

Differential Privacy Approaches in a Clinical Trial

Martin Leuckert

Faculty of Computer Science,
Otto-von-Guericke University
Magdeburg, Germany
e-mail: martin@leuckert.de

Antao Ming

Clinic of Nephrology, Hypertension, Diabetes and
Endocrinology, Otto-von-Guericke University Magdeburg
Magdeburg, Germany
e-mail: antao.ming@med.ovgu.de

Abstract— Clinical trials are essential for advancements in the medical field. The study subjects of clinical trials agree that the data may be used within the scope of the clinical trial and they trust the study center to not misuse the data. Limiting access and anonymizing the data is usually the only way of offering privacy to the subjects. Currently, the collected data may only be used within the scope of the respective study, and in the case of external entities evaluating the data, potential privacy risks occur. To improve the situation, we investigated the applicability of Differential Privacy approaches for clinical trials by looking into differentially private queries as well as differentially private Machine-Learning approaches. Different configurations have been tested for two Differential Privacy mechanisms. The Laplacian Mechanism is much more influenced by the chosen epsilon compared to the Functional Mechanism implemented in this study. However, both mechanisms trade accuracy for privacy. In summary, both queries and Machine Learning can be made secure by applying differential privacy approaches, but the implementation and configuration overhead is still likely to exceed the capacity of clinical trials, especially the smaller ones.

Keywords-Differential Privacy; Clinical Trial; Sensor Data; Machine Learning; Privacy Preservation; Data Security.

I. INTRODUCTION

There is an increasing number of companies collecting massive amounts of data about virtually every aspect of our lives. The availability of big data can be useful for many reasons, for instance, to gain statistical insights or to build Machine-Learning (ML) models. When it comes to confidential data, we expect entities that we trust our data with to release information only as long as privacy is maintained. Participants in medical trials expect their data to be handled with confidentiality, but, on the other hand, having as much available data collected as possible can be key to new scientific insights in medical trials.

In many cases, often including medical trials, the assumption is that anonymizing data suits this need. Often, it is considered safe to use pseudonyms and not release other identifying data, such as phone numbers and addresses. However, the Netflix prize dataset linkage attack performed by Narayanan and Shmatikov [1] in 2007 using the Internet Movie Database (IMDb) to successfully identify users is a good example of why pseudonymization and anonymization as the only means of privacy-preservation are insufficient.

The advances in privacy-preserving approaches are released proportionally to the increasing importance and awareness of privacy. The clinical implementation of privacy-preserving mechanisms, on the other hand, is often lagging many years behind because of the previously described misconception; and the data protection laws either do not require the implementation of advanced security functions or have, according to Koch et al. [3], insufficient requirements. On the basis of a real clinical study, we discuss an approach to improve the situation. This work focuses on the applicability of Differential Privacy (DP) in a specific medical trial scenario rather than surveying or evaluating different DP mechanisms to find the most suitable mechanism. However, the outcome of relevant surveys of DP ML in practice, such as Jayaraman and Evans [14], has been considered.

A. Problem Definition

Initially, we explain the setup of a real-world scientific study to illustrate the privacy problem and how specific privacy-preserving mechanisms can be used to solve them. The research was carried out in the context of a clinical trial that studied ulcer prevention using a smart insole. The study, which is based on the findings in Armstrong et al. [2], found that the temperature at the affected foot regions increases weeks before the inflammation. The study, conducted in [6], aimed at providing 300 diabetics who suffer from comorbidities like nerve damage and are at risk of developing ulcers with a smart insole in order to intervene in time. The insole has multiple temperature sensors and transfers the measurement data to a smartphone app which then forwards the data to an electronic trials system located at the research facility.

Researchers then analyze the data to learn about potential diseases like ulcers, gout, or peripheral arterial occlusive disease that can be detected by continuously measuring the foot temperature. Further research intends to find automated alarm signals by using ML algorithms to identify arising ulcers early. In order to benefit the most from the data, it makes sense to involve third-party scientists specialized in data mining and ML.

