Vision: A Noisy Picture or a Picker Wheel to Spin? Exploring Suitable Metaphors for Differentially Private Data Analyses

Farzaneh Karegar
Karlstad University
Sweden
farzaneh.karegar@kau.se

Simone Fischer-Hübner
Karlstad University
Sweden
simone.fischer-huebner@kau.se

ABSTRACT

Personal data analyses, for instance, in the area of eHealth, can provide many benefits while posing privacy challenges at the same time. Applying differentially private mechanisms have become one of the dominant approaches for the protection of personal data in statistical analyses. Transparency of the privacy functionality of differentially private mechanisms can facilitate informed decision-making for using differentially private systems and understanding the privacy consequences of such decisions. However, differential privacy is a complex concept that makes it a challenge to explain the privacy functionality it comprises to lay users. Our research outlined in this vision paper aims to address this challenge in three phases by creating and analysing metaphors for conveying the functionality of differential privacy to lay data subjects who should decide about sharing their data in the context of differentially private data analysis. In this paper, we report the results of the first two phases of our study for extracting the metaphors and adapting and extending them based on two rounds of analytical evaluations and feedback from privacy experts. Further, we briefly discuss how, in the third phase, we want to move forward and empirically assess the resulted metaphors from previous steps by involving lay users.

CCS CONCEPTS

• Security and privacy → Human and societal aspects of security and privacy; **• Human-centered computing → Empirical studies in HCI**.

KEYWORDS

Differential Privacy, Metaphors, Privacy protection, User understanding

ACM Reference Format:
Farzaneh Karegar and Simone Fischer-Hübner. 2021. Vision: A Noisy Picture or a Picker Wheel to Spin? Exploring Suitable Metaphors for Differentially Private Data Analyses. In *European Symposium on Usable Security 2021 (EuroUSEC '21), October 11–12, 2021, Karlsruhe, Germany.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3481357.3481525

EuroUSEC '21, October 11–12, 2021, Karlsruhe, Germany

ACM ISBN 978-1-4503-8423-0/21/10.
https://doi.org/10.1145/3481357.3481525

1 INTRODUCTION

Although the analysis of a vast amount of sensitive data can be beneficial for both users and services it involves severe privacy risks. Differential privacy has emerged as a relatively new privacy-preserving mechanism for protecting personal data in statistical data analyses. Initially, in 2006, differential privacy was defined by Cynthia Dwork [6] for calculating statistics of samples of individuals in a statistical database while protecting the privacy of the individuals in the samples. Since then, different variants and extensions of differential privacy were proposed for other types of data analysis scenarios, such as local differential privacy or differential privacy for collaborative (federated) learning. Differential privacy [6, 7] is a rigorous definition of what it means to have privacy with statistical privacy guarantees and is a property that an algorithm may have, and not an algorithm itself. Using differentially private mechanisms, data is perturbed in a controlled manner that allows quantifying privacy through a privacy loss parameter ϵ. Lower values of privacy loss parameters provide more privacy but affect the accuracy of the results more negatively. Therefore, there is a trade-off between privacy and accuracy in differentially private data analyses. Informally, for each individual who contributes her data in a differentially private data analysis differential privacy guarantees that the output of such analysis will be roughly the same, whether or not the individual contributes her data to the data sample to be analysed.

In practice, differential privacy has been deployed by giant IT companies, including Apple, Google, Microsoft, Uber, and LinkedIn. With the growing application of differentially private mechanisms, we need to address the human aspects of implementing such complex mechanisms, including the potential issues surrounding users' understanding and mental models of differentially private data analyses, their perceptions of the privacy provided by them, and their trust. Further, usable transparency about the functionality of the underlying differentially private mechanisms can help users form correct mental models of how their data is protected and make informed decisions. Therefore, in this work, our objective is to contribute to the body of knowledge on usable differential privacy and investigate the effective communication of the underlying differentially private data analyses to lay users to help them make informed decisions on whether to share their data as input to differentially private data analyses.

