A Practical and Scalable Privacy-preserving Framework

Nikos Avgerinos
Diadikasia Business Consulting S.A
Athens, Greece
navgerinos@diadikasia.gr

Salvatore D'Antonio
TRUSTUP
Naples, Italy
salvatore.dantonio@trustup.it

Irene Kamara
Tilburg Institute for Law, Technology, and Society
Tilburg Law School
Tilburg, The Netherlands
i.kamara@tilburguniversity.edu

Christos Kotselidis
Department of Computer Science
The University of Manchester
Manchester, United Kingdom
christos.kotselidis@manchester.ac.uk

Ioannis Lazarou
EXUS
Athens, Greece
g.lazarou@exus.ai

Teresa Mannarino
Department of Advanced Biomedical Sciences
University of Naples Federico II
Naples, Italy
teresa.mannarino@unina.it

Georgios Meditskos
School of Informatics
Aristotle University of Thessaloniki
Thessaloniki, Greece
gmeditsk@csd.auth.gr

Konstantina Papachristopoulou
Eight Bells Ltd
Athens, Greece
konstantina.papachristopoulou@8bellsresearch.com

Angelos Papoutsis
Information Technologies Institute, CERTH
Thessaloniki, Greece
apapoutsis@iti.gr

Paolo Roccetti
R&D Labs
Engineering Ingegneria Informatica
Rome, Italy
paolo.roccetti@eng.it

Martin Zuber
Universite Paris-Saclay, CEA, List,
Palaiseau, France
martin.zuber@cea.fr

***Abstract*— ENCRYPT is an EU funded research initiative, working towards the development of a scalable, practical, adaptable privacy-preserving framework, allowing researchers and developers to process data stored in federated cross-border data spaces in a GDPR-compliant way. ENCRYPT proposes an intelligent and user-centric platform for the confidential processing of privacy-sensitive data via configurable, optimizable, and verifiable privacy-preserving techniques. Hence, ENCRYPT builds on top of cutting-edge technologies such as Fully Homomorphic Encryption, Secure Multi-Party Computation, Differential Privacy, Trusted Execution Environment, GPU acceleration, knowledge graphs, and AI-based recommendation systems, making them configurable in terms of security and, most importantly, performance. The ENCRYPT framework is being designed taking into consideration the needs and preferences of relevant actors and will be validated in realistic use cases provided by consortium partners in three sectors, namely healthcare (oncology domain), fintech, and cyber threat intelligence domain. This position paper provides an overview of ENCRYPT by presenting project objectives, use cases, and technology pillars.**

Keywords*— *Differential Privacy, Fully Homomorphic Encryption, Trusted Execution Environment

I. Introduction

Several technologies and tools exist in literature to allow for privacy-preserving data processing. However, they are not largely used yet in application domains due to some limitations and constraints. Fully homomorphic encryption, for example, despite being versatile in allowing various computations over federated sets of encrypted data, it suffers from a significant performance degradation as the amount of data to be processed increases, while complex calculations in large-scale deployments take a significant toll when multi-party computation methods are used. Another reason for the lack of uptake of such technologies is related to their user-friendliness, both for researchers and service providers, as well as for data owners. The type and configuration of the privacy-preserving technology to be used, as well as the level of privacy required for a given dataset and a given output, is often unclear to all parties involved. This is further exacerbated by the fact that not all relevant actors are aware of the legal and technical terms used in guidelines related to the privacy requirements of certain types of data.

In this paper we present how the ENCRYPT research project aims at addressing the challenge of maximizing the exploitation of big data available in several sectors, such as health, communication, finance, while preserving privacy, since those data could contain sensitive information and are subject to several data protection laws.