First, the data subject must give explicit consent to all of the primary (article 6 (1)(a) of General Data Protection Regulation (GDPR)) and secondary research activities (article 6 (1)(b) of GDPR) involving their personal data:

“Personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a

manner that is incompatible with those purposes; further processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes shall, per Article 89(1), not be considered to be incompatible with the initial purposes ('purpose limitation');"

This clinical trial setup relies on third parties to analyze the acquired data. Under the assumption that all requirements of the GDPR, including the explicit consent, are met, the privacy of the participants is at risk: Both the data queries and the ML models reveal data about the study participants. According to Jagannathan et al. [15]: "The difficulty of individual privacy is compounded by the availability of auxiliary information, which renders straightforward approaches based on anonymization or data masking unsuitable."

Significant progress was made when Cynthia Dwork [4] defined DP as retrieving useful information while maintaining privacy. Pre-eminently, DP uses randomized noise to protect individuals in a data set. The required range of noise that needs to be added to a query depends on the sensitivity of the respective function. The sensitivity describes the maximum difference between two queries on an underlying data set and is therefore proportional to the magnitude of the required noise to maintain privacy. Depending on the underlying data set, the amount of required noise can be very high if the global sensitivity is high. There are investigations to still achieve DP in these cases; Lundmark and Dahlman [5], for instance, address the issue of applying noise based on global sensitivity to reduce the required noise.

B. Goals

First, this work will demonstrate why the security regulations required by European law and their national implementations are insufficient in the context of preserving the participant's privacy. This includes the General Data Protection regulation (art.70.1.b of the GDPR) and the Clinical Trials Regulation (CTR).

Second, we will demonstrate that it is possible to implement DP in the context of the clinical trial described in the problem definition to improve privacy without significantly affecting the usefulness of the results (utility). This is possible for both queries and ML operations. We will conclude this paper with a subjective assessment of the results.

C. Setup

Implementing privacy-preserving mechanisms extending further than pseudonymization or anonymization might be hard to sell to physicians. They potentially fear for the usability of their data if encryption or noise of some sort is implemented. In the same vein, looking into the field of homomorphic encryption reveals many cases of rejection due to performance concerns [25]. Among other reasons, this is why most clinical trials implement legally required privacy measures without questioning them.

The open-label, prospective, and single-blinded study recruited participants with diabetes mellitus type I or II who are randomly assigned to the control (n=150) or the intervention group (n=150). All study participants are diagnosed with severe peripheral neuropathy (e.g., vibration perception $\leq 4/8$).

The study participant provides data by regularly measuring their foot temperature using smart insoles and a mobile application. The application uploads the raw data. Data analysts perform queries on the data with the goal of finding patterns that could help in developing and improving automatic ulcer detection algorithms. The analysts apply both Data Mining as well as ML approaches to make sense of the collected data and to predict future behavior (see examples described in Section IV).

Section II addresses related work that is the foundation of this study. Next, Section III describes possible attack scenarios. Section IV and section V describe DP queries and DP ML. The article closes with section VI summarizing the results and providing an outlook.

II. RELATED WORK

ML models are commonly used in the health care field. For instance, Orfanoudaki et al. [17] identify a non-linear Framingham stroke risk score using Optimal Classification Trees. With regard to the subject matter at hand, Tabaei et al. [18] use a logistic regression model to predict the likelihood of study subjects suffering from diabetes. Maniruzzaman et al. [19] expand on the aforementioned studies by addressing the impact of missing values and outliers and verified their results in different scenarios by testing six feature selection techniques and ten different qualifiers with Random Forest-based models showing the best performance. The given example and many more studies aim to create or improve their models and databases. Moreover, other studies focus on identifying various approaches to making ML algorithms privacy-preserving. This may partly be the case because the nature of the underlying ML algorithms is substantially different, but it is also driven by the system design and data flow. There are, among others, supervised, unsupervised, and reinforced ML algorithms that require different types of data and produce different types of results. Furthermore, the system can follow a local or a global privacy approach. Local privacy can be achieved by perturbing the individual input. Global privacy can be achieved by cost function or output perturbation, which will be explained in detail in Section IV and V. Privacy Preservation can be further expanded to other fields, like, for instance, Deep Learning. Phan et al. [9] proposed an adaptive Laplacian mechanism that can be used in a Deep Learning setting.