Design models for explaining a system to users can employ accustomed metaphors [3]. Metaphors are the means to present new ideas through the use of more familiar ones [1]. In other words, metaphors are simply a transfer or a mapping of meaning between different domains [8]. In this paper, we do not make a distinction between metaphor-based or analogical-based descriptions. We imply

analogies as well when we refer to metaphors. However, some researchers argue that in metaphor-based interfaces, the mapping of ideas involves transformation. Otherwise, it is analogy or juxtaposition when we make a direct transferal between existing knowledge and a novel situation [1]. Therefore, to reach our objective, as the first stepping stone, we would like to explore suitable metaphors of differential privacy and address the following research question. Our focus is on graphical metaphors that are elaborated with short simple accompanying information conveyed as text in its simplest form.

RQ: What are the suitable metaphors for conveying the concept of differential privacy to lay users in the context of various differentially private data analyses to help them make informed decisions about sharing their data?

To address our research question, similar to Demjaha et al. [4] who benefited from the framework proposed by Alty et al. [1] to generate and evaluate their explanatory metaphors for E2E encryption, we employed and adapted the framework to fit our objective. Our approach consists of three phases: 1) metaphor generation, 2) metaphor analytical evaluations based on expert analyses, and 3) metaphor empirical evaluations involving lay users.

In this paper, we first present background and related work in Section 2 and our methodology in Section 3. Section 4 presents the preliminary results of our first two research phases. Finally, Section 5 concludes the paper and outlines the third phase of our research, which will involve interviews to empirically evaluate the resulted metaphors from the second phase and receive lay users' feedback.

2 BACKGROUND AND RELATED WORK

2.1 Differential Privacy and Data Analysis Scenarios

As defined by Dwork et al. [7], a randomized mechanism A is ϵ-differentially private, where $0 \leq \epsilon$, iff for any two data sets D and D' that differ in at most one record, and any set R of possible outputs of A, we have $Pr[A(D) \in R] \leq e^{\epsilon} * Pr[A(D') \in R]$. The definition prevents an attacker who knows all but one record in a database from inferring the last one after viewing the output. Differentially private mechanisms can be generally divided into local and centralized (aggregate-level) models. In the centralized model of differential privacy, the aggregator has access to the actual information of users who should rely on the trustworthiness of the aggregator. The trusted aggregator gathers data from individual users and processes the data in a way that satisfies differential privacy before it publishes the results. However, when using local differential privacy, the aggregator does not see the actual data of an individual and the data is perturbed before being shared with the aggregator. Therefore, users do not need to trust the aggregator.

In our research, instead of focusing on conveying differential privacy as a notion of privacy and as a property that algorithms may have without considering the context, we define three distinct scenarios of differentially private data analyses in the context of eHealth as depicted in Figure 1 for which we will explore metaphors. One scenario is related to the local model of differential privacy (Figure 1a) and two are the scenarios related to the centralized model. For the latter, in one scenario we have one data aggregator (Figure 1b), i.e. a health company that conducts differentially private data analyses on actual information it collects and combines from its users. In another aggregate-level scenario, we have several data aggregators (different health companies), depicted in Figure 1c, that collaboratively work together to make an improved machine learning model, i.e to train a model collaboratively with the help of a cloud-based analyser while preserving the privacy of their users. The third scenario is related to differentially private federated learning.

2.2 Related work

There are a few studies in which researchers took first steps forward to achieve usable differential privacy [2, 10]. For instance, Bullek et al. used the randomized response technique (RRT) to describe a variant of differentially private mechanism using a spinner, i.e. a picker wheel [2]. They examined whether users trust the RRT mechanism which proposes to ensure their privacy and if they adjust their privacy decisions when they see more details of the privacy promises made by RRT. Xiong et al. [10] analysed the effects of using different approaches to verbally communicate differentially private techniques to laypersons in a health app data collection setting. Across different approaches of short textual descriptions, their results show that descriptions explaining implications, i.e. what happens if the aggregator's database is compromised can facilitate people's data-sharing decisions and their comprehension of the local and centralized techniques [10]. To the best of our knowledge, no attempts have yet been made to generate, test, and compare metaphors for conveying the underlying differentially private data analysis to lay users.

