposes of Regulation 2016/679, `https://ec.europa.eu/newsroom/article29/item-detail.cfm?item_id=611236`. *Article 29*, 2017.

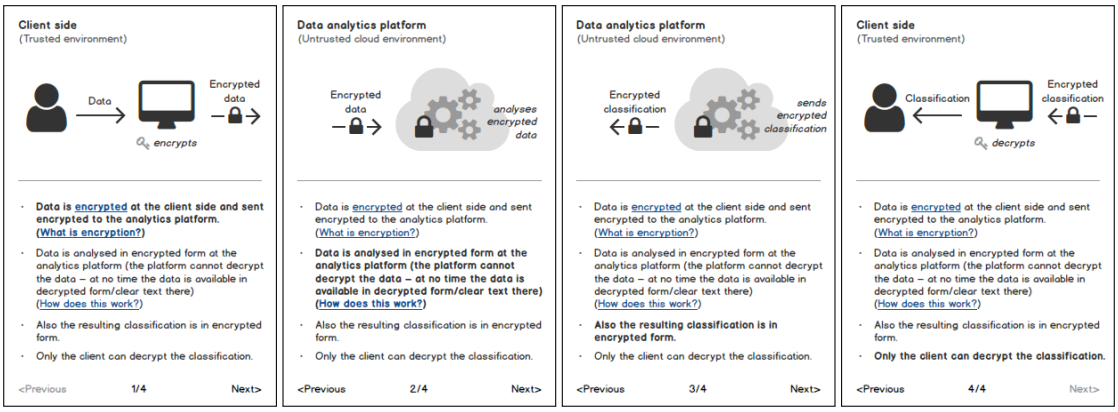

(a) Screen 1 (b) Screen 2 (c) Screen 3 (d) Screen 4

Figure 31: Prototype for UC 1/iteration 1: analytics of encrypted data: graphics change while accompanying text at the bottom remains static.

A UIs Iterations for Privacy Preserving Neural Networks

This appendix holds supplementary material related to Section 3, in particular earlier UI iterations, meant to illustrate the progression of the UI design over time.

A.1 Explaining Classification based on Homomorphic Encryption

Figs. 31, 33 and 34 represent the first, second and third iteration of UC 1 (Section 3.1: Homomorphic encryption), respectively. Fig. 32 shows the visualisation of how encryption in general, and homomorphic encryption in particular were explained.

A.2 Explaining Collaborative Training with Differential Privacy

Figs. 35–37 show the mockups that served as a template for the design (Fig. 14).

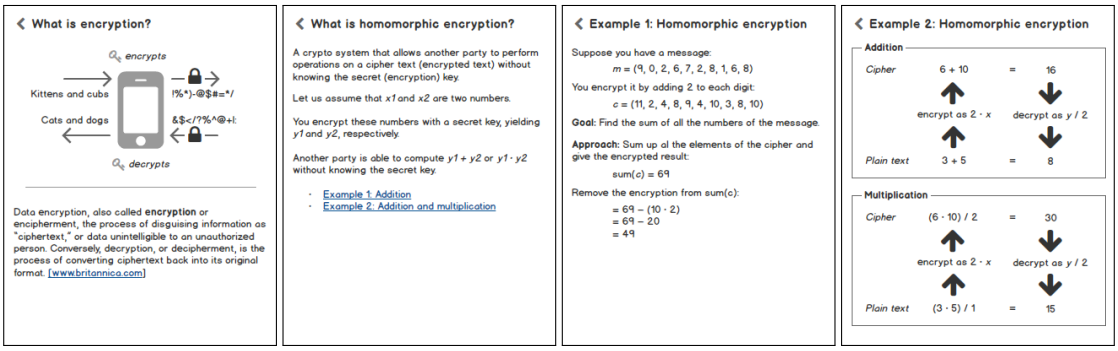

(a) Encryption (b) Homomorphic encrypt. (c) Example 1 (d) Example 2

Figure 32: Prototype for UC 1/iteration 1: secondary information on (a) encryption, and on (b–d) homomorphic encryption.

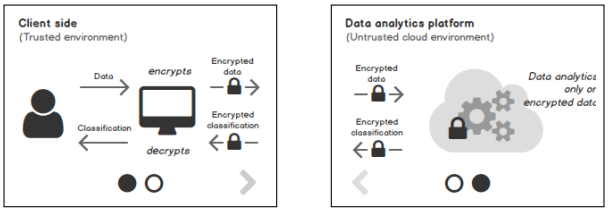

(a) Screen 1 (b) Screen 2

Figure 33: Prototype for UC 1/iteration 2: analytics of encrypted data: graphics change while accompanying text remains static.