In [16], Bos et al. provide a good introduction to the topic of publicly available databases as well as privately compiled databases containing medical records. The authors expand on the point made in [15] that masked data records state privacy concerns when publicly available. Respectively, according to Bos et al. [16], publicly available databases provide the most benefit while also "creating the steepest privacy challenge". They first compared "conventional encryption" to

homomorphic encryption, concluding that both encryption approaches can be used to assure privacy, but homomorphic encryption provides more operations on the encrypted data without the need for a decryption key. Second, they describe possible scenarios to conduct predictive analysis privately. In their outlook section, Bos et al. describe the need for performance improvements, which remains an issue with homomorphic encryption.

DP mechanisms use different ways of data perturbation to protect the privacy of individuals in a data set. Local DP approaches perturb the data on input time while global DP approaches do so when the data is queried by an adversary. The DP mechanisms range from applying random noise (e.g., coin toss) to more advanced systems using Laplacian noise [8]. Fundamental work and surveys by Dwork et al. can be found in [4], [8], [20], [21]. DP can be applied both to queries and ML approaches. For instance, Cheu et al. [22] introduce a system that works with sensitive data in a distributed setting and applies DP via shuffling.

Other contributions discuss the application fields of DP and that it has been successfully applied. Nguyen et al. [26] stated in 2013 that DP “[...] has become the de facto principle for privacy-preserving data analysis tasks”. The application of DP on medical data is actively researched: Lee and Chung [24], for instance, propose “Informative attribute preserving anonymization” (IPA), which is further discussed in Section IV.

III. ATTACKS ON DATA RECORDS AND MODELS

This section goes into detail about why and how the previously described medical trial raises privacy concerns for participants even though it acts within the legal requirements. The study participants agreed that their data may be shared with data analysts. Data analysts can access the masked data via a query interface using a secure channel, which allows for a similar linkage attack as described in Section I. Data analysts can query personal information like a subject’s birthday, sex, diabetes type, and other known information regarding medication or medical anamnesis. The combination of the information becomes a quasi-identifier, rendering the pseudonymization meaningless.

A. Membership inference

In medical trials, the ML models are trained on highly sensitive data of real persons and could potentially leak information about them. Membership inference attacks aim to prove the existence of a data record in a data set. According to [13], this is done by training an attack model which intends to distinguish the behavior based on input that was part of the training and input that was not. Publicly available ML models are usually block-boxes with unknown structures and parameters. Shokri et al. [13] propose multiple generic techniques to tackle this problem. For instance, they introduced “shadow training”. Shadow training creates multiple models that imitate the original ML model’s behavior with known training data.

B. Attribute Inference

Attribute inference attacks are based on publicly available information about a person that is either provided directly by the user or gathered indirectly via their connections (“friends”) on their social media accounts. The combined knowledge can then be used to infer or validate further information about an individual. Jayaraman and Evans [14] describe attacks on social network profiles of users and infer data about individuals by creating “social-behavior-attribute networks” and run different mechanisms like, e.g., “friends-based attack” on them.

The work of Shokri et al. [13] and Jayaraman and Evans [14] are examples of privacy breaches while potentially fulfilling the requirements of GDPR and CTR (see first goal in Section I), but both attacks can be mitigated by DP because there is plausible deniability or reasonable doubt about the presence or authenticity of data.

IV. DIFFERENTIAL PRIVATE QUERIES

There are two stakeholders performing queries on the data set: the trial staff located at the study center observing the study data to intervene if necessary and the data analysts. Data analysts can be understood as adversaries in this setup and should be prohibited from finding sensitive information about individuals.

The original data may not be changed, which is why a preceding data perturbation is not a suitable solution but can be done on intermediate data sets. Purely syntactic approaches are also not suitable because the use case can be understood as a data mining problem rather than a data publication problem. Other means of anonymization and DP are mandatory to protect the privacy of individuals. This section describes different approaches to create differentially private versions of the queries. Transforming the queries into differentially private queries has an impact on the usefulness of the result due to reduced accuracy. We decided to go with the rather straightforward and well-known approaches to show their applicability in a real-world telemedical use case and do not focus on maximizing privacy or improving existing DP approaches.

