

BUNET: Blind Medical Image Segmentation Based on Secure UNET

Song Bian¹, Xiaowei Xu², Weiwen Jiang³, and Yiyu Shi³ Takashi Sato¹

¹ Kyoto University

{sbian, takashi}@easter.kuee.kyoto-u.ac.jp

² Guangdong Provincial People's Hospital

xiao.wei.xu@foxmail.com

³ University of Notre Dame

{wjiang2, yshi4}@nd.edu

Abstract. The strict security requirements placed on medical records by various privacy regulations become major obstacles in the age of big data. To ensure efficient machine learning as a service schemes while protecting data confidentiality, in this work, we propose blind UNET (BUNET), a secure protocol that implements privacy-preserving medical image segmentation based on the UNET architecture. In BUNET, we efficiently utilize cryptographic primitives such as homomorphic encryption and garbled circuits (GC) to design a complete secure protocol for the UNET neural architecture. In addition, we perform extensive architectural search in reducing the computational bottleneck of GC-based secure activation protocols with high-dimensional input data. In the experiment, we thoroughly examine the parameter space of our protocol, and show that we can achieve up to 14x inference time reduction compared to the-state-of-the-art secure inference technique on a baseline architecture with negligible accuracy degradation.

1 Introduction

The use of neural-network (NN) based machine learning (ML) algorithms in aiding medical diagnosis, especially in the field of medical image computing, appears to be extremely successful in terms of its prediction accuracy. However, the security regulations over medical records contradicts the use of big data in the age of ML. Highly sensitive patient records are protected under the Health Insurance Portability and Accountability Act (HIPAA), where strong protection measures need to be taken over all of the electronic Protected Health Information (ePHI) possessed by a patient. In particular, access controls and client-side encryption are mandated for the distribution of all ePHI records over public networks [12, 14]. In addition, while qualified professionals are allowed to handle ePHI, the data exposure is required to be kept minimal [10, 13], i.e., just enough to accomplish the necessary professional judgements.

A central question to the real-world deployment of NN-based ML techniques in medical image processing is how the related data transfer and computations

can be handled securely and efficiently. Previous security measures on medical data generally involved physical means (e.g., physically disconnected from the internet), and these techniques clearly cannot benefit from the large-scale distributed computing networks available for solving ML tasks. Recent advances in cryptography and multi-party secure computing seek alternatives to address the security concerns. In particular, the concept of secure inference (SI) is proposed, where Alice as a client wishes to inference on some of her inputs with the machine learning models provided by the server, called Bob. The security requirement is that no one, including Bob, learns anything about the inputs from Alice, while Alice also learns nothing about the models from Bob. Over the past few years, prior arts on SI flourished [6,15,17,18,22,24], where secure protocols targeted on general learning problems were proposed. In addition, we also observe protocol- and system-level optimizations [2,22] on SI. Unfortunately, most existing works mentioned above do not have a clear application in mind. Thus, the utilized network architectures and datasets (e.g., MNIST, CIFAR-10) are usually generic, without immediate practical implications.

In this work, we propose BUNET, a secure protocol for the UNET architecture [23] that enables input-hiding segmentation on medical images. In the proposed protocol, we use a combination of cryptographic building blocks to ensure that client-side encryption is enforced on all data related to the patients, and that practical inference time can also be achieved. As a result, medical institutions can take advantage of third-party machine learning service providers without violating privacy regulations. The main contributions of this work are summarized as follows.

- **Privacy-Preserving Image Segmentation:** To the best of our knowledge, we are the first to propose a secure protocol for image segmentation.
- **Architectural Exploration for Secure UNET:** We perform a search on the possible alternative UNET architectures to reduce the amount of computations (in terms of cryptographic realizations) in SI.
- **Thorough Empirical Evaluations:** By performing architectural-protocol co-design, we achieved 8x–14x inference time reduction with negligible accuracy degradation.