(a) Start page

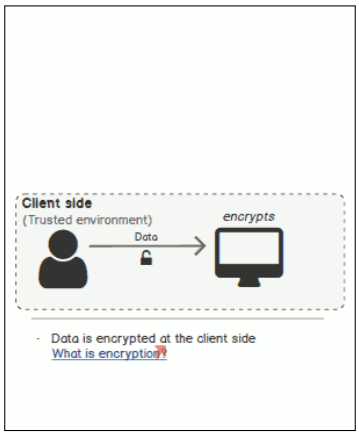

(b) Screen 1/5

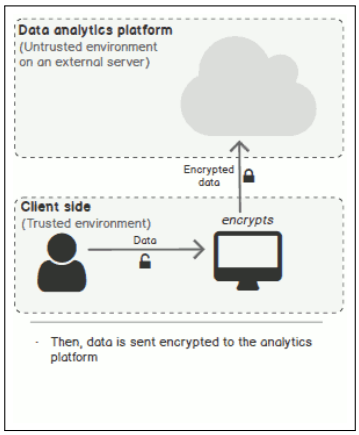

(c) Screen 2/5

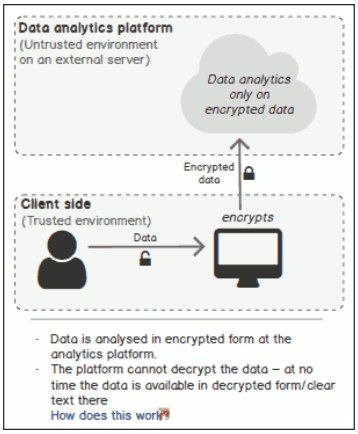

(d) Screen 3/5

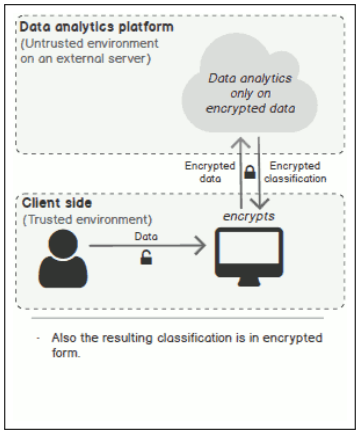

(e) Screen 4/5

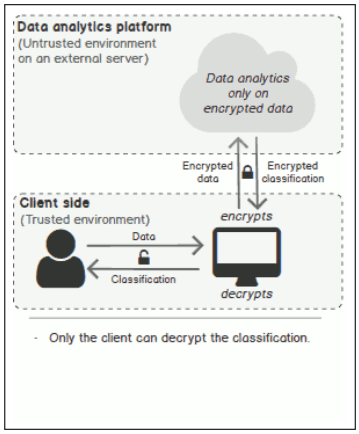

(f) Screen 5/5

Figure 34: Prototype for UC 1/iteration 3: analytics of encrypted data: graphics and accompanying text at the bottom both change.

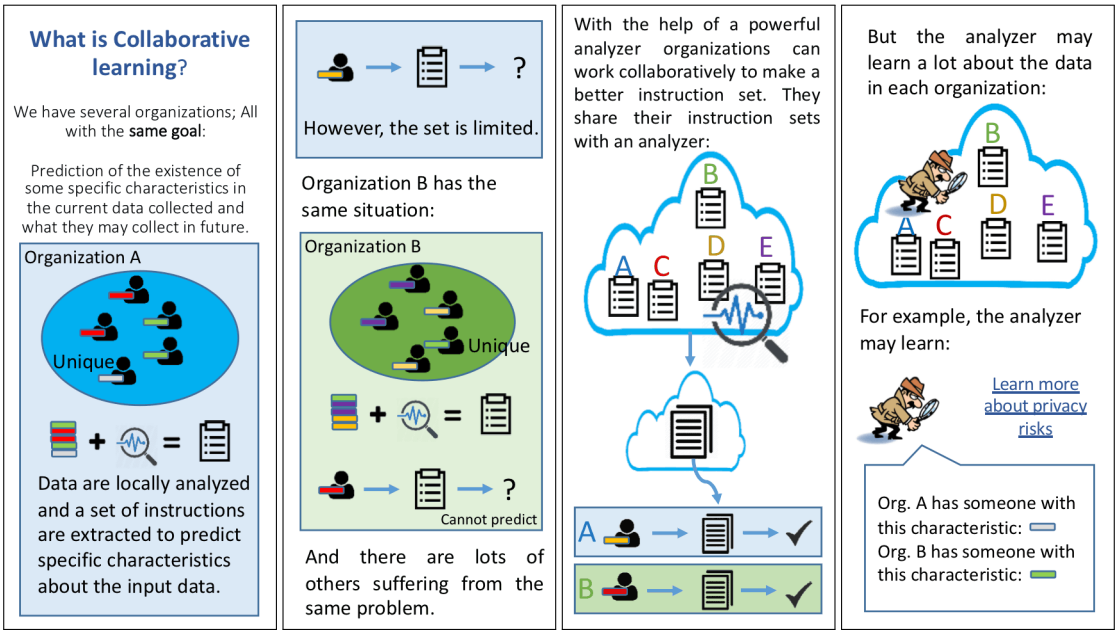

(a) Introduction 1/3 (b) Introduction 2/3 (c) Introduction 3/3 (d) Risks

Figure 35: Mockups for UC 2: explaining collaborative learning.