A. Basic DP query mechanisms

To evaluate a selection of differentially private analyses, the following example query will be used:

“Did study participant x have a foot ulcer in the past?”

This is revealing information and, therefore, worthy of being protected. Across all study participants, the percentage to answer the query with “yes” lies at $p=0.3$. Each query on the medical database, including DP queries, reveals information about a patient and causes a certain amount of privacy loss. The privacy loss is defined by the parameter ϵ . The closer ϵ gets to 0, the smaller the privacy loss will be for each query. However, smaller ϵ also decreases the accuracy of the result due to the increased noise level. Fig. 1 illustrates how the usability increases when a larger n is available.

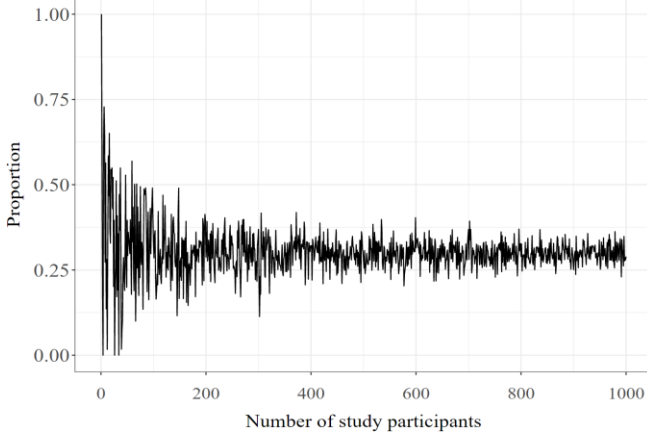


Figure 1: Laplacian mechanism’s decreasing usability impact with larger n

The first global DP method is an ϵ -DP mechanism called the Laplacian method which is popular for numeric functions. Dwork [4] and Dwork and Roth [8] introduced solutions to (ϵ, δ) -DP by applying Gaussian noise to query results. In [4] the summand

$$\sqrt{2 \ln \frac{2}{\delta} \times \frac{\Delta f}{\epsilon}}, \quad (1)$$

which was changed in [8] to

$$\sqrt{2 \ln \frac{1.25}{\delta} \times \frac{\Delta f}{\epsilon}}, \quad (2)$$

is added independently to each query answer for a query with L2 sensitivity Δf . The Gaussian mechanism does not satisfy ϵ -DP but achieves (ϵ, δ) -DP for some $\delta \in [0, 1]$ while Laplacian achieves ϵ -DP.

Consequently, the Laplacian mechanism works best with low sensitivity and smaller amounts of queries. Vice versa, large amounts of queries require a larger ϵ , which produces less accurate results. The relaxed (ϵ, δ) -DP definition and the smaller accuracy of the Gaussian mechanism turn out useful for vector-valued functions. The Laplacian mechanism requires the use of L1 sensitivity, while the Gaussian mechanism supports both L1 and L2 sensitivity. There are extensions and improvements to Gaussian and Laplacian mechanisms available with higher privacy results described, among others, in [7].

The Gaussian and Laplacian mechanisms are both focused on numerical queries. However, McSherry and Talwar [23] proposed a mechanism that is able to solve different types of problems that require retrieving a certain element of an existing set R that fits a query. A simple example could be: “What is the most common comorbidity of diabetic foot neuropathy?” from a set that could be

$R = \{“Ulcer”, “Gait”, “Macroangiopathy”, “Fasciitis”, “Angiopathy”, “Arthrosis”\}$.

B. Medical DP queries

Naturally, there are more complex queries than the query used for 5.1. Likewise, the requirements for DP queries exceed the possibilities of the basic mechanisms. It becomes both interesting and complex when different approaches are combined, may they be sequential or parallel compositions of DP functions.