II. Rationale and Contribution

Existing privacy-preserving technologies, such as Homomorphic Encryption (HE), Secure Multi-Party-Computation (SMPC), Trusted Execution Environment (TEE)

or Differential Privacy (DP) even if promising at a small-scale level still need to overcome several limitations in order to become mainstream security solutions. Moreover, none of the aforementioned privacy-preserving techniques can be used as a single standalone security mechanism. In most cases a combination of them has to be deployed to cover the full spectrum of possible cyber-threats, while taking into account the regulations needed by the end-users and the established infrastructures. On the other hand, a large amount of big data is available nowadays to be used for addressing new challenges and developing better research and digital services. However, the major impediment in the processing of these data, that usually contain sensitive or personal information, lies in the risk of cyber-security attacks and/or in data breaches and misuses. The regulations on data protection and the EU's high norms and laws on personal data, such as the General Data Protection Regulation [2], pose additional obligations and safeguards to be taken into consideration while storing and processing personal data. In order to address the above issues, the advanced privacy-preserving computation technologies, such as FHE, SMPC or DP can provide valid GDPR-compliant solutions once they become more scalable and reliable - i.e., ready for realistic scenarios.

In the following we illustrate the main limitations affecting different privacy-preserving methods and describe how the project intends to go beyond the state-of-the-art to make them applicable in real-world use cases involving a high volume of sensitive data. Both the FHE and SMPC solutions for privacy-preserving of data in use have scalability issues when dealing with a lot of data. While FHE has a high computational overhead to treat the encrypted data, SMPC requires a high communication cost for the secret sharing. Another common limitation is that their integration with the existing networking infrastructure and security protocols is a neglected aspect of the ongoing research. The DP technique requires a predefined privacy budget that linearly depends on a fixed number of queries. This can impact its utility in practice, thereby making it complex to apply DP in adaptive settings. Finally, SGX-based TEE provides secure computing, but for small workloads. This weakness can make TEE difficult to apply to large-scale aggregated computations that involve the input of many users (large overhead due to the limited paging).

The ENCRYPT project goes beyond the state-of-the-art to overcome the limitations of these privacy-preserving technologies in several aspects. First of all, it addresses the scalability issue by going beyond the single-key FHE paradigm and explore the application and the practicability of new multi-key and threshold FHE schemes especially in a federated context. Second, to address the drawbacks of each technology in terms of the covered cybersecurity threats and performance, ENCRYPT investigates the combinations of several of these privacy-preserving methods. Third, ENCRYPT addresses the slow computation times associated with the existing solutions for privacy-preserving technologies based on HE or SMPC, by providing hardware acceleration in a user-friendly way. Fourth, ENCRYPT looks at the necessary methods in order to make these advanced privacy-preserving data processing technologies more suitable to interoperate with existing infrastructures and traditional security mechanisms. In particular, it will investigate the use and the application of the transciphering method for the FHE, allowing to switch from "traditional" symmetric encryption to a homomorphic one, without the need to decrypt the sensitive data. This powerful method will permit not only to keep standard symmetric cryptography on the clients' terminals, but also to reduce the bandwidth requirements for exchanging encrypted data, thus ultimately addressing scalability issues. Since a major impediment in the adoption of these privacy-preserving data processing technologies is the lack of "user-friendliness", ENCRYPT also provides (a) a privacy risk assessment methodology supporting adopters in evaluating privacy risks, and to link them to cyber vulnerabilities they often depends on; (b) an AI-based recommendation system allowing to choose one or a combination of those technologies and to configure them in order to meet system requirements and the identified needs in terms of protection of the users' and personal data and of performance. Finally, the proposed solutions are being developed and will be validated in several settings and real-world use cases including the challenging cross-border federated processing of large datasets.

III. ENCRYPT USE CASES

A. *Medical Use Case*

Cooperative oncology involves different specialists from different medical disciplines evaluating and analyzing the same patient from different perspectives. This leads to large amounts of medical data being shared in real time and across different hospitals, leading to data protection and privacy preservation issues.

Patients are registered to the hospital system, clinical data are collected and imaging procedures are performed. Once the images and the clinical data have been evaluated, the physician compiles a report. The physician might need a consultation with a colleague for a second opinion, or a consultation among different specialists might be useful to determine the best therapeutic option for the patient. The different professional figures involved may not be in the same hospital/institution. Similarly, in Radiotherapy Unit, different professional figures have access to patient data and imaging files. Constant data exchange is essential to plan the treatment and verify that the scheduled treatment is appropriately delivered.

