Simplifying Differential Privacy for Non-Experts: The ENCRYPT Project Approach

Stelios Erotokritou
Eight Bells Ltd
Nicosia, Cyprus
stelios.erotokritou@8bellsresearch.com

Ioannis Giannoulakis
Eight Bells Ltd
Athens, Greece
giannoul@8bellsresearch.com

Emmanouil Kafetzakis
Eight Bells Ltd
Athens, Greece
mkafetz@8bellsresearch.com

Konstantinos Kaltakis
Eight Bells Ltd
Athens, Greece
konstantinos.kaltakis@8bellsresearch.com

Abstract
The ENCRYPT project, funded under the Horizon Europe Framework, aims to advance privacy-preserving technologies across various sectors, ensuring robust data protection while maintaining utility. By providing users with core methodologies including the Fully Homomorphic Encryption, Trusted Execution Environments, Differential Privacy and advanced Hybrid Protection Services, ENCRYPT seeks to address the challenge of ensuring data privacy and utility, across federated data spaces within the EU. Differential privacy is an important approach within ENCRYPT, in protecting individual privacy in the growing landscape of digital data. In this paper, we provide an overview of fundamental concepts and present an overview of differential privacy foundations, examining its theoretical underpinnings and practical implementations. We also provide an insight into how it will be applied within the ENCRYPT project. Experiments carried out demonstrate that differential privacy can maintain high data accuracy despite the addition of noise, and we will describe how the ENCRYPT platform simplifies the use of this privacy-preserving technology for non-expert users by automating privacy parameter selection and model optimization. This approach enhances data security, efficiency and accessibility, helping to develop a more privacy-conscious environment for data analysis to carry out research and innovation in a secure and private manner. We will also explore potential future developments and applications of differential privacy within various industries and sectors.

Keywords - Differential Privacy, Data Security, ENCRYPT Project, Privacy-Preserving Technology, Data Accuracy, Automated Model Optimization, User-centric Privacy Solutions

I. INTRODUCTION

In today's digital era, the abundance of data presents significant opportunities for addressing emerging challenges, advancing research and fostering innovative services. One notable application is the enhancement of machine learning models and data analytics through federated learning on extensive datasets [1].

In light of this, the primary challenge in handling such data - which often contains sensitive or personal information, is the potential threat of cybersecurity attacks and the risk of disclosure or misuse of private data. Adhering to data protection regulations and the strict privacy standards set by the EU on personal data - such as the General Data Protection Regulation (GDPR) further complicates the management of such information [2].

Furthermore, increasing public awareness and concerns with regards to data privacy requires strict measures to be followed to ensure data protection. Importantly, these need to be readily and easily available to an increasing proportion of the population. Addressing these concerns is critical to maintaining public trust and ensuring the continued availability of data and their security, for beneficial uses.

In this landscape, the ENCRYPT project [3] is an initiative under the Horizon Europe Framework. This project focuses on integrating cutting-edge privacy-preserving technologies to protect sensitive data across federated data spaces, enhancing security and compliance with the GDPR. Through its innovative platform, ENCRYPT facilitates secure and efficient data processing, addressing critical needs in various sectors such as finance, healthcare, and cybersecurity.

ENCRYPT leverages advanced Privacy-Preserving Technologies (PPTs) such as Fully Homomorphic Encryption (FHE) [4], Trusted Execution Environment (TEE) [5, 6] and Differential Privacy (DP) which offer potential GDPR-compliant solutions, and these aim to achieve applicability and reliability for real-world applications.

Despite their promise, existing PPTs face several limitations before they can become widely adopted security solutions. As an example, FHE struggles with scalability when processing large amounts of data due to its high computational overhead. Additionally, many PPTs lack integration with existing networking infrastructure and security protocols. This poses a challenge for ongoing research. Such difficulties highlight the importance for continued innovation and collaboration in the field of PPTs to create more efficient, accessible and user-friendly solutions. Innovations in this area can unlock the potential for secure data use across various sectors and new practical applications and solutions for various fields.

