Secure Multicenter Medical Model Inference from Homomorphic Encryption

Jingwei Chen[1,2,3], Chen Yang[1,2,3], Yuwen Chen[1,2,5(✉)], Kunhua Zhong[1,2,5], Wenqiang Yang[1,2,3], Jiang Liu[1,2,5], Wenyuan Wu[1,2,3], and Bin Yi[4,5]

[1] Chongqing Institute of Green and Intelligent Technology, CAS, Chongqing 400714, China
chenyuwen@cigit.ac.cn
[2] Chongqing School, University of Chinese Academy of Sciences, Chongqing 400714, China
[3] Chongqing Key Laboratory of Secure Computing for Biology, Chongqing 400714, China
[4] Southwest Hospital, Army Medical University, Chongqing 400038, China
[5] Chongqing Perioperative Medical Big Data Laboratory, Chongqing 400038, China

Abstract. To address the privacy requirements of sensitive medical data, we present a multicenter diagnostic inference framework based on homomorphic encryption (HE). Leveraging the CKKS scheme implemented in the Lattigo library with 128-bit security, our method enables efficient and privacy-preserving inference for diseases such as bladder cancer, breast cancer, and sepsis. By exploiting the sparsity of LASSO parameters, we significantly reduce the computational overhead of encrypted inference. Notably, our LASSO-based analysis reveals a potential therapeutic target for bladder cancer. Experimental results across multiple datasets show that encrypted inference achieves performance comparable to plaintext inference, demonstrating that strong privacy can be preserved without sacrificing diagnostic performance.

Keywords: Disease Diagnosis · Homomorphic Encryption · Medical AI Models · Secure Inference

1 Introduction

Recent advances in artificial intelligence (AI) have revolutionized disease diagnosis by improving accuracy and efficiency. AI models, applied in areas like cancer detection [1] and genomic analysis [2], ease physician workloads and provide faster, more accurate diagnoses, shifting from expert-driven to data-driven systems. However, privacy and security concerns hinder AI's widespread use in medical diagnostics. Medical data, often distributed across institutions, contains sensitive patient information, making it difficult to share and collaborate. Secure multicenter medical model inference enables AI model owner to provide services across multicenter without compromising privacy. Implementing this approach is challenging due to the need for protocols that ensure data confidentiality and interoperability. Privacy-preserving techniques like secure multi-party computation (MPC) [3], differential privacy [4], homomorphic encryption (HE)

D.-S. Huang et al. (Eds.): ICIC 2025, LNBI 15868, pp. 135–147, 2025.
https://doi.org/10.1007/978-981-95-0033-8_12

[5], and trusted execution environments [6] enable us to do so. HE, which allows computations on encrypted data, mitigates exposure risks. In HE- based frameworks, data remains encrypted throughout, reducing unauthorized access concerns. This makes HE suitable for cloud environments and multicenter medical model inference. Moreover, HE schemes like BGV [7], BFV [8, 9], FHEW [10], TFHE [11], and CKKS [12] are quantum-resistant, offering long-term security.