The following somewhat simplified query is a realistic example that was run on the data in a similar fashion:

```
SELECT AgeGroup, Disease, COUNT(*)
FROM (
  SELECT FLOOR (Age/5) * 5 as AgeGroup, *
  FROM Patients
  WHERE Sex = ‘male’ AND DiabetesType = 1
) GroupedResults
GROUP BY AgeGroup, Disease
```

The query goes through the study subjects and divides them into age groups and diseases. A possible way of applying compositions of DP functions is the IPA approach proposed in Lee and Chung [24]. The IPA approach goes through a processing pipeline as illustrated in a simplified version in Fig. 2. The authors of [24] classify the data perturbation into different methods: generalization, suppression, and insertion. Each method achieves a different goal, such as reducing the number of counterfeit records or reducing information loss.

V. DIFFERENTIALLY PRIVATE MACHINE LEARNING

In the previous sections, we have introduced basic DP mechanisms. Now, we go a step further by implementing privacy-preserving ML using the same clinical trial as our use case. In contrast to the more theoretical Section IV, this section is more detailed and looks into the trade-off between the privacy parameter ϵ and the prediction quality of the ML model.

ML allows more stages to perturb data to make ML DP. We will consider output and cost function perturbation.

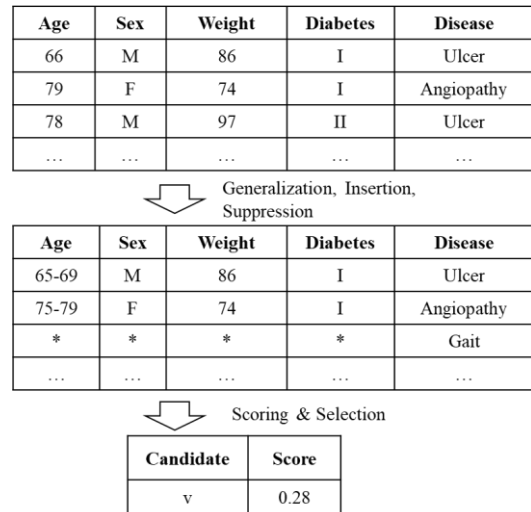


Figure 2: Simplified IPA model by Lee et al. [24]

This section will describe how output and objective perturbation have been applied to linear regression in a real-world application.

A. Differentially Private Linear Regression

In our use case, we gather significant amounts of data from many different patients. One of our goals is to build a predictive model to identify inflammations or other diseases at an early stage and maybe even predict them before they occur. Using ML on the data sets has the potential to improve the accuracy of our prediction. However, this first example takes a step back and provides a forecast of the temperature development.

Let

$$y = f(\vec{x}, \vec{w}) = w_0 + w_1x_1 + \dots + w_Dx_D, \quad (3)$$

where $\vec{x} = (x_1, \dots, x_D)^T$ and w_i are weights. With N data records, X has dimension $(N \times D)$, which will become $(N+1 \times D)$ -dimensional matrix \vec{X} when accounting for w_0 . Then $\vec{y} = \vec{X}\vec{w}$. When training a model from a data set, \vec{y} can then be used to evaluate the chosen \vec{w} . A popular cost function can be the Mean Squared Error (MSE)

$$MSE = \frac{1}{N} \sum_{i=1}^N (y_i - y_i^*)^2. \quad (4)$$

To make the linear regression DP, we can add noise at several stages in the process including the dataset, the cost function, and the prediction output as shown in Fig. 3. As mentioned before, we will not alter the original datasets because the trial staff must have access to an immaculate dataset. Instead, we could create a secondary synthetic data set from the original data set that can be used to achieve a DP Linear Regression [10]. However, creating a synthetic data set was not part of this work. The linear regression is executed on a dataset of feet temperature measurements as described in Section I.

Several features are collected during the clinical trial as described in Ming et al. [6] and Section I. For simplicity reasons, no thought-out feature selection has been performed, but the features have been reduced to the available temperature data. The trained model has an MSE of 1.27.

The first example will add Laplacian noise to the prediction output of the linear regression. In order to do that, we need to calculate the sensitivity l_1 . According to [8] the sensitivity l_1 is determined by finding Δf of a function $f : N^{|x|} \rightarrow R^k$ over all pairs of neighboring databases.