The ENCRYPT project aims to address these challenges by providing researchers and service providers who handle personal and sensitive data with a scalable, practical and adaptable privacy-preserving framework. The ENCRYPT platform facilitates GDPR-compliant processing of data stored in federated cross-border data spaces.

By developing and integrating a number of PPTs, the ENCRYPT project aims to make privacy-preserving data analysis accessible and practical for a wide range of users. In this way, it will help promote a culture of data security, private and anonymous computation and compliance to security and privacy standards. The ultimate goal is to enable a more secure data environment which encourages innovation and research.

This paper focuses on DP and how it will be developed and deployed within the ENCRYPT project to allow use by non-expert users through the ENCRYPT platform and its recommendation engine, which proposes privacy parameter selection and automates model optimization.

The rest of this paper is organized as follows. Section II discusses the theoretical foundations of DP. Section III details DP's application within the ENCRYPT project and its impact on privacy solutions. Section IV describes the methodologies and platform design of ENCRYPT, followed by Section V which presents experimental results validating the effectiveness of DP. Section VI highlights future research directions in privacy-preserving technologies. The paper concludes in Section VII, summarizing the ENCRYPT project's contributions to data privacy.

II. Background

A. Introduction to Differential Privacy

DP is a PPT which prevents gaining knowledge of a dataset and so does not compromise the privacy of individuals within the dataset. This security property is very important in the field of data privacy within data analytics, where the challenge lies in balancing the utility of data analysis with the need to protect individual privacy.

Understanding the mathematical foundations of DP is important to appreciating its effectiveness and potential applications. DP was formulated in 2006 in the work by Cynthia Dwork [7], and since then is used as a privacy preserving data analytics technology. DP ensures that the outcome of any statistical analysis is indistinguishable whether a single individual's data is included in the dataset or not. This is achieved by altering data (or the results of queries made on a dataset) with a controlled amount of random noise. The addition of noise is one of the most important aspect of DP, as it obfuscates the contribution of individual data points while maintaining data utility [8].

Addition of noise can be drawn from a Laplace distribution, Gaussian distribution or other mechanisms. The security parameter of DP is the privacy loss parameter - denoted by ε (epsilon), which quantifies the trade-off between privacy and accuracy. Smaller values of ε introduce greater amounts of noise to a dataset, achieve stronger privacy guarantees, but can potentially lead to less accurate data analytic results. On the other hand, greater values of ε add reduced amounts of noise to datasets, enhancing accuracy but potentially compromising privacy. This highlights the importance of selecting appropriate ε values. The ENCRYPT platform makes this a transparent process, by suggesting ε values to users based on their security requirements and application.

DP has been used in different practical applications in various fields, such as in the U.S. Census Bureau's adoption of DP techniques for the 2020 Census to protect respondents' data [9, 10], while still providing useful statistical information. Technology companies including Apple and Google [11] have also implemented differential privacy in their data collection processes to enhance user privacy. These real-world examples demonstrate the versatility and effectiveness of DP in various sectors and applications, from government data collection to consumer technology. These implementations showcase the growing recognition and importance of DP in protecting user privacy across different domains.

DP is based on a rigorous mathematical foundation, which provides strong privacy guarantees, and its flexibility allows it to be applied and adopted across various data types and use cases. As concerns over data privacy grow, DP offers a robust solution to enable the beneficial use of data while safeguarding individual privacy. The mathematical foundations of DP ensure that privacy guarantees can be formally proved and trusted, making it a reliable choice for sensitive data applications. This mathematical foundation is crucial in establishing trust and reliability in DP solutions.

B. Local vs Global Differential Privacy

DP can be implemented using one of two main approaches - local differential privacy (LDP) and global differential privacy (GDP) [12]. Understanding the differences between these approaches is important in selecting the appropriate model for specific use cases. Both approaches offer unique advantages and are suitable for different scenarios, enhancing the versatility of DP implementations.