3 METHODS

Figure 2 shows a general view of our approach, based on the extended and adapted version of Alty et al.'s framework. [1], to address our research question for the three data analysis scenarios depicted in Figure 1. The framework provides a practical way for the application of metaphor in the design of interactive systems [1]. However, due to contextual differences, we applied the adapted version of the framework. We excluded the step related to the integration of metaphors into the user interfaces of real systems and included two rounds of analytical evaluations. After the first round of analytical evaluation, we received feedback from experts, adapted our metaphors, and analytically evaluated our adapted metaphors before we empirically evaluate them in our ongoing user studies.

3.1 Phase 1: Metaphor generation

Alty et al. proposed a few methods to generate metaphors for new interactive systems, including *extension* and *design metaphors*. The *design metaphor* technique accents the crucial role of prospective users of a system to derive metaphors and is based on careful monitoring of users' language when they discuss their requirements and understanding of a system. The *extension* technique supports the idea of recycling the metaphors that are currently used by similar interactive systems and extending them in a way that suits the metaphor for a new system. Therefore, to begin with, we reviewed literature and media outlets (see Section 4.1.1) to see how others conveyed the concept of differential privacy to users using

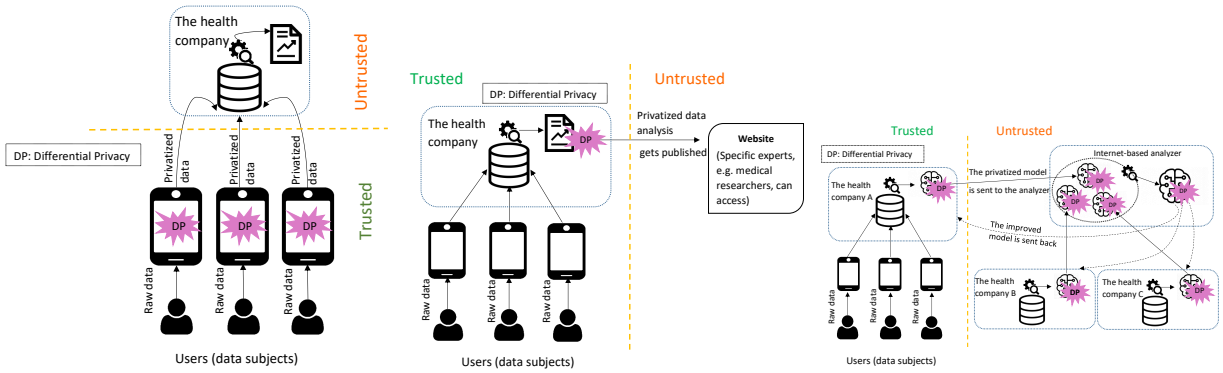

(a) Scenario 1 - local (b) Scenario 2 - centralized (c) Scenario 3 - centralized

Figure 1: Data analysis scenarios.

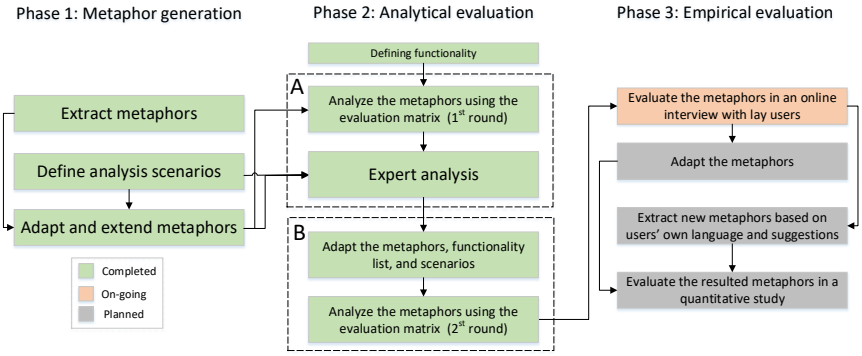

Figure 2: The steps we follow to address our research question.

metaphors or analogies. We used the results of our investigation to extend and adapt the metaphors for our context of use as reported in Section 4.1.2. In the third phase, we intend to analyse and monitor users' language when they are talking about their understanding of differentially private data analyses to derive new metaphors.

