Towards Privacy-Preserving Deep Learning based Medical Imaging Applications

Anamaria Vizitiu*†, Cosmin Ioan Niță*†, Andrei Puiu*†, Constantin Suciu*†, Lucian Mihai Itu*†
*Department of Automation and Information Technology
Transilvania University of Braşov, Braşov, Romania
†Corporate Technology
Siemens SRL, Braşov, Romania

***Abstract*—Motivated by state-of-the-art performances across a wide variety of areas, over the last few years Machine Learning has drawn a significant amount of attention from the healthcare domain. Despite their potential in enabling personalized medicine applications, the adoption of Deep Learning based solutions in clinical workflows has been hindered in many cases by the strict regulations concerning the privacy of patient health data. We propose a solution that relies on Fully Homomorphic Encryption, particularly on the MORE scheme, as a mechanism for enabling computations on sensitive health data, without revealing the underlying data. The chosen variant of the encryption scheme allows for the computations in the Neural Network model to be directly performed on floating point numbers, while incurring a reasonably small computational overhead. For feasibility evaluation, we demonstrate on the MNIST digit recognition task that Deep Learning can be performed on encrypted data without compromising the accuracy. We then address a more complex task by training a model on encrypted data to classify X-ray coronary angiography views. These results underline the potential of the proposed approach to outperform current solutions by delivering comparable results to the unencrypted Deep Learning based solutions, in a reasonable amount of time. Lastly, the security aspects of the encryption scheme are analyzed, and we show that, even though the chosen encryption scheme favors performance and utility at the cost of weaker security, it can still be used in certain practical applications.**

***Index Terms*—Homomorphic encryption, Deep Learning, medical data, privacy**

I. Introduction

In recent years machine learning algorithms, and specifically Deep Neural Networks, have shown promising results in delivering personalized medicine, allowing for tailored diagnosis, treatment planning and disease prevention [1]. Since Deep Neural Networks have the ability to learn from past observations, they represent an attractive solution for integrating the knowledge and experience of medical experts into Computer-aided Detection (CADe) solutions.

Machine learning relies extensively on existing and future patient data to deliver accurate and reliable results [2]. However, among all types of data associated with an individual, medical data has some of the highest privacy requirements. Thus, as access to sensitive plaintext data is required in deep learning based applications, privacy and security concerns have been raised [3]. Moreover, the currently adopted regulations towards confidentiality guarantees for personal data manipulation (e.g. GDPR in EU, HIPAA in USA) urges for the adoption of more effective privacy-preserving techniques.

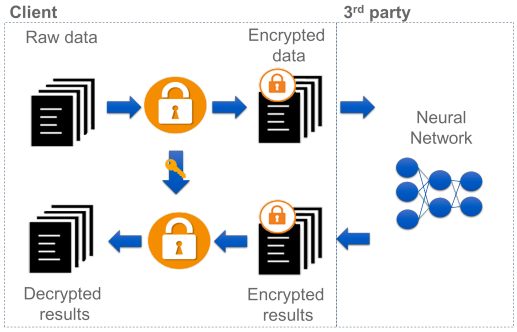

Fig. 1. Workflow of a privacy-preserving deep learning based application relying on homomorphic encryption.

Typically, to export sensitive data without compromising privacy, proper anonymization has to be performed. Thus, some of the data properties are modified, leading to a trade-off between privacy and utility. To address this limitation, herein we rely on a specific form of encryption, called homomorphic encryption, which represents a promising solution for guaranteeing privacy while still maintaining full utility. Specifically, the chosen homomorphic encryption scheme (MORE) [4] enables a limited set of operations to be performed directly on encrypted data, without having to reveal the underlying data or the encryption key. This ensures that both data and predictions remain private and data is analyzed in its encrypted from.