2 Preliminaries

2.1 Cryptographic Primitives

In this work, we mainly consider the four types of cryptographic primitives: a packed additive homomorphic encryption (PAHE) scheme based on the ring learning with errors (RLWE) problem [4,5,8,11], additive secret sharing (ASS) [9], garbled circuits (GC) [26], and multiplication triples (MT) [1,16]. In what follows, we provide a brief overview for each primitive.

PAHE: A PAHE is a cryptosystem, where the encryption (Enc) and decryption (Dec) functions act as group (additive) homomorphisms between the plaintext and ciphertext spaces. Except for the normal Enc and Dec, a PAHE

scheme is equipped with the following three abstract operators. We use $[\mathbf{x}]$ to denote the encrypted ciphertext of $\mathbf{x} \in \mathbb{Z}^n$, and $n \in \mathbb{Z}$ the maximum number of plaintext integers that can be held in a single ciphertext.

- Homomorphic addition (\boxplus): for $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^n$, $\text{Dec}([\mathbf{x}] \boxplus [\mathbf{y}]) = \mathbf{x} + \mathbf{y}$. Note we can also perform homomorphic subtraction \boxminus , where $\text{Dec}([\mathbf{x}] \boxminus [\mathbf{y}]) = \mathbf{x} - \mathbf{y}$.
- Homomorphic Hadamard product (\boxtimes): for $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^n$, $\text{Dec}([\mathbf{x}] \boxtimes [\mathbf{y}]) = \mathbf{x} \circ \mathbf{y}$, where \circ is the element-wise multiplication operator.
- Homomorphic rotation (rot): for $\mathbf{x} \in \mathbb{Z}^n$, let $\mathbf{x} = (x_0, x_1, \dots, x_{n-1})$, $\text{rot}([\mathbf{x}], k) = (x_k, x_{k+1}, \dots, x_{n-1}, x_0, \dots, x_{k-1})$ for $k \in \{0, \dots, n-1\}$.

ASS and Homomorphic Secret Sharing: A two-party ASS scheme consists of two operators, (Share, Rec), and some prime modulus $p_A \in \mathbb{Z}$. Each operator takes two inputs, where we have $\mathbf{s}_A = \text{Share}(\mathbf{x}, \mathbf{s}_B) = (\mathbf{x} - \mathbf{s}_B) \bmod p_A$ and $\mathbf{x} = \text{Rec}(\mathbf{s}_A, \mathbf{s}_B) = (\mathbf{s}_A + \mathbf{s}_B) \bmod p_A$. In [15], homomorphic secret sharing (HSS) is adopted, where ASS operates over an encrypted \mathbf{x} . For HSS, we have that

$$[\mathbf{s}_A] = \text{Share}([\mathbf{x}], \mathbf{s}_B) = ([\mathbf{x}] \boxminus \mathbf{s}_B) \bmod p_A, \text{ and} \quad (1)$$

$$[\mathbf{x}] = \text{Rec}([\mathbf{s}_A], \mathbf{s}_B) = (\mathbf{s}_A \boxplus \mathbf{s}_B) \bmod p_A. \quad (2)$$

GC: GC can be considered as a more general form of HE. In particular, the circuit garbler, Alice, “encrypts” some function f along with her input x to Bob, the circuit evaluator. Bob evaluates $f(x, y)$ using his encrypted input y that is received from Alice obliviously, and obtains the encrypted outputs. Alice and Bob jointly “decrypt” the output of the function $f(x, y)$ and one of the two parties learns the result.