Personal and sensitive data are currently accessible to all the health-care professionals and researchers involved. At the state of the art, data protection measures include removing sensitive data by implementing a non-standardized anonymization procedure. Moreover, currently the only available option to share patients' data and images is data transfer, considering that hospital systems are not accessible by external subjects.

The ENCRYPT platform will allow radiologists to protect data communication between different health care professionals in other departments/hospitals, while also preserving data integrity and patient privacy. The ENCRYPT framework is expected to enable the processing of the data prior to the sharing with different parties, ensuring the preservation of data confidentiality and integrity, and compliance with data protection law. This will permit the use of the data in the subsequent steps of the diagnostic/therapeutic pathway and during patient's follow-up consultations. In the Radiotherapy

setting, using the ENCRYPT platform, the integrity of treatment planning data and images metadata will be ensured. Moreover, all the data shared to any other stakeholder through this service/tool will comply with GDPR, without the risk that such data can lead to the actual identification of a real patient.

B. Fintech Use case

Just like any other emerging technological domains, data security and privacy preservation mechanisms ensuring the anonymity of the sensitive data are of paramount importance to the Fintech domain as well. In the ENCRYPT's Fintech use case, two main challenges will be tackled during the project's pilot activities, implemented in two different sub-use cases.

Sub-use case 1 – Security and Impact Assessment of data owned by financial institutions: The first Fintech sub-use case tackles the need of the financial institutions and banks to be sure about the security and privacy levels required to ensure the anonymity of their clients not only internally to their organization, but also when they share their data to potential 3rd parties to perform data analysis and/or to deliver tailored software solutions for the bank's activities. More specifically, in the majority small scale financial institutions and banks, there is a need to assess the security and privacy levels of the data they own with respect to their clients. These data comprise among others: a) personal data of the client itself including sensitive data (name, ID, social security address, income, bank accounts, etc.), b) history of actions made by the client (transactions, calls, payments, timings in payments, etc.), c) actions made by the bank towards the client (phone call history, transactions, etc.) and d) specific services packages tailored to each client. All of these data tend to increase over time and more attributes linked to each client are added so that more sophisticated product services are offered to the clients by the bank. Especially for small size financial institutions it is of a paramount importance their Data Protection Officers (DPOs) to be able to assess the current security and privacy level of the data they own in order to perform an assessment of potential risks and impact to the organization itself in case of an incident that might compromise the security and anonymity of the clients' sensitive data. Moreover, in case these data are either processed internally or shared to external 3rd parties for further processing, compliance only to GDPR currently applied in these organizations is not enough, since information linked with a client might potentially lead to the actual identification of person or/and sensitive information linked with that person. In this scenario the DPO of EPIBANK will exploit the functionalities of the ENCRYPT recommendation system, where it will be informed about the type of data owned by its organization and their level of security and privacy. The DPO will be also informed about the potential security impact it might occur to its organization in case there is a security issue on its owned data. Additionally, the ENCRYPT recommendation system will inform the DPO about the available privacy-preserving technologies that can be applied to its use case presenting also the potential trade-offs in terms of system's performance. Finally, the overall ENCRYPT platform will provide the option to select one of the ENCRYPT's privacy-preserving method to be applied to the EPIBANK's data to be shared with external 3rd party entities so that they can perform data processing and deliver to the EPIBANK tailored software solutions that will help the bank on its dept collection policies.

Sub-use case 2 – Training of tailored-made AI models with data shared by financial organization to 3rd parties: The second sub-use case describes the situation when the 3rd party/entity receives these data from the financial institutions/banks and wishes to perform AI-driven data analysis in order to deliver tailored software solutions serving the strategies and policies of the bank in specific business portfolios. More specifically, in the ENCRYPT project EXUS, as an AI-software house delivering solutions to financial organizations to manage their debt collection portfolio, will develop tailored AI-models using the EPIBANK's clients data in order to perform among others client stratification, behaviour forecasting and overall scoring of the bank agents responsible for handling each client. Since these functionalities are tailored specifically to the needs of EPIBANK, clients' historical data over a long period of 0.5-1 year have to be delivered to EXUS, so that EXUS will be able to develop and train its AI-models.

