

Measuring Computational Universality of Fully Homomorphic Encryption

Jiaqi Xue¹ Xin Xin¹ Wei Zhang¹ Mengxin Zheng¹ Qianqian Song²
 Minxuan Zhou³ Yushun Dong⁴ Dongjie Wang⁵ Xun Chen⁶ Jiafeng Xie⁷
 Liqiang Wang¹ David Mohaisen¹ Hongyi Wu⁸ Qian Lou^{1*}

¹University of Central Florida ²University of Florida ³Illinois Institute of Technology ⁴Florida State University
⁵The University of Kansas ⁶Samsung Research America ⁷Villanova University ⁸University of Arizona

ABSTRACT

Many real-world applications, such as machine learning and graph analytics, involve combinations of linear and non-linear operations. As these applications increasingly handle sensitive data, there is a significant demand for privacy-preserving computation techniques capable of efficiently supporting both types of operations—a property we define as “computational universality.” Fully Homomorphic Encryption (FHE) has emerged as a powerful approach to perform computations directly on encrypted data. In this paper, we systematically evaluate and measure whether existing FHE methods achieve computational universality or primarily favor either linear or non-linear operations, especially in non-interactive settings. We evaluate FHE universality in three stages. First, we categorize existing FHE methods into ten distinct approaches and analyze their theoretical complexities, selecting the three most promising universal candidates. Next, we perform measurements on representative workloads that combine linear and non-linear operations in various sequences, assessing performance across different bit lengths and with SIMD parallelization enabled or disabled. Finally, we empirically evaluate these candidates on five real-world, privacy-sensitive applications, where each involving arithmetic (linear) and comparison-like (non-linear) operations. Our findings indicate significant overheads in current universal FHE solutions, with efficiency strongly influenced by SIMD parallelism, word-wise versus bit-wise operations, and the trade-off between approximate and exact computations. Additionally, our analysis provides practical guidance to help practitioners select the most suitable universal FHE schemes and algorithms based on specific application requirements. The code is publicly available at <https://github.com/UCF-Lou-Lab-PET/FHE-Universality>.

CCS CONCEPTS

• Security and privacy → Privacy-preserving protocols; • General and reference → Measurement;

KEYWORDS

Fully Homomorphic Encryption, Universality Measurement

1 INTRODUCTION

FHE is a cryptographic technique that allows computations to be performed directly on encrypted data, eliminating the need for

decryption and ensuring sensitive information remains confidential throughout processing. Because of its strong privacy guarantees, FHE has become a cornerstone of secure data analysis in numerous domains, including privacy-preserving machine learning [70, 95, 109], genomics [51, 84, 103], and finance [11, 45]. It enables researchers and practitioners to extract valuable insights from confidential data while rigorously safeguarding user privacy.

Since Gentry’s first FHE scheme in 2009 [40], numerous constructions [9, 19, 20, 26, 31, 34, 36, 37] have been proposed, primarily focusing on enabling either linear or non-linear computations. Existing FHE methods can generally be divided into two categories based on their computational strengths. The first category is word-wise FHE, such as BGV [20], BFV [19, 37], and CKKS [26]. These schemes excel at performing efficient word-wise linear operations, such as matrix multiplication, by leveraging Single Instruction Multiple Data (SIMD) [87] to pack multiple plaintexts into a single ciphertext. The second category is bit-wise FHE, including FHEW [36] and TFHE [31], which specialize in non-linear logic operations by supporting arbitrary functions through Boolean circuit evaluations. Bit-wise schemes encrypt each plaintext bit individually, allowing for flexible computation over complex non-linear functions.

However, little attention has been given to systematically enabling mixed linear and non-linear operations in a non-interactive manner. We define this capability as *universality*. Universality is crucial for many real-world applications that inherently involve a combination of linear and non-linear computations. For example, neural networks [41, 72] perform linear matrix multiplications followed by non-linear activation functions such as ReLU, while graph algorithms like Floyd-Warshall [38] often involve linear aggregation and non-linear comparisons.

In addition, given the diverse landscape of FHE methods that are specifically optimized for different applications and workloads, evaluating the FHE’s universality is challenging. Many FHE methods are theoretically capable of universal operations but remain impractical due to significant computational complexity. For instance, directly using bit-wise FHE for universal operations is impractical because of its extremely high computational complexity for linear computation [74], i.e., multiplying two 16-bit integers in TFHE can take up to 30 seconds. On the other hand, using word-wise FHE for non-linear functions with linear approximation appears to be a potential direction for handling universal operations, such as polynomial approximation [27, 28, 60] for CKKS. However, these methods only work on a small interval, such as $[-1, 1]$, inevitably introducing errors around 0. Such errors render approximation-based methods unsuitable for applications where even small errors are intolerable, such as genomics [51, 84, 103] and finance [11, 45].

*Corresponding Author Email: qian.lou@ucf.edu.

Another line of work for universal operations involves scheme switching [18, 74] between word-wise and bit-wise FHE, allowing for the evaluation of the linear operations in word-wise and the non-linear operations in the bit-wise FHE without decryption. While this leverages the strengths of both types, switching cost is prohibitively high, especially for input with large bit-width [29]. Most importantly, additional bootstrapping might be needed after scheme switching [7], which is one of the most computationally expensive operations in the word-wise FHE and is often avoided in practice [48].

Also, many new advancements have emerged and claimed to support non-linear operations using word-wise FHE, but they are not verified in universal computation measurements. For example, word-wise BGV/BFV can perform exact non-linear computations without approximation errors by leveraging polynomial interpolation over the base field \mathbb{Z}_p [48, 79, 90, 94, 96, 102, 105]. To enable them to support universal operations in a non-interactive manner, recent work [102] proposed an encoding switching method, HEBridge. Recently, several functional bootstrapping methods [9, 57, 65] have been proposed to allow arbitrary function computation during the bootstrapping process. Functional bootstrapping refreshes ciphertexts and applies a chosen function simultaneously, effectively removing noise while completing a specific Look-Up Table (LUT) evaluation. Specifically, Lee et al. [57] proposed a functional bootstrapping technique that takes an RLWE ciphertext, i.e., CKKS or BFV, as input and outputs the refreshed BFV ciphertext, building upon BFV-style bootstrapping. However, such BFV-style bootstrapping is far from efficient; for instance, it can take about 3 minutes for a single evaluation due to the intrinsically inefficient slot utilization of BFV-style bootstrapping [9]. To address this issue, Alexandru et al. [9] presented a general functional bootstrapping technique based on CKKS-style bootstrapping, which has the best throughput among all FHE methods, to improve amortized performance. However, this method suffers from high computational complexity for large input spaces. Their experimental results are limited to 12-bit LUTs because the scaling factor nears the limit for 64-bit modular operations, and the complexity of polynomial evaluation increases significantly, i.e., approximately 1 minute per evaluation for 9-bit LUTs, but about 10 minutes for 12-bit LUTs.