LDP ensures the privacy of an individual's data *at the source*, before any data collection or analysis takes place. Noise is therefore added to user data locally on their device, and only noisy data is sent to a central server. LDP is particularly useful when users do not fully trust a data collector. An example of LDP is Apple's implementation in iOS, where user data such as typing habits are anonymized locally before being sent to Apple's servers for analysis [13]. LDP is considered to provide strong privacy guarantees since the data is anonymized through the addition of noise before it leaves the user's device, ensuring privacy even if the central server is compromised. The LDP approach allows users to maintain control over their data privacy, enhancing security and trust in secure services used.

GDP on the other hand, applies the privacy mechanism on a centralized database/server side. In this model, a dataset is collected, stored and processed by a trusted entity and noise is added to the outputs of queries made on the database. This approach relies on the trustworthiness of the data collector to apply the noise and maintain privacy and anonymity properties. The U.S. Census Bureau's use of DP for the 2020 Census is an example of GDP, where noise is added to published statistics to protect individual respondents data. GDP allows for more accurate aggregate analysis since the noise is applied only once, after data collection, but it requires that users trust the central entity to handle their raw data securely. GDP can be the best option for large-scale data analytics, provided that the central entity is reliable, secure, cannot be compromised and more importantly can be trusted.

The two approaches of LDP and GDP can be visualised below in Figure 1.

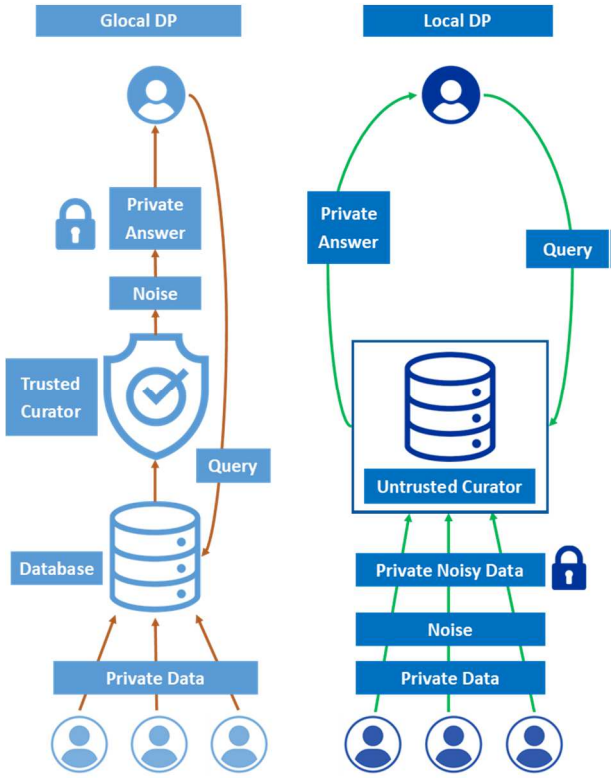

Figure 1: Global and Local Differential Privacy

In the ENCRYPT project, the LDP model is followed. This choice is motivated by the project's goal to empower users with control over their data privacy while ensuring robust protection against potential breaches. This choice is key in ENCRYPT's commitment to user-centric privacy solutions.

C. *Related Work*

ENCRYPT's use of DP - as will be presented in this paper, is a practical implementation of the foundational theoretical models pioneered by Cynthia Dwork and others. ENCRYPT's methodology is to incorporate theory and extend its applicability to practical, real-world scenarios, for various industrial sectors, and importantly for users of varying technical expertise.

In this section we explore related work, namely how other companies have implemented DP in their services.

Apple: Apple was among the first of the big technology companies to use DP for their products and services on a large scale [13]. DP is used by Apple to collect data from user devices, mainly for them to be able to analyse this data towards an improved user experience [14]. Such data stem from keyboards – mainly from smartphones, watches, fitness trackers and others [15].