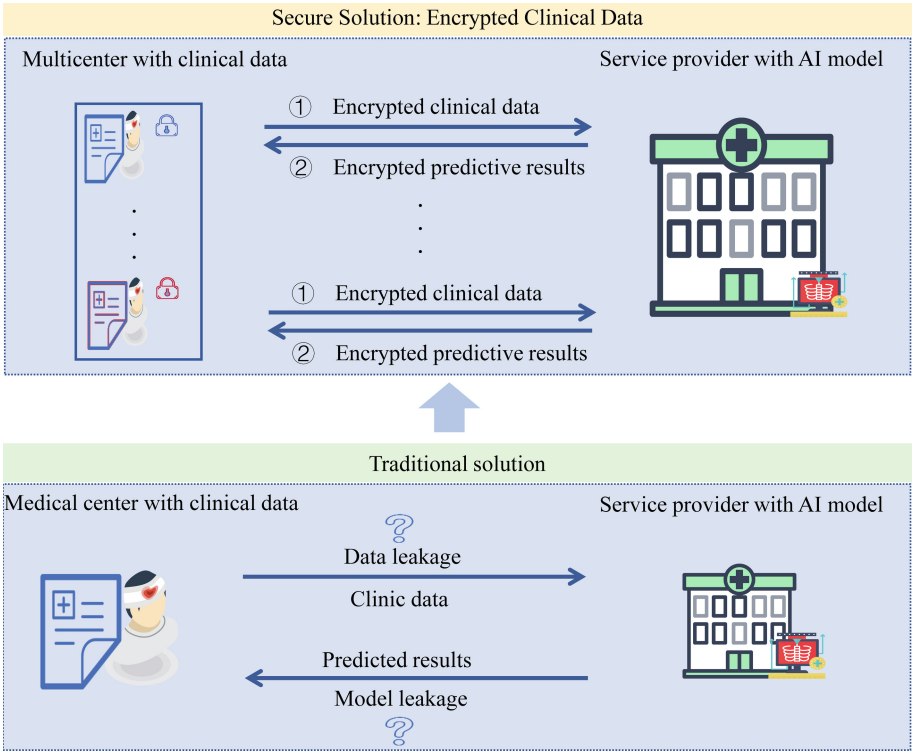

Fig. 1. HE-based multicenter medical model inference. In traditional solutions (below), there is a risk of data leakage from both the clinical data of medical centers and the model data of service providers. In our framework (above), the clinical data from each medical center is independently encrypted using its own HE public key. The service provider then performs encrypted inference based on HE and returns the encrypted results to the corresponding medical center. The medical center decrypts the results using its own HE private key to obtain the final diagnostic prediction. Throughout this process, the clinical data of each medical center remains encrypted, and the service provider's model never leaves the domain. This ensures data security for all parties while completing the model inference.

In this paper, we propose a multicenter medical model inference framework for diseases diagnosis based on HE, as shown in Fig. 1. In this framework, we implement efficient, accurate, and secure model inference for various diseases (bladder cancer, breast cancer, sepsis) across different datasets (see Sect. 3.1) using both polynomial disease

prediction models (LASSO, least absolute shrinkage and selection operator [13], polynomial dendritic neural (PDN) [14]) and non-polynomial ones (MLP [15], KAN [16]). In particular, for the LASSO-based prediction model for bladder cancer, we enhance the efficiency of model inference on encrypted data by leveraging the sparsity of model parameters, thus enabling efficient predictions for medical forecasting. Furthermore, based on the corresponding dataset, our experiments identify a potential new therapeutic target for bladder cancer, which could contribute to clinical diagnosis and treatment; see Sect. 5 for more details.

We implemented the above ciphertext inference using the CKKS scheme [12] from the HE library Lattigo [17]. Experimental results show that, under setups for 128-bit security, our proposed framework efficiently completes all encrypted inference tasks, whether using polynomial or non-polynomial models. Furthermore, the performance metrics (such as accuracy, precision, recall, F1 score, and AUC) are almost identical to those of plaintext model inference; see Sect. 4.

2 Related Work

Generally, plenty of AI models for inference over encrypted data have been investigated based on HE schemes, including decision tree [18], naive Bayes [19], k-means [20], CNN [21], etc. Some of them are hybrid, i.e., combining HE with some other techniques, which typically needs interactions between participants with communication overhead. The datasets used to validate privacy-preserving AI inference are typically the most commonly used ones, such as MNIST [22], CIFAR [23], etc.

For healthcare or medical applications, privacy-preserving machine learning (PPML) is a highly active research field; we refer readers to recent surveys [24, 25] and references therein for more details. Here, we focus specifically on privacy-preserving inference based on HE for medical datasets. Vizitiu et al. [26] explored the application of privacy-preserving technologies in medical images, proposing the use of federated learning, HE, and MPC to protect sensitive patient data. Yue et al. [27] introduced a hybrid deep learning framework combining HE and differential privacy, aiming to conduct privacy-preserving analyses of time-series medical images. T'Jonck et al. [28] proposed a framework for medical data classification using HE, allowing machine learning inference on encrypted data while maintaining data encryption throughout the classification process, ensuring patient privacy. Sarkar et al. [29] investigated how HE can be employed to predict cancer types using encrypted genomic data. These studies collectively demonstrate the potential and security of HE in medical data analysis.

3 Methods

3.1 Datasets

Bladder Cancer Dataset. The dataset called "bc total" is collected from The Cancer Genome Atlas (TCGA) [30] and the Gene Expression Omnibus (GEO) [31] database, aiming at identifying the bladder cancer (BCa)-specific DNA methylation markers. This dataset includes 1, 069 cases, and each has 385, 456 markers. After removing NAN sites and T-test, 342, 669 markers remain.

Breast Cancer Dataset. The well-known Wisconsin Breast Cancer (WBC) data set in UCI Machine Learning Repository [32]. Features are computed from a digitized image of a fine needle aspirate (FNA) of a breast mass. They describe the characteristics of the cell nuclei present in the image. Each sample has 9 attributes. Class distribution: 458 benign, 241 malignant.

