Towards Private Medical Data Donations by Using Privacy Preserving Technologies

Arno Appenzeller
Karlsruhe Institute of Technology
Karlsruhe, Germany
arno.appenzeller@kit.edu

Erik Krempel
Fraunhofer IOSB, Fraunhofer Institute of Optronics,
System Technologies and Image Exploitation
Karlsruhe, Germany
erik.krempel@iosb.fraunhofer.de

Nick Terzer
Karlsruhe Institute of Technology
Karlsruhe, Germany
unhfb@student.kit.edu

Jürgen Beyerer
Fraunhofer IOSB, Fraunhofer Institute of Optronics,
System Technologies and Image Exploitation
Karlsruhe, Germany
juergen.beyerer@iosb.fraunhofer.de

ABSTRACT

Through the growing amount of personal health data collected by the individual itself digital data donations become more and more attractive. Wearables like Apple Watch or Fitbit trackers make tracking of heart rate, daily step counts and other lifestyle data easier than ever. While this data is collected on the dedicated device, it can help research in many promising ways. Even if the potential benefit of this data is very clear, there are open questions regarding privacy. Traditional privatization measures like anonymization and pseudonymization can only provide limited privacy guarantees especially with the growing amount of personalized data. To mitigate those risks privacy enhancing technologies like differential privacy can be used. While the theoretical foundation of such technologies is strong, only limited data is available about their practical use in large scale applications and the trade-off between privacy and utility. In this paper we will present a data donation scenario that is inspired by a real-world use case using lifestyle data for its analyses. We will apply the local differential privacy technology "RAPPOR" to improve the privacy protection for the data donors and evaluate the impact of this technique to the data utility.

CCS CONCEPTS

• **Security and privacy → Domain-specific security and privacy architectures**; **Privacy protections**; *Social aspects of security and privacy*; • **Social and professional topics → Patient privacy**; *Health information exchanges*; Medical records; Personal health records.

KEYWORDS

Differential Privacy, Data donation, Privacy Enhancing Technology, Medical data protection

PETRA '22, June 29-July 1, 2022, Corfu, Greece
© 2022 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9631-8/22/06.
https://doi.org/10.1145/3529190.3534768

ACM Reference Format:
Arno Appenzeller, Erik Krempel, Nick Terzer, and Jürgen Beyerer. 2022. Towards Private Medical Data Donations by Using Privacy Preserving Technologies. In *The15th International Conference on PErvasive Technologies Related to Assistive Environments (PETRA '22), June 29-July 1, 2022, Corfu, Greece.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3529190.3534768

1 INTRODUCTION

It can be seen as yesterday's news that every year people generate more data about themselves and their environment.

One of the latest additions into this collection of data are body worn fitness trackers. Their capabilities strongly vary between models. Simple versions can count steps, measure the wearers blood pressure, and track sleep. Newer developments monitor blood glucose levels and blood oxygen levels. Starting with Series 4 the Apple Watch is capable to make an ECG and detect heart rhythm anomalies which are a sign of atrial fibrillation[1]. While this might look like a toy example, the sensor and software behind it have a high enough quality to reach FDA certification[2].

Not only Apple, Facebook and Google are interested in these kinds of data. Researchers see a high potential in monitoring the public health and advancing medical research with it or using general activity data to improve care of individuals. A recent example where data from activity trackers is donated to research is the ”Corona-Datenspende-App” (Covid data donation app) by the German federal agency for disease control and prevention Robert Koch Institute (RKI). The researchers use the data of more than 500.000 participants to predict the course of the COVID-19 pandemic[3].

Working with medical data of volunteers sets high demands on data protection for the researchers. Legal regulations such as the GDPR and more specific regulations for medical data must be fulfilled. Technical and organizational measures need to be in place to ensure a level of security appropriate to the risk of the data processing. Additionally, the volunteers have a high demand for

[1] https://www.cnet.com/health/apple-watch-ecg-app-what-cardiologists-want-you-to-know/ [Accessed: 25th April 2022]

[2] https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN180044.pdf [Accessed: 25th April 2022]

[3] https://corona-datenspende.de/science/en/reports/nowcast/ [Accessed: 25th April 2022]

data protection[4]. Errors in handling the data can have a vast impact on the participants' privacy. This was already shown by Latanya Sweeney when she re-identified allegedly anonymized medical data containing information about Massachusetts governor, William Weld [16].

Looking at terms like "anonymity" and "de-identification" it gets clear that these terms are hard to define and even harder to achieve. In our scenario of research on medical data from volunteers we are working with personal identifiable information (PII). The question remains, whether it is possible to remove all elements that identify a person behind the data and still work with it. On this topic Cynthia Dwork is quoted with "De-identified data isn't". You either alter the data so much that it becomes useless, or you alter it so little that it will be possible to re-identify persons in your data.

To help to handle medical data and its privacy implication in research the German "TMF – Technologie- und Methodenplattform für die vernetzte medizinische Forschung e.V." a registered association for connected medical research [5] separates medical data into two categories:

- IDAT (identity data): Data that is primarily used for identification of a person (e.g., Name, birthday, address)
- MDAT (medical data): Data that is primarily used for the treatment of a person (e.g., blood pressure, medication plan)

IDAT and MDAT are not disjoint sets, the birthday of a patient might be used in treatment and in identification. Nonetheless separation is useful as it allows for a better understanding. Many operations can be used to alter IDAT to protect privacy. For example, replacing all IDAT with a pseudonym and limiting access to the table where IDAT and pseudonym is stored. Or altering IDAT in a way that makes it harder for an attacker to re-identify a person, i.e., removing the last 3 digits of a ZIP code.

Altering MDAT is more challenging. It is highly use case-specific which operations can be used to protect data subjects without rendering the data useless. This needs strong cooperation between medical domain expert and privacy experts.