3.2 Phase 2: Analytical evaluation

3.2.1 Metaphor evaluation matrix. Two of the important steps in the framework proposed by Alty et al. [1] are *identifying system functionality* and *analysing metaphor-system pairings*. To analyse the metaphors of E2E encryption, Demjaha et al [4] benefited from a metaphor evaluation matrix. The template of the metaphor evaluation matrix we use for analysing metaphor-system pairings is shown in Table 1 and is adapted from [4]. As we have three different scenarios, we adapted the template table to evaluate whether a metaphor is suitable for each of our scenarios.

Before analysing metaphor-system pairings, we need to identify system functionality. The system in our work refers to general features of differentially private analyses. We made a balance between a functionality list that is detailed enough to cover the main purposes and characteristics of differentially private analyses and sufficiently simple for conducting analytic assessment and finding suitable metaphors. Thus, the list excluded any details about the implementation of specific mechanisms in different contexts, for example, the exact details about perturbation. In addition, the list excludes post-processing, group privacy, composability, and the feature related to the privacy budget, i.e. the threshold on the level of privacy loss when making multiple queries on the same data. Excluding these features does not imply that they are not regarded as relevant for making informed decisions. Nonetheless, as differential privacy is a complex concept, especially for lay users, first we aim to succeed at effective communication of the core features including the concept of perturbation, its relation with privacy, and its effects on accuracy. Then we can extend our work to convey other features such as privacy budget. Note that the functionality list should always be adapted based on the requirements and tasks of the target group in a system. Section 4.2.1 presents our functionality list.

3.2.2 Expert analysis. Eight privacy experts, knowledgeable about differential privacy, both from academia and industry reviewed our materials in step A of phase 2 (see Figure 2), including our description of scenarios, our original functionality list, the resulted metaphors in phase 1 and the first round of our analytical evaluation. We reached the experts through personal contacts and ongoing collaborations within joint projects. Section 4.2.2 provides the results of

Table 1: The metaphor evaluation matrix (DP: differential privacy, M: metaphor) adapted from Demjaha et al.'s work [4]

	M+	M-
DP+	Desirable: features provided by DP and supported by the metaphor	Undesirable: features provided by DP and not supported by the metaphor
DP-	Very undesirable: features implied by the metaphor and not supported by DP: conceptual baggage	Not important: features not implied by the metaphor and not supported by DP.
Is metaphor suitable for each of our scenarios:	Yes?	No?

the first analytical evaluation and expert analysis and Section 4.2.3 provides the results of our second analytical evaluation.

4 PRELIMINARY RESULTS

4.1 Results of Phase 1

4.1.1 Deriving metaphors from literature and media outlets. Warner [9] for the first time proposed randomization of responses by a spinner for improving the reliability of responses to sensitive questions. Our literature review revealed that the spinner metaphor was used by Bullek et al. [2]. We however adapted it after our first round of analytical evaluation. The spinner has been also used in media outlets to convey how differential privacy works [1].

We investigated the media outlets[2] and companies applying differential privacy to see how differential privacy is conveyed to users. We found that differential privacy is conveyed to people using an example of tossing a coin for changing responses to sensitive questions,[3] noisy sound waves of radio channels,[4] and a noisy portrait[5] from the media outlets. Exploring how companies described differential privacy to their users did not result in any further metaphors which we could analyse and use in our study. For example, Apple mentions the term differential privacy when describing how the analytics collected about users and their devices would be protected. Apple describes that this type of analysis would help improve Apple's products and reduce problems like apps crashing. Apple also conveys that the collected information cannot be associated with the user and the user's account.[6] Note that Apple also published white papers (e.g. [5]) presenting more details on how they are leveraging local differential privacy.