To ensure that there is not a possibility that the actual identity of the person or sensitive data related to this person could lead to its actual identification and since simple pseudonymization techniques as mentioned above are not sufficient, EPIBANK will deliver the data needed, after one of the ENCRYPT's privacy-preserving methods will be applied to the data. Since the privacy-preserving method selected will have a direct impact on the AI model's development and training phases as well as to the performance/accuracy of the models developed, all ENCRYPT's privacy-preserving methods have to be tested and validated so that the optimum solution will be opted.

C. Cybersecurity Use Case

Over time, cybercriminals are constantly improving their techniques and strategies for launching cyber-attacks, becoming more advanced and sophisticated. In addition, it is well known that the ever-evolving cyber-attack landscape leads to an array of new and varied attacks. To defend against such threats, organizations rely solely on their data (e.g internal logs), which may not be sufficient in detecting and responding to diverse cyber threats on time. Cyber Threat Intelligence (CTI) can gather information about these attacks and sharing this knowledge can improve the understanding of potential threats and strengthen defence strategies for individuals and organizations. However, organizations are reluctant to share information due to concerns about exposing confidential data. To this extent, the CTI use case will tackle the above objectives and concerns. A CTI gathering, extraction and sharing tool will be used to collect, combine, and correlate data from various internal (i.e., data from ENCRYPTs' end-users) and external sources, such as social media platforms and vulnerability databases. During the use case, different techniques will be used such as extraction of Indicators of Compromise (IoCs), correlation using the MISP (Malware Information Sharing Platform) correlation engine [1] (simple correlation), Exploratory Data Analysis (EDA) and Machine Learning analysis for advanced correlation (using internal and external data). Thanks to the CTI tool, personal information such is automatically rectified during data gathering. The CTI Extraction component minimizes and anonymizes the data, while the CTI correlation component pseudonymizes the data from ENCRYPT's data providers and external sources. The CTI sharing component provides a secure platform for sharing information with interested parties while maintaining data owner

privacy. However, regarding the previous techniques, there is still a risk of data privacy violation during the data gathering and CTI sharing processes. Nevertheless, implementing strong privacy-preserving techniques can address these concerns by securing the privacy of data holders. During the CTI use case scenario, data providers will use ENCRYPT's privacy-preserving techniques to anonymize their data before sending it to the data processor for the extraction of CTI.

IV. Legal Considerations

ENCRYPT delivers a privacy-preserving framework, which aims at incorporating an ethically and legally aware design. Data protection by design is a legal obligation - Art. 25 GDPR [2], which entails that organizations and entities implement technical and organizational measures to integrate data protection principles, such as data minimization and purpose limitation, in the manipulation and processing of personal data. In addition, security by design is a key component of the new proposed law, the Cyber Resilience Act (CRA). Both design obligations contribute to enhancing trust and security, allowing for the processing of data, while respecting privacy and data protection. The ENCRYPT user-centric framework allows data owners to comply with those obligations, by processing data, without sharing sensitive or other information to unauthorized parties. ENCRYPT offers a scalable privacy solution by allowing the users to determine the purposes of the processing and their needs, which is in line with Art. 5(2) GDPR and the accountability principle. In addition, the ENCRYPT framework, and its recommendation engine will integrate by default legal and ethical risks, alongside risks of technical nature, allowing for a warning system providing the trade-offs between privacy and data exploitation, when envisaged uses of the data, are not meeting the thresholds and requirements imposed by the GDPR.

V. Framework Architecture

ENCRYPT proposes an intelligent and user-centric framework (Fig. 1) for the confidential processing of privacy-sensitive data via configurable, optimizable, and verifiable privacy- preserving techniques and its overall architecture is given in the figure below.