Google: Google uses DP in various applications, including Google Maps and the Chrome browser [16]. For example, Google Maps uses DP to gather data - whilst still protecting the privacy of users, to provide information such as how busy a business is over the course of a day or how popular a particular restaurant's dish is in Google Maps.

Microsoft: Microsoft has incorporated DP into its products and services – such as the Azure platform and Microsoft 365 suite [17], to enhance data privacy and be compliant with regulations. The "SmartNoise" tool of Microsoft Azure allows developers to develop applications using DP.

Uber: Uber's "Elastic Sensitivity" tool implements DP in the data analytics process carried out by the company [18]. Protecting the privacy of individual trip records, it still allows for aggregated data analysis to be carried out to improve services and operations.

Commercially, while these technology giants have used DP, these applications mainly remain internal for their products and services. ENCRYPT on the other hand uses DP for it to be applied by researchers in various data oriented applications, in sectors such as healthcare and finance. This highlights ENCRYPT's focus on versatility of its solutions while of course complying to regulations such as GDPR.

III. ENCRYPT Use Case for Differential Privacy

The ENCRYPT project centers around three use cases - in the healthcare domain, cyber threat domain and fintech domain.

DP will mostly focus on the fintech use case which is a real-world application specifically chosen to demonstrate how PPTs can be effectively used in the financial sector, ensuring data privacy and compliance with stringent regulations such as GDPR. In this scenario, EXUS [19] serves as the service provider and data processor, while EPIBANK [20] functions as the data steward, managing data from its customers who are the data owners. This collaboration between service providers and data stewards exemplifies the practical application of DP in securing sensitive financial data.

The primary objective of the fintech use case is to demonstrate how ENCRYPT's platform can enable financial institutions to share and analyse their data securely, protecting sensitive customer information and proprietary financial models through the use of PPTs. These ensure that data remains encrypted or anonymized while allowing for valuable insights to be derived without exposing the underlying sensitive information. This enhances operational efficiency and service delivery and also builds customer trust by demonstrating a strong commitment to data privacy and security. By implementing DP, financial institutions can achieve a balance between data utility and privacy, facilitating informed decision-making without compromising customer trust.

In this use case, DP plays an important role in safeguarding the privacy of sensitive financial data. EPIBANK shares randomised and anonymized data with EXUS to train machine learning models that optimize debt collection services. This data includes detailed customer profiles, transaction histories, and demographic information. DP is therefore applied to ensure the anonymity and confidentiality of customer information when sharing data between external parties while still allowing for data analysis to take place.

Beyond this use case, the DP component of the ENCRYPT platform will also allow researchers and innovators from other fields to use DP in their work. Following a user-centric approach, the ENCRYPT recommendation engine will be able

to propose when DP is suitable for a user to use, and also suggest the value of ε to be applied. The ENCRYPT platform also provides an interface for randomness to be applied locally to a dataset before it can be uploaded to the ENCRYPT platform for training and testing of machine learning models. These features enable DP to be more usable and user friendly for non-experts and facilitates its adoption across diverse industry sectors and academic fields.

DP will also be tested on the healthcare use case of ENCRYPT, to explore if this technology is suitable for their datasets.

Potential enhancements for DP within ENCRYPT are presented in Section VI, where an exploration of future research directions are described.

IV. ENCRYPT PLATFORM

The ENCRYPT platform is a system designed to facilitate secure, efficient and scalable handling of sensitive information. It brings together several distinct software components, each utilizing different technologies to implement specific functionalities independently.

Specifically, it hosts all PPT technologies, other supporting technologies developed in the ENCRYPT project and provides the framework for these components to work together. This modular approach allows for flexible integration, where development cycles are divided into smaller modules, tested incrementally and integrated to deploy a cohesive platform. The modular design also ensures that the platform can adapt to evolving privacy requirements and technological advancements.