To evaluate the computational universality of FHE, we proceed in three stages, as shown in Figure 1. First, we survey and analyze the theoretical complexities of current FHE methods—grouped into 10 distinct directions—to identify three strong candidates for universal computation. Next, we conduct micro-benchmarks on three representative workloads involving mixed linear and non-linear operations: (1) linear followed by non-linear, (2) non-linear followed by linear, and (3) a sequence of linear, non-linear, and linear operations. Finally, we extend our evaluation to five real-world, privacy-sensitive applications that inherently require universal FHE capabilities, including Floyd–Warshall graph analysis, decision tree inference, sorting, database aggregation, and neural network inference. These applications span domains such as finance and healthcare, where protecting sensitive information during complex computations is critical. By integrating theoretical analysis with empirical runtime measurements, we deliver a comprehensive, application-driven benchmark of universal FHE solutions.

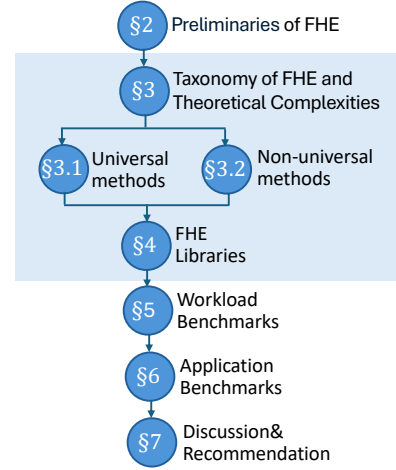


Figure 1: Overview of Our FHE Universality Measurements

2 PRELIMINARIES

2.1 Fully Homomorphic Encryption

Fully Homomorphic Encryption (FHE) is an encryption technique that enables computations to be carried out directly on encrypted data, ensuring confidentiality throughout processing. Formally, if $y = f(x)$ represents an arithmetic function on plaintext x , there exists a function $g(\cdot)$ operating in the encrypted domain such that $y = f(x) \Leftrightarrow y = \text{Dec}(g(\text{Enc}(x)))$. $\text{Enc}(\cdot)$ and $\text{Dec}(\cdot)$ denote encryption and decryption, respectively, while $g(\cdot)$ mirrors $f(\cdot)$ in the ciphertext space. FHE methods introduce noise into ciphertexts to ensure security, typically following the Learning With Errors (LWE) problem. The noise prevents cryptanalysis but must remain below a threshold to ensure successful decryption. The security of modern FHE methods relies on the hardness of LWE and its ring-based variant, Ring LWE (RLWE). In RLWE, the goal is to find the secret $s \in R_p$ in the equation $b = a \cdot s + e$, where $b, a \in R_p$ are known, and e is an error sampled from a distribution over R_p . The ring R_p is defined as $R_p = \mathbb{Z}_p[x] / \langle x^n + 1 \rangle$, with n a power of 2 and $p \equiv 1 \pmod{2n}$. In practice, ciphertexts are encrypted as tuples of polynomials in this ring, taken modulo an irreducible cyclotomic polynomial whose roots are the N^{th} primitive roots of unity. The cyclotomic order (degree of the ciphertext polynomials) is typically between 2^{10} and 2^{15} , and the coefficient modulus q is often a product of primes, reaching several hundred bits.

Most current FHE methods can be categorized according to their capability to support different types of operations. Word-wise FHE methods (such as BFV [37], BGV [20], and CKKS [26]) excel at linear operations such as matrix multiplication, while bit-wise FHE methods (such as TFHE [31] and FHEW [36]) are better suited for non-linear operations such as comparison.

Bit-wise FHE methods. The first category includes bit-wise FHE methods, such as TFHE [31] and FHEW [36], which encrypt individual bits into separate ciphertexts, allowing for a wide range of encrypted bitwise manipulations and logic gate evaluations. These schemes support the construction of Boolean circuits composed of encrypted logic gates, facilitating efficient operations directly on

encrypted data. TFHE, in particular, is known for its fast bootstrapping, which is essential for evaluating homomorphic logic gates. Bit-wise FHE methods enable developers to leverage decades of digital circuit design research, enabling many applications such as privacy-preserving machine learning and relevant TFHE accelerators [46, 49, 55, 67–69, 107, 108, 110]. However, they face challenges with addition and multiplication circuits [69, 74], especially when dealing with large circuit depths and fan-in bits. For example, TFHE [31] takes about 30 seconds to multiply two encrypted 16-bit integers. Also, the expansion ratio of bit-wise FHEs is usually several orders of magnitude larger than that of word-wise FHEs, leading to higher communication costs.

Word-wise FHE methods. Word-wise FHEs can be classified into integer schemes [19, 20, 37] and floating-point schemes [26]. The two most important FHE schemes that encrypt integers modulo a user-determined modulus p are BGV [20] and BFV [37]. FHE addition and multiplication correspond to modular addition and multiplication over the plaintext values. It is also possible to emulate Boolean circuits in integer schemes by setting $p = 2$, which causes addition to behave as an XOR gate and multiplication to work as an AND gate. Unlike bit-wise FHE schemes, both BGV and BFV have considerably slower bootstrapping; depending on parameter choices (such as the ring dimension), a single bootstrapping can take anywhere from several seconds to several hours [39]. Thus, most implementations of BGV and BFV do not include bootstrapping and are used exclusively in LHE mode [42].

For the floating-point schemes, they use a plaintext type of floating-point numbers, and the most popular scheme in this category is CKKS [26]. It is desirable for certain classes of algorithms, such as machine learning [95, 104, 111]. In practice, the core operations of floating-point schemes parallel those of integer schemes, with only a few differences. Similarly, the floating-point schemes have slow bootstrapping procedures and are predominantly used in LHE mode. The key difference is the need to keep track of the *scale* factor that is multiplied by plaintext values during encoding, which determines the bit-precision associated with each ciphertext. The scale doubles when two ciphertexts are multiplied, which may result in overflow as the scale becomes exponentially larger. Thus, *rescaling* must be invoked to preserve the original scale after multiplications, which is similar to modulus switching and serves a similar role in reducing ciphertext noise.

The key advantage of word-wise FHE over bit-wise methods is its support for batching. Single Instruction Multiple Data (SIMD) enables encrypting a vector of plaintexts into a single ciphertext [26, 87]. Each plaintext occupies a *slot*, with the number of slots determined by encryption parameters. Batched ciphertexts support slot-wise addition and multiplication. A key limitation is slot dependency—if operations require interaction across slots (e.g., summing all slots), slot-wise rotations are needed, incurring significant memory and runtime overhead.

Despite their efficiency in linear operations, word-wise FHE methods struggle with non-linear functions like comparison [71, 73], prompting efforts to support such functions within word-wise schemes [48, 90]. CKKS supports approximate computation on

floating-point numbers, and the most prevailing method to perform the non-linear functions is through polynomial approximation [27, 28, 60]. These methods use a composition of minimax approximation polynomials [60] to approximate the sign function with errors. These methods primarily work on only small intervals, such as $[-1, -\epsilon] \cup [\epsilon, 1]$ with errors. The smaller the errors are, the higher the degree of polynomial required. Thus, more multiplication and multiplicative depth are needed. Most importantly, the approximation can never be accurate around 0. Such errors limit the application of approximation-based non-linear where accuracy is important, such as genomics [51, 84, 103] and finance [11, 45]. On the other hand, since computations are exact in BFV and BGV, the non-linear functions can be implemented in BFV and BGV without approximation error. There has been a series of works [48, 79, 90] exploring the efficient comparison in BFV and BGV. On a high level, they compute the non-linear functions by evaluating an interpolation polynomial over the base field. However, the degree of the interpolation polynomial grows exponentially with the bit-width of the input [90] or relies on a special plaintext space [48] that does not support seamless computation of the linear and non-linear operations.