MT: Beaver’s MT [1] is a technique that performs multiplication on a pair of secret-shared vectors $\mathbf{x} = \text{Rec}(\mathbf{s}_{A,x}, \mathbf{s}_{B,x})$ and $\mathbf{y} = \text{Rec}(\mathbf{s}_{A,y} + \mathbf{s}_{B,y})$ between Alice and Bob. Here, we take computations performed by Alice as an example, and only note that the exact same procedure is also executed by Bob on his shares of secrets. To compute $\mathbf{x} \circ \mathbf{y}$, Alice and Bob first pre-share a set of respective multiplication triples $(\mathbf{a}_A, \mathbf{b}_A, \mathbf{c}_A)$ and $(\mathbf{a}_B, \mathbf{b}_B, \mathbf{c}_B)$, where we have $\text{Rec}(\mathbf{a}_A, \mathbf{a}_B) \circ \text{Rec}(\mathbf{b}_A, \mathbf{b}_B) = \text{Rec}(\mathbf{c}_A, \mathbf{c}_B)$. Alice locally calculates $\mathbf{d}_A = \mathbf{s}_{A,x} \circ \mathbf{a}_A \bmod p_A$, $\mathbf{e}_A = \mathbf{b}_A \circ \mathbf{s}_{A,y} \bmod p_A$. Then, Alice and Bob publish their results $\mathbf{d}_A, \mathbf{d}_B$ and $\mathbf{e}_A, \mathbf{e}_B$. Finally, Alice obtains

$$\mathbf{g}_A = (\mathbf{c}_A + (\mathbf{e} \circ \mathbf{s}_{A,0,i}) + (\mathbf{d} \circ \mathbf{s}_{A,1,i}) - \mathbf{d} \circ \mathbf{e}) \bmod p_A, \quad (3)$$

where $\mathbf{d} = \text{Rec}(\mathbf{d}_A, \mathbf{d}_B)$ and similarly for \mathbf{e} . MT guarantees that $\text{Rec}(\mathbf{g}_A, \mathbf{g}_B) = \mathbf{x} \circ \mathbf{y}$, where \mathbf{g}_B is the MT results computed by Bob.

2.2 Related Works on Secure Neural Network Inference

While a limited number of pioneering works have been proposed for secure inference and training with neural networks [15, 17, 18, 24], it was not until recently that such protocols carried practical significance. For example, in [17], an inference with a single CIFAR-10 image takes more than 500 seconds to complete.

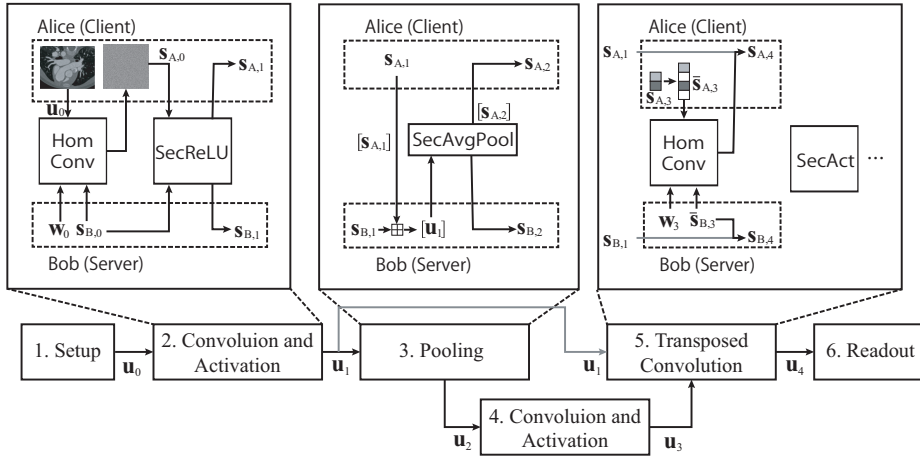


Fig. 1. The overview of the BUNET protocol.

Using the same neural architecture, the performance was improved to less than 8 seconds in one of the most recent arts on SI, ENSEI [15]. Unfortunately, as shown in Section 4, without UNET-specific optimizations and protocol designs, existing approaches carry significant performance overhead, especially on the 3D images (e.g., CT scans) used in medical applications. Hence, in this work, we establish our protocol based on the ENSEI construction, and explore UNET-specific optimizations and cryptographic protocol designs to improve the practicality of secure inference in medical segmentation.

3 Secure UNET for Blind Segmentation

3.1 BUNET: The Protocol

Fig. 1 shows an example of the BUNET protocol structured as a UNET architecture, where the input image goes through four steps. The operators used in each step will be discussed in detail later in Section 3.2.