Main features of the ENCRYPT platform include:

- **The front-end service** - which hosts the user interface
- **Hosting and deployment** – upon Microsoft Azure Cloud and utilizing its SGX-enabled VMs for secure enclaves and other features needed within ENCRYPT
- **Interconnections and communication** – leveraging Azure Virtual Network for secure communication across the platform
- **APIs and message specification** – for inter-component communication, ensuring validation, consistency and reliability

These features, interfaced together create a robust and secure environment for data processing, ensuring that privacy-preserving measures are consistently applied and maintained.

A. ENCRYPT Recommendation Engine

The recommendation engine is one of the PPT supporting technologies within ENCRYPT, and it is a novel tool which has been developed to help users with the PPT technologies provided by ENCRYPT. Specifically, it can be used to recommend to users the PPT they should use for their data processing scenarios – based on a set of criteria. The recommendation engine therefore represents a novel feature which makes complex PPTs accessible and usable for a broader set of users.

The tool uses a sophisticated algorithm which leverages Artificial Intelligence (AI) to analyze the characteristics of the user's data, the intended processing activities and the associated privacy requirements. It considers factors such data sensitivity, data size, computational intensity, performance constraints, time constraints and computational constraints to make PPT recommendations that balance data utility with privacy and security. This AI analysis ensures that recommendations are suitable to the specific needs and constraints of each user, while considering optimization, performance and data utility.

The recommendation system is also designed to continuously update its knowledge base with the latest research findings, technological advancements and regulatory changes. This ensures that the recommendations consider current state of the art but are also forward-looking, which are able to adapt to future developments in the privacy domain. This updating capability allows the ENCRYPT platform to remain relevant and effective in light of rapid technological and regulatory changes.

The recommendation system is also designed in a user-centric approach, providing a justification for its recommendation tailored to the knowledge the user, which helps build trust in the system. By explaining its recommendations, the system enables user understanding and confidence, encouraging wider use of the platform and PPTs. This makes it more usable for non-experts and facilitates its adoption across diverse sectors, beyond fields of the ENCRYPT use cases.

Specifically, for the ENCRYPT DP component, upon users providing relative details to the recommendation engine, it will inform them of how appropriate DP is for them to use in their setting and will also suggest a value of ε to be used.

B. User Interface

The user interface of ENCRYPT, enables users to interact with the platform seamlessly. This interface supports functionalities such as data pre-processing, user authentication, and role management. Users can upload data files, which are processed through a pre-processing pipeline with metadata stored in a database for further use. The user-friendly design of the interface enables users with limited technical expertise to effectively utilize the platform's features.

Importantly for the ENCRYPT DP component, the user interface allows for users to add noise to their dataset before these are uploaded to the ENCRYPT platform. Specifically, users are able to select their datasets via the user interface and identify the ε amount of noise that should be added to the data.

It should be noted that the addition of noise takes place locally, with computation carried out on the user's device. Once this is complete, users are then able to upload their noised datasets to the ENCRYPT platform for training/testing of AI models or for carrying out analysis on a dataset on an already saved AI trained DP model. This local computation of noise addition enhances security by ensuring that raw data never leaves the user's device, further safeguarding privacy and trust in the system.

V. ENCRYPT DP DEPLOYMENT

The ENCRYPT DP component aims to train and test various AI models with a variety of optimizations for each, upon a training/testing dataset provided by a user. Details of the model which is found to be most accurate will be saved and closely tied to the user and their account. When the user wants to carry out data analytics on new datasets, the ENCRYPT platform will use the saved details of the most accurate model for data analytics purposes. This approach ensures that users can consistently achieve high-quality analysis results while achieving high security and privacy data-protection.

A. Initial Experiments and Results

In the first DP module implementation, we have conducted experiments applying the developed functionalities in the fintech use case, with synthetically generated data provided by EXUS.