4.1.2 Metaphors generated in phase 1. The metaphors we extracted from media outlets and literature, however, are not necessarily suitable for conveying differentially private data analyses in all scenarios. Randomized response techniques, the coin flip and spinner examples, are only suitable for local differential privacy. In addition, a noisy picture may not be suitable for centralized differential privacy because it does not convey that perturbation happens on the aggregate level and not on single records from individuals. We adapted and extended the metaphor of a noisy picture by adding a different amount of noise to a picture which is a *combination* of several portraits, which better reflects that noise is added to the calculated aggregate (result) and not to the data itself.

Spinning a picker wheel and flipping a coin to perturb responses are similar examples. To convey the trade-off between privacy and accuracy, we wanted to show at least two different cases in which the probability of landing on yes/heads and no/tails would differ. However, we assumed that it would be harder for users to think of a deformed coin that would make it more probable, for example, to have tails rather than heads compared to spinners with different probabilities. Consequently, we excluded the coin metaphor. The picker wheel, both variants of noisy pictures, and noisy broadcasts of a radio channel served as the input to our first analytical evaluation.

4.2 Results of Phase 2

4.2.1 Functionality list. Here we provide the functionality list that we used for our second round of analytical evaluation (step B in Figure 2) that is the adapted list after receiving feedback from experts.

(F1) A differentially private analysis[7] bounds and quantifies the probability of additional privacy risk that any individual would face because of her/his participation in a data analysis.
(F2) The privacy of a differentially private analysis is controlled by tuning a privacy loss parameter.
(F3) The smaller the value of the privacy loss parameter, the better the privacy guarantee for an individual.
(F4) The smaller the value of the privacy loss parameter, the less accurate the results of data analysis are.
(F5) A differentially private analysis randomly perturbs data on an aggregate level (i.e. the results of the analysis) or individual level (i.e. the input data) dependent on the context.
(F6) The amount of perturbation is controlled by the underlying differentially private analysis.[8].
(F7) A differentially private analysis is resistant to privacy attacks based on auxiliary information, i.e. any past, present, and future information that an attacker may have.
(F8) A differentially private analysis does not promise unconditional freedom from privacy risks[9].

The first feature (F1) in the list can be interpreted in different ways. For example, F1 should convey for the centralised model that the results of a differentially private data analysis do not significantly depend on any particular individual's data so an individual will not be affected, adversely or otherwise, by allowing her data to

[1] An example of using spinner by Mark Hansen:https://accuracyandprivacy.substack.com/
[2] We googled *differential privacy* keyword and also its combination with *users, people, definition,* and *introduction* keywords and investigated each of the first five pages of the results to find videos or pages which conveyed the concept at a high level rather than with technical details.
[3] Simply Explained: https://www.youtube.com/watch?v=gI0wk1CXlsQ
[4] National Institute of Standards and Technology: https://www.youtube.com/watch?v=-JRURYTfBXQ
[5] Nikolas Sartor at Aircloak blog: https://aircloak.com/explaining-differential-privacy/
[6] For example, see the description under the title Analytics: https://www.apple.com/privacy/control/
[7] A differentially private analysis is often called a mechanism
[8] To have a differentially private data analysis we should know what to perturb and to what extent.
[9] Note that any useful data analysis carries the risks that it will reveal information about individuals

be used in the analysis. F1 can also be rephrased in terms of plausible deniability: a differentially private analysis gives individuals who contribute their data to the analysis plausible deniability. For example, in local differential privacy, an individual can deny that a particular data record is her true data and in a centralized differentially private analysis individuals can even deny that they have participated in the analysis. When we do the analytical evaluation, we consider different interpretations of F1. Although a metaphor may not directly convey F1, it may imply different interpretations of it.