2.2 Universality Required Privacy-sensitive Applications

Over the past decade, FHE has evolved from a purely theoretical construct into a practical solution for today’s data-centric challenges [42]. In finance, Intesa Sanpaolo’s collaboration with IBM leverages FHE secure asset transactions [47], allowing for both confidential database querying and the execution of complex machine learning algorithms on encrypted data. Similarly, in the medical domain, Roche’s partnership with ETH investigates the use of FHE for secure graph analyses of genomic information [66, 100]. As industries increasingly adopt advanced analytics that require both linear and non-linear operations, FHE’s *universal capacity* to handle complex computations on encrypted data makes it an indispensable tool for ensuring data confidentiality and regulatory compliance.

In Section 3, we analyze the universality of various FHE methods from a theoretical perspective. Section 5 presents experimental evaluations assessing their performance across diverse workloads. These workloads include sequences of operations such as non-linear applied after linear operations, mirroring the activation functions in neural networks; as well as linear following non-linear operations, which form the basic unit of fundamental $\text{Max/Min}(A, B)$ operations, expressed as $\text{Compare}(A, B) \times A + (1 - \text{Compare}(A, B)) \times B$. Finally, in Section 6, we evaluate privacy-sensitive applications built on these workloads, demonstrating the practical universality of FHE methods. The applications under study include graph algorithms, decision trees, database aggregation, neural-network inference, and sorting, each representing a critical function in real-world research and commercial settings. Benchmarking different FHE methods on these tasks provides insight into their real-world performance in universal scenarios.

Table 1: Universality comparison of FHE methods under b -bit precision and 2^λ security level. N is the ring dimension. For a polynomial approximation of CKKS, the n denotes the degree of evaluation polynomial.

Methods	Non-linear Operation (i.e., Comparison)					Linear Operation (i.e., Multiplication)				Universal
	Depth	Complexity	SIMD	Exact	Efficiency	Depth	Complexity	SIMD	Efficiency	
Word-wise BGV/BFV [19, 20]	-	-	-	-	-	1	$O(\lambda^2)$	✓	✓	✗
Word-wise CKKS [26]	-	-	-	-	-	1	$O(N \log N)$	✓	✓	✗
Poly. Approx. (CKKS) [27, 60]	$O(\log_2 n)$	$O(\sqrt{n})$	✓	✗	✓	1	$O(N \log N)$	✓	✓	✗
Poly. Interp. (BGV/BFV) [80]	$O(b)$	$O(\sqrt{2^b})$	✓	✓	✗	1	$O(\lambda^2)$	✓	✓	✗
Decomp. (BGV/BFV) [48]	$O(\log_2 b)$	$O(b \log_2 b)$	✓	✓	✓	-	-	-	-	✗
XCMP (BGV/BFV) [75]	$O(1)$	$O(1)$	✓	✓	✓	-	-	-	-	✗
General Functional Bootstrapping [9]	$O(b)$	$O(\sqrt{2^b} + b)$	✓	✓	✗	1	$O(N \log N)$	✓	✓	✗
Bit-wise TFHE [31]	$O(b)$	$O(b)$	✗	✓	✓	$O(b^2)$	$O(b^2 p)$ CMUX + $O(bp)$ KS	✗	✗	✓
Scheme Switching [18, 74]	-	$\Omega(2^b)$	✗	✓	✗	1	$O(\lambda^2)$	✓	✓	✓
Encoding Switching [102]	$O(\log_2 d + d \log_2 p)$	$O(d^2 \sqrt{p})$	✓	✓	✓	$O(1)$	$O(\lambda^2)$	✓	✓	✓

3 FHE UNIVERSALITY TAXONOMY AND COMPLEXITY ANALYSIS

Universality is an important and general feature of current real-world privacy-sensitive applications, even though there is little research investigating the universality of various FHE methods. Specifically, a universal FHE method must satisfy two conditions:

Mixed Linear and Non-Linear Operations: The universal FHE method must support both linear and non-linear operations on ciphertexts and allow them to be composed in any sequence. This enables complex computations that alternate between operation types, such as machine learning algorithms that use linear operations (e.g., matrix multiplications) followed by non-linear activation functions (e.g., ReLU). By seamlessly mixing linear and non-linear operations, universal FHE can accommodate various applications. **Non-Interactive Computation:** A universal FHE method should support mixed operations in a non-interactive manner—allowing computations on encrypted data without further communication between the data owner and the computing party. This non-interactivity is essential for scalability and efficiency, especially in distributed or cloud settings where reducing communication overhead is critical.

In this section, we evaluate the universality of prevalent FHE methods. We first classify them into two categories: universal FHE methods supporting both linear and non-linear operations, and non-universal FHE methods, which do not. We then provide a detailed analysis of each category in Sections 3.1 and 3.2, respectively, and summarize their universality properties and complexities in Table 1.

3.1 Analysis of universal FHE methods

This section analyzes three universal FHE methods that support both linear and non-linear operations in a non-interactive manner. The first is bit-wise FHE [31, 36], which inherently supports both operation types by constructing arithmetic and Boolean circuits. The second is scheme switching [18, 74], which combines bit-wise and word-wise FHE to leverage the strengths of both. The third is encoding switching (i.e., HEBridge [102]), which integrates linear and non-linear operations within word-wise BGV/BFV methods by bridging distinct encoding spaces using proposed homomorphic reduction and lifting functions.

Bit-wise FHE. TFHE [31] is a bit-wise FHE method that supports gate-level operations on encrypted data, enabling fast non-linear operations due to its efficient bit-wise processing. For a ciphertext

with b -bit precision, TFHE can perform non-linear functions with a depth of $O(b)$ and complexity of $O(\sqrt{2^b})$. However, TFHE’s linear operations are comparatively slower due to its bit-wise design. Multiplication on ciphertexts requires a depth of $O(b^2)$, along with $O(b^2 p)$ CMUX operations and $O(bp)$ key-switching. Additionally, TFHE lacks support for batch processing, i.e., SIMD, making linear computations less efficient than in word-wise FHEs, which can parallelize linear operations.

Scheme Switching. Scheme switching [18, 74] between bit-wise and word-wise ciphertexts enables universal FHE operations by leveraging the strengths of both schemes: linear functions are evaluated in word-wise FHE, and non-linear functions in bit-wise FHE. However, the complexity of scheme switching grows at least exponentially with input bit-width [16, 85], i.e., $\Omega(2^b)$. For example, evaluating 6-bit inputs takes 43.8 seconds, while 8-bit inputs require 162.7 seconds. Consequently, scheme switching becomes impractical for larger inputs.

Encoding Switching. Prior studies [48, 90] have shown that word-wise BGV/BFV support efficient and precise non-linear operations via polynomial interpolation. However, these methods cannot be seamlessly incorporated with linear operations because the plaintext space for linear operations in FV over \mathbb{Z}_{p^d} is different that used for non-linear operations in base-encoded FV (beFV) over \mathbb{Z}_p . To enable continuous evaluation of both operation types, Zhang et al. [102] introduced conversion techniques that employ a reduction function, which homomorphically converts the plaintext space from FV to beFV, allowing non-linear operations to follow linear operations, and a lifting function that restores the result from beFV back to the original FV space for subsequent linear operations. In terms of complexity, given that the FV plaintext modulus is p^d , the reduction function requires $O(d^2 \sqrt{p})$ multiplications and $O(d \log_2 p)$ multiplicative depth, while the lifting function requires $O(\sqrt{dp})$ multiplications and a depth of $O(\log_2 d + \log_2 p)$.