Several AI algorithms have been tested, including random forests, decision tree classifiers and artificial networks. These experiments aim to evaluate the effectiveness of DP across different types of machine learning models.

The performance of DP models – where randomness is added to datasets, has been found to be high at around 88% which is just 4% lower than models upon datasets on their original state (where no noise was added). This also occurred when high levels of noise were added to datasets (low value of ε). These results demonstrate the potential of DP to maintain data utility even with significant privacy protections in place.

These experiments demonstrate that upon the dataset provided by our use case, DP can provide a high level of utility of data analysis while still being able to protect individual privacy. These results highlight the practicality of DP in real-world scenarios, validating its use in the ENCRYPT platform.

B. Workflow

When the DP component is fully developed and integrated to the ENCRYPT platform, we foresee the following workflow which users will follow:

a. The user will interact with the ENCRYPT recommendation engine and provide details of their computational setting

b. The recommendation engine will propose DP to be used (where appropriate of course) and will also suggest the ε value to be used on the user dataset(s)

c. The user will use the ENCRYPT platform user interface and add noise to the dataset – by specifying the proposed ε-value

   - It is reminded that this noise is added to the dataset locally on the user's system

d. The user will upload the noisy datasets to the ENCRYPT platform – though the user interface

e. The DP component will use these datasets and train various AI models – each with different parameters

   - Details of the most accurate model will be saved and associated with the user

f. At any future time, the user will apply the same amount of noise (ε-value) to any other dataset upon which they may want to carry out data analytics on

g. The DP component will use the saved configuration of the most accurate model upon the provided dataset and provide results to the user.

This workflow ensures a seamless integration of DP into users data analysis processes, enhancing both privacy and usability.

VI. CHALLENGES AND FUTURE DIRECTIONS

While DP offers robust privacy guarantees, its practical and scalable implementation still has a number of challenges. A significant issue is the trade-off between data privacy and data utility. As discussed, lower ε values with higher privacy guarantees often result in reduced data accuracy, which can be problematic for certain applications requiring high precision. Future research could focus on optimizing this trade-off to achieve better balance.

Another challenge is the scalability of DP deployments. As datasets grow larger and more complex, the computational requirements associated with adding noise and maintaining privacy can become prohibitive to use. Developing more efficient algorithms and use of advanced computational techniques including parallel programming and hardware acceleration could help address issues associated with scalability, energy and running time.

Furthermore, more user-friendly tools and interfaces that simplify the use of DP for non-experts are required. While the ENCRYPT platform makes significant contributions in this direction, further advancements and innovations are necessary to ensure that a wider range of users can easily use and apply these technologies in their work.

Interoperability is also another factor that needs to be considered. DP tools will need to seamlessly integrate with various data storage, processing, analysis systems and workflows. Ensuring this will enable for widespread adoption. Developing standard protocols and APIs can facilitate this integration and promote consistent implementation across different platforms.

As privacy regulations and standards evolve, DP must also adapt to meet new requirements. Ongoing collaboration with policymakers and regulatory bodies will be essential to ensure that PPTs remain compliant and relevant.

Future research could also explore novel applications of DP beyond traditional data analytics. For example, its use in emerging fields such as quantum computing, blockchain technology and the Internet of Things presents exciting possibilities. Investigating these applications can open up new avenues for protecting privacy in diverse technological landscapes.

While DP has made significant advancements, addressing these challenges and exploring future directions will be critical for its continued evolution and adoption. By overcoming these hurdles, we can unlock the full potential of DP to safeguard individual privacy in an increasingly data-driven world.

VII. Conclusion

Currently, the ENCRYPT DP component is still under development and integration to the ENCRYPT platform, so the described workflow can be made available to users.

Early experiments have shown that DP can be used with high amounts of noise added to dataset (low ε-values), thus ensuring data privacy while still maintaining data utility – as shown by the high accuracy of our initial results.

These promising initial results suggest that DP can effectively balance privacy and accuracy, making it a valuable tool for a wide range of applications.