This property is particularly useful in the context of deep learning solutions. As outlined in Figure 1, privacy is preserved at three levels: (i) during training, when the external party (e.g. a cloud or processor) processes directly ciphertexts, (ii) during inference, when the patient's data remains confidential: algorithm receives as input ciphertexts and outputs ciphertexts, which are revealed only on the client side after performing the decryption, and (iii) the external party's deep learning model remains confidential. Consequently, the secure processing of medical data is performed in such a way that the external party cannot derive knowledge from the data, and the user is unable to obtain information regarding the machine learning model.

Driven by the difficulties that arise in practice when em-

ploying deep learning over encrypted data and also by the inefficiency of current solutions, herein we propose a method that increases the efficiency of the encrypted models in real-world applications by enabling: (i) computations over rational numbers, (ii) faster operations and, (iii) results comparable to those obtained with the unencrypted model. We assess the feasibility of the proposed solution for delivering reliable results, and show that performance is not lost when deep neural networks operate on data encrypted using the MORE homomorphic encryption scheme. We evaluate the privacy-preserving deep learning algorithms on the classic benchmarking application of digit classification, and on a personalized medicine application.

II. Related work

Recent advances in homomorphic encryption have lead to several encryption schemes, with different properties and constraints. The most notable drawback of the majority of fully homomorphic encryption (FHE) schemes is that each operation adds noise to the underlying message, therefore limiting the overall number of operations that can be performed without losing too much accuracy. Furthermore, to the best of our knowledge, there is no currently available partially or fully homomorphic encryption scheme that can process rational numbers (only integer numbers are supported). As a consequence, a variant of a matrix-based method, called MORE (Matrix Operation for Randomization or Encryption) [4] was adapted in the current work. Compared to currently studied schemes, in the context of privacy-preserving networks [5], [6], [7], MORE is noise free (unlimited number of operations can be performed on ciphertext data) and non-deterministic (multiple encryptions of the same message and with the same key result in different ciphertexts). Moreover, both division and multiplication operations can be performed over encrypted data.

While fully homomorphic encryption seems to offer a high level solution for privacy-preserving computations with deep learning models, there are still important practical challenges that urge for stronger security, faster running time, and improved generalization performance [8].

To empower privacy-preserving computations in the context of deep learning, it is crucial for the encryption scheme to be applicable to rational numbers. Previously reported approaches for handling this aspect rely on the encoding of rational numbers as integers or as a sequence of integers [9]. Such an approach has limited usability since it does not allow for any operation to be performed on the encoded form. Moreover, adopting an encoding strategy as a way of enabling computations to be performed on real-data introduces not only a clear limitation in its utility but also directly affects the outcome of the computations. To address this limitation, the MORE encryption scheme was adapted to directly support floating point arithmetic. A more detailed description is provided in section III-A.

III. Matrix based data randomization

With Gentry's first introduction of a fully homomorphic encryption schemes [10], numerous variations of the original strategy were proposed in literature [11]. While most of the schemes were shown to be secure, they suffer from very poor performance, being several orders of magnitude slower than the plaintext computations. Alternatively to the original fully homomorphic encryption schemes, some simpler methods which are based on linear transformations emerged. Although criticized due to weaker security [12], [13], this class of methods appears to be currently the only practical approach for performing privacy-preserving computations in real-world applications.

Herein we have employed a variant of the MORE encryption scheme. The MORE scheme relies on matrix algebra and can be used to encrypt a numerical value as a matrix. Therefore, operations performed on encrypted values will turn into matrix operations, e.g. addition of unencrypted scalars will result in addition of encrypted matrices. The MORE encryption scheme is defined as follows. It can be directly generalized to n by n matrices, however, for simplicity, herein we present only the 2 by 2 setup:

1) Let m be the scalar value to be encrypted
2) Let S be a 2 by 2 invertible matrix, representing the encryption and decryption key
3) m is mapped to a 2 by 2 matrix M as follows: $M = \begin{pmatrix} m & 0 \\ 0 & r \end{pmatrix}$ where r is a random parameter
4) Encryption: $C = SMS^{-1}$, C is the encrypted matrix
5) Decryption $M = S^{-1}MS$, the element on the first row and column represent the plaintext value