F6 can be elaborated further by including the parameters which affect the amount of perturbation in an analysis. For example, the amount of perturbation, for most of the differentially private mechanisms, is dependent on the privacy loss parameter and the sensitivity of the analysis function. However, we avoided including further details to keep the feature simple for to convey and evaluate. To communicate the underlying differentially private data analysis to lay users we did not focus on the privacy loss parameter but on the role of perturbation in providing privacy and the effects of perturbation on the accuracy of the results. Therefore, if a metaphor conveys that more perturbation leads to better privacy but less accuracy we assume it covers F3 and F4.

A single graphical metaphor with textual elements may not convey all the features without accompanying clarifying information, for example, in text format. Thus, it is important to find the gaps and cover the functionality not conveyed by the metaphor in an accompanying text which presents the metaphor to users. However, we should concurrently take users' needs of what should be conveyed and how to convey it, their understanding, and their satisfaction with the metaphors into account. It highlights the importance of empirical evaluations of metaphors by involving users and exposing them to the metaphors to gauge their needs and understandings that we planned to do in the third phase.

4.2.2 The results of the 1st round of analytical evaluation and expert analysis. As a result of the expert analysis, we excluded the metaphor of noisy sound waves of a radio channel because it has a highly undesirable feature (conceptual baggage). The metaphor implies that the original data collected by the aggregator can be heard by anyone who listens to the radio channel at the right frequencies. Nonetheless, those who should receive the sound waves without noise (those who should have access to the original data) are either the data subjects or trusted data aggregators. Further, the metaphor conveys that anyone who receives the data, i.e. listen to the radio channel can decide on the amount of noise. However, an adversary, as an example, does not decide on how much noise should be added to the results of an analysis. In addition, our experts advised that if we have public information, e.g. radio broadcast through an FM radio channel, it does not make sense to apply differentially private mechanisms. Despite its problems, this metaphor, as also confirmed

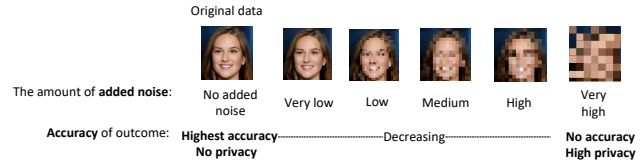

(a) Noisy picture metaphor - local model

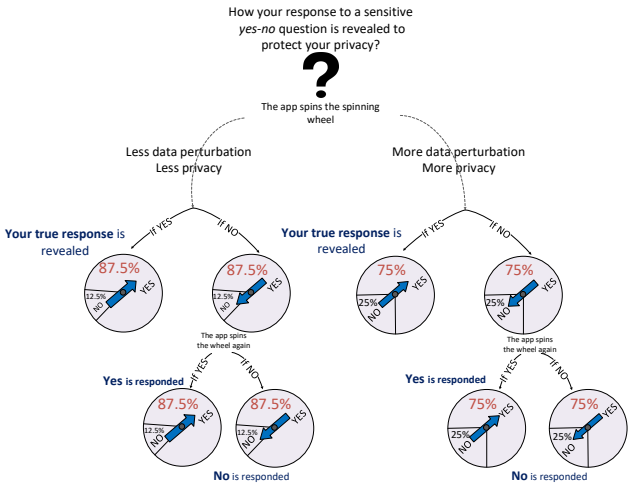

(b) The picker wheel (spinner) metaphor

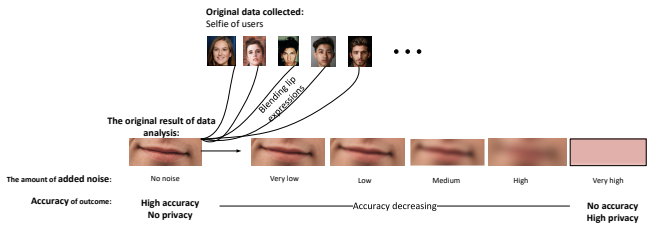

(c) Noisy (combined) picture metaphor - centralized model

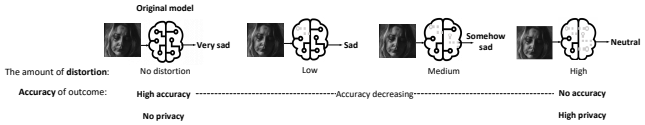

(d) A cut of the noisy brain model metaphor - centralized model

Figure 3: The adapted metaphors we aim to empirically test in our planned user studies.

by our experts, is suitable if we want to highlight the importance of tuning the privacy loss parameter, for example, by an aggregator rather than if want to convey that there is a trade-off in every differentially private system.