However, the pairs can only be found by making many

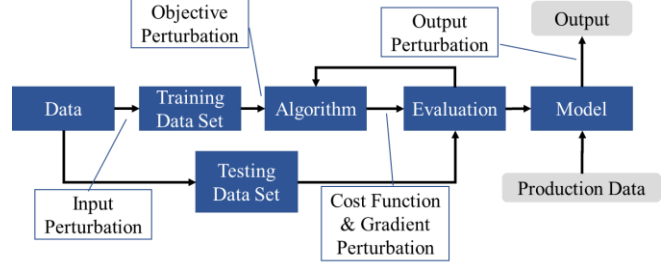


Figure 3: Perturbation approaches

assumptions about, for instance, the highest and lowest possible temperatures. Alternatively, we follow the approach proposed in Ji et al. [11] to find neighboring databases by deleting an element rather than changing it. Finding the element with the biggest impact on the model is still a challenging task; particularly if large amounts of data are gathered. Our use case allows applying a brute force approach because we have a maximum of 1,424 data records per study participant. We were able to identify a neighboring database with the highest difference in the MSE by deleting the element with the largest impact at index 64. Now that we have our original database and the one with the most differing outputs, the difference between their MSE can be used to find an approximated sensitivity of 0.62.

With the sensitivity value, we can now apply the following $\text{Lap}(0, 0.62/\epsilon)$ with ϵ being the selected security parameter. If ϵ is very small, e.g., 0.01, the noise addition will be very high, and the usability of the data gets very low due to a high mean error rate. With a higher value for ϵ , the error rate decreases but so does the privacy gained by the noise addition. Dwork [4] and Dwork and Roth [8] proposed a range between 0.01 and $\log 3$. Finding the “best value” for ϵ is not a trivial task and always needs to be a compromise between usability and privacy, depending on the requirements. If we use $\log 2$, for instance, and repeat the test using 10-fold cross-validation, we get an MSE of 5.74. The average processing time increased from approximately 7ms to 44ms on a machine with Intel Core i7-8665U with 48GB RAM.

To perturb the cost function of linear regression, we have to preprocess our data because it needs to be in the range $[-1, 1]$. This was achieved by scaling it using a min-max normalization

$$x' = \frac{x - x_{\min}}{x_{\max} - x_{\min}}. \quad (5)$$

Following the approach described by Zhang et al. in [12], we are not only able to perturb the cost function but also the function itself. This can be done by adding noise directly to the cost functions. Here, again, we used the noise from the Laplacian distribution. Following the definitions from [12], we define our problem to have a set of features x_1 to x_n resulting from the temperature measurements and a Boolean y indicating whether the participant has developed a disease.

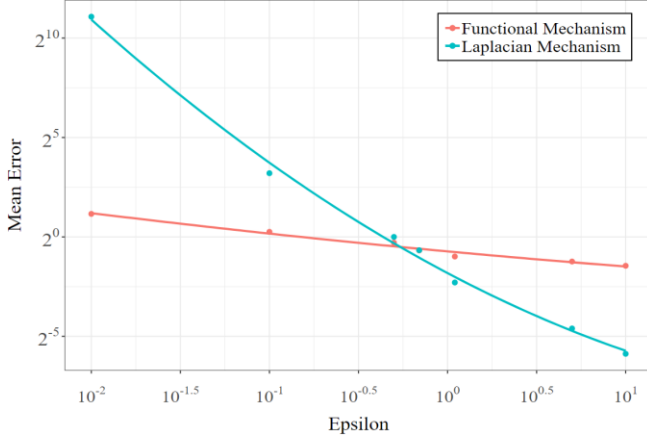


Figure 4: Laplacian and Functional Mechanism in comparison

This leaves us with a prediction function to predict $y = 1$ with probability:

$$p(x_i) = \frac{\exp(x_i^T \omega^*)}{(1 + \exp(x_i^T \omega^*))}, \quad (6)$$

Zhang et al. [7] describe ω^* as a vector of d real numbers that minimize a cost function:

$$\omega^* = \arg \min \sum_{i=1}^n \left(\log(1 + \exp(x_i^T \omega)) - y_i x_i^T \omega \right). \quad (7)$$

Using this function, a logistic regression on our dataset will be able to return a probability of a participant having an inflammation. To achieve DP by perturbing this function, we use the functional mechanism and the polynomial extension to this mechanism from [12], which have been proven to achieve DP for logistic regressions. The functional mechanism averaged at approximately 15ms on the same machine.