In this paper we will investigate the usage of differential privacy (DP) as a method to preserve the privacy of an individual in a scenario of a medical data donation. Secure, private, and easily accessible data donations could have a large impact in the domain of Active Assisted Living (AAL) to collect long term data about affected individuals. This data could then be used to improve care and health conditions of the individuals through Big Data research. While the benefit of a data donation seems clear, open questions in terms of privacy remain. Because privacy is also a huge impact factor for trust and acceptance of data sharing, the here presented approach offers solutions to improve this factor.

There is a multitude of related work (e.g., [9, 11, 12, 19]) that also investigate the topic of privacy preserving technologies for medical data usage but none of them use a local differential privacy (LDP) approach in the context of a data donation like our work does by using the RAPPOR technology. This has a lot of advantages as our approach does not require a trusted third party and is not limited to numerical values that could be a limitation of other LDP technologies.

Additionally, we want to clarify, that we do not claim that DP in this case will allow us to completely anonymize the data. Just as one example our setup could be vulnerable to side-channel attacks that allow an attacker to learn something about a participant and re-identify him or her. We use DP as a privacy enhancing technology (PET) that makes it harder for an attacker to re-identify data and thus improve privacy. Most likely additional technical and organizational measures such as encryption or access control need to be in place to protect privacy and built a solid system.

Section 2 of this paper will start by introducing the concept of DP, its extension LDP and the implementation framework RAPPOR. In section 3 we introduce our scenario in detail and present our proof-of-concept implementation before evaluation in section 4. As usual we end our paper with some additional thought on future work and a conclusion in section 5.

2 PRIVACY MECHANISMS

This chapter will give a crash course for the used privacy mechanisms without claiming to offer enough detail for a full understanding of it.

2.1 Differential Privacy

Differential privacy (DP) goes back to Cynthia Dwork and the year 2006 [4]. It is a mechanism to publish information about a dataset and its patterns without publishing information about the individuals in it. The first formal definition by Dwork is as follows:

$$\Pr[\mathcal{K}(D_1) \in S] \leq \exp(\epsilon) \times \Pr[\mathcal{K}(D_2) \in S]. \tag{1}$$

This definition means, that DP gives us guarantees when comparing the results of two datasets D_1 and D_2 that differ in at least a single entry, i.e., they differ in the data of a single person. To evaluate the datasets, we use the *randomized* evaluation mechanism $\mathcal{K}$ with $S \subseteq \mathit{Range}(\mathcal{K})$. This mechanism $\mathcal{K}$ is called ϵ-DP if the results of $\mathcal{K}$ in the two datasets differ at most in $\exp(\epsilon)$ with ϵ a real number. ϵ is also called privacy budget.

The most crucial part of every DP system is the choice of the randomized mechanism $\mathcal{K}$. The Laplace mechanism is used when looking at means, min or max values in a dataset and works by adding random noise on the data . An exponential mechanism can be used to have private auctions and special mechanisms can be used in the training of neural networks [1, 5].

2.2 Randomized response

In our work we use *randomized response* (RR) as random mechanism $\mathcal{K}$ in DP. RR itself is even older than DP. It was first introduced by Stanley Warner in 1965 to allow participants in a study to have plausible deniability when answering sensitive questions (e.g., drug use, sexuality) [20]. Figure 1 shows the process of a survey using RR. Before answering a question, the participants throw a coin that nobody else can observe. If the coins lands heads they should answer truthfully. If it comes up tails, they should throw it again. If it comes up tails again, they answer with "yes", if it comes up heads they should answer with "no". This mechanism adds a random noise to the result. Looking at the result of a single participant it remains

[4] See https://www.tmf-ev.de/DesktopModules/Bring2mind/DMX/Download.aspx?Method=attachment&Command=Core_Download&EntryId=32187&PortalId=0 [Accessed: 25th April 2022]

[5] https://www.tmf-ev.de/EnglishSite/Home.aspx [Accessed: 25th April 2022]

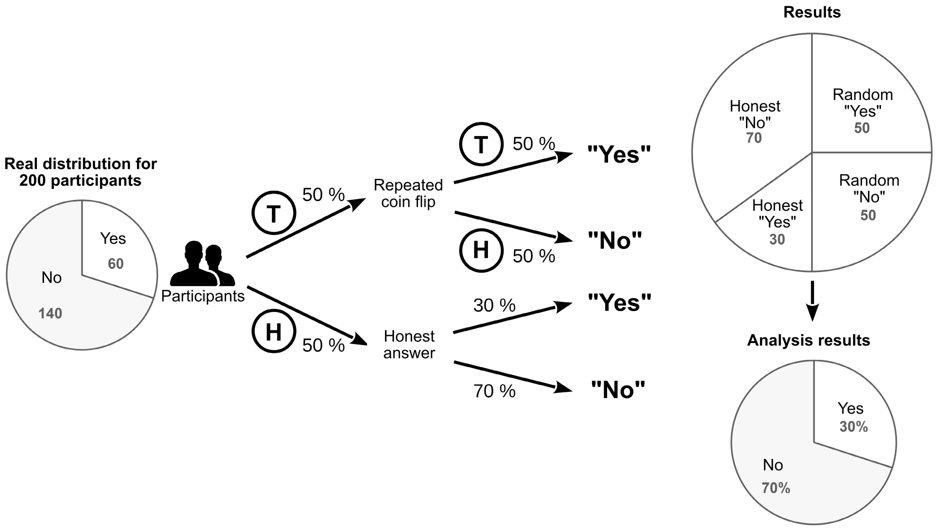

Figure 1: Visualization of the Randomized Response process using a double coin flip for a survey.

unclear if the answer was truthfully or the result of the coin flip. But looking at the results of a large population of participants, it is possible to calculate the real percentage. RR with a fair coin's flip fulfills the ϵ-DP requirements for $\epsilon \approx 1,09$.