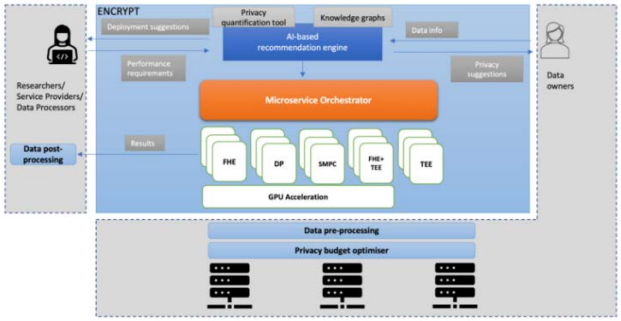

Fig. 1: The ENCRYPT framework architecture

ENCRYPT leverages, improves, and complements technologies and cryptographic schemes that represent the current state-of-the-art in the field of data-in-use protection, that is: Fully Homomorphic Encryption, Secure Multi-Party Computation, Differential Privacy, Trusted Execution Environment. ENCRYPT builds on top of these techniques making them configurable in terms of security and performance. It intelligently uses them by also combining their intrinsic security mechanisms to mitigate limitations and take advantage of their benefits. Since in most cases, performance represents the Achilles' heel of solutions for privacy-preserving computing, ENCRYPT provides a transparent to the user GPU-based acceleration service that is capable of being used in conjunction with each technology or cryptographic scheme that can exploit parallel processing of confidential data.

The platform also encloses intelligent units to adjust itself to users' and data requirements. More specifically, it comes with an AI-based recommendation system, which provides data owners with recommendations on the privacy level necessitated for their data, and data processors with suggestions on the deployment and configuration of privacy-preserving technologies based on the type of data they want to process. Taking into account the lack of familiarization with privacy-preserving technologies, the ENCRYPT recommendation system will provide intelligible suggestions tailored to the needs of different types of users. ENCRYPT also provides data owners with a prototype that streamlines the continuous assessment of privacy risks for a given personal data processing.

The knowledge graph building tool assists data owners and data stewards in standardising their datasets, so that their processing is performed in a GDPR-compliant way, based on the assessment of the ENCRYPT recommendation system. On the data processors side, a privacy quantification tool assists data processors in identifying the level of privacy offered by specific configurations of privacy-preserving technologies, as well as the level of privacy required for the data they want to process. As a result, the recommendation system will match specific privacy requirements imposed by the types of data to be processed, with a specific deployment and configuration plan for the type of processing the data processor wants to perform.

For the optimised deployment of privacy-preserving technologies, a data pre-processing module developed within ENCRYPT will process data and configure datasets in a way that facilitates the execution of the privacy-preserving technology selected. Automation of this process will ease the burden of configuration from the data owners or data stewards. For the HE case, this component will analyse the data/features to identify which contribute more to the objective of a given analysis and encrypt only those. This will allow users to experience less computational overhead, improving scalability. Similarly, for the case of DP, a privacy budget optimiser will analyse data to identify the privacy protection – data utility curve for each specific dataset. This will allow data owners and stewards to make an informed decision on the level of noise added in their datasets to ensure adequate levels of privacy, while minimising detrimental effects to data utility. Similarly, on the data processor side, and in order to facilitate the deployment of privacy-preserving technologies, ENCRYPT will develop a lightweight data post-processing module in order to allow the easier recuperation and exploitation of the results of execution for the selected privacy-preserving technologies. In an indicative ENCRYPT user journey, the data owner or data steward provides the recommendation system with information about the

data they have, with the support of the Knowledge Graph tool that supports interoperability between different data types. The recommendation system responds with suggestions on the level of privacy needed, and technologies supporting this level of privacy. The service provider/data processor provides other requirements on the expected performance and functionality. This information is sent via dedicated ENCRYPT UIs to the AI-based recommendation system, which selects the most appropriate technology (or combination of technologies), as well as the configuration of privacy parameters. The Microservice orchestrator (the project will consider Kubernetes and OpenStack for the orchestration) – based on the inputs coming from the recommendation system – deploys the hosting machines with related microservices to support the selected ENCRYPT solutions, and to meet possible constraints posed by the data owner and data processor. On the data owner side, the ENCRYPT data pre-processing service is executed to format and configure data appropriately for the selected privacy-preserving technology to be used. When DP is selected as the privacy-preserving technology to be deployed for a given application, the privacy budget optimiser is responsible for configuring the noise levels, as calculated by the AI-based recommendation engine. Similarly, for SMPC, this service splits data in pieces to be stored in distributed servers, and in FHE it homomorphically encrypts only the necessary data. During runtime, the orchestrator monitors performance metrics, and spawns new services to ensure that the Service Level Agreements (SLAs) with regard to delay are met. The data post-processing service, deployed on the side of the entities interested by the outputs of the selected privacy-preserving technologies (either researchers, service providers or data processors) helps them recuperate the results in an easily understandable and readable form (e.g., if FHE was selected, a decryption and decoding of the result is needed; for SMPC a reconstruction of the final result is required, etc.)