In Fig. 4 it can be seen that for the smallest $\epsilon=0.01$ the Laplacian approach reaches a mean error of around 211 where the functional approach only reached 2.2, making the latter significantly more suitable. However, with decreasing ϵ the mean error also drops exponentially, eventually falling below the mean error of the Functional Mechanism. This explanation lies in the nature of the Laplacian algorithm which adds noise based on the underlying distribution. If all samples are very close together, it is much simpler to hide the original values but with strong outliers, much more noise needs to be added. The Functional Mechanism is better suited for smaller ϵ because it provides more accurate results than the Laplacian Mechanism.

VI. CONCLUSION

ML problems can have different data types which are more or less suitable for the previously described DP mechanisms. If the data is strongly correlated, it gets even

worse. Eventually, the practicability of the DP mechanisms remains dependent on the application. Sections IV and V have shown that it is a possible but not a trivial task to select the correct DP mechanism, since it requires a deep understanding of the (ML) task as well as DP. The exemplary privacy breaches from Section III can be prevented by choosing the right trade-off between usability and privacy.

The authors of this work are not aware of any openly available DP libraries which can be used for ML tasks, but existing open-source libraries can be integrated into, e.g., Microsoft’s “ML.net” framework, which was one of the chosen approaches for this paper. Hence, each clinical study faces the problem of finding the correct DP approach to their individual ML tasks. Because of the unavailability of out-of-the-box solutions, smaller scale studies like our use case from Section I using DP correctly likely exceeds their possibilities and could be solved differently. However, when creating large databases like a diabetes register of a state with thousands of entries that could be used by multiple studies at once, DP becomes a more realistic approach.

Assuming the masked dataset is publicly available, it would allow for creating a huge learning data set. On the other hand, the public availability would pose a great privacy challenge. The privacy challenge can be addressed by applying differential privacy-preserving techniques, which enables users to query for approximate answers based on trained models. The suitability of DP techniques that build on ML training models requires further investigation [16]. Furthermore, the question remains whether this can be applied efficiently in the encrypted domain. Syntactic approaches were also not considered in this work and may be a valid solution for certain problems.

Privately compiled databases are a more typical scenario to handle patient data because companies and hospitals usually do not disclose their data freely. Regulations and applicable laws bind stakeholders to not only handle data confidentially but to use them for predetermined purposes. Nevertheless, both clinics and companies wish to learn as much as possible from their data and, consequently, to improve their work. Another approach to overcome this dilemma could be using a homomorphic encryption function. It may be possible to outsource the homomorphically encrypted storage and prediction model building and still maintain confidentiality.

The previously described use case is not as time-sensitive as, for example, an ECG evaluation. Nevertheless, a fast and efficient implementation is always desirable with respect to cost-efficiency. On the other hand, this approach may be adapted in a different medical use case that works with continuous data flows.

REFERENCES

- [1] A. Narayanan, V. Shmatikov, “How To Break Anonymity of the Netflix Prize Dataset,” Arxiv, 2006.
- [2] D. G. Armstrong, A. J.M. Boulton, and S. A. Bus, “Diabetic Foot Ulcers and Their Recurrence,” *N Engl J Med.*, 376(24) pp. 2367-2375, 2017.