2.3 Local Differential Privacy

RR has an additional characteristic. Many DP concepts require that we first transmit data to a central repository and later evaluate them with a random mechanism $\mathcal{K}$ to achieve DP. While the results of the evaluation are protected by DP the central repository holds all sensitive information and represents an interesting target for attackers. With RR we do not have this problem. The random mechanism $\mathcal{K}$ is executed by the participant, even before the answer is transmitted to the interviewer. This characteristic, called *local differential privacy* (LDP) and can strongly add into protecting participants' privacy.

2.4 RAPPOR

With LDP and RR we have all the building block we need to understand the inner workings of RAPPOR – Randomized Aggregatable Privacy-Preserving Ordinal Response. RAPPOR was developed by Erlingsson et. al at Google research [6]. It is be one of the first DP mechanisms that is used in a commercial setting, namely, to submit security critical parts of the Google Chrome settings for monitoring.

RAPPOR was designed to work with unpredictable string values as input. To achieve this, so-called Bloom filters are used to map strings into a bit array. Figure 2 gives an insight into the necessary steps. Input values are first hashed and then placed into a bit array B. Now we do RR for every individual bit in the bit array. This will generate B' for us. RAPPOR offers different parameters to influence its behavior. Here we explain the three most important ones:

- k : The size of the bit arrays B and B'. A higher value will prevent hash collisions and overall improve the precision of the design. At the same time a higher value will decrease performance.
- h : Bloom filters can be used with one or multiple hash algorithms. Using multiple algorithms will offer some protection against hash collisions while decreasing performance.
- f : The probability that we alter the value of B during the RR. To simulate a fair coin toss, this value should be set to 0.5. This parameter can be altered to increase precision by reducing privacy and vice versa. The value of f is connected to ϵ.

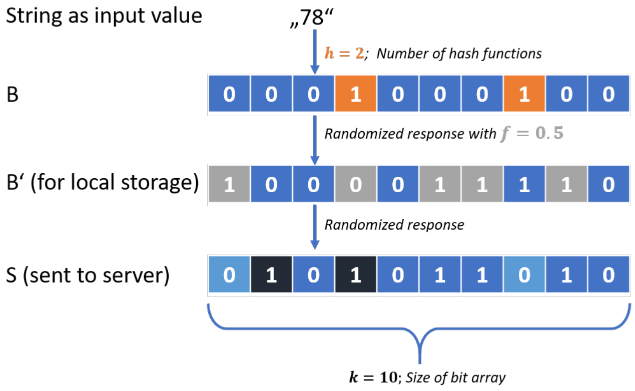

Figure 2: The RAPPOR system adding the string "78" to the Bloom filter B before altering it using RR

If we want to use RAPPOR to send a set of values a single time we can stop here. If we want to send the same values over a longer period, we must do one additional step to protect the users against information leak and fingerprinting. To do so we store the value of B' after first creating it. Every time we want to send it, we first do another round of RR on B' to generate a new S. This S is send to the server and then deleted. This process is also called Instantaneous Randomized Response (IRR).

3 PRIVACY PRESERVING DATA DONATION

In this section our proposal for a privacy preserving data donation is described. At first, we introduce the usage scenario and will show the potential privacy risk for an individual by presenting an attacker for this scenario. After this we will explain the prototypical implementation of our workflow.

3.1 Scenario

We use a population wide fever monitor as use case for our implementation. This scenario is inspired by the so called "Corona-Datenspende-App"[6] (Covid data donation app) by the German federal agency for disease control and prevention Robert Koch Institute (RKI). This app collects step counts and heart rate measurements from fitness trackers like Apple Watch and asks the users to donate this data. According to the responsible researchers those measurements are used to model a correlation between reduced activity and fever. This could help to predict the course of the COVID-19 pandemic. The project itself is a huge success. Over 500.000 individuals participated and donated their health data. Recent findings[7] of the project showed that it enabled a relatively precise prediction of the virus spread in Germany. Unfortunately, the raw data and processing algorithms of the "Corona-Datenspende-App" are not available publicly. Nevertheless, the use case still shows the importance of the data donation scenario even when the actual data and algorithms are replaced with our own technologies that try to mimic this behavior.
For our scenario we want to use this data analysis setting and see how good it works and how accurate the data is when Privacy Enhancing Technologies (PETs) like RAPPOR are applied to them, so that an individual cannot be linked to its donated data. It remains to be noted that recent work about data privatization often used synthetic data generation (See for example [10, 13, 14]). However, we choose RAPPOR for this mechanism because synthetic data generation requires a decent amount of data to produce reasonable results. This would not be possible in a scenario of a data donation with LDP where the real data should only be on the affected persons device. Synthetic data would require a large database of raw data which would impose an additional privacy risk. Therefore, we choose the more traditional technique of LDP with RAPPOR. Since the original dataset is not publicly available, we use a crowd-sourced Fitbit datasets which has similar properties [7]. Table 1 shows some sample entries from the dataset. It consists out of heart rate measurements (5 seconds interval), step counts (60 second interval) and duration of sleep from 30 participants over a a time span of a month making it approximately 15 million data points.

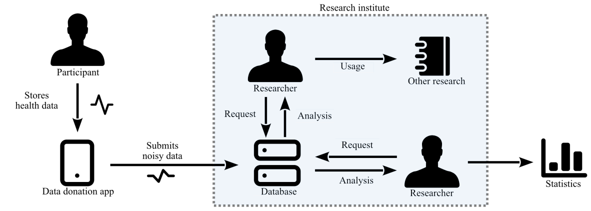

Figure 3: Visualization of the data donation scenario

With those 30 participants we simulate the scenario that every participant sends his data to a central point which can be a research institute that analyzes the data. Figure 3 shows a visualization of the use case and components that are used for the implementation. The data is sent by the user with a smartphone app or some other interface which is not further discussed in this paper. We use the collected data to perform some typical data evaluation steps. One of those analyses is a resting heart rate curve which can be calculated by using the heart rate and step count of an individual. Instead of sending the raw data we use LDP with the RAPPOR Algorithm to increase the privacy protection for the data donors. The RAPPOR parameters are set by the research institute which also hosts the database where the noisy data is stored. This database can then be used by researchers for the previously described or other analyses.