1. **Setup:** As our protocol takes both 2D and 3D images as inputs, the input image is of dimension $\dim(U) = c \times w \times h \times d$ ($d = 1$ for 2D inputs). Alice first raster-scans her input image into a one-dimensional vector \mathbf{u} of length $\ell = c \cdot w \cdot h \cdot d$. Bob does a similar transformation on his filter weights to obtain the one-dimensional vector \mathbf{w} . We use \mathbf{u}_0 to denote the input from Alice at the 0-th layer, and \mathbf{w}_0 for that of Bob.
2. **Convolution and Activation:** We conduct a standard (input-hiding) convolution on the input image with activation functions followed. For each convolution layer with activation, we run

$$\mathbf{v}_0 = \text{HomConv}(\mathbf{u}_0, \mathbf{w}_0), \text{ and } \mathbf{u}_1 = f_a(\mathbf{v}_0) = \text{SecAct}(\mathbf{v}_0). \quad (4)$$

for some weight vector \mathbf{w}_0 . Here, f_a is some abstract activation function (e.g., ReLU or square activation). The output $\mathbf{u}_1 = f_a(\mathbf{v}_0)$ will become the input

to the next layer. We point out that the activation function has a significant impact on both the accuracy and the inference time of UNET. Therefore, we propose a hybrid UNET architecture, where both ReLU and square activations are used. Here, the security guarantee is that Alice obtains no knowledge on \mathbf{w}_0 , and Bob knows nothing about \mathbf{u}_0 and \mathbf{v}_0 after the protocol execution.

3. **Pooling:** While the standard UNET architecture employs max pooling [23] as the downsampling method, in the experiment, we show a large performance difference between max pooling and average pooling protocols, due to the change of underlying protocol. We also demonstrate in Table 1 that the two pooling methods result in marginal accuracy differences. Hence, we modified the UNET architecture to employ only average pooling. Consequently, the proposed protocol executes $\mathbf{u}_2 = \text{SecAvgPool}(\mathbf{u}_1)$. For an input of length ℓ and pooling size $\zeta_p = z_{w,p} \times z_{h,p} \times z_{d,p}$, the pooled output have a dimension of ℓ/ζ_p .
4. **Bottom-Level Convolution and Activation:** Here, Step 2 is repeated, and we get \mathbf{u}_3 as output. Note that the input and output dimension is reduced by Step 3, so Step 4 is computationally lighter than Step 2.
5. **Transposed Convolution and Concatenation:** While the arithmetic procedures for transposed convolution is essentially the same as a normal convolution, protocol-level modifications are required for the image concatenation and padding operations. Concretely, after obtaining input from the previous layer, e.g., \mathbf{u}_3 , Alice needs to zero-pad \mathbf{u}_3 in an interleaving manner, according to some stride size $\zeta_t = z_{w,t} \times z_{h,t} \times z_{d,t}$. The padded result, $\bar{\mathbf{u}}_3$, will have a length of $\ell \cdot \zeta_t / \zeta_p$. Subsequently, Alice uses $\bar{\mathbf{u}}_3$ as input to execute the following protocols.

$$\mathbf{v}_4 = \text{HomConv}(\bar{\mathbf{u}}_3, \mathbf{w}_3), \text{ and } \mathbf{u}_4 = \mathbf{v}_4 || \mathbf{u}_1. \quad (5)$$

One subtlety is that, the output from the 1-st layer, \mathbf{u}_1 , needs to be concatenated with the output from the transposed convolution layer for the rest of the normal convolutions. However, \mathbf{u}_1 and \mathbf{v}_4 will actually be encrypted under different keys. Thus, both results need to be decrypted and concatenated by Alice. The concatenated result, $\mathbf{v}_1 || \mathbf{u}_4$, will become the inputs to later layers.