3.2 Analysis of un-Universal FHE methods

Word-wise FHEs, including BGV, BFV, and CKKS, naively support efficient linear operations with SIMD. However, computing non-linear functions in SIMD-enabled word-wise FHE is non-trivial. As a result, many efforts have been made to enable non-linear functions within word-wise FHE to leverage SIMD. For the CKKS scheme, the most common approach is polynomial approximation [27, 28, 60].

On the other hand, for integer-based BGV and BFV, Narumanchi et al. [80] proposed evaluating non-linear functions via polynomial interpolation over the base field.

CKKS with Polynomial Approximation. CKKS is a word-wise FHE method capable of encrypting real numbers. Since it supports approximate computation on floating-point numbers, the most common method for performing non-linear functions is polynomial approximation [61, 62]. Specifically, these methods use compositions of minimax approximation polynomials to approximate the sign function over a small interval, typically $[-1, -\epsilon] \cup [\epsilon, 1]$, with errors. Given an n -degree approximation polynomial, the complexity is $O(\sqrt{n})$ and the multiplicative depth is $\log_2 n$. The larger the interval and the smaller the error, the higher the polynomial degree required—thus increasing the number of multiplications and multiplicative depth. Most importantly, approximation can never be accurate around 0. These errors limit the applicability in domains where accuracy is critical, such as genomics [51, 84, 103] and finance [11, 45].

Overall, polynomial approximation strikes a balance between computational efficiency and accuracy, making it well-suited for applications such as neural network inference, where approximate computations are acceptable.

BGB/BFV with Polynomial Interpolation. In contrast to CKKS, computations in BFV/BGV are exact. Accordingly, non-linear functions can be implemented in BFV/BGV without approximation error. Narumanchi et al. [80] proposed computing non-linear functions via polynomial interpolation.

Take the non-linear comparison function as an example, since the comparison is a fundamental operation for many logic functions and is key to evaluating various non-linear operations [102]. The core idea is to convert the comparison between two encrypted integers, a and b , into a comparison between their difference, $z = b - a$, and zero. This is achieved by constructing a polynomial $P(x)$ with the property:

$$P(x) = \begin{cases} 1, & \text{if } x < 0; \\ 0, & \text{if } x \geq 0; \end{cases} \quad (1)$$

$P(x)$ are expressed using Lagrange interpolation:

$$P(x) = \sum_{i=0}^p \left(\prod_{j=0, j \neq i}^p \frac{x - x_j}{x_i - x_j} \right) \cdot y_i \mod p, \quad (2)$$

where x_i and y_i are chosen such that $y_i = 1$ for $x_i < 0$ and $y_i = 0$ for $x_i \geq 0$. Once the polynomial is constructed, it is evaluated homomorphically on the encrypted difference $\text{Enc}(a - b)$:

$$\text{Enc}(z) = P(\text{Enc}(a - b)). \quad (3)$$

Upon decryption, $\text{Dec}(\text{Enc}(z)) = 1$ indicates $a < b$, while $\text{Dec}(\text{Enc}(z)) = 0$ indicates $a \geq b$.

The polynomial interpolation has notable limitations. Specifically, the standard BGV and BFV support efficient linear operations over \mathbb{Z}_{p^r} [24, 44]. However, polynomial interpolation in \mathbb{Z}_{p^r} is impractical, as it leads to polynomials of degree $O(p^r)$, which grow exponentially with the input bit-width. Evaluating such high-degree polynomials is prohibitive in practice [48, 90]. Therefore, non-linear operations are typically performed in the base field \mathbb{Z}_p . The degree of interpolation polynomials equals p , requiring $O(\sqrt{p})$

non-scalar multiplications and $O(\log_2 p)$ multiplicative depth using the Paterson-Stockmeyer algorithm. To enable efficient evaluation, the plaintext space is typically confined to a small prime p , making it unsuitable for large inputs. The value of p ranges from $2 \sim 7$ in [90], is at most 173 in [48] and at most 257 in [79]. In conclusion, this limitation makes it impractical for real-world universal applications.

Polynomial Interpolation with Special Encoding. Aiming to enhance the capability of polynomial interpolation to support non-linear operations on large inputs, several special encoding-based methods have been proposed [48, 90], which encode a large integer as a vector of base- p digits: $a = (a_0, a_1, \dots, a_{r-1})$, where each $a_i \in \mathbb{Z}_p$. Consequently, a non-linear operation, e.g., a comparison between two integers a and b , is transformed into a lexicographical comparison of their respective digit vectors. Formally, the process begins with the decomposition of the input integers a and b into their digit representations:

$$\begin{aligned} \text{Enc}(a) &= (\text{Enc}(a_0), \text{Enc}(a_1), \dots, \text{Enc}(a_{r-1})), \\ \text{Enc}(b) &= (\text{Enc}(b_0), \text{Enc}(b_1), \dots, \text{Enc}(b_{r-1})). \end{aligned} \quad (4)$$

For each digit pair a_i and b_i , a polynomial $P(\cdot)$ is constructed to evaluate the comparison $a_i < b_i$ as Equation 1. This polynomial is then applied homomorphically:

$$\text{Enc}(z_i) = P(\text{Enc}(a_i - b_i)). \quad (5)$$

The overall comparison is obtained by combining these digit-wise comparisons lexicographically:

$$\text{Enc}(z) = \text{Enc}(z_0) + \sum_{i=1}^{r-1} \text{Enc}(z_i) \prod_{j=0}^{i-1} (1 - \text{Enc}(z_j)), \quad (6)$$

where $\prod_{j=0}^{i-1} (1 - \text{Enc}(z_j))$ ensures that less significant digits only contribute when more significant digits are equal.

By decomposing the large field into smaller subfields and performing comparisons at the digit level, this method achieves a significantly lower circuit depth. The depth of the circuit evaluating the comparison of two b -bit integers is given by:

$$\log_2 \log_p 2^b + \log_2(p - 1) + 4. \quad (7)$$

In contrast, directly interpolation on \mathbb{Z}_{p^r} has a multiplicative depth of $O(\log_2 p^r)$. This reduced depth makes the method more practical for larger fields. By addressing the challenge of deep circuit depth in polynomial interpolation, [48] enables faster homomorphic non-linear operations for practical applications in BGV/BFV. Despite its significant efficiency improvement, the special encoding-based ciphertexts do not support linear operations since the number is expressed in the vector form, making it not a universal scheme.

Similarly, other special encoding-based approaches have also been proposed to enhance the efficiency of non-linear operations in word-wise FHE. Unfortunately, these methods sacrifice universality, as the specialized ciphertext formats do not support linear operations. Below, we present a representative example.

Exponential Encoding (XCMP). The exponential encoding method for FHE private comparison, named XCMP, was first proposed by Lu et al. [75]. The core idea is based on the fact that the message space in HE is a polynomial ring, i.e., $\mathbb{Z}_p[x]/(x^n + 1)$, and that ciphertext multiplication corresponds to polynomial multiplication.

This enables the design of a special encoding scheme that encodes the values to be compared into the polynomial's degree coefficients.