3.2 Potential Attackers

Our system should protect the privacy of all participants in the present of an attacker. A recent study also underlines that such privacy and security issues are relevant for data donation scenarios [8]. A potential attacker wants to re-identify participants by a combination of the transmitted data and publicly available information like voter registers or location data. Additionally, the attacker could observe individuals to collect information about their activities. Many studies show the potential re-identification risk of linking such datasets [2, 16]. The potential privacy impact of a successful attack varies with the goals of the attacker. The used heart rate data can be used to gain knowledge about the general health of the data donor. Additionally, pattern from activity timestamps and sleep data could be created, to reconstruct the day of an individual (e.g., when he goes to bed, to work, workout schedule). For our scenario we ignore the risk of an attacker gaining access to the device of the data donor or attackers intercepting the communication between participant and server. We consider those attacks vectors to be handled by using security best practices like a secure encryption of the local data. We are focusing on attackers that gain access to the data stored on the research database, this attacker also models a malicious researcher that missuses his access to the database.

If attackers get access to the database, they can gain an apparent knowledge about certain properties of the user like described above. However, using LDP with the noisy data the attacker cannot be sure what the actual value is thus reducing the risk. Additionally, an attacker that has access to the whole database can reconstruct profiles about an individual even if the data is anonymized (so the actual sender is not known). If the attacker has additional knowledge about individuals in the database, he can even link the data for a so called *Attribute disclosure*. This can also be mitigated by using LDP. In this case not even the central database has the raw data making re-identification a lot harder. Even if participant IDs,

[6] https://corona-datenspende.de [In German; Accessed: 25th April 2022]
[7] https://corona-datenspende.de/science/reports/nowcast/ [In German; Accessed: 25th April 2022]

Table 1: Exemplary entries of an individual in the Fitbit dataset

Id	Time	Heartrate	Steps	SleepStage	Calories
2022484408	2016-04-12 07:21:35	101	17	0	3
2022484408	2016-04-12 07:21:40	87	9	0	3
2022484408	2016-04-12 07:21:45	58	0	0	1

IP addresses or real timestamps (IDAT) must be stored, LDP ensures that the attacker can not gain solid knowledge regarding the medical data (MDAT) of an individual in the database.

3.3 Implementation

Instead of implementing the whole workflow of a data donation with a dedicated app and a research infrastructure, we build only the necessary parts to demonstrate our use case. For RAPPOR we use the open-source library *Pure-LDP*[8] which implements a RAPPOR client besides different PETs. The library is enhanced by an implementation of the IRR and some other additions to adapt it to our use case. We provide a prototypical client and server application written in Python[9]. The server creates the RAPPOR interface and sets all the required parameters for RAPPOR. In addition, the server collects all data sent by the exemplary data donors and runs the analysis on the data. The client simulates the data donors and reads the data from the Fitbit dataset. For every user of the dataset a RAPPOR client is created which receives the parameters from the RAPPOR server. With those parameters every user can run the RAPPOR LDP on its data and send them to the server. With this we simulate an *n* parties data donation scenario.

The server runs different analyses on the received data. A simple example is to measure the frequency count of each heart rate in the received dataset as an average or median operation. This can be also done for body weight, sleep duration and step counts. More use case specific analyses following the example of the "Corona-Datenspende-App" to create a fever monitor can be also done by the server. Our model for fever is a higher average resting heart rate than usual (the higher the average the more likely the population has a higher body temperature than usual [15]). The resting heart rate estimation by the "Corona-Datenspende-App" is shown in Figure 4. On the right side of this figure there is our analysis of the Fitbit data. To estimate the average resting heart rate, we combine the step counts with the average heart rate as an association analysis. Since the step counts are also noisy through the RAPPOR processing we arrange the step count values in four activity classes: no, low, medium, and high activity. This also improves accuracy. Several studies show that useful ranges for the step counts are 0, 70 and 140 steps per minute [3, 17, 18]. Combining those values results in an average heart rate subdivided in activity categories per minute. To get the resting heart rate the values of the no activity class can be used which results in a curve analogously to the left one in Figure 4. Please note that the data used for our plot is from 2016 and from only 30 participants. So, we do not have a way to directly compare the quality of our system to the Corona-Datenspende-App. In the next section we will discuss some important elements of the system and what we can say about the quality of the results.

4 EVALUATION

In this section the results of the analysis with our implementation of a data donation with RAPPOR will be presented and discussed.

4.1 Results

We separate the discussion of the results in multiple steps. First, we look at the overall performance and then go into more specific parameters and their effect on the evaluation.

4.1.1 Data histograms. Figure 5 shows a simple histogram of all heart rates of all participants in the dataset. For every evaluation mechanism we run tests in advance to find the best parameters for the evaluation. We display the used RAPPOR parameters f, h and k in the figures as well. For the histogram of the heart rate, we choose to use a single hash function ($h = 1$) and a bit array size of 220. The size of the bit array allows us to have a single bit for every value we expect. The selected value for the privacy budget ($\epsilon = 3$), results in $f = 0.365$.

4.1.2 Offset of mean values. Figure 6 shows the average heart rate. Here we compute the average of all the received heart rates of all participants per day. The figure shows that the curve of the real values and the data from RAPPOR look similar, but the curve of the RAPPOR values is floating above the real values. This is due to the fact, that our value range is purely positive. As the RAPPOR algorithm will add additional positive measurements as noise but never negative ones, the average is increased. A strategy to deal with this offset would be to calibrate the offset with an expected average for the resting heart rate. Because the relative change is considered more relevant than in the actual values, we do not include this correction step in this work.

Going into more detail one can observe that the precision of RAPPOR or the downstream evaluations heavily depend on the value range of the input data. The larger the value range, the less precise the estimation will be.