Sepsis Dataset. The sepsis experimental dataset is sourced from multiple tier-one hospitals. There are a total of 689 cases, with 227 positive cases and 462 negative cases. Among the positive patients, there are 175 cases from the Southwest Hospital, 28 cases from West China Hospital, and 24 cases from Xuanwu Hospital. Among the negative patients, there are 378 cases from the Southwest Hospital and 84 cases from West China Hospital. Each patient is with 36 attributes.

3.2 Plaintext AI Models

Artificial intelligence models are capable of handling tasks such as classification and regression. In this paper, we focus solely on the former. Assume that after model training, a classification AI model f is obtained. For a given input instance $\boldsymbol{x}$, the model outputs a predicted label $y = f(\boldsymbol{x})$. Different AI model architectures will result in different models f.

We will use different AI models to deal with the three datasets introduced above, including Least Absolute Shrinkage and Selection Operator (LASSO) [13], MultiLayer Perceptron (MLP) [15], exponential and asymptotic polynomial dendritic neural (EPDN and APDN) [14], and Kolmogorov-Arnold Networks (KAN) [16].

Bladder Cancer. Recall that the objective of LASSO is to solve

$$\min_{\boldsymbol{\beta}} \frac{||\boldsymbol{y} - \boldsymbol{X}\boldsymbol{\beta}||_2^2}{N} + \alpha||\boldsymbol{\beta}|| \tag{1}$$

where $\boldsymbol{X}$ is the covariate matrix, N is the number of samples, and $\boldsymbol{y}$ is the label. Here, α is a hyperparameter that determines the degree of regularization. Using the LASSO method on the public dataset, training was conducted with different values of α in (1), resulting in different linear models with the same form $f(\boldsymbol{x}) = \langle \boldsymbol{w}, \boldsymbol{x} \rangle + b$, where $\boldsymbol{x} \in \mathbb{R}^{342,669}$ is the input sample and $\boldsymbol{w}$ is the model parameter that is sparse.

Breast Cancer and Sepsis. For the bladder cancer dataset, LASSO outperformed models such as decision trees, random forests, and MLP, and thus, no further consideration will be given to this dataset. In this section, we primarily focus on using various AI models for diagnostic classification on the breast cancer and sepsis datasets.

Model Structure and Model Training. The structure of each network is described as in Table 1. Each model and dataset is trained for 100 epochs with a batch size of 32, using a learning rate of 0.001. The models are optimized using the Adam optimizer based on the backpropagation algorithm. The training: validation: test ratio for each dataset is 6:2:2. A 5-fold cross-validation was performed respectively. We record all model parameters for encrypted inference.

3.3 A Unified Protocol

Homomorphic Encryption. Homomorphic encryption (HE) enables computation on ciphertexts, producing correct results upon decryption—a concept dating back decades

Table 1. Description of Network Structures for Each Model.

Network	Network Structure	Structure Description
Selu	2-6-4-2	Input dimension is 2, with 6 hidden layers and 4 nodes, output dimension is 2, each layer is based on the Selu activation function
APDN	2-6-4-2	Input dimension is 2, with hidden layers of 6 and 4 nodes, output dimension is 2, each layer utilizes a multiform quadratic function $(\boldsymbol{w}\cdot\boldsymbol{x}+b)\cdot(\boldsymbol{w}'\cdot\boldsymbol{x}+b')$
EPDN	2-6-4-2	Input dimension is 2, with hidden layers of 6 and 4 nodes, output dimension is 2, each layer utilizes a multiform quadratic function $(\boldsymbol{w}\cdot\boldsymbol{x}+b)\cdot(\boldsymbol{w}'\cdot\boldsymbol{x}+b')$
KAN forbreast cancer	9-36-4-2	Input dimension is 9, with hidden layers of 36 and 4 nodes, output dimension is 2, each layer utilizes some activation functions, such as *sin*, *tanh*, *abs*, etc
KAN for sepsis	37-1	Input dimension is 37, output dimension 1, each layer utilizes activation functions, including *sin*, *tanh*, *abs*, etc

[33]. Gentry's 2009 breakthrough [5] led to practical FHE schemes such as BGV [7], BFV [8, 9], CKKS [12], GSW [34], TFHE [11], and FHEW [10]. Among them, GSW, TFHE, and FHEW support binary logic operations, BGV and BFV support integer arithmetic, and CKKS supports real or complex number operations. FHE schemes are typically IND-CPA secure, with security grounded in LWE [35], RLWE [36], or their variants. In practice, minimizing multiplicative depth is crucial, as deeper circuits incur higher computational costs.