6. **Readout:** Here, we applies the **SecArgmax** function over the label dimensions. It is noted that, since the **Softmax** operator is monotonic and is only required in the learning process, we avoid using a separate protocol for **Softmax**, and directly perform a secure **Argmax**. Since the **Argmax** function is a pixel-wise comparison function across the label dimension, it can be implemented using a simple GC protocol similar to the secure ReLU protocol, and we omit a formal presentation.

Threat Model and Security The threat model for BUNET is that both Bob and Alice are semi-honest, in the sense that both parties follow the protocol prescribed above (e.g., encrypting real data with **Enc**, etc.), but want to learn as much information as possible from the other party. Our protocol guarantees that Alice learns only the segmentation results, while Bob learns nothing about the

inputs from Alice. We provide a figure in the appendix to illustrate the security relationship between BUNET and the previous protocols [3, 15].

3.2 The Cryptographic Building Blocks

Here, we discuss each of the cryptographic primitives used in the previous section in details.

HomConv: The HomConv operator obviously convolve two vectors \mathbf{u} and \mathbf{w} . A very recent work [3] discovered that, instead of the complex rotate-and-accumulate approach proposed by previous works [15], homomorphic convolution can be performed in the frequency domain, where the only computation needed in the homomorphic domain is the \boxtimes (homomorphic Hadamard product) operator. Therefore, the homomorphic convolution protocol proceeds as follows.

1. First, Alice performs an integer discrete Fourier transform (DFT) (i.e., number theoretic transform in [3]) on \mathbf{u} and obtain its frequency-domain representation, $\hat{\mathbf{u}}$. She simply encrypts this input array into a ciphertext $[\hat{\mathbf{u}}]$ (when $\ell > n$, the vector is encrypted into multiple ciphertexts) by running $[\hat{\mathbf{u}}] = \text{Enc}_{\mathcal{K}}(\hat{\mathbf{u}})$, where \mathcal{K} is the encryption key. The resulting ciphertext is transferred to Bob.
2. Before receiving any input from Alice, Bob applies DFT on his filter weights to obtain $\hat{\mathbf{w}}$. In this process, the size of the filter will be padded to be ℓ . Upon receiving the inputs $[\hat{\mathbf{u}}]$ from Alice, Bob computes

$$[\hat{\mathbf{v}}] = [\hat{\mathbf{u}} \circ \hat{\mathbf{w}}] = [\hat{\mathbf{u}}] \boxtimes \hat{\mathbf{w}} \quad (6)$$

for all ciphertexts $\hat{\mathbf{u}}$.

3. Finally, since $\hat{\mathbf{v}}$ contains information of the weights from Bob, Bob applies HSS as $[\hat{\mathbf{s}}_A] = \text{Share}([\hat{\mathbf{v}}], \hat{\mathbf{s}}_B)$. Bob keeps $\hat{\mathbf{s}}_B$ and returns $[\hat{\mathbf{s}}_A]$ to Alice, where Alice decrypts and obtain $\hat{\mathbf{s}}_A$. Both Alice and Bob run inverse DFT on their shares of secrets (i.e., $\hat{\mathbf{s}}_A$ and $\hat{\mathbf{s}}_B$) and obtain \mathbf{s}_A and \mathbf{s}_B , respectively, completing the protocol.

SecAct: The SecAct protocols are summarized as follows.

- **ReLU:** We follow the construction in [15] based on the GC protocol. Alice first garbles the circuit with her share of secret \mathbf{s}_A . The garbled circuit obviously computes the following function

$$\mathbf{v} = \text{Rec}(\mathbf{s}_A, \mathbf{s}_B), \text{ and } \bar{\mathbf{s}}_A = \text{Share}(\text{ReLU}(\mathbf{v}), \bar{\mathbf{s}}_B) \quad (7)$$

where $\bar{\mathbf{s}}_B$ is a freshly generated share of secret from Bob. After protocol execution, Alice obtains $\bar{\mathbf{s}}_A$, which contains $\text{ReLU}(\mathbf{v})$ in an oblivious manner.