4.1.3 Effect of data range on precision. Figure 7 shows a comparison of the LDP estimation and the raw data heart rate distribution and the same for the average calories burn. It can be noticed that the heart rate estimation alters more from the raw data curve than the average calories burn curve does. This can be explained by the larger range of possible data values. While heart rate can be between 0 and 220 the average calories burn is between 0 and 20. As the number of participants is the same in both experiments,

[8] https://pypi.org/project/pure-ldp/ [Accessed: 25th April 2022]

[9] https://www.python.org [Accessed: 25th April 2022]

[10] Source: https://corona-datenspende.de/science/reports/pulse/ [In German; Accessed: 25th April 2022]

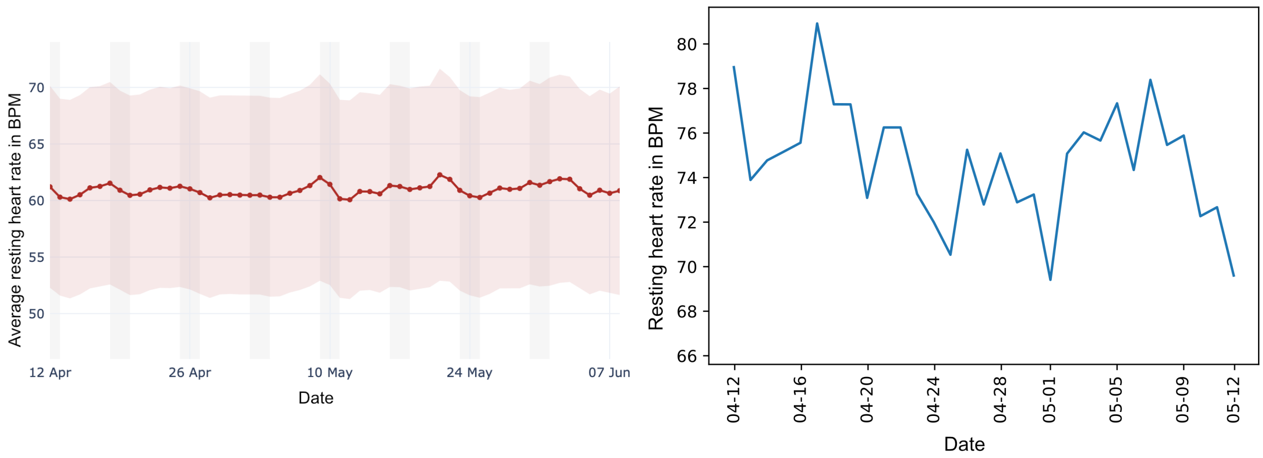

Figure 4: Comparison of the resting heart rate curve from the "Corona-Datenspende-App"^{10} on the left and our approach with RAPPOR on the right

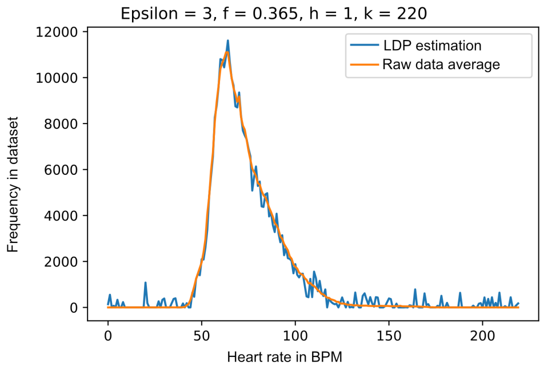

Figure 5: Frequency count of the heart rates from the dataset

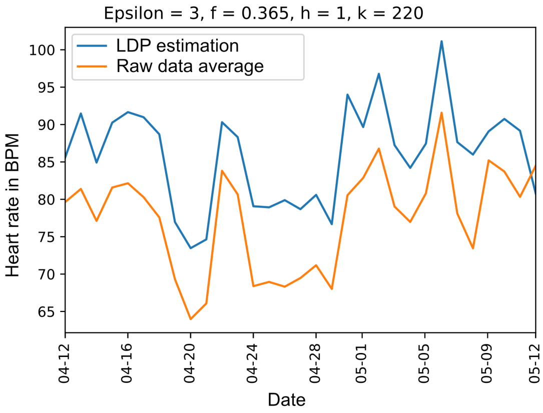

Figure 6: Daily average heart rate from the dataset

there are less data points for every value of the heart rate than for the calorie burn. So, the RAPPOR algorithm has a smaller range of values and will therefore create less outliers that will influence the results.

4.1.4 Single user evaluation vs. group evaluation. Our experiments show the desired feature of LDP is achieved. Data about an individual is disguised while we have usable results for analyses of all participants. On the left side of Figure 8 you can see the daily average heart rate of a single data donor. The RAPPOR estimation differs a lot from the raw data average. On the right side of the same figure you can see the average heart rate of all participants. The estimated curve and the real curve converge much more. This is exactly the behavior we want to have. The privacy of the individual is protected, while the data is still useful to make assertions about all participants.

4.1.5 The effect of different privacy budgets ϵ. Figure 9 shows the daily average heart rate for all data donors with two different values for the privacy budget ϵ. Again, you can see an offset between the real values and the estimation after the data was altered with RAPPOR. It can be noticed, that this offset depends on the privacy budget ϵ. The smaller ϵ the large the offset will be. This again can be explained by the inner workings of RAPPOR. A smaller value of privacy budget will result in more random altered bits in the bit array. As all our values are positive, the offset of the average will bet bigger the more random noise you add. While you could get rid of the offset with calibration as mentioned earlier, a lower value for the privacy budget ϵ (without increasing the number of participants) will result in less precise evaluations.