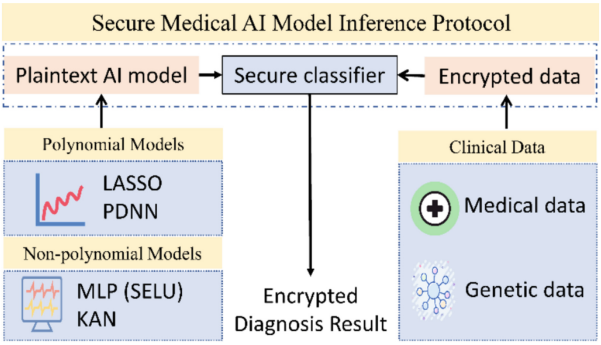

Fig. 2. The workflow of secure medical AI model inference.

The Protocol. Following the workflow in Fig. 2 and the framework in Fig. 1, the medical center first sends the encrypted clinical data to the service provider for a diagnosis. The service provider runs the classifier with its plaintext model together with the received encrypted clinical data as input, and results in an encrypted diagnosis result, and sends it to the corresponding medical center. The medical center decrypts the received ciphertext to disclose the result of the diagnosis. This describes the following general protocol for an arbitrary model f. We will discuss the details in Sect. 3.4 for different AI models f.

Protocol 1 Secure model inference from FHE

Input of the client (C): A sample $\boldsymbol{x}$, the secret and public key sk and pk of an FHE scheme.

Input of the service/model provider (S): The model f.

Client:

1: Encrypt the input data: $\mathrm{ct}.\boldsymbol{x} \leftarrow Enc_{pk}(\boldsymbol{x})$.

2: Send $ct.\boldsymbol{x}$ to the service/model provider.

Service/model provider:

3: Compute $\mathrm{ct}.\boldsymbol{y} \leftarrow \mathrm{Eval}(f, \boldsymbol{x})$.

4: Send $ct.\boldsymbol{y}$ to the client.

Client:

5: Decrypt $\boldsymbol{y} \leftarrow Dec_{sk}(\mathrm{ct}.\boldsymbol{y})$.

6: Output $\boldsymbol{y}$.

Adversarial model. Within the framework in Fig. 2, the protocol only involves the client and the server, respectively. We consider the honest-but-curious (semi-honest or passive) model as described in, e.g., [37, Sec. 7.2].

Security. By employing the classical simulation-based proof method in secure multiparty computation, the security of our protocol can be reduced to the security of the HE scheme, since the protocol only requires a single round of communication.

3.4 Encrypted Inference

We discuss the encrypted AI model inference by categorizing the models into polynomial (LASSO and PDNN) and non-polynomial models (MLP and KAN).

Polynomial Models. For inference using a linear model, $f(x) = \langle w, x \rangle + b$, where $\boldsymbol{w}$ is the coefficient vector and b the bias term, the inference essentially involves computing an inner product between the vector w (in plaintext form) and the vector x (in encrypted form). While this is a relatively simple computational task, the presence of the L_1 regularization term in the LASSO model results in w being a highly sparse vector. This sparsity improves model interpretability and enhances computational efficiency. Fortunately, under the framework illustrated in Fig. 2, this sparsity can be fully leveraged. Specifically, assuming the dimension d of the vector x is extremely large, exceeding the maximum dimension ℓ that a single ciphertext can represent, the vector x needs to be encrypted as $n = \lceil \frac{d}{\ell} \rceil$ ciphertexts. Correspondingly, the vector w must be divided

into n segments as well. Given that w is sparse, let us assume only $n_1 \ll n$ segments are non-zero. In this case, we only need to compute n_1 plaintext-ciphertext inner products of dimension ℓ instead of n. For polynomial models with degrees larger than one, such as PDNN, $f(\boldsymbol{x})$ is essentially a polynomial function. To homomorphically evaluate a polynomial of degree d by using $O(d)$ non-scalar multiplications, one can use the Paterson-Stockmeyer algorithm [38]. It costs $\lceil \log d \rceil$ multiplicative depth.