- **Square:** Since the square activation (i.e., $y = x^2$) is essentially evaluating a polynomial over the inputs, the computationally-light MT can be used instead of GC. To use MT, we first share the secret \mathbf{v} twice among Alice and Bob. Then, the MT protocol outlined in Section 2.1 can be executed, where both \mathbf{x} and \mathbf{y} equal \mathbf{v} . After the protocol execution, Alice and Bob respectively obtain \mathbf{g}_A and \mathbf{g}_B where $\text{Rec}(\mathbf{g}_A, \mathbf{g}_B) = \mathbf{v}^2$. The main observation here is

Table 1. The Dice Accuracy on the HVSMR Dataset for Different Neural Architectures

	ReLU+Max Float	ReLU Float	ReLU+Max 32-bit (Baseline)	ReLU 16-bit	Hybrid 20-bit	Square 16-bit
HVSMR Myo. Dice	0.74±0.04	0.73±0.04	0.73±0.05	0.73±0.05	0.71±0.05	0.52±0.14
HVSMR BP Dice	0.87±0.04	0.87±0.04	0.87±0.04	0.87±0.04	0.86±0.05	0.83±0.05
HVSMR Time (s)	-	-	42616	14205	3054.2	1118.2
EM Dice	0.9405	-	0.9411	0.9398	0.9385	0.8767
EM Time (s)	-	-	8838.1	2968.6	1077.9	227.16

that, all computations are coefficient-wise multiplications and additions over $\lg p_A$ -bit integers. As a result, MT-based square activation is much faster than GC-based ReLU activation.

SecPool: As mentioned above, two types of secure pooling can be implemented, the **SecMaxPool** and the **SecAvgPool** operator. In [15], it is shown that max-pooling can be implemented using the GC protocol as in Eq. (7), where we replace the ReLU operator with the MaxPool operator. Meanwhile, for secure average pooling, we can use a simple protocol that is purely based on PAHE. Specifically, we can compute the window-wise sum of some vector \mathbf{v} by calculating $\text{SecAvgPool}(\mathbf{v}) = \sum_{i=0}^{\zeta_p-1} \text{rot}([\mathbf{v}], i)$, where ζ_p is the pooling window size. Since both homomorphic rotations and additions are light operations compared to GC, **SecAvgPool** is much faster than **SecMaxPool**.

4 Accuracy Experiments and Protocol Instantiation

4.1 Experiment Setup

Due to the lack of immediate existing works, we compare BUNET with the standard UNET architecture implemented by the ENSEI [3] protocol, which is the best performing protocol on secure multi-class inference. Here, the standard UNET architecture only utilizes max-pooling for pooling layers, and ReLU for activation layers. We denote this architecture as the baseline architecture. As shown in the appendix, the baseline architecture consists 19 convolution layers including three transposed convolution layers, 14 activation layers, three average pooling layers, and a readout layer implementing the **SecArgmax** function.

Our experiments are conducted on three datasets, GM [21], EM [7], and HVSMR [20]. Due to the space constraint, we only present the accuracy and performance results on the EM (two dimensional) and HVSMR (three dimensional) datasets (GM will be added to the appendix).

The cryptographic performance is characterized on an Intel i5-9400 2.9 GHz CPU, and the accuracy results are obtained with an NVIDIA P100 GPU. The adopted PAHE library is SEAL version 3.3.0 [25] (we set q, p to be 60-bit and 20-bit integers, respectively, and $n = 2048$, ensuring a 128-bit security) and MT/GC protocols are implemented using LowGear provided by MP-SPDZ [16, 19].

4.2 Accuracy and Performance Results

In this section, we explore how the neural architecture of 3D-UNET impact on the segmentation performance. In particular, it is important to see if the the

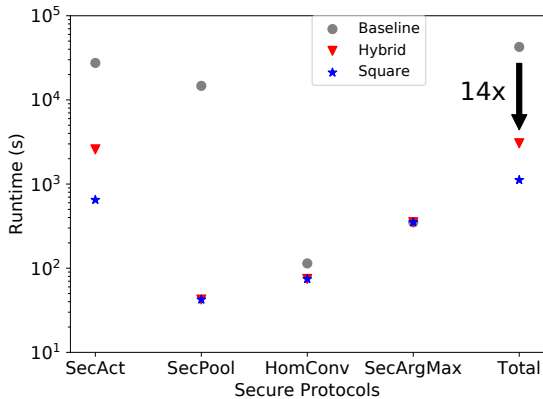


Fig. 2. Runtime distribution for different cryptographic building blocks in a single run of secure segmentation.