4.1.6 The effect of local aggregation on result quality. Figure 10 shows the resting heart rate for the fever monitor scenario. As described in Section 3.3 an association analysis from the step count and heart rate is used. At first it can be noticed that through the combination of two different values the number of outliers is rising so that there is a noticeable offset. In this analysis the small number of data points is even more reduced as we only look at the heart rate if the participant is currently in a no activity phase. This can be

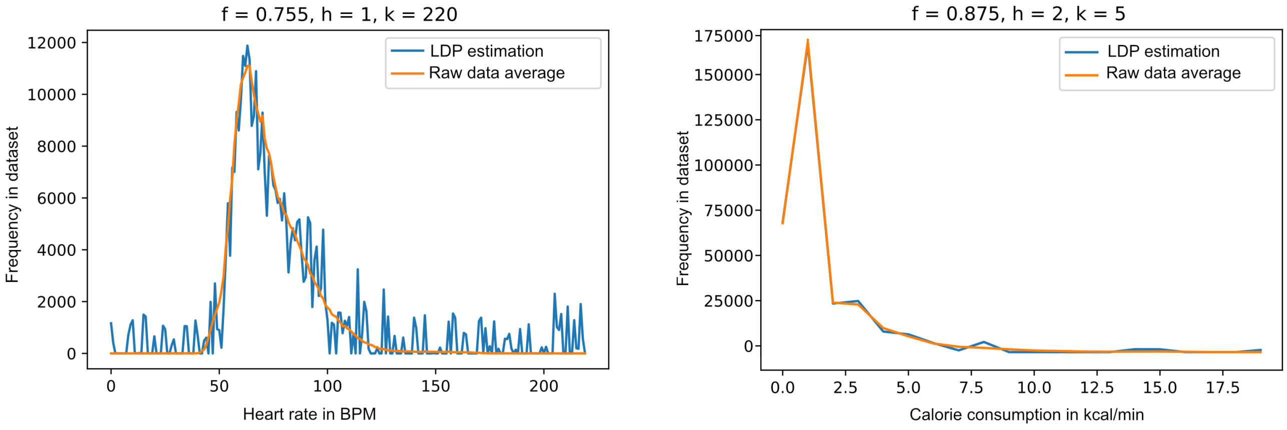

Figure 7: Distribution of heart rate and calorie consumption per minute (Privacy budget $\epsilon = 1$)

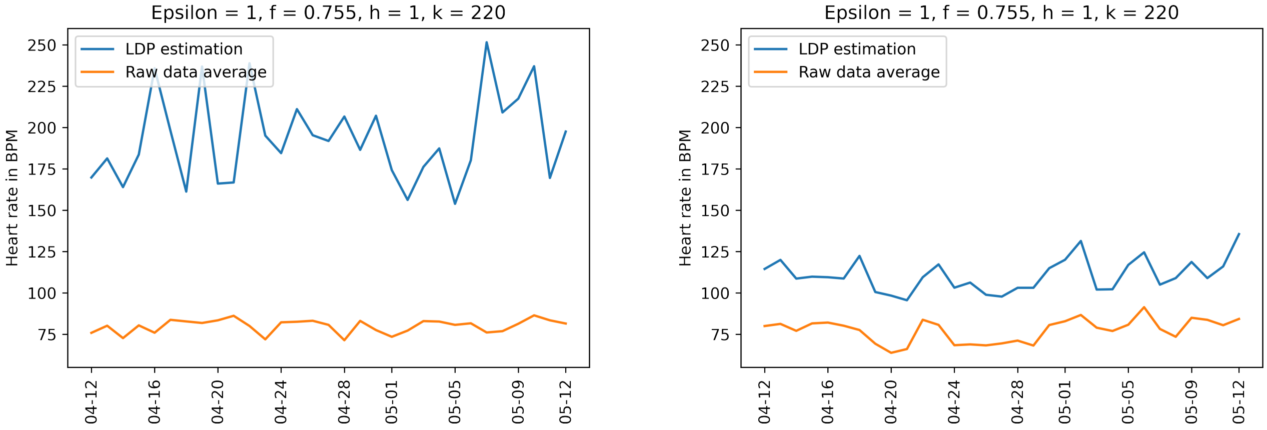

Figure 8: Daily average heart rate of single data donor vs. the average heart rate of all participants

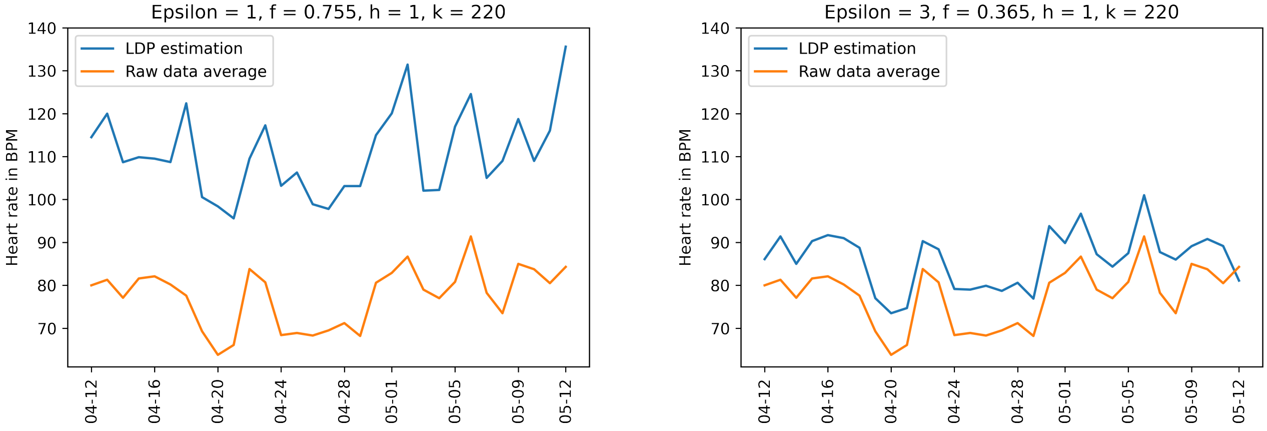

Figure 9: Daily average heart rate of all data donors (Privacy budget left: $\epsilon = 1$ right: $\epsilon = 3$)

seen on the LDP estimation curve on the left-hand side, which does not preserve the course of the raw data curve while the right side does. Using the data as it is, it results in nearly useless evaluations, as can be seen on the left side.