The MORE scheme is fully homomorphic with respect to algebraic operations, i.e. given two encrypted matrices $C_1 = SM_1S^{-1}$ and $C_2 = SM_2S^{-1}$, for multiplication $C_1C_2 = SM_1S^{-1}SM_2S^{-1} = SM_1M_2S^{-1}$, which is the encryption of the multiplication M_1M_2, and for addition $M_1 + M_2 = SM_1S^{-1} + SM_2S^{-1} = S(M_1 + M_2)S^{-1}$. Similarly, this applies also for subtraction and division, and even for operations involving unencrypted scalars.

A. Encryption of rational numbers

The original MORE scheme, as described by Kipnis et al. [4], applies the encryption to positive integer numbers modulo N, and all the operations are performed modulo N. This is a typical characteristic of fully homomorphic or partially homomorphic encryption schemes. Typical approaches for extending the methodology to rational numbers consist in employing an encoding operation. More specifically rational numbers are first encoded as integers, or as a sequence of integer numbers, and then the encryption is applied on the resulting encoded form. In essence, this is a straightforward problem but, no solution has been found to date in the context of homomorphic encryption. One approach consists in encoding rational numbers as continued fractions [9], however it is difficult to perform operations on numbers represented under this form. Another approach consists in turning rational

numbers into an integer by multiplying with a large scale factor. Unfortunately this approach will not allow for divisions as it will cause the large scale factor to be reduced.

One of the most important advantages of the MORE encryption scheme is that it can also be directly applied on rational numbers without the need of an encoding operation. The drawback is that the method becomes vulnerable to known ciphertext attacks, as described in Section V-C.

B. Performing operations over encrypted data

It was shown previously that the MORE method is fully homomorphic with respect to algebraic operations. In real world applications, a broader spectrum of operations need to be performed, e.g. non-linear (exponential, logarithmic, square root, etc), comparison operations, etc. Typical approaches for performing non-linear operations consist in approximating the given functions as finite polynomial series (e.g. truncated Taylor series), therefore relying only on algebraic operations. The MORE scheme allows for a simple approach for performing such operations.

Knowing that operations performed on encrypted values turn into matrix operations, an intuitive approach is to compute most of the non-linear functions used in neural networks as matrix functions. However, a second approach can also be derived using a property of the MORE scheme: the secret message will always be one of the eigenvalues of the encrypted matrix, e.g. for the 2x2 case, the encrypted matrix C will have two eigenvalues: m and r corresponding to the message and the random secret. If the random secret r is chosen to be statistically indistinguishable from the message, it is impossible to separate the two without knowing the decryption matrix S. Therefore, given an encrypted matrix C, and knowing that m is one of the eigenvalues of C, one can perform eigen decomposition, and then evaluate the given non-linear function directly on the eigenvalues of C. More specifically, given the eigen decomposition VLV^{-1} where V is the eigenvector matrix, and L is the diagonal matrix containing the eigenvalues, one can evaluate any unary function by performing the evaluation separately for each eigenvalue $L_1, L_2, \ldots, L_n$ and then reconstructing the new encrypted matrix as $C_f = Vf(L)V^{-1}$. This approach can even be used for comparing an encrypted matrix C with a plain scalar s. Non-linear binary operations involving two encrypted values cannot be performed, but these types of operations can be avoided in deep learning based applications.

IV. Deep Neural Networks over Encrypted Data

In this section we evaluate the proposed encryption scheme in two types of deep learning applications: binary and multi-class classification. We first address a well known benchmarking application (digit classification), and then focus on training a neural network model on encrypted data to determine coronary angiography views. Experiments demonstrate that we can ensure data security and, at the same time, efficiently perform deep learning based data analysis.