Non-polynomial Models. We observe that the primary difference between non-polynomial models and polynomial models is that the former incorporates non-polynomial functions into the model to enhance its expressive capability. In MLP, a non-polynomial activation function, such as Sigmoid or Selu, is applied between different layers. KAN takes this further by performing a non-polynomial transformation at each node, subsequently combining these results to construct a more complex functional structure. The main challenge in applying encrypted inference with these models lies in how to evaluate these non-polynomial functions homomorphically on encrypted data. HE schemes only support addition and multiplication on ciphertexts. Thus, non-polynomial functions should be approximated by polynomials, which consists of only additions and multiplications, to make it compatible with HE schemes. And we use Chebyshev polynomials to approximate those non-polynomial functions. Compared with Taylor approximation, Chebyshev approximation is more accurate and numerically stable. Furthermore, the Paterson-Stockmeyer algorithm can also be applied to Chebyshev approximation; see, e.g., [39].

4 Results

4.1 Plainxtext Model Inference

LASSO for BCa. Here, we only describe the results of LASSO method for the bladder cancer datasets. Other models for other datasets are stored and used to test the performance of our secure model inference protocol. Recall the hyperparameter α in (1). For $\alpha = 0.1, 0.02, 0.01$, we obtain three different models with w has 1, 5, and 9 non-zero entries, respectively, which correspond to different markers that may be related to BCa.

- $\alpha = 0.1$
  - Model: w = (0.553), b = 0.200180702.
  - Marker: cg26112797.
- $\alpha = 0.02$
  - Model: w = (0.164, 0.0989, 0.374,0.0235, 0.475), and b = 0.013430789.
  - Markers: cg26112797, cg06153925, cg10351284, cg10590292, cg23359665.
- $\alpha = 0.05$
  - Model: w = (0.0319, 0.225, 0.00128, 0.251, 0.38, 0.0673, 0.0801, 0.0261, 0.192), and b = 0.006486069.

- Markers: cg26112797, cg06153925, cg10351284, cg10590292, cg23359665, cg00025044, cg06390079, cg11436362, cg24757533.

In the training of the plaintext model, the "bc total" dataset was split into a 7:3 ratio for training and test. During model training, we employed 10-fold cross-validation. On the test dataset, AUC equals 1.0 for all three models.

KAN for Sepsis. Although the KAN model performs a bit worse than EPDN, APDN and Selu on the breast cancer dataset, it outperforms EPDN and APDN on the sepsis dataset; See Fig. 3 and 4.

4.2 Secure Model Inference

In this section, we will present the performance of the HE version of our model. The clinical data is encrypted and sent to the classifier, who owns the model, to compute the result in encrypted state. Then, the encrypted result is sent back and decrypted by the medical center; see Fig. 1. In this way, the clinical data is protected from the classifier since the data is encrypted, and the model owned by the service provider is also protected from the medical center since the medical center can only access the final result. So, the data leakage and model leakage are both eliminated.

Setup. The library we use is Lattigo (v5.0.2) [17], and the HE scheme we choose is CKKS [12]. Table 2 describes the parameters, ciphertext modulus log q and degree of ring N for RLWE, that we used in experiments. The column entitled "depth" indicates the multiplicative depth required by evaluating the corresponding model inference. The depth required by these models is supported by the setup parameters without needing bootstrapping. Under this setup, we can achieve a 128-bit security according to a draft security standard [40] and the latest lattice estimator [41]. The code for secure model inference, together with all model parameters, are available at https://github.com/JohnJimAir/SecureMulticenterMedicalModel_viaHE.

Table 2. Parameter setup for encrypted model inference.

Model	depth	log$q$	N
LASSO	2	160	2^{13}
EPDN/APDN	6	386	2^{14}
Selu	13	721	2^{15}
KAN for Sepsis	17	881	2^{15}
KAN for Breast Cancer	24	1356	2^{16}

Performance. As mentioned previously, under the framework of Protocol 1, we lever-age the sparsity of the model to perform encrypted inference on bladder cancer data using the LASSO linear model. This approach enables the classification of gene data of dimension 342, 669 within one second, with the intermediate and resulting ciphertexts not exceeding 20 MB. For the two polynomial models, EPDN and APDN, since they

are constructed through polynomials, which is very friendly for HE schemes, the result from HE is the same as in plaintexts. While the Selu model consists of exp(x), we use polynomials to approximate it in the intervals observed from the process of plaintexts. The interval is [4], and the degree of the approximate polynomial is 15. Although the polynomial approximation of exp(x) introduces numerical errors, the difference in final result accuracy between ciphertexts and plaintexts on the breast cancer dataset is less than 2%. And from our experiments, the Selu model preform not well on the spesis dataset, so we omit to report its performance on the sepsis dataset here.