Looking at our data, we have the heart rate in 5 seconds intervals, while we have activity only in 60 seconds intervals. To try to improve our results we interpolated our data and created additional activity measurements with the existing information. This was done on the client before sending the data to the repository and it resulted in a significant more accurate LDP estimation. This can be seen on the right side of Figure 10. The improved diagram shows us two things. First, as with all data evaluation, the more data the more accurate our privatized estimation will be. Second, with RAPPOR or other LPD approaches, it might be a good choice to do data aggregation later in the processing. While data aggregation offers many benefits for privacy and performance, we think that RAPPOR offers even stronger privacy guarantees and the performance cost should be balanced by the greater usability.

As final remark we want to point out that there are only around 30 participants in our dataset. The "Corona-Datenspende-App" - a real world example - had around 500.000 participants, so this accuracy issue would be less relevant in the real world.

4.2 Discussion

The results show that there must be a general knowledge of the RAPPOR and LDP technology to interpret the privatized data. This helps to understand the occurrence of outliers or different levels of accuracy when using data that have different value ranges. Furthermore, the analysis methods used should have some error tolerance to generate clear results. When working with the data, there must be also some use case depending on decisions how to treat the data or how to define possible value ranges. A good example from the fever monitor use case is the activity categorization for the step counts per minute. In addition, there are also some considerations that need to be done before a data donation with RAPPOR is started. This applies especially to the RAPPOR parameters. Besides that, they need to be defined on the central server so that the reports from the individual donors remain compatible, most of them can also not be changed when the data collection is already running. If the privacy budget is badly chosen and the results have a bad accuracy, the change cannot be undone without losing the reports with the old privacy budget. On the other hand, RAPPOR can handle a growing number of data donors without any issues and does not require a static number of donors. The relevance of the number of donors could also be seen in the results of this use case. The evaluation shows that the small number of participants can lead to imprecise results. Additionally, data quality is another issue. This especially applies to data that is collected by wearables that are not worn regularly like in our use case. Irregular gaps in the datasets of the individuals can lead to inaccurate results for the whole analysis. In practice this can be balanced by a larger number of participants. For rather simple analysis like averages of heart frequencies or calorie burn a small number of participants is no big issue. Our evaluation shows that even the small number of the Fitbit dataset led to usable results. But the utility limit was reached with the more complex association analysis that was done for the resting heart rate analysis. Besides this the LDP technique shows its strength when trying to analyze the data of single individuals. With the privatized data of a single data donor a reconstruction of any pattern in the data was not possible anymore. This can be ensured by a well-chosen privacy budget. Another consideration when using RAPPOR is the long-term privacy protection of an ongoing data donor. If a single participant is asked multiple times the uncertainty about his randomized response which is used for the privatization sinks. RAPPOR already has some techniques like the IRR that helps to mitigate this risk and ensures long term protection. Additional measures and a wise choice of the RAPPOR parameters should be considered to improve the long-term privacy protection. Finally, it should be noted that RAPPOR or LDP is no silver bullet for privacy protection but is a very useful and usable technique that helps to improve privacy in the scenario of data donation. It provides a solid protection and provides accurate although for non-technicians a bit cryptic privacy protection guarantee.

5 FUTURE WORK & CONCLUSION

The authors are confident that data donations to researchers will be an important topic in the future. With upcoming devices and sensors, a broader variety of medical information can be collected. While this might be good news for medical research and the scientific results one can expect it seems plausible that the sensitivity of the data will also increase. Protecting participants rights and freedoms while allowing for data evaluation will be very important for acceptance.

We have shown that using RAPPOR as a DP mechanism is a promising PET to modify data, even before it leaves the participants devices. If we have enough participants, the precision of the data should be sufficient. As we assume that in the future more and more people will start tracking their health data, we see no problem in requiring additional participants to strongly improve privacy.

Ironically advertising that a study is protecting the participants' privacy with DP might be one of the hardest parts. While participants want to have privacy protection and DP can offer it, it can get hard to explain. Even experts struggle with giving a meaning to ϵ values. So, in the future researchers might chose not to mention it to not irritate their participants.

Of course, the quality of the results depends on the concrete evaluation that researchers plan to perform. So, we do not demand that every research project should use DP. But the question if data can be protected by DP could be a logical extension of the question which data needs to be collected in the first place. However, our results shows that PETs have the potential to improve privacy and still produce usable results for data donations. Other PETs besides DP should be looked at and evaluated against different data donations scenarios.

The system that was presented in this paper was a first prototype. While all the data modification with RAPPOR and the evaluation are present we have not yet built the part of the user app. But we are confident that large challenges are present to test in a real-world example. Speaking of real-world examples, it would be interesting to test the system with larger datasets. Even with our relatively small number of participants some performance problems occurred in the beginning of the work. They could be fixed by switching

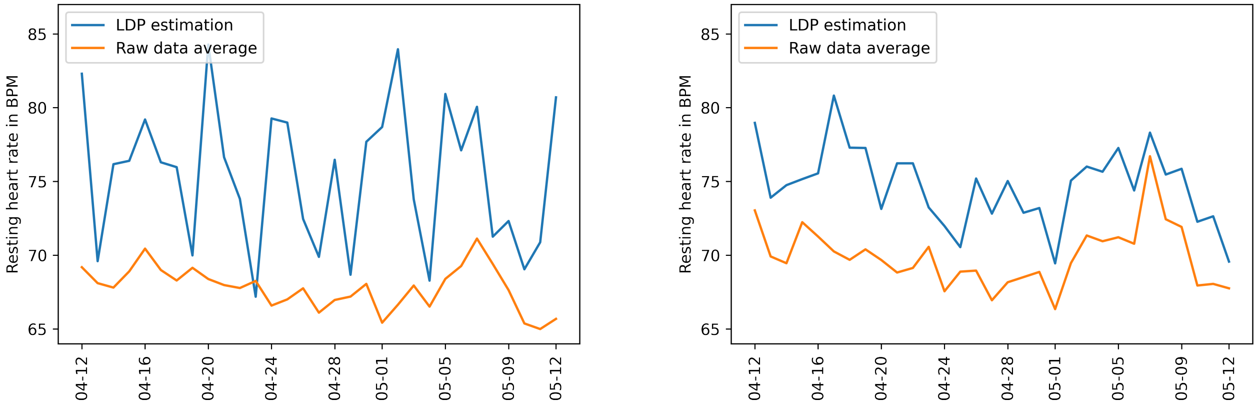

Figure 10: Daily average resting heart rate of all data donors. The left side was created with heart rate per minute. The right side with heart rate per five seconds.