The comparison process of integers a and b begins with encoding a and $-b$ as polynomials X^a and X^{-b} , respectively. These polynomials are encrypted into ciphertexts $\text{Enc}(X^a)$ and $\text{Enc}(X^{-b})$, which are used to perform:

$$\text{Enc}(C(X)) = T(X) \times \text{Enc}(X^a) \times \text{Enc}(X^{-b}) \mod (X^n + 1) \quad (8)$$

where $T(X) = 1 + X + \dots + X^{n-1}$. Note that the polynomial with a negative degree $X^{-b} \equiv -X^{n-b} \mod (X^n + 1)$.

Finally, decrypt $\text{Enc}(C(X))$ and evaluates the 0-th coefficient $C(X)[0]$, which indicates the comparison result:

- If $a \leq b$, X^{b-a} from $T(X)$ aligns with X^{a-b} , yielding $C(X)[0] = 1$.
- If $a > b$, the $(n - (a - b))$ -th term of $T(X)$ aligns with X^{a-b} , resulting in $C(X)[0] = -1$ due to the wrap-around: $X^{n-(a-b)} \times X^{a-b} \equiv -1 \mod (X^n + 1)$.

In other words, $C(X)[0] = 1$ if $a \leq b$, and $C(X)[0] = -1$ if $a > b$.

XCMP has a multiplicative depth and complexity of 1 due to its use of a single ciphertext multiplication. However, it is limited by a small input field, requiring both a and b to be smaller than the polynomial degree n . This constraint reduces efficiency for large inputs, as a larger n is needed [75]. Although there have been discussions about extending the domain to \mathbb{F}_{n^2} by decomposing numbers in base n and processing digits separately:

$$\mathbb{I}[a \leq b] = \mathbb{I}[a_1 \leq b_1] + \mathbb{I}[a_1 = b_1] \cdot \mathbb{I}[a_0 \leq b_0], \quad (9)$$

where the equality check result $\mathbb{I}[a_1 = b_1]$ must be converted to the XCMP format for multiplication with $\mathbb{I}[a_0 \leq b_0]$. This conversion involves Fermat's little theorem, resulting in a large multiplicative depth. As the result, XCMP performs well for small domains (less than 16 bits). For larger domains, performance significantly decreases [75].

General Functional Bootstrapping. Functional bootstrapping [9, 57, 65] extends traditional bootstrapping techniques by not only refreshing ciphertexts but also by enabling the homomorphic evaluation of arbitrary functions on encrypted data. The key idea is to express a target function as a look-up table (LUT) and then approximate this LUT via a constructed interpolation polynomial. Lee et al. [57] proposed functional bootstrapping for BFV, building upon regular BFV-style bootstrapping [39]. However, BFV-style bootstrapping efficiently supports only a small number of slots. To improve efficiency, Alexandru et al. [9] employed CKKS-style bootstrapping to increase slot utilization. Specifically, the target function is first encoded as a LUT over a finite base- p field \mathbb{Z}_p , and then approximated using a trigonometric Hermite interpolation polynomial $R(x)$. For instance, in the first-order case, the interpolation is constructed so that:

$$R\left(\frac{k}{p}\right) = f(k) \text{ for all } k \in \{0, 1, \dots, p-1\} \text{ and } R'\left(\frac{k}{p}\right) = 0 \quad (10)$$

This polynomial is then evaluated homomorphically using an efficient technique such as the Paterson-Stockmeyer algorithm. The complexity of evaluating a polynomial of degree d using this method is roughly proportional to $\sqrt{2d} + \log d$. For a first-order interpolation, where $d = p - 1$, the complexity increases approximately

as \sqrt{p} . Moreover, if further noise reduction is required, higher-order interpolations can be used (e.g., second-order or third-order), which increase the polynomial degree (to about $\frac{3p}{2}$ or $2p - 1$, respectively) and thus the overall cost. Moreover, extending the input space from \mathbb{Z}_p to \mathbb{Z}_{p^r} requires representing each integer as a vector of base- p digits, similar to special encoding-based polynomial interpolation in BFV/BGV [48, 102], which precludes the continuous evaluation of linear and non-linear functions on large inputs. Experiments in [9] demonstrate that general functional bootstrapping remains practical only for LUT evaluations up to 12-bit inputs.

4 FHE LIBRARIES SUPPORTING UNIVERSALITY

Various open-source FHE libraries implement the aforementioned FHE methods and provide high-level APIs that allow users to select encryption parameters tailored to their applications. In this section, we list six widely used libraries and summarize their support for these FHE methods.

Table 2 summarizes FHE libraries alongside the encryption schemes they implement and the types of supported operations. In general, HELib, Lattigo, and SEAL offer word-wise FHEs by supporting BFV, BGV and CKKS. In contrast, TFHE and Zama's TFHE-rs are designed primarily for bit-wise FHEs, offering Boolean gate-level primitives that can be composed into more complex functions. Notably, OpenFHE and Zama extend this universality further by facilitating Scheme Switching or hybrid approaches that bridge bit-wise and word-wise computations. Additionally, methods Polynomial Interpolation [80] and Polynomial Interpolation with Special Encoding [48] are open-sourced based on HELib, while Encoding Switching [102] is open-sourced using HELib. Functional bootstrapping techniques [9, 57] are implemented on Lattigo and OpenFHE but have not yet been open-sourced.

HELlib. The Homomorphic Encryption Library (HELlib) was introduced in 2013 by IBM and supports the BGV scheme (with bootstrapping), as well as CKKS. Polynomial Interpolation [48, 80] with BGV and Encoding Switching [102] are implemented on HELlib and have been open-sourced. HELlib is written in C++17 and uses the NTL mathematical library.

Lattigo. The lattice-based multiparty homomorphic encryption library in Go was first developed by the Laboratory for Data Security (LDS) at EPFL and is currently maintained by Tune Insight. It supports BFV, BGV, and CKKS. Lattigo enables scheme switching to compute non-linear functions. Functional FV bootstrapping [57] is built on the lattigo, but their implementation is not open-sourced yet.

SEAL. The Simple Encrypted Arithmetic Library (SEAL) was developed by Microsoft Research and was first released in 2015 [3]. SEAL supports leveled BFV, BGV, and CKKS.

TFHE. The Fast Fully Homomorphic Encryption Library over the Torus (TFHE) was released in 2016 by Chillotti et al. [31] and proposes the CGGI cryptosystem. The library exposes homomorphic Boolean gates such as AND and XOR but does not build complex functional units (e.g., adders, multipliers, and comparators) and leaves that to the developer.

OpenFHE. OpenFHE [7] is developed by Duality, NJIT, MIT, and other organizations. It supports a wide range of FHE methods,

including BGV, BFV, and CKKS with approximate bootstrapping, as well as DM/FHEW, CGGI/TFHE, and LMKCDEY for evaluating Boolean circuits. OpenFHE enables scheme switching between CKKS and between CKKS and FHEW/TFHE to evaluate non-smooth functions, e.g., comparison, using FHEW/TFHE functional bootstrapping. Recently, Alexandru et al. [9] leveraged OpenFHE to build general functional bootstrapping to enable the non-linear function for any RLWE ciphertexts, but it is not open-sourced yet. **Zama.** Zama [97] is an open-source cryptography company developing state-of-the-art FHE solutions for blockchain and AI. Its products include TFHE-rs [101] for Boolean and small integer arithmetic, Concrete [99] for compiling Python to FHE with LLVM, Concrete ML [98] for encrypted machine learning, and fhEVM for confidential smart contracts in Solidity.