Table 2. Number of Activation Functions per Layer Batch

1st	2nd	3rd	4th	5th	6th	7th
33554432	8388608	2097152	524288	2097152	8388608	33554432

proposed architectural modifications for UNET result in satisfactory prediction accuracy while accelerating the network inference time. We downsampled the images in the HVSMR [20] dataset to a dimension of $64 \times 64 \times 64$ containing three class labels, i.e., background, myocardium (myo.) and blood pool (BP). The images in the EM dataset has 200×200 input for binary segmentation. Table 1 summarizes the dice scores and runtime under various architectural settings (HVSMR with 10-fold cross validation). Here, the pooling function is average pooling unless otherwise stated. Hybrid refers to the neural architecture where the first and the last layer batches use square activation, while all other layer batches adopt ReLU activation. Here, a layer batch denotes two convolution and activation layers with the same output feature dimensions and channels. We have three main observations. First, the use of average pooling instead of max pooling results in negligible accuracy degradation, on a level that can likely be compensated by parameter tuning. Second, the UNET architecture is robust in low-quantization environment, where we see little accuracy difference between floating point, 32-bit and 16-bit quantization factors, especially on the EM dataset. Lastly, replacing all ReLU activations with squares results in significant accuracy degradation for the segmentation of myocardium. However, the BP prediction can be acceptable for a quick run of cheap evaluations, and the hybrid architecture successfully achieves a good balance between accuracy and performance.

We record the total runtime for a single blind image segmentation with respect to different 3-D UNET architectures. In Table 2 and Fig. 2, we illustrate the per-layer-batch number of neuron activations and the runtime distribution for different secure protocols in BUNET. As expected, the first and last two layers contain the most amount of activations, and replacing the GC-based heavy

ReLU activation with square activation results in an immediate 5x total runtime reduction. In addition, it is observed that the runtime for ReLU activation functions dominate the total runtime across architectures, while square activation is as light as a frequency-domain homomorphic convolution operation.

Compared to the baseline 32-bit ReLU architecture, we obtain 8x–14x runtime reduction with the reasonably accurate Hybrid architecture, and up to 39x reduction with the cheapest (all-square) UNET implementation on EM. Finally, we note that most NN operations can be parallelized as well as the cryptographic building blocks. Therefore, since our performance is recorded on a single-thread CPU, we expect further runtime reduction for BUNET on parallel computing architectures.

5 Conclusion

In this work, we propose BUNET to perform blind medical image segmentation on encrypted medical images. The observation we make is that protocol and network designs need to be jointly performed to achieve the best accuracy-performance trade-off. By designing UNET-specific protocols and optimizing the UNET architecture, we show that up to 8x–14x inference time reduction can be achieved with negligible accuracy degradation on several medical datasets.

Acknowledgment

This work was partially supported by JSPS KAKENHI Grant No. 20K19799, 20H04156, Edgecortex Inc, the Science and Technology Planning Project of Guangdong Province under Grant No. 2017A070701013, 2017B090904034, 2017B030314109, 2018B090944002, and 2019B020230003, Guangdong peak project under Grant No. DFJH201802, the National Key Research and Development Program under Grant No. 2018YFC1002600, the Natural Science Foundation of Guangdong Province under Grant No. 2018A030313785.