A. MNIST: a typical dataset for neural networks

The MNIST (Modified National Institute of Standards and Technology database) dataset [14] contains images representing handwritten digits, and is ly used as reference for benchmarking image classification algorithms. The training dataset consists of 60,000 grayscale images, of relatively small dimension, 28x28, each image being labeled with the digit it depicts.

To address the challenge of privacy-preserving computations and evaluate the use of deep neural network models over encrypted data, the focus lies on solving the classification problem using a convolutional neural network (CNN) employed on encrypted input-output value pairs. Therefore, with a message $m \in \mathbb{R}$ encoded as a matrix $M \in \mathbb{R}^{2x2}$, for a training example, both the input image and the associated label vector are now represented as ciphertexts in the $\mathbb{R}^{28x28x2x2}$, and $\mathbb{R}^{10x2x2}$ domains. By leveraging the homomorphic property of the scheme, and with the direct support for floating-point arithmetic, training can be performed in a straightforward way.

The trained network has 6 layers, organized as follows: conv-pool-conv-pool-fc-fc. The first convolutional layer has 8 filters, the second one 16 filters, and both layers handle kernels of size 3x3. The pooling layer downsamples the images by a factor of two through averaging. The last two fully connected layers cover 100 and, respectively 10 nodes and all the activation functions employed in the network, except for the last layer, are sigmoid functions. The network was trained using stochastic gradient descent (SGD) to minimize a cross entropy loss between encrypted targets and encrypted predictions. A learning rate of 0.01 was considered, and training was performed in batches of 32 images for a number of 100 epochs. This network leads to an accuracy of 98.3% on the test dataset.

B. View classification in X-ray coronary angiographies

Invasive X-ray coronary angiography (ICA) is a diagnostic imaging procedure that provides important information on the structure and function of the heart, and represents the gold standard in Coronary Artery Disease (CAD) imaging [15]. During a coronary angiogram, radio opaque dye is injected into the coronary arteries and an X-ray scanner rapidly takes a series of images, offering a detailed overview of the coronary arteries. ICA enables the assessment of the anatomical severity of coronary stenoses either visually or by computer-assisted quantitative coronary angiography (QCA) [16].

In view of the limitations of the pure anatomical evaluation of CAD, the functional index of Fractional Flow Reserve (FFR) has been introduced as an alternative [17], and recent technological advances also allow for image-based functional assessment of coronary stenoses based on ICA [18], [19], [20]. Coronary angiographies are recorded separately and sequentially for the right coronary artery (RCA) and the left coronary artery (LCA) (Figure 2).

An important research area in CAD is the fully automated post-processing of coronary angiographies [21], having as objectives:

- Anatomical assessment: automatically determining the anatomical severity of stenoses.
- Non-invasive functional assessment: automatically computing functional diagnostic indices. [19], [20].
- Reporting: composing medical reports automatically based on the findings in the coronary angiographies.

In this and other clinical settings based on the use of ICA, automatic LCA / RCA view classification represents an important pre-processing step. In the following section we describe our approach for automatic coronary angiography view classification. We considered a dataset composed from 3378 coronary angiographies, which were manually annotated as displaying the LCA or the RCA. For each angiography we extracted automatically one frame, in which the arteries were well visible. The dataset was split into 1996 samples for training and 680 for validation, while 702 images were kept for the final testing of the trained model (splitting was performed at patient level, i.e. ensuring that all coronary angiographies of a patient are put in the same dataset - train, validation or test). All 3 datasets were balanced, with a 1:1 prevalence for LCA and RCA cases.

Moreover, to limit overfitting, aside from the regularization added into the network we also performed an offline augmentation, increasing the size of the training dataset by a factor of 4. As augmentation strategies we adopted transformations involving rotating the images by ± 10 degrees, shifting and zooming. For runtime efficiency, we down sampled the coronary angiography images by a factor of 2, resulting in a 256x256 pixel resolution. We have conducted multiple experiments and concluded that for coronary angiography view classification, images having the original 512x512 resolution do not improve the classification accuracy.