The ENCRYPT platform with its recommendation engine and user interface greatly simplifies the process of using differential privacy, thus opening up the use of advanced security technologies to non-expert users for academic research and industrial innovation purposes.

By democratizing access to PPTs, the ENCRYPT project paves the way for more secure and privacy-conscious data analysis practices across various sectors and for them to be used by a wider range of users.

Acknowledgment

This work has received funding from the European Union's Horizon 2020 research and innovation programme - ENCRYPT project under Grant Agreement no. 101070670.

References

[1] Reza Shokri and Vitaly Shmatikov. Privacy-preserving deep learning. In Proc SIGSAC conference on computer and communications security, 2015.

[2] European Parliament and Council of the European Union. (2016). General Data Protection Regulation (GDPR). Official Journal of the European Union, L119, 1-88.

[3] ENCRYPT Project: https://encrypt-project.eu/

[4] Craig Gentry. Fully homomorphic encryption using ideal lattices. In STOC, volume 9, pages 169–178, 2009.

[5] James Menetrey, Marcelo Pasin, Pascal Felber, Valerio Schiavoni, Giovanni Mazzeo, Arne Hollum, and Darshan Vaydia. 2023. A Comprehensive Trusted Runtime for WebAssembly with Intel SGX. IEEE Transactions on Dependable and Secure Compu-ting (2023).

[6] Coppolino, L., D'Antonio, S., Mazzeo, G., Romano, L., & Sgaglione, L. (2022). PriSI-EM: Enabling privacypreserving Managed Security Services. Journal of Network and Computer Applications, 203, 103397.

[7] C. Dwork. Differential privacy. In Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)(2), pages 1–12. 2006.

[8] Ú. Erlingsson, V. Pihur, and A. Korolova, ''RAPPOR: Randomized aggregatable privacy-preserving ordinal response,'' in Proc. ACM SIGSAC Conf. Comput. Commun. Secur., Nov. 2014, pp. 1054–1067.

[9] Cynthia Dwork. Differential privacy and the us census. In Proceedings of the 38th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, pages 1–1, 2019.

[10] Christopher T Kenny, Shiro Kuriwaki, Cory McCartan, Evan TR Rosenman, Tyler Simko, and Kosuke Imai. 2021. The use of differential privacy for census data and its impact on redistricting: The case of the 2020 US Census. Science Advances 7, 41 (2021), eabk3283.

[11] Differential Privacy Codelabs (accessed 31 May 2024). https://gsec-onair.withgoogle.com/events/codelab

[12] Shiva Prasad Kasiviswanathan, Homin K. Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. 2011. What can we learn privately? SIAM Journal on Computing 40, 3 (2011), 793–826.

[13] Differential Privacy Team, Apple. Learning with Privacy at Scale. https://docs-assets.developer.apple.com/ml-research/papers/learning-with-privacy-at-scale.pdf

[14] Apple. Differential Privacy Overview. https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

[15] Apple in collaboration with Stanford University. Protection Against Reconstruction and Its Applications in Private Federated Learning. https://machinelearning.apple.com/research/protection-against-reconstruction-and-its-applications-in-private-federated-learning

[16] Google. Enabling Developers and Organizations to Use Differential privacy. https://opensource.googleblog.com/2019/09/enabling-developers-and-organizations.html

[17] Microsoft. How differential privacy enhances Microsoft's privacy and security tools: SmartNoise Early Adopter Acceleration Program Launched. https://blogs.microsoft.com/on-the-issues/2020/12/10/differential-privacy-smartnoise-early-adopter-acceleration-program/

[18] Joe Near. Differential privacy at scale: Uber and berkeley collaboration. In Enigma 2018 (Enigma 2018), Santa Clara, CA, 2018. USENIX Association.

[19] EXUS AI Labs, https://www.exus.ai/

[20] Cooperative Bank of Epirus. https://www.epirusbank.com/