- [3] H. Koch, B. Schütze, G. Spyra, and M. Wefer, "Datenschutzrechtliche Anforderungen an die medizinische Forschung unter Berücksichtigung der EU Datenschutz-Grundverordnung (DS-GVO)," GMDS e.V., Köln, 2017.
- [4] C. Dwork, "Differential Privacy," International Colloquium on Automata, Languages and Programming, part II (ICALP 2006), vol. 4052, pp. 1-12, 2006.
- [5] M. Lundmark, C. Dahlman "Differential privacy and machine learning: Calculating sensitivity with generated data sets," KTH, Stockholm, 2017.
- [6] A. Ming, I. Walter, A. Alhajjar, M. Leuckert, and P. R. Mertens, "Study protocol for a randomized controlled trial to test for preventive effects of diabetic foot ulceration by telemedicine that includes sensor-equipped insoles combined with photo documentation," *Trials*, vol. 20, 521, 2019.
- [7] J. Zhao, "Reviewing and Improving the Gaussian Mechanism for Differential Privacy," arXiv:1911.12060, 2019.
- [8] C. Dwork and A. Roth, "The Algorithmic Foundations of Differential Privacy," *Foundations and Trends in Theoretical Computer Science (FnT-TCS)*, vol. 9, no. 3-4, pp. 211-407, 2014.
- [9] N. Phan, X. Wu, H. Hu, and D. Dou "Adaptive Laplace Mechanism: Differential Privacy Preservation in Deep Learning," arXiv:1709.05750, 2017.
- [10] J. Lei, "Differentially Private M-Estimators," *Advances in Neural Information Processing Systems*, pp. 361-369, 2011.
- [11] Z. Ji, Z. C. Lipton, and C. Elkan, "Differentially Privacy and Machine Learning: A Survey and Review," arXiv:1412.7584, 2014.
- [12] J. Zhang, Z. Zhang, X. Xiao, Y. Yang, and M. Winslett, "Functional Mechanism: Regression Analysis under Differential Privacy," arXiv:1208.0219, 2012.
- [13] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, "Membership Inference Attacks Against Machine Learning Models," 2017 IEEE Symposium on Security and Privacy (SP), San Jose, CA, pp. 3-18, doi: 10.1109/SP.2017.41, 2017.
- [14] B. Jayaraman and D. Evans, "Evaluating Differentially Private Machine Learning in Practice," 28th USENIX Security Symposium, pp. 1895-1912, 2019.
- [15] G. Jagannathan, K. Pillaipakkamatt, and R. N. Wright, "A Practical Differentially Private Random Decision Tree Classifier," *Transactions on Data Privacy*, vol. 5, pp. 273-295, 2012.
- [16] J. W. Bos, K. Lauter, and M. Naehrig, "Private Predictive analysis on encrypted medical data," *Journal of Biomedical Informatics*, vol. 50, pp. 234-243, 2014.
- [17] A. Orfanoudaki et al., "Machine learning provides evidence that stroke risk is not linear: The non-linear Framingham stroke risk score," *PLoS ONE*, 15(5), pp. 1-20, 2020.
- [18] B. P. Tabaei and W. H. Herman, "A multivariate logistic regression equation to screen for diabetes: development and validation," *Diabetes Care*, 25(11):1999-2003, 2002.
- [19] Md. Maniruzzaman et al., "Accurate Diabetes Risk Stratification Using Machine Learning: Role of Missing Value and Outliers," *J Med Syst* 42, 92, 2018.
- [20] C. Dwork, "Differential Privacy: A survey of results," TAMC: Theory and Applications of Models of Computation, 5th International Conference, pp. 1-19, 2008.
- [21] C. Dwork, "Differential Privacy in new settings," *SODA*, pp. 174-183, 2010.
- [22] A. Cheu, A. Smith, J. Ullman, D. Zeber, and M. Zhilyaev, "Distributed Differential Privacy via Shuffling," *Advances in Cryptology – EUROCRYPT 2019*, 2019.
- [23] F. McSherry and K. Talwar, "Mechanism design via differential privacy," *FOCS*, volume 7, pp. 94-103, 2007.
- [24] H. Lee and Y. D. Chung, "Differentially private release of medical microdata: an efficient and practical approach for preserving informative attribute values," *BMC Medical Informatics and Decision Making* 20, pp. 1-15, 2020.
- [25] M. E. Y. Geng, "Homomorphic Encryption Technology for Cloud Computing", *Procedia Computer Science*, pp. 73-83, 2019.
- [26] H. Nguyen, J. Kim, and Y. Kim, "Differential Privacy in Practice", *J. of Computing Science and Engineering*, pp. 177-186, 2013.