Since we are dealing with sensitive data, we focused on training the CNN network on encrypted data. We chose to encrypt only the input data, i.e. the coronary angiography images, and leave the target, i.e. binary label 0 or 1, as plaintext to show that training can as well be performed if labels are kept unencrypted. Note that training can be performed by also encrypting the target, as shown in the multi-class classification problem on the MNIST dataset.

For classifying X-ray coronary angiographies we adopted the following topology of the CNN network:

- Convolutional layer (4 filters of size 3x3, stride of 1x1).
- Sigmoid activation layer.
- Average pooling layer (stride of 2x2).
- Convolutional layer (8 filters of size 3x3, stride of 1x1).
- Tanh activation layer.
- Average pooling layer (stride of 2x2).
- Convolutional layer (16 filters of size 3x3, stride of 1x1).
- Tanh activation layer.
- Average pooling layer (stride of 2x2).
- Convolutional layer (32 filters of size 3x3, stride of 1x1).
- Tanh activation layer.
- Average pooling layer (stride of 2x2).
- Fully connected layer (64 units) and Tanh activation function.

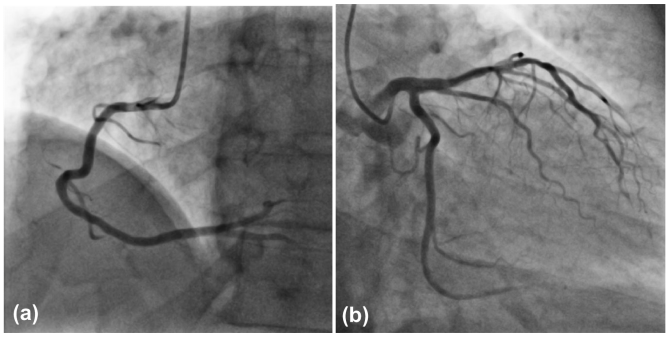

Fig. 2. (a) Right coronary artery, (b) Left coronary artery.

- Dropout layer (dropout rate set to 25%).
- Fully connected layer (1 unit) and Sigmoid activation function.

We have set the learning rate to 0.01 and then trained the network on mini-batches of 16 images, over 100 training epochs, to solve a binary classification problem which minimizes a cross entropy loss. Once the training is finalized, the encrypted form of the model can be employed to predict new encrypted instances, where angiographic images are encrypted with the same key as the ones used during the training phase.

V. Results

A. Performance

A first goal was to verify the correctness of the computations. Hence, in the following we present results obtained by running the algorithms with unencrypted data (plaintext) and encrypted data (ciphertext). Note that for consistency, and for enabling a fair comparison, the same hyper-parameters and random initializations were adopted.

A common question which is raised while training neural networks is when to stop the training to achieve the optimal performance. While an insufficient training may result in non-optimal results by underfitting the data, a too long training phase may lead to overfitting, which again can result in poor performance on the unseen dataset. A typical strategy is to closely monitoring both the training and the validation losses and to stop the training when the first trend of overfitting is observed. Alternatively, the number of epochs may be set to an arbitrary large number and the training is stopped if the validation loss does not improve for a certain number of epochs. While both strategies are straightforward to implement during training on plaintext data, they becomes impractical when dealing with ciphertext data. In the latter case the loss becomes encrypted, and if two encrypted numbers are compared, the result is also a ciphertext, which cannot be used inside a conditional statement. The inconvenience of not seeing the actual loss value forces the training to take place for a predefined number of epochs.