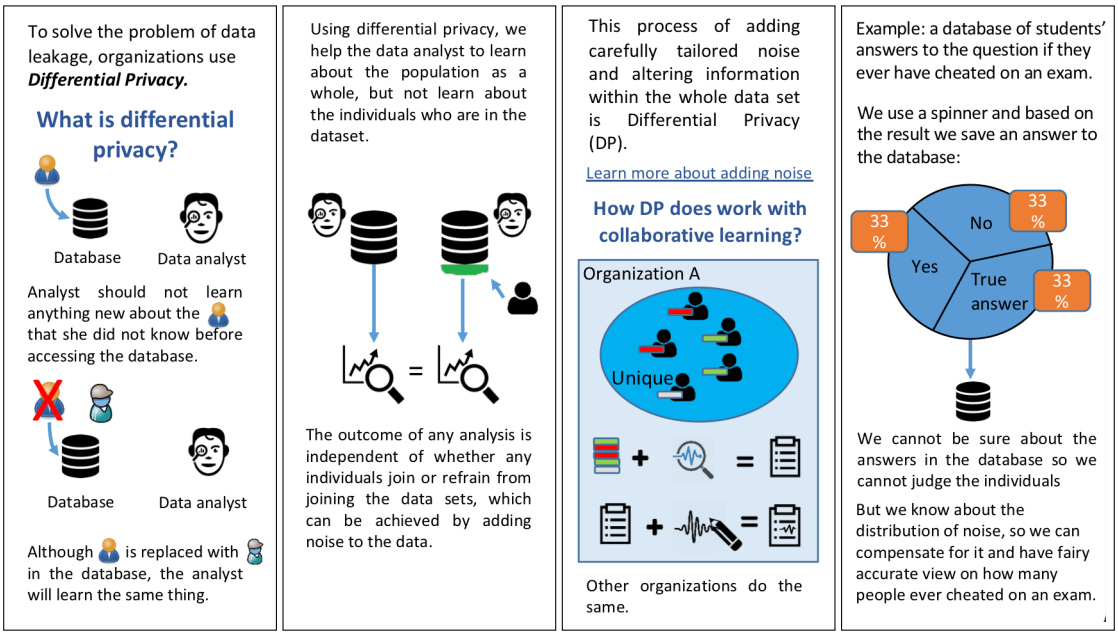

(a) Diff. privacy 1/3 (b) Diff. privacy 2/3 (c) Diff. privacy 3/3 (d) Example

Figure 36: Mockups for UC 2: Explaining differential privacy.

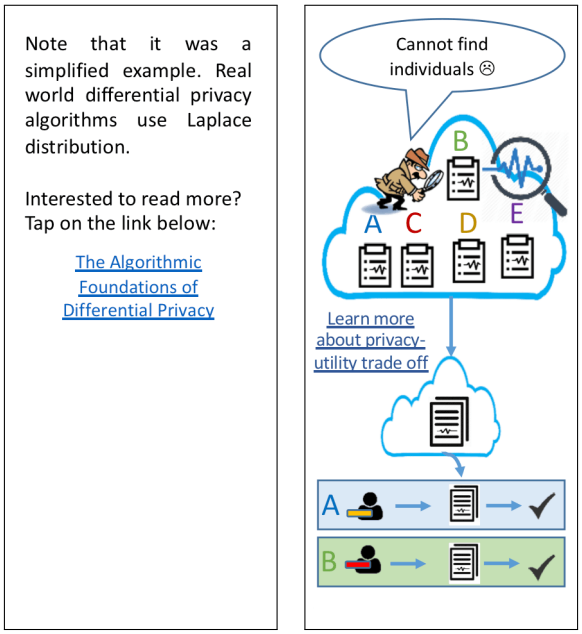

(a) Further information (b) Application

Figure 37: Mockups for UC 2: applying differential to collaborative learning.

B Questionnaire Used for UC 1 (Homomorphic Encryption)

The following questionnaire was used during the user tests to follow-up the interactive part of the user test conducted for UC 1 (Homomorphic encryption, Section 3.1). All answers were collected via free form text fields. After the questionnaire was completed, data collection continued by discussions around walkthroughs of the UI.

1. Please briefly describe what would happen according to the mobile app you just had a look at.
2. Was there anything that was unclear? (Regarding the scenario, the encryption, the user interface, the icon, pictures, terms.)
3. Are the data secure during the data analysis?
   ☐ Yes, ☐ No, ☐ I don't know
   If, why? If not, why?
4. Who can decrypt the data and the derived classification?
5. What can the data analytics platform learn about the user?
6. Would you trust analysis being done in this way?
   ☐ Yes, ☐ No, ☐ I don't know
   Please comment.
7. Would you mind your data being treated this way?
   ☐ Yes, ☐ No, ☐ I don't know
   Please comment.