from a small laptop to a more capable server. But so far, we cannot predict how system will perform with large participants numbers (>10.000) and if the performance problems are due to not optimized code or the complexity of the problem. So, a larger test with more participants seems to be the next logical step for our research. In addition, a real-world setup like the "Corona-Datenspende-App" is needed to benchmark the approach in a more realistic scenario.

REFERENCES

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. 308–318.

[2] Yves-Alexandre de Montjoye et al. 2013. Unique in the Crowd: The privacy bounds of human mobility. *Scientific Reports* 3 (2013).

[3] S. W. Ducharme, D. S. Turner, J. D. Pleuss, C. C. Moore, J. M. Schuna, C. Tudor-Locke, and E. J. Aguiar. 2021. Using Cadence to Predict the Walk-to-Run Transition in Children and Adolescents: A Logistic Regression Approach. *J Sports Sci* 39, 9 (May 2021), 1039–1045.

[4] 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.

[5] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* 9, 3-4 (2014), 211–407.

[6] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security* (2014). 1054–1067.

[7] Robert Furberg, Julia Brinton, Michael Keating, and Alexa Ortiz. 2016. *Crowd-sourced Fitbit datasets 03.12.2016-05.12.2016*. https://doi.org/10.5281/zenodo.53894

[8] Sandra Gabriele and Sonia Chiasson. 2020. Understanding Fitness Tracker Users' Security and Privacy Knowledge, Attitudes and Behaviours. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–12. https://doi.org/10.1145/3313831.3376651

[9] Raquel Hill. 2015. *Evaluating the Utility of Differential Privacy: A Use Case Study of a Behavioral Science Dataset*. Springer International Publishing, Cham, 59–82. https://doi.org/10.1007/978-3-319-23633-9_4

[10] Hao Jin, Yan Luo, Peilong Li, and Jomol Mathew. 2019. A Review of Secure and Privacy-Preserving Medical Data Sharing. *IEEE Access* 7 (2019), 61656–61669. https://doi.org/10.1109/ACCESS.2019.2916503

[11] Jong Wook Kim, Beakcheol Jang, and Hoon Yoo. 2018. Privacy-preserving aggregation of personal health data streams. *PLOS ONE* 13, 11 (11 2018), 1–15. https://doi.org/10.1371/journal.pone.0207639

[12] Chi Lin, Zihao Song, Houbing Song, Yanhong Zhou, Yi Wang, and Guowei Wu. 2016. Differential privacy preserving in big data analytics for connected health. *J. Med. Syst.* 40, 4 (April 2016), 97.

[13] Yunhui Long, Suxin Lin, Zhuolin Yang, Carl A. Gunter, and Bo Li. 2019. Scalable Differentially Private Generative Student Model via PATE. *CoRR* abs/1906.09338 (2019). arXiv:1906.09338 http://arxiv.org/abs/1906.09338

[14] C. Ma, L. Yuan, L. Han, M. Ding, R. Bhaskar, and J. Li. 5555. Data Level Privacy Preserving: A Stochastic Perturbation Approach based on Differential Privacy. *IEEE Transactions on Knowledge and Data Engineering* 01 (dec 5555), 1–1. https://doi.org/10.1109/TKDE.2021.3137047

[15] Jennifer M Radin, Nathan E Wineinger, Eric J Topol, and Steven R Steinhubl. 2020. Harnessing wearable device data to improve state-level real-time surveillance of influenza-like illness in the USA: a population-based study. *The Lancet Digital Health* 2, 2 (Feb. 2020), e85–e93. https://doi.org/10.1016/s2589-7500(19)30222-5

[16] Latanya Sweeney. 2013. Matching known patients to health records in Washington State data. (2013). https://doi.org/10.2139/ssrn.2289850

[17] Catrine Tudor-Locke, Ho Han, Elroy J Aguiar, Tiago V Barreira, John M Schuna Jr, Minsoo Kang, and David A Rowe. 2018. How fast is fast enough? Walking cadence (steps/min) as a practical estimate of intensity in adults: a narrative review. *British Journal of Sports Medicine* 52, 12 (2018), 776–788. https://doi.org/10.1136/bjsports-2017-097628 arXiv:https://bjsm.bmj.com/content/52/12/776.full.pdf

[18] C. Tudor-Locke, S. B. Sisson, T. Collova, S. M. Lee, and P. D. Swan. 2005. Pedometer-determined step count guidelines for classifying walking intensity in a young ostensibly healthy population. *Can J Appl Physiol* 30, 6 (Dec 2005), 666–676.

[19] Zhiqiang Wang, Pingchuan Ma, Ruming Wang, Jianyi Zhang, Yaping Chi, Yanzhe Ma, and Tao Yang. 2018. Secure Medical Data Collection via Local Differential Privacy. In *2018 IEEE 4th International Conference on Computer and Communications (ICCC)*. 2446–2450. https://doi.org/10.1109/CompComm.2018.8780925

[20] Stanley L Warner. 1965. Randomized response: A survey technique for eliminating evasive answer bias. 60, 309 (1965), 63–69.