VI. Technology Pillars

This section provides an overview of the ENCRYPT privacy-preserving technologies.

A. Virtual Secure Enclave

Trusted Execution Environment and Homomorphic Encryption are widely accepted technologies for privacy-preserving data computing. However, they present some drawbacks. The capability of HE to perform operations on encrypted data does not come for free. Not only HE results in an increase of the execution time, but it is also affected by the Ciphertext Expansion (CTE) phenomenon, which is not negligible. Another drawback of HE is its poor scalability with respect to the multiplicative depth of the circuit being evaluated. We call this problem the *noise explosion* problem. As the number of successive multiplications over the same ciphertext grows, so does a noise inside the ciphertext that was introduced at encryption time for security purposes. When this noise increases beyond a certain point, the ciphertext is undecipherable. This forces the client to use higher parameters which in turn have a non-linear, negative impact on performance. As far as TEE is concerned, the limited memory space available in secure enclaves (e.g., 128MB for Intel SGX). The Limited Memory Size (LMS) issue might result in a serious limitation for the data processing requirements of many memory-intensive applications. ENCRYPT proposes the concept of Virtual Secure Enclave that combines TEE and HE in an effective way, providing protection from privileged-user memory-access violations, while also limiting (i.e., as compared to their use in isolation) the inherent drawbacks of the two techniques [3] [4]. In this approach, the computing activity to implement HE is moved from the customer's local host to the TEE of the cloud provider: the server's TEE receives data from the field –no more affected by CTE– only after a TLS-secured channel is created. Virtual Secure Enclave leverages remote attestation features of TEEs to establish a trusted communication channel from the client to the host of the server farm acting as an Ingress node to the hosting platform. Once data arrives at the trusted host, it is homomorphically encrypted. Then, it is moved outside the TEE for further processing. This faces the LMS issue since processing of homomorphically encrypted data is performed on off-premises machines with no TEE and with full memory resources available. Finally, since the HE key is kept secret within the TEE, we can refresh the noise of homomorphic ciphertexts within the TEE of the attested server. In order to protect the data that is decrypted by the TEE, the server can add a random mask to the data before sending it over to the TEE for processing. This mask can then be removed homomorphically afterward. The advantage of such a solution is twofold: i) it allows to refresh ciphertext noise within the TEE of the attested server, thus mitigating the HE noise explosion problem due to frequent client-assisted decryptions; and ii) it allows to manage the TEE limited memory size since HE data processing is performed on the untrusted platform with no TEE and with full memory resources available.

B. Homomorphic Encryption

Homomorphic encryption has been longtime considered the Holy Grail of modern cryptography since it allows performing computation directly over encrypted data. Nowadays, the research field has become more mature and we dispose of several stable homomorphic schemes (BGV [5], BFV [6], CKKS [7], TFHE [8]) each one with its strengths and weaknesses, a theoretical framework (Chimera [9]) to switch between these cryptosystems, different available open-source libraries (e.g. OpenFHE [1], SEAL [2]) as well as dedicated optimization techniques and compilers (e.g. Cingulata [3]). However, in order to take it a step longer and deploy the homomorphic encryption as a complete security solution for real-work application, there is a need to further research for better scalability and performances. In parallel with the investigation of the combination of the homomorphic encryption with TEE, one of the objectives of ENCRYPT project is to conceive and implement protocols for federated

[1] https://github.com/openfheorg/
[2] https://github.com/microsoft/SEAL
[3] https://github.com/CEA-LIST/Cingulata

data processing based on homomorphic encryption with support for multi-users (e.g. threshold homomorphic encryption). Another step forward to the deployment of HE at a larger scale requires efficient ways of transciphering, i.e. combine homomorphic computation with exiting, lighter symmetric encryption. This will be another direction which will be investigated by ENCRYPT team, since this allows to improve the scalability of the overall system by diminishing the memory required for the encryption and the transmission of the private, sensitive data to the processing server. Another line of research consists in increasing the performances of non-linear homomorphic operators by using the functional bootstrapping, an interesting method offered by TFHE cryptosystem.