Table 2: Compatibility of libraries with various FHE methods.

Methods	HElib	Lattigo	SEAL	TFHE	OpenFHE	Zama
Word-wise BGV	✓	✓	✓	✗	✓	✗
Word-wise BFV	✗	✓	✓	✗	✓	✗
Word-wise CKKS	✓	✓	✓	✗	✓	✗
Bit-wise TFHE	✗	✗	✗	✓	✓	✓
Poly. Interp.	✓	✗	✗	✗	✗	✗
Poly. Approx.	✓	✓	✓	✗	✓	✗
Scheme Switching	✗	✓	✗	✗	✓	✗
Encoding Switching	✓	✗	✗	✗	✗	✗
General Functional Bootstrapping	✗	✗	✗	✗	✗	✗

5 UNIVERSALITY REQUIRED PRIVACY-SENSITIVE WORKLOADS

5.1 Design Principle

To comprehensively evaluate and compare different universal FHE methods, we designed three workloads requiring universality, each representing basic computational units in privacy-sensitive applications, covering diverse scenarios:

- **Workload-1:** Compare $(\text{Enc}(A) \times \text{Enc}(B), \text{Enc}(C))$, a non-linear operation following a linear operation that serves as a basic unit for database queries. For instance, it can be used to determine if the product of two encrypted values exceeds a given threshold.
- **Workload-2:** Compare $(\text{Enc}(A), \text{Enc}(B)) \times \text{Enc}(C)$, a linear operation following a non-linear operation that forms a basic unit in decision tree algorithms. For example, this unit can evaluate a comparison between two encrypted feature values and then multiply the result by another encrypted value to determine the weight or selection criteria at a decision node.
- **Workload-3:** Compare $(\text{Enc}(A) \times \text{Enc}(B), \text{Enc}(C)) \times \text{Enc}(D)$, a composite sequence of linear, non-linear, and linear operations. It is a basic component of a neural network, where it can model the execution of a convolution layer, followed by a ReLU activation, and then another convolution layer.

5.2 Experimental Setup

To evaluate the performance of universal FHE methods, we conducted experiments on the above workloads using input vectors of 100 integers. Comparisons were performed on pairs of these vectors

to yield an output vector of 100 Boolean values, while element-wise multiplications on two vectors produced a vector of 100 integers.

System Setup. The experiments are conducted on a server equipped with an AMD Ryzen Threadripper PRO 3955WX (2.2 GHz) and 125 GB of RAM. All tests are run in single-thread mode for a fair comparison. For Encoding Switching, we use the official open-sourced implementation from HEBridge [102] built on HELib. Scheme Switching is adopted from OpenPEGASUS [74] built on Microsoft SEAL, and TFHE-rs [101] are used for the bit-wise TFHE method. All encryption parameters are configured to maintain a security level above 128 bits, following the "BKZ-beta" classical cost model from the LWE estimator [8] unless stated otherwise.

Parameter Setup. For both BGV/BFV-based methods, i.e., Encoding Switching and Scheme Switching, we set the secret distribution to be a Hamming weight distribution over the set of ternary polynomials with coefficients in $\{0, 1, -1\}$, such that each secret has exactly $h = 64$ nonzero entries.

For Encoding Switching, which is built on HELib, the plaintext modulus in the FV space is given by p^r . We choose the parameters p and r based on the input bit-width. Specifically, we use the pairs $\{(4, 4), (5, 4), (7, 5), (17, 4)\}$ as (p, r) for input bit-widths of 6, 8, 12, and 16, respectively. The multiplicative depth is determined by the ciphertext capacity, defined as $\log_2 \frac{q}{\eta}$, where q is the ciphertext modulus and η is the current noise bound. Accordingly, we set $\log_2 q$ to $\{256, 320, 488, 648\}$ for 6, 8, 12, and 16 bits of input. After choosing p , r , and q , we select an appropriate degree for the polynomial modulus to ensure that the security level satisfies $\lambda > 120$. In this context, the cyclotomic order of the polynomial ring m is chosen so that the ring degree n and the order of the base prime, $d = \text{Ord}(p)$, meet the following values: $\{(m, n) = (13201, 12852), (16151, 15600), (25301, 25300), (31621, 31212)\}$, corresponding to 6, 8, 12, and 16 bits of input, respectively.

For Scheme Switching, which is implemented using Microsoft SEAL, we set the parameters in a way that mimics the above setup to ensure a fair comparison. Specifically, we explicitly configure the plaintext modulus p^r using the identical (p, r) pairs, and we select the ciphertext modulus q with the same $\log_2 q$. Then, we choose a polynomial modulus degree n (which in SEAL must be a power of two) such that the overall security level satisfies $\lambda > 120$. Although Microsoft SEAL does not allow explicit setting of the cyclotomic order m and the order of p (denoted as d), these are implicitly determined by our choices of n and p (with m being typically interpreted as $2n$ for the cyclotomic polynomial $x^n + 1$, and $d = \text{Ord}(p)$ accordingly).

For the bit-wise TFHE method, TFHE-rs library provides objects `FheUint6`, `FheUint8`, `FheUint12` and `FheUint16` to represent different input bit precision.

Our benchmarks focus exclusively on server-side FHE operations, including all evaluation tasks except key generation and encryption/decryption, as these are one-time user costs. No custom parallelization is applied to HELib, OpenFHE, or TFHE-rs, as each defaults to single-core execution.

5.3 Experimental Results

The experimental results are presented in Figure 2. The TFHE [31], which evaluates both linear and non-linear operations in a bit-wise

manner, handles non-linear operations efficiently but performs sub-optimally for linear operations, especially for larger input bit lengths, due to increased circuit depth. More importantly, TFHE lacks SIMD support, a key advantage of word-wise FHE methods that can process multiple ciphertexts simultaneously.

Scheme Switching [18], on the other hand, evaluates linear operations using word-wise FHE and non-linear operations using bit-wise FHE. However, the switching process’s complexity grows exponentially with the input bit-width. For instance, on workload-1, evaluating 6-bit inputs takes 9.87 seconds, while 8-bit inputs require 32.1 seconds. This exponential growth makes scheme switching impractical for larger inputs, as similar prohibitive runtimes were estimated for 12-bit and 16-bit inputs.

In contrast, Encoding Switching [102] avoids the overhead of expensive scheme switching and leverages SIMD capabilities by performing both linear and non-linear operations in word-wise FHE. For 8-bit inputs, evaluating workload-1 takes only 15.5 seconds. More importantly, Encoding Switching scales well to larger inputs, taking 23.0 seconds for 12-bit inputs and 46.9 seconds for 16-bit inputs, yielding 14.1 \times and 25.4 \times improvements over the bit-wise TFHE and Scheme Switching, respectively.




Figure 2: Compare all comparisons on three workloads.

6 UNIVERSALITY REQUIRED PRIVACY-SENSITIVE APPLICATIONS

6.1 Design Principle

In this section, we evaluate universal FHE methods on five representative privacy-sensitive applications that require universal capabilities: private graph algorithms, private decision tree inference, private database aggregation, private neural network inference, and private sorting. These applications cover a diverse range of real-world use cases across domains such as healthcare, finance, and other industrial sectors. Each application incorporates both linear and non-linear operations and is built on the workloads analyzed in Section 5.