As the overall goal of the study is to assess the feasibility of the deep neural network to operate directly on ciphertext data, i.e. showing that the performance does not drop compared to the plaintext setting, we have chosen an arbitrarily large

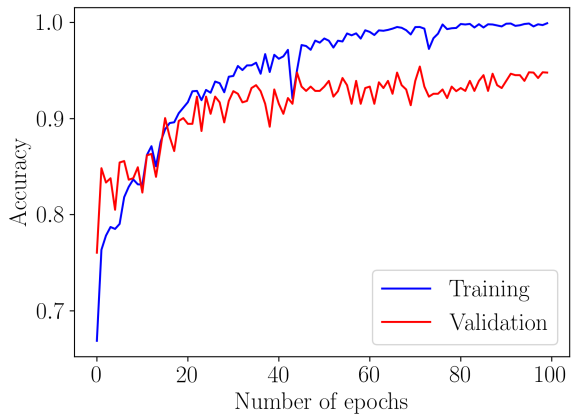

Fig. 3. Accuracy evolution for the network trained on ciphertext data.

number of epochs to conduct the experiments and report the performance.

All experiments indicated that the training progresses similarly in both the encrypted and the unencrypted use cases, as is outlined in the following.

1) MNIST classification: The most important metric is the absolute accuracy of the classification models, i.e. the percentage of correctly labeled digit images. To compute the metric, the outputs of the model outputting ciphertext results are decrypted with the symmetric key. The unencrypted network achieved a classification accuracy of 98.3% on the testing dataset, which is preserved by the encrypted network.

While 98.3% is a marginally acceptable accuracy on the MNIST dataset, it is still relatively far away from 99.77%, declared as the state of the art accuracy for the digit recognition task. However, this is not surprising, as the network proposed to solve the classification task was chosen not with the intention of improving recognition accuracy, but rather to validate privacy preserving computations in the context of neural network models. The accuracy of any predictive model generally improves with more favorable activation functions and optimization algorithms.

2) X-ray coronary angiographies classification: We validated the encrypted model at two levels: (i) at training level, in terms of its capability of preserving the correctness of the computations, and (ii) at inference level, where the focus lies on the overall capability to classify the X-ray coronary angiographies.

To show the ability of the network to learn from ciphertext data, the training and validation accuracy, as resulted after decryption, are depicted in Figure 3.

Regarding the classification accuracy, the CNN network trained on ciphertext data achieves 96.2% of correctly classified samples, when evaluated on unseen encrypted angiographies. When compared to the unencrypted model, accuracy was identical.

TABLE I
RUNTIME ANALYSIS OF THE ENCRYPTED AND PLAINTEXT CNNS FOR MNIST DIGIT RECOGNITION.

Operation	Runtime [s] on cipertext data	Runtime [s] on plaintext data	Encrypted - Unencrypted ratio
Data encryption and key generation	2.44±0.016	-	-
Training (1 epoch)	444.59±8.53	12.98±1.17	34.25
Data encryption	0.39±0.009	-	-
Inference (10K images)	20.42±0.32	0.54±0.08	37.81
Data decryption	0.001±0.0005	-	-

TABLE II
RUNTIME ANALYSIS OF THE ENCRYPTED AND PLAINTEXT CNNS FOR ANGIOGRAPHIC VIEW CLASSIFICATION.

Pearson correlation	Runtime [s] on cipertext data	Runtime [s] on plaintext data	Encrypted - Unencrypted ratio
Training (1 epoch)	1075.47±45.54	34.48±1.12	31.19
Inference (702 images)	26.36±1.98	0.8±0.06	32.95

B. Execution time

All runtimes reported in the current section were measured on a machine equipped with an Intel(R) Xeon(R) CPU running at 2.10GHz. The deep learning library which integrates the MORE encryption scheme was written in C++. The library is still under active development, with minimal multi-threading support.

A detailed comparison of the runtime for each of the applications is given in Table I and Table II. Note that all results were reported under the assumption of employing data parallelism (8 threads) at training and inference level.