We adapted and extended our preliminary spinner metaphor to better communicate F3, F4, and F6. Figure 3b shows the resulted metaphor. We purposefully made it prominent that the probability of landing on Yes and No is different in the two cases shown which affects the accuracy of the results and the individuals' privacy in different ways. Based on our experts' feedback we confirm that the spinner metaphor is only suitable for the local differential privacy model. The spinner metaphor may also suffer from a rather undesirable feature if not properly introduced to users. Bullek et al. [2] reported that some participants preferred the most truthful spinner because they associated the perturbation they made on their answers to lying to the entity asking questions. To avoid it, we convey the message that users do not perturb their data but, for example, a device on their phone does the perturbation before sharing the data with a remote server. We assume that it would at least help absolve users from the feeling of deceiving the party that requests their responses.

Initially, for the metaphors based on noisy figures, we used the portraits of famous people. However, based on our experts' comments, we avoided using photos of famous and well-known people, although it does not solve the undesirable features of these metaphors. Unless distorted with a high amount of noise, a noisy picture may still reveal the identity of the person in the picture and several of his or her facial characteristics so it may not be possible for that person to deny that a noisy figure is her true figure or to deny that her data was used to conduct the analysis, e.g. to derive the combined picture. Therefore, metaphors based on the noisy pictures do not necessarily cover F1 and F7. Based on our experts' feedback we confirm that a noisy picture when we have a single picture, as depicted in Figure 3a, is only suitable for the local models and a noisy picture which is a combination of several other pictures, as depicted in Figure 3c, is more suitable for centralized models. In our third scenario, instead of using a noisy picture, we used a distorted brain as a metaphor of a differentially private trained model as depicted in Figure 3d. Note that Figure 3d only shows a cut of the original metaphor for conveying the concept in which users' pictures are shown to be used for training a brain model to recognize emotions based on facial expressions. The brain model is then distorted.

Our expert analysis revealed that the usefulness of the results of the analysis might be too abstract for users to comprehend. For example, for the spinner metaphor, users should understand that the entity who asks questions still can benefit from the perturbed responses to fulfil its intention of data analysis. Thus, the accompanying information can help to convey the usefulness of perturbed results of analysis and can serve as a means to make a connection between the type of data in the metaphor, e.g. the picture or the yes-no questions and the type of data in the real scenarios.

4.2.3 The results of the 2nd round of analytical evaluation. Table 2 shows whether each of our adapted metaphors (shown in Figure 3) conveys or implies the features in the functionality list, although it is subjected to be validated by users studies, and shows for which scenario it can be used. What we assume could be understood from a metaphor is different from what lay users grasp. The Y (Yes) in Table 2 means that the metaphor has the potential to convey or implies the feature without clarification by further information, for example, in the form of an accompanying text. Features F3 to F6 are conveyed by all four metaphors. F1 and F8 are implied by the spinner metaphor. F8 is implied by the other metaphors as well. Until completely distorted, e.g. when the area for YES is zero in the spinner or the amount of noise is very high for the figures, we can still have a useful analysis that may carry a risk of revealing information about individuals. The noisy picture metaphor for the local model (see Figure 3a) does not cover F1 and F7. The noisy combined picture metaphor may convey F1 and F7. However, whether it really covers F1 and F7 is pretty much dependent on the combination of all pictures selected for that metaphor. In addition, users' understanding and perception of, for example, how much the aggregate-level picture might be revealing and if and how the added noise can circumvent privacy leakage from a combined/aggregated picture play a significant role. The distorted brain metaphor shown in Figure 3d is quite abstract and whether it conveys or implies F1 and F7 is very much dependent on what users know or understand from the concept of a model.