6.2 Experimental Setup

Generally, we evaluate various applications using Scheme Switching, Encoding Switching, and bit-wise TFHE, with experimental setup details identical to those described in Section 5.2. For private graph evaluation, a graph with n nodes is represented as an $n \times n$ adjacency matrix. In the case of bit-wise TFHE, each element of the matrix is individually encrypted into multiple ciphertexts (for multiple bits), whereas for SIMD-enabled methods, each row is batched

into a single ciphertext. For private decision tree evaluation, each node is encrypted as multiple ciphertexts for bit-wise FHE, while for SIMD-enabled methods, an entire layer is batched into one ciphertext. In private sorting, each element is encrypted into a single ciphertext as described in [48]. For private database aggregation, every element is encrypted into multiple ciphertexts when using bit-wise TFHE, whereas for SIMD-enabled methods, each column is batched into a single ciphertext. Finally, for private neural network inference, each parameter is encrypted into multiple ciphertexts with bit-wise TFHE, while for SIMD-enabled methods, each filter or weight matrix is batched into a single ciphertext.

6.3 Private Graph

Private graph algorithms [35, 64, 76, 92] enable secure computations on graph-structured data while preserving individual privacy [5, 77, 92]. In this threat model, an untrusted server performs computations without learning any details about the input graph. A common operation in graph algorithms is selecting the shorter path from paths P_1 and P_2 , which can be implemented as: $(P_1 > P_2) \times P_2 + (1 - (P_1 > P_2)) \times P_1$; this computation exemplifies workload-2, where a linear operation follows a non-linear operation. In this section, we demonstrate the performance of universal FHE methods on the private Floyd-Warshall algorithm [38].

Floyd-Warshall on Plaintext. The Floyd-Warshall algorithm is a classic algorithm for finding the shortest paths between all pairs of nodes in a graph G . It returns on a distance matrix D , where $D[i, j]$ represents the shortest known distance from node i to node j .

As shown in Algorithm 1, the algorithm iteratively updates D by considering each node as an intermediate node in potential paths. During each iteration, it updates the distance $D[i, j]$ for each pair of nodes (i, j) by comparing it with $D[i, k] + D[k, j]$, which represents a path from i to j via k . This process is repeated for all nodes k , ensuring all potential intermediate nodes are considered. After all iterations, the D contains the shortest paths between all pairs of nodes in the graph. The algorithm has a time complexity of $O(|V|^3)$, where $|V|$ is the nodes number.

Algorithm 1 Floyd-Warshall Algorithm on Plaintext

```

1: Input:  $G$ : adjacency matrix
2: Output:  $D$ : matrix of shortest path distances between each pair of nodes,  $P$ : matrix of predecessors for each node
3: # Initialize distance array  $D$  and predecessor array  $P$ 
4: for each vertex  $v$  in  $G$  do
5:    $D[v] = G[v]$ ,  $P[v] = \text{null}$ 
6: # Iteratively update the distances
7: for each vertex  $k$  in  $G$  do
8:   for each vertex  $i$  in  $G$  do
9:     for each vertex  $j$  in  $G$  do
10:      if  $D[i, k] + D[k, j] < D[i, j]$  then
11:         $D[i, j] = D[i, k] + D[k, j]$ 
12:         $P[i, j] = P[k, j]$ 
13: return  $D$ ,  $P$ 

```

Bit-wise FHE Private Floyd-Warshall. To implement a private Floyd-Warshall algorithm, a naive method [35] encrypts each element of the adjacency matrix G to ciphertext. The server then conducts the same procedure on the encrypted G as delineated in Algorithm 1. For if-else operations based on comparison results, the server uses HE multiplication to implement the conditional logic.

Specifically, lines 10-11 of Algorithm 1 can be computed in HE as:

$$D[i, j] = M \cdot (D[i, k] + D[k, j]) + (1 - M) \cdot D[i, j], \quad (11)$$

where M is the encrypted comparison result. This naive private Floyd-Warshall requires $|V|^3$ private comparisons and $4|V|^3$ private multiplications, resulting in significant computational overhead.

Word-wise FHE Private Floyd-Warshall with SIMD. The cubic complexity of bit-wise naive private Floyd-Warshall stems from the three nested loops. However, we can optimize this process by leveraging the SIMD mechanism of word-wise FHE with the algorithm's inherent parallelism.

After fixing the middle node k , updates for paths from node i to other nodes via k can be computed simultaneously. This parallelism allows SIMD by encoding the adjacency matrix row by row.




Figure 3: Toy example of SIMD update distances between node 1 to others through node 4.

Figure 3 illustrates this SIMD optimization for a graph with 4 nodes, focusing on updating distances from node 1 to all others through node 4. The process is: ① Generate ciphertext $D_{1 \rightarrow 4}$ by multiplying $[D_{1 \rightarrow 1}, D_{1 \rightarrow 2}, D_{1 \rightarrow 3}, D_{1 \rightarrow 4}]$ with plaintext $[0, 0, 0, 1]$, then replicating $D_{1 \rightarrow 4}$ across all slots by rotations and addition. ② Compute $D_{1 \rightarrow 4 \rightarrow all}$ by adding $D_{1 \rightarrow 4}$ to $D_{4 \rightarrow all}$, representing new potential distances via node 4. ③ Compare $D_{1 \rightarrow 4 \rightarrow all}$ and $D_{1 \rightarrow all}$, yielding mask M (1 where new distance is shorter, 0 otherwise). ④ Update distances by:

$$D_{1 \rightarrow all} = M \cdot D_{1 \rightarrow 4 \rightarrow all} + (1 - M) \cdot D_{1 \rightarrow all}. \quad (12)$$

only paths via node 4 is shorter than original one will be updated.

This SIMD-enabled approach, as shown in Algorithm 2, allows each row of the adjacency matrix to be encrypted as a single ciphertext. Consequently, paths from node i to all other nodes via node k can be compared in a single operation, reducing the total number of HE comparisons from $O(|V|^3)$ to $O(|V|^2)$, and HE multiplications from $O(4|V|^3)$ to $O(4|V|^2)$.

Experimental Analysis. To assess the performance of different universal FHE methods for the Private Floyd-Warshall algorithm, we conducted experiments across various input bit widths and graph sizes. Figure 4 (a) compares the performance for a graph with 32 nodes under different input bit lengths, where Encoding Switching exhibits the best performance and Scheme Switching the worst, particularly for larger bit widths. Figure 4 (b) shows results for graphs of varying sizes using an 8-bit input. As noted earlier, the runtime of bit-wise TFHE grows approximately as $O(|V|^3)$ due to the lack of SIMD support, while the complexities for both Scheme Switching and Encoding Switching are reduced to roughly $O(|V|^2)$. For 8-bit inputs, TFHE takes 13.3 seconds, and Encoding Switching takes 17.3 seconds on a 16-node graph; however, for a 128-node

Algorithm 2 Floyd-Warshall Algorithm on Ciphertext Using SIMD

```

1: Input:  $G$ : adjacency matrix (encrypted row-by-row)
2: Output:  $D$ : matrix of shortest path distances between each pair of nodes (encrypted),  $P$ : matrix of predecessors for each node (encrypted)
3: # Initialize distance array  $D$  and predecessor array  $P$ 
4: for each vertex  $v$  in  $G$  do
5:    $D[v] = G[v]$ ,  $P[v] = \text{null}$ 
6: # Iteratively update the distances using SIMD
7: for each vertex  $k$  in  $G$  do
8:   for each vertex  $i$  in  $G$  do
9:      $D_{i \rightarrow k} = \text{SIMD}(D[i, k])$  # Broadcast  $D[i, k]$  to all slots
10:     $D_{k \rightarrow all} = D[k, :]$ 
11:     $D_{new} = D_{i \rightarrow k} + D_{k \rightarrow all}$ 
12:     $M = D_{new} < D[i, :]$  # Mask for minimum
13:     $D[i, :] = M \cdot D_{new} + (1 - M) \cdot D[i, :]$ 
14:     $P[i, :] = M \cdot P[k, :] + (1 - M) \cdot P[i, :]$ 
15: return  $D, P$ 

```

graph, TFHE's runtime escalates dramatically to 4,933 seconds, whereas Encoding Switching completes in 632 seconds.