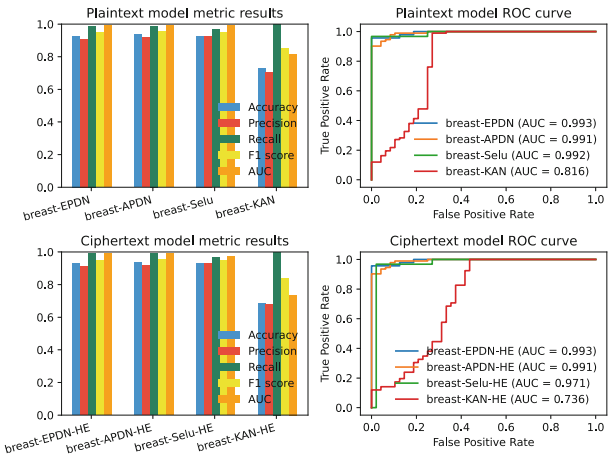

Fig. 3. Plaintext and encrypted inference for breast cancer.

Implementing encrypted KAN is more challenging than the Selu because almost all the nodes in KAN need to be approximated by polynomials. For sin, tanh and abs, we approximate them using polynomials in interval of [16] and with degree of 31; for sqrt and log, using polynomials in interval of [0, 16] and with degree of 31. Fortunately, the difference in final result accuracy between ciphertexts and plaintexts is almost the same, with about 4% at the breast dataset and 0% at the sepsis dataset. Figure 3 and 4 demonstrate the accuracy of plaintext and ciphertext inference on test data. From our experiments, the homomorphic evaluation of EPDN, APDN, Selu, and KAN on the real datasets "breast" and "sepsis" shows little decline in accuracy.

The machine we use to test our implementation is equipped with Intel Xeon Gold 6248R (3.00 GHz, 24Core) and 128G (32G $\times$ 4) memory. We present the cost of encrypted inference for different datasets with different models in Table 3. Experiments demonstrate that in terms of both time and storage overhead, encrypted AI model inference has approached a highly practical level, particularly in fields like disease diagnosis, where real-time performance is not as critical.

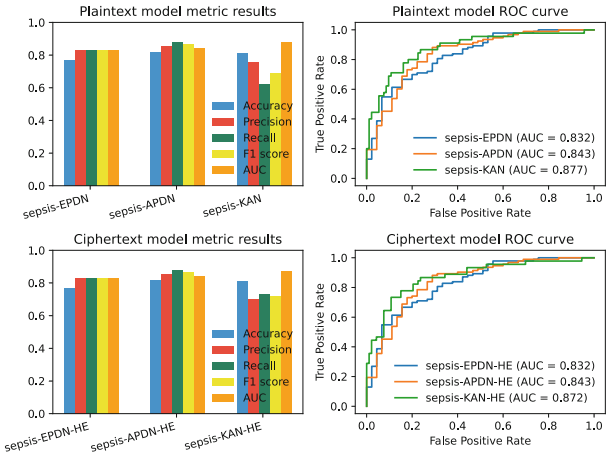

Fig. 4. Plaintext and encrypted inference for breast sepsis.

Table 3. Time and memory cost for secure model inference

Method	Bladder Cancer	Breast Cancer	Sepsis
LASSO	1 s, 20 MB	-	-
EPDN	-	2 s, 196 MB	20 s, 642 MB
APDN	-	2 s, 202 MB	22 s, 658 MB
Selu	-	7 s, 548 MB	-
KAN	-	204 s, 4343 MB	61 s, 1904 MB

5 Discussion and Conclusion

Many BCa markers identified by our LASSO analysis are not reported in prior studies (e.g., [2]), suggesting potential value for early bladder cancer detection after further bioinformatics and clinical validation.

Unlike general-purpose secure inference systems [18–21], our framework relies solely on HE (Fig. 1), requiring only a single communication round between centers and the provider. We are the first to apply HE to privacy-preserving PDN [14] inference. While KAN inference has only one general-purpose HE-based solution [42], we adapt it for medical datasets here. Prior work often uses public datasets like MNIST [22] or CIFAR [23]; we instead evaluate real-world datasets for bladder cancer, breast cancer, and sepsis.