C. Security concerns

While the MORE design is simple and clean, with homomorphic properties tailored to privacy-preserving deep neural network, the linear transformations used as the only component of the encryption algorithm limits the security. As stated in [13], [12], the scheme is vulnerable to the chosen plaintext attacks. In particular, if an attacker has access to a large enough number of pairs of encrypted and unencrypted messages, it is possible to compute the secret key by formulating and solving a numerical optimization problem, i.e. by finding the best fit of a matrix S such that $(S^{-1}C_iS)_{1,1} = m_i$ for each known pair (C_i, m_i). This key search attack cannot be applied on the original MORE scheme (on integers modulo N) because the modulo operation is nonlinear.

Although this methodology has weaker security than other homomorphic encryption schemes, it can still be used in applications where the key is never disclosed, e.g. a hospital encrypts the data and then uploads encrypted data to an external computing service. Similarly, it can be employed in a case where encryption is performed per patient, e.g. an application where one can upload personal medical data to

a service that provides a personalized risk factor or other relevant health indices.

VI. Discussion and Conclusions

In the past few years, the raised concern for protecting the privacy of sensitive medical data while still encouraging the delivery of personalized medicine solutions, increased the focus on enabling privacy-preserving computations inside Deep Neural Networks.

The proposed solution aims at ensuring the privacy by incorporating a data encryption mechanism and delivering reliable results, to be used in clinical workflows. We have showcased the applicability of incorporating the MORE encryption scheme into Deep Learning models by tackling two different problems: digit recognition and coronary angiography view classification. We have addressed both the training and the inference phase, and showed that both can be performed on encrypted data. We demonstrated that the accuracy of the encrypted model is statistically not discernable from the unencrypted model, and that, by following the proposed strategy, computations over ciphertext data are only slightly more costly than the ones performed on plaintext data.

In conclusion, we showed that by employing the MORE fully homomorphic encryption scheme as a privacy preserving mechanism, we enabled the application of Deep Learning models on encrypted data without compromising the accuracy at all. Although the runtime increased by more than one order of magnitude, the encrypted models are still outputting results in a reasonable amount of time. With its direct support for computations over rational numbers, and the ability to perform operations without adding noise, the scheme becomes eligible for more complex models from the realm of Deep Learning.

Although the MORE encryption scheme is an attractive choice due to its unbiased advantages in terms of performance and usability, we acknowledge that it offers a lower security compared to standard schemes, and it is by no means a definitive option for problems requiring homomorphic encryption. Improving the security while maintaining the performance and potential to be used in real-world applications represents our main future work direction.

VII. Acknowledgment

The research leading to these results has received funding from the European Union's Horizon 2020 Programme (H2020/2014-2020) under grant agreement no. 732907.

This work was supported by a grant of the Romanian National Authority for Scientific Research and Innovation, CCCDI – UEFISCDI, project number ERANET-FLAG - CONVERGENCE (2), within PNCDI III.

References

[1] R. Miotto, F. Wang, S. Wang, X. Jiang, and J. T. Dudley, "Deep learning for healthcare: review, opportunities and challenges," *Briefings in Bioinformatics*, vol. 19, no. 6, pp. 1236–1246, 05 2017. [Online]. Available: https://dx.doi.org/10.1093/bib/bbx044

[2] Z. Obermeyer and E. J. Emanuel, "Predicting the future - big data, machine learning, and clinical medicine." *The New England journal of medicine*, vol. 375 13, pp. 1216–9, 2016.

[3] R. Shokri and V. Shmatikov, "Privacy-preserving deep learning," in *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '15. New York, NY, USA: ACM, 2015, pp. 1310–1321. [Online]. Available: http://doi.acm.org/10.1145/2810103.2813687

[4] A. Kipnis and E. Hibshoosh, "Efficient methods for practical fully homomorphic symmetric-key encrypton, randomization and verification,"

[5] rig, and

[6] Cryptodl: Deep neural networks over encrypted data," *CoRR*, vol. abs/1711.05189, 2017. [Online]. Available: http://arxiv.org/abs/1711.05189

[7] H. Chabanne, A. de Wargny, J. Milgram, C. Morel, and E. Prouff, "Privacy-preserving classification on deep neural network," *IACR Cryptology ePrint Archive*, vol. 2017, p. 35, 2017.