References

1. Beaver, D.: Efficient multiparty protocols using circuit randomization. In: Annual International Cryptology Conference. pp. 420–432. Springer (1991)
2. Bian, S., Jiang, W., Lu, Q., Shi, Y., Sato, T.: NASS: Optimizing secure inference via neural architecture search. arXiv preprint arXiv:2001.11854 (2020)
3. Bian, S., Wang, T., Hiromoto, M., Shi, Y., Sato, T.: ENSEI: Efficient secure inference via frequency-domain homomorphic convolution for privacy-preserving visual recognition (2020)
4. Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical GapSVP. In: Advances in Cryptology–CRYPTO 2012, pp. 868–886. Springer (2012)
5. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. ACM Transactions on Computation Theory (TOCT) **6**(3), 13 (2014)

6. Brutzkus, A., Gilad-Bachrach, R., Elisha, O.: Low latency privacy preserving inference. In: International Conference on Machine Learning. pp. 812–821 (2019)
7. Cardona, A., Saalfeld, S., Preibisch, S., Schmid, B., Cheng, A., Pulokas, J., Tomancak, P., Hartenstein, V.: An integrated micro-and macroarchitectural analysis of the drosophila brain by computer-assisted serial section electron microscopy. *PLoS biology* **8**(10) (2010)
8. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 409–437. Springer (2017)
9. Damgård, I., Nielsen, J.B., Polychroniadou, A., Raskin, M.: On the communication required for unconditionally secure multiplication. In: Annual International Cryptology Conference. pp. 459–488. Springer (2016)
10. Drolet, B.C., Marwaha, J.S., Hyatt, B., Blazar, P.E., Lifchez, S.D.: Electronic communication of protected health information: privacy, security, and hipaa compliance. *The Journal of hand surgery* **42**(6), 411–416 (2017)
11. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. *IACR Cryptology ePrint Archive* **2012**, 144 (2012)
12. HHS.gov: <https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html> (2009), accessed: 2020-03-04
13. HHS.gov: <https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/minimum-necessary-requirement/index.html> (2009), accessed: 2020-03-04
14. Hoffman, S., Podgurski, A.: Securing the hipaa security rule. *Journal of Internet Law*, Spring pp. 06–26 (2007)
15. Juvekar, C., et al.: Gazelle: A low latency framework for secure neural network inference. arXiv preprint arXiv:1801.05507 (2018)
16. Keller, M., Pastro, V., Rotaru, D.: Overdrive: making spdz great again. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 158–189. Springer (2018)
17. Liu, J., et al.: Oblivious neural network predictions via MinioNN transformations. In: Proc. of ACM SIGSAC Conference on Computer and Communications Security. pp. 619–631. ACM (2017)
18. Mohassel, P., et al.: Secureml: A system for scalable privacy-preserving machine learning. In: Proc. of Security and Privacy (SP). pp. 19–38. IEEE (2017)
19. MP-SPDZ: <https://github.com/data61/MP-SPDZ/> (2018), accessed: 2020-03-10
20. Pace, D.F., Dalca, A.V., Geva, T., Powell, A.J., Moghari, M.H., Golland, P.: Interactive whole-heart segmentation in congenital heart disease. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 80–88. Springer (2015)
21. Prados, F., Ashburner, J., Blaiotta, C., Brosch, T., Carballido-Gamio, J., Cardoso, M.J., Conrad, B.N., Datta, E., Dávid, G., De Leener, B., et al.: Spinal cord grey matter segmentation challenge. *Neuroimage* **152**, 312–329 (2017)
22. Riazi, M.S., Samragh, M., Chen, H., Laine, K., Lauter, K.E., Koushanfar, F.: Xonn: Xnor-based oblivious deep neural network inference. *IACR Cryptology ePrint Archive* **2019**, 171 (2019)
23. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
24. Rouhani, B.D., et al.: Deepsecure: Scalable provably-secure deep learning. In: Proc. of DAC. pp. 1–6. IEEE (2018)
25. Microsoft SEAL (release 3.3). <https://github.com/Microsoft/SEAL> (June 2019), microsoft Research, Redmond, WA.

26. Yao, A.C.: Protocols for secure computations. In: Foundations of Computer Science, 1982. SFCS'08. 23rd Annual Symposium on. pp. 160–164. IEEE (1982)