Compared to MPC-based medical inference systems [26], our solution—like [27–29]—relies entirely on HE for security, and demonstrates near-practical performance across diverse models and datasets. Figure 1 assumes encrypted clinical data is sent to the provider, which may not be feasible due to data-sharing restrictions. Figure 5 presents an alternative: the provider encrypts the model and deploys it locally, allowing each

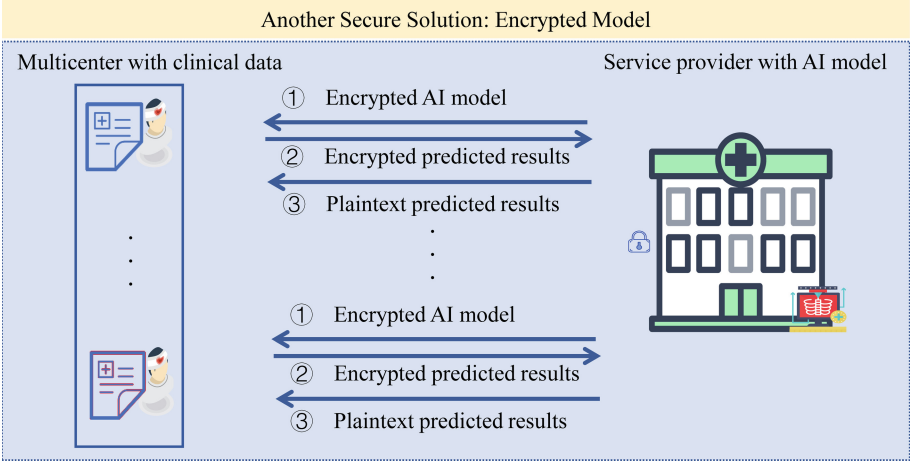

Fig. 5. Another framework for secure model inference.

center to run inference and return encrypted results. The provider decrypts and delivers the plaintext diagnosis. This preserves both model and data privacy, though outputs are exposed. Finally, both frameworks assume semi-honest adversaries. However, in Fig. 5, the provider controls the encrypted computation and medical centers lack transparency over the process. Malicious providers could embed trapdoors to extract private data. Extending the framework to resist malicious behavior is an important direction.

Acknowledgments. This work was supported partly by National Key Research and Development Project of China (2020YFA0712303), Natural Science Foundation of China (62371438), Natural Science Foundation of Chongqing (2022yszx-jcx0011cstb, cstb2023yszx-jcx0008, cstb2024nscq-msx1043) and the Light of West China Program of CAS.

References

1. Elmarakeby, H.A., et al.: Biologically informed deep neural network for prostate cancer discovery. Nature **598**(7880), 348–352 (2021)
2. Chen, X., et al.: Urine DNA methylation assay enables early detection and recurrence monitoring for bladder cancer. J. Clin. Investig. **130**(12), 6278–6289 (2020)
3. Yao, A.C.: Theory and application of trapdoor functions. In: FOCS 1982, pp. 80–91 (1982)
4. Dwork, C.: Differential privacy. In ICALP 2006, 2006, LNCS 4052, pp. 1–12 (2006)
5. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: STOC 2009, pp. 169–178. ACM, New York (2009)
6. Sabt, M., et al.: Trusted execution environment: what it is, and what it is not. In: 2015 IEEE Trust-Com/BigDataSE/ISPA, vol. 1, pp. 57–64. IEEE (2015)
7. Brakerski, Z., Gentry, C., Vaikuntanathan V.: (Leveled) fully homomorphic encryption without bootstrapping. In: ITCS2012, pp. 309–325. ACM, New York (2012)
8. Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical GapSVP. In: CRYPTO 2012, LNCS 7417, pp. 868–886. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_50

9. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. Cryptology ePrint Archive https://eprint.iacr.org/2012/144 (2012)
10. Ducas, L., Micciancio, D.: FHEW: bootstrapping homomorphic encryption in less than a second. In: EUROCRYPT 2015, LNCS 9056, pp. 617–640. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_24
11. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: fast fully homomorphic encryption over the torus. J. Cryptol. **33**(1), 34–91 (2020)
12. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017, pp. 409–437. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-319-70694-8_15
13. Tibshirani, R.: Regression shrinkage and selection via the lasso. J. Royal Stat. Soc. Ser. B (Methodol.) **58**(1), 267–288 (1996)
14. Chen, Y., Liu, J.: Polynomial dendritic neural networks. Neural Comput. Appl. **34**(14), 11571–11588 (2022)
15. Rosenblatt, F.: The perceptron: a probabilistic model for information storage and organization in the brain. Psychol. Rev. **65**(6), 386–408 (1958)
16. Liu, Z., et al.: KAN: Kolmogorov-Arnold networks. arXiv:2404.19756 (2024)
17. EPFL-LDS and Tune-Insight. Lattigo v5.0.2. https://github.com/tuneinsight/lattigo
18. Akhavan Mahdavi, R., et al.: Level Up: Private non-interactive decision tree evaluation using levelled homomorphic encryption. In: CCS 2023, pp. 2945–2958. ACM, New York (2023)
19. Chen, J., Feng, Y., et al.: Non-interactive privacy-preserving naïve Bayes classifier using homomorphic encryption. In: Shi, W., Chen, X., Choo, KK.R. (eds.) EAI SPNCE 2021, pp. 192–203. Springer, Cham (2022). https://doi.org/10.1007/978-3-030-96791-8_14
20. Yang, C., Chen, J., Wu, W., Feng, Y.: Optimized privacy-preserving clustering with fully homomorphic encryption. In: SecureComm 2024, Springer (2024)
21. Kim, D., Guyot, C.: Optimized privacy-preserving CNN inference with fully homomorphic encryption. IEEE Trans. Info Forensic Secur. **18**, 2175–2187 (2023)
22. LeCun, Y., Cortes, C., Burges, C.J.C.: The MNIST database of handwritten digit (1998). http://yann.lecun.com/exdb/mnist/
23. Krizhevsky, A., Nair, V., Hinton, G.: The CIFAR-10 dataset (2009). https://www.cs.toronto.edu/~kriz/cifar.html
24. Guerra-Manzanares, A., et al.: Privacy-preserving machine learning for healthcare: open challenges and future perspectives. In: Chen, H., Luo, L. (eds.) TML4H 2023, pp. 25–40. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-39539-0_3
25. Naresh, V.S., et al.: Privacy-preserving deep learning in medical informatics: applications, challenges, and solutions. Artif. Intell. Rev. **56**(1), 1199–1241 (2023)
26. Anamaria Vizitiu, et al. Towards privacy-preserving deep learning based medical imaging applications. In 2019 MeMeA, pages 1–6. IEEE, 2019
27. Yue, Z., et al.: Privacy-preserving time-series medical images analysis using a hybrid deep learning framework. ACM Trans. Internet Technol. **21**(3), 57:1–21, (2021)
28. T'Jonck, K., et al.: Privacy preserving classification via machine learning model inference on homomorphic encrypted medical data. In: ET, pp. 1–6. IEEE (2022)
29. Sarkar, E., et al.: Privacy-preserving cancer type prediction with homomorphic encryption. Sci. Rep. **13**(1), 1661 (2023)
30. Gao, G.F., et al.: Before and after: Comparison of legacy and harmonized TCGA genomic data commons' data. Cell Syst. **9**(1), 24-34.e10 (2019)
31. Ron Edgar, M. Domrachev, Lash, A.E.: Gene expression omnibus: NCBI gene expression and hybridization array data repository. Nucleic Acids Res. **30**(1), 207–210 (2002)
32. Dua, D., Graff, C.: UCI machine learning repository (2017). http://archive.ics.uci.edu/ml
33. Rivest, R., Adleman, L., Dertouzos, M.: On data banks and privacy homomorphisms. In Foundations of Secure Computation, pp. 165–179. Academic Press, Atlanta (1978)

34. Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: conceptually-simpler, asymptotically-faster, attribute-based. In CRYPTO 2013, LNCS 8042, pp. 75–92. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40041-4_5
35. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. J. ACM **56**(6), 34:1–40 (2009)
36. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. J. ACM, **60**(6), 43:1–35 (2013)
37. Goldreich, O.: Foundations of Cryptography-Volume II. Cambridge University Press (2004)
38. Paterson, M.S., Stockmeyer, L.J.: On the number of nonscalar multiplications necessary to evaluate polynomials. SIAM J. Comput. **2**(1), 60–66 (1973)
39. Chen, H., Chillotti, I., Song, Y.: Improved bootstrapping for approximate homomorphic encryption. In: EUROCRYPT 2019, LNCS 11477, pp. 34–54. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17656-3_2
40. Albrecht, M., et al.: Homomorphic encryption security standard. Technical report, HomomorphicEncryption.org, Toronto, Canada (2018)
41. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of learning with errors. J. Math. Cryptol. **9**(3), 169–203 (2015)
42. Lai, Z., Zhou, Y., Zheng, P., Chen, L.: Efficient privacy-preserving KAN inference using homomorphic encryption (2024). https://arxiv.org/abs/2409.07751