C. Differential Privacy

Differential Privacy [10] has several properties that make it particularly useful in applications such as those envisioned by the ENCRYPT project: composability, group privacy, and robustness to auxiliary information. Composability enables modular design of mechanisms: if all the components of a mechanism are differentially private, then the same holds for their composition. Group privacy implies graceful degradation of privacy guarantees if datasets contain correlated inputs, such as the ones contributed by the same individual, while robustness to auxiliary information means that privacy guarantees are not affected by any side information available to the adversary. In principle, during the iterative method for optimizing the objective functions in AI systems, this methodology averages together multiple updates induced by training data examples, clips each of these updates, and adds some kind of noise (e.g., Gaussian) to the final average. Based on the above, ENCRYPT will offer the expansion of privacy-preserving libraries for Deep Learning with additional guarantees that can usefully strengthen the protections offered by other privacy techniques. ENCRYPT will offer to the user a tool for the optimization and tuning of the "privacy budget" parameter, which comes in differential privacy techniques to quantitatively dial up or down the privacy guarantee [11]. Additionally, ENCRYPT will make good use of the "strength" that differential privacy is parametric: privacy does not come for free. Asking for more privacy will have the cost of either diminished data utility or the need to collect more data.

ACKNOWLEDGMENT

This work is supported by the European Union's Horizon Europe programme under grant agreement No 101070670 (ENCRYPT). In addition, this work is funded by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee (10039809).

REFERENCES

[1] Wagner, C., Dulaunoy, A., Wagener, G., and Iklody, A., "MISP: The Design and Implementation of a Collaborative Threat Intelligence Sharing Platform", Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security, Vienna, Austria, 2016, pp. 49–56.
[2] GDPR: Regulation 679/2016 OJ L 119, 4.5.2016, p. 1–88
Cyber Resilience Act: Proposal for a Regulation COM(2022) 454 final EDPB: Guidelines 4/2019 v. 2 (Oct. 2020).
[3] Coppolino, L., D'Antonio, S., Formicola, V., Mazzeo, G., Romano, L., "VISE: Combining Intel SGX and Homomorphic Encryption for Cloud Industrial Control Systems", IEEE Transactions on Computers, Volume 70, Issue 5.
[4] Coppolino, L., D'Antonio, S., Mazzeo, G., Romano, L., Sgaglione, L., "Securing the weak link of federated systems via trusted execution: A case study from the eHealth domain", International Journal of Critical Computer-Based SystemsVolume 9, Issue 4, Pages 293 – 3172019.
[5] Brakerski, Z., Gentry, C., and Vaikuntanathan, V., "(Leveled) Fully Homomorphic Encryption Without Bootstrapping", Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12. 2012. 309–325.
[6] Fan, J. and Vercauteren, F., "Somewhat practical fully homomorphic encryption", IACR Cryptology ePrint Archive, 2012:144.
[7] Cheon, J. H., Kim, A., Kim, M., and Song, Y., "Homomorphic encryption for arithmetic of approximate numbers", Cryptology ePrint Archive, Rapport 2016/421.
[8] Chillotti, I., Gama, N., Georgieva, M., and Izabachène, M., "Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds", Advances in Cryptology–ASIACRYPT 2016. Springer., 2016. Proceedings, Part I 22, pages 3–33.
[9] Boura, C., Gama, N., Georgieva, M. and Jetchev, D., "CHIMERA: Combining Ring-LWE based Fully Homomorphic Schemes" https://eprint.iacr.org/2018/758
[10] Dwork, C., "Differential privacy" in Automata, Languages and Programming, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Springer Berlin Heidelberg, 2006, pp. 1–12.
[11] Van der Veen, K. L., Seggers, R., Bloem, P., and Patrini, G., "Three tools for practical differential privacy", 2018. arXiv preprint arXiv:1812.02890. [Online]. Available: https://arxiv.org/abs/1812.02890