[8] J. Mancuso. Privacy-preserving machine learning 2018: A year in review. [Online]. Available: https://medium.com/dropoutlabs/privacy-preserving-machine-learning-2018-a-year-in-review-b6345a95ae0f

[9] H. Chung and M. Kim, "Encoding rational numbers for fhe-based applications." *IACR Cryptology ePrint Archive*, vol. 2016, p. 344, 2016.

[10] C. Gentry and S. Halevi, "Implementing gentry's fully-homomorphic encryption scheme," in *CRYPTO 2011*, 2010.

[11] A. El-Yahyaoui and M. D. Elkettani, "Fully homomorphic encryption: state of art and comparison," *International Journal of Computer Science and Information Security*, vol. 14, no. 4, p. 159, 2016.

[12] D. Vizár and S. Vaudenay, "Cryptanalysis of chosen symmetric homomorphic schemes," in *CRYPTO 2014*, 2014.

[13] B. Tsaban and N. Lifshitz, "Cryptanalysis of the more symmetric key fully homomorphic encryption scheme," *J. Mathematical Cryptology*, vol. 9, pp. 75–78, 2014.

[14] L. Deng, "The mnist database of handwritten digit images for machine learning research [best of the web]," *IEEE Signal Processing Magazine*, vol. 29, pp. 141–142, 2012.

[15] T. J. Ryan, "The coronary angiogram and its seminal contributions to cardiovascular medicine over five decades." *Circulation*, vol. 106 6, pp. 752–6, 2002.

[16] V. G. Ng and A. J. Lansky, "Novel qca methodologies and angiographic scores," *The International Journal of Cardiovascular Imaging*, vol. 27, pp. 157–165, 2010.

[17] P. A. L. Tonino, B. D. Bruyne, N. H. J. Pijls, U. Siebert, F. Ikeno, M. van' t Veer, V. Klauss, G. Manoharan, T. Engstrom, K. G. Oldroyd, P. N. V. Lee, P. A. Maccarthy, and W. F. Fearon, "Fractional flow reserve versus angiography for guiding percutaneous coronary intervention." *The New England journal of medicine*, vol. 360 3, pp. 213–24, 2009.

[18] S. Tu, E. Barbato, Z. Köszegi, J. Yang, Z. Sun, N. R. Holm, B. Tar, Y. Li, D. Ruşinaru, W. Wijns, and J. H. C. Reiber, "Fractional flow reserve calculation from 3-dimensional quantitative coronary angiography and timi frame count: A fast computer model to quantify the functional significance of moderately obstructed coronary arteries," *JACC. Cardiovascular interventions*, vol. 7 7, pp. 768–77, 2014.

[19] M. Tröbs, S. Achenbach, J. Roether, T. Redel, M. Scheuering, D. Winneberger, K. Klingenbeck, L. M. Itu, T. Passerini, A. Kamen, P. Sharma, D. Comaniciu, and C. Schlundt, "Comparison of fractional flow reserve based on computational fluid dynamics modeling using coronary angiographic vessel morphology versus invasively measured fractional flow reserve." *The American journal of cardiology*, vol. 117 1, pp. 29–35, 2016.

[20] L. M. Itu, S. Rapaka, T. Passerini, B. Georgescu, C. Schwemmer, M. Schoebinger, T. G. Flohr, P. Sharma, and D. Comaniciu, "A machine-learning approach for computation of fractional flow reserve from coronary computed tomography." *Journal of applied physiology*, vol. 121 1, pp. 42–52, 2016.

[21] A. Arbab-Zadeh, "What will it take to retire invasive coronary angiography," *JACC. Cardiovascular imaging*, vol. 9 5, pp. 565–7, 2016.