Table 2: Features of functionality list covered (and not covered) by each metaphor.

Metaphor\feature	F1	F3	F4	F5	F6	F7	F8	Context
Picker wheel	Y	Y	Y	Y	Y	Y	Y	Scenario 1
Noisy single picture	N	Y	Y	Y	Y	N	Y	Scenario 1
Noisy picture - combined	Y	Y	Y	Y	Y	Y	Y	Scenario 2
Distorted brain model	Y	Y	Y	Y	Y	Y	Y	Scenario 3

5 CONCLUSION AND MOVING FORWARD

We presented a three-phase approach towards effective communication of the underlying differentially private analysis to users. The completed first two phases resulted in adapted metaphors, their analytical evaluations, and a functionality list that can be used to analytically evaluate the suitability of other metaphors for conveying privacy functionality of differentially private data analysis to users. In the third phase, first, we investigate to what extent the metaphors convey the features in the functionality list in lay users' opinions conducting online interviews which are currently ongoing. We also explore the factors contributing to their trust in a differentially private system and their willingness to share their data with such systems. We recruit our interviewees through the Prolific platform. We aim for interviewing people who have not heard about differential privacy and are not knowledgeable about any other privacy protection mechanisms although they might have used some privacy protection tools or heard about other privacy techniques. The preliminary insights from three pilots show that our metaphors can successfully convey the trade-off between privacy and accuracy. However, achieving more privacy at the expense of losing accuracy may affect users' trust in the results of differentially private data analysis and lead to the preference of having less amount of perturbation applied. The results of our interviews may lead to further tailored metaphors and the extraction of new

metaphors based on users' own language. Finally, at the last step, we plan to evaluate our metaphors in a quantitative study.

ACKNOWLEDGMENTS

This work was funded by the H2020 Framework of the European Commission under Grant Agreement No. 786767 (PAPAYA project) and by the Swedish Knowledge Foundation (TRUEdig project).

REFERENCES

[1] James L. Alty, Roger P. Knott, Ben Anderson, and Michael Smyth. 2000. A framework for engineering metaphor at the user interface. *Interacting with computers* 13, 2 (2000), 301–322.

[2] Brooke Bullek, Stephanie Garboski, Darakhshan J. Mir, and Evan M. Peck. 2017. *Towards Understanding Differential Privacy: When Do People Trust Randomized Response Technique?* ACM, 3833–3837.

[3] Louise Clark and M. Angela Sasse. 1997. Conceptual Design Reconsidered: The Case of the Internet Session Directory Tool. In *People and Computers XII*, Harold Thimbleby, Brid O'Conaill, and Peter J. Thomas (Eds.). Springer, 67–84.

[4] Albese Demjaha, Jonathan M Spring, Ingolf Becker, Simon Parkin, and M Angela Sasse. 2018. Metaphors considered harmful? An exploratory study of the effectiveness of functional metaphors for end-to-end encryption. In *Proc. USEC*, Vol. 2018.

[5] Apple Differential Privacy Team. 2017. Learning with Privacy at Scale. https://docs-assets.developer.apple.com/ml-research/papers/learning-with-privacy-at-scale.pdf [Online; accessed 1-June-2021].

[6] Cynthia Dwork. 2006. Differential Privacy. In *Automata, Languages and Programming*, Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener (Eds.). Springer, 1–12.

[7] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*. Springer, 265–284.

[8] Arnold H Modell. 2009. Metaphor—The bridge between feelings and knowledge. *Psychoanalytic Inquiry* 29, 1 (2009), 6–11.

[9] Stanley L Warner. 1965. Randomized response: A survey technique for eliminating evasive answer bias. *J. Amer. Statist. Assoc.* 60, 309 (1965), 63–69.

[10] Aiping Xiong, Tianhao Wang, Ninghui Li, and Somesh Jha. 2020. Towards Effective Differential Privacy Communication for Users' Data Sharing Decision and Comprehension. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 392–410.