Figure 4: Homomorphic Floyd-Warshall. (a) Execution time for a 32-node graph with varying input bit lengths. (b) Execution time for graphs of different sizes with 8-bit input.

6.4 Private Decision Tree Evaluation

Private Decision Tree Evaluation (PDTE) [4, 6, 13, 52] allows a server to provide predictions using a private decision tree on a client's confidential input, preserving both input privacy (the server learns nothing about the client's data) and model privacy (the client only obtains the inference result without learning any details of the decision tree).

Decision trees primarily involve comparisons and traversals [88]. In these trees, comparisons are made between the client's inputs and decision thresholds, yielding ciphertext values $c_i \in [0, 1]$ at each node, while traversals determine the active path by computing the product $c_1 \cdot c_2 \cdots c_d$, where c_i is the comparison result at the i -th level and d is the tree's depth. The correct decision is identified by the leaf node that decrypts to 1, with all other leaves decrypting to 0. Since private decision tree evaluation involves both linear and non-linear operations, it requires universal FHE methods. Specifically, linear operations following non-linear comparisons correspond to workload-2 defined in Section 5. Notably, the traversal operation $c_1 \cdot c_2 \cdots c_d$ can be reformulated [6] as:

$$c_1 \cdot c_2 \cdots c_d = \overline{c_1 + c_2 + \cdots + c_d} \quad (13)$$

which significantly reduces the required multiplicative depth, further improving the efficiency, especially for decision trees with larger depths.

Experimental Analysis. In Figure 5 (a), we evaluate the performance of various universal FHE methods on a complete decision tree with a depth of 6 across different input bit lengths. Figure 5 (b) shows the performance for complete decision trees of varying depths using an 8-bit input. The results indicate that for deeper trees, the word-wise Encoding Switching method performs more efficiently than others since it enables SIMD so that can compare the same input with multiple thresholds simultaneously while avoiding the significant cost of scheme switching between bit-wise and word-wise ciphertexts. On the other hand, the word-wise Scheme Switching perform worst because the significant frequent of switching since every non-linear operation following a linear operation.




Figure 5: Private Decision Tree. (a) Execution time of inference for a depth-6 private decision tree under varying input bit precision; (b) Execution time of private decision trees with different depths using 8-bit input.

6.5 Private Sorting

Private sorting is a protocol that enables secure ordering of encrypted data without revealing the underlying values or the sorting process. In this scenario, a client encrypts its data and sends to a server, which then performs the sorting directly on the ciphertexts. Private sorting is critical and commonly used in privacy sensitive area such as financial. For instance, several financial institutions might want to merge and sort transaction data to detect market trends without exposing individual client details.

The most efficient homomorphic sorting algorithm in terms of running time is the direct sorting algorithm proposed by Çetin, et al. [23]. This approach directly determines final position of each element by counting the number of elements less than it. Specifically, for an array of size m , the server performs $m(m-1)/2$ homomorphic less-than operations and an equal number of homomorphic additions on the less-than operations' result for computing the Hamming weight, which yields the correct indices I for each element.

Once the indices are established, the elements are repositioned in one efficient step using a conditional multiplexer (CMUX)-like operation, as follows:

$$S[i] = \sum_{j=1}^m X[j] \cdot EQ(i, I[j]), \quad (14)$$

where $S[i]$ is the element at the position I of sorted array S and X is the input array. $EQ(i, I[j])$ is 1 if the i equals the index $I[j]$ corresponding to $X[j]$, and 0 otherwise. Constructing one position

$S[i]$ requires m equality evaluation and m multiplication, so sorting the entire encrypted array requires $O(m^2)$ non-linear operations (less-than and equality evaluation) and $O(m^2)$ linear operations (additions and multiplications).

Thus, homomorphically sorting an encrypted array of m elements demands universal FHE methods, particularly as the linear operations are performed on the results of non-linear operations, corresponding to our workload-2.

Figure 6 shows an example to private sort a three-dim array $X = [a, b, c]$, the server first computes a less-than matrix L by:

$$L_{ij} = \begin{cases} LT(x_i, x_j) & \text{if } i < j, \\ 0 & \text{if } i = j, \\ 1 - LT(x_j, x_i) & \text{if } i > j. \end{cases} \quad (15)$$

Next, the server computes the Hamming weight of each row, i.e., i -th row, to obtain how many elements are larger than x_i . The Hamming weight, i.e., $I[a]$, $I[b]$ and $I[c]$, denote the indices of a , b and c in the sorted array, respectively. Finally, the sorted array S is constructed by selecting elements $X[j]$ where the index matches the hamming weight, with Equation 14.




Figure 6: Toy example of direct sorting on the array $[a, b, c]$. (a) demonstrates the computation of the Comparison Matrix L and the derivation of indices from the hamming weight of L , where each element's index corresponds to the count of elements greater than it. (b) shows the sorting mechanism where each element is multiplied by a binary value: only the element with a matching index retains its value (multiplied by 1), while others are set to zero (multiplied by 0).

Experimental Analysis. We conduct experiments using various universal FHE methods on the private sorting tasks. Figure 7 (a) presents the performance of different FHE methods under different input bit precisions, and Figure 7 (b) shows the performance under different array dimensions. For instance, sorting an 8-element array requires 54.4 seconds with TFHE, compared to 645 seconds for Encoding Switching and 7,062 seconds for Scheme Switching. This efficiency advantage is due to two key factors. First, unlike other applications, private direct sorting on a single array cannot leverage SIMD [48], which diminishes the benefits of word-wise methods like Scheme Switching and Encoding Switching. Second, in this context, one of the operands in the linear operations is the output of a non-linear operation (either 0 or 1), resulting in a simpler circuit compared to linear operations on arbitrary integers. Conversely, Scheme Switching shows the lowest efficiency because its frequent switching, triggered by each non-linear operation following a linear one, adds significant overhead.




Figure 7: Homomorphic Sorting. (a) Execution time of sorting an 8-element array under varying input bit precision; (b) Execution time of sorting arrays of different lengths with 8-bit input.

The reasons are two-fold; firstly, unlike other applications, private direct sorting an encrypted array cannot leverage the SIMD [48], which disables the advantages of word-wise Scheme Switching and Encoding Switching. Secondly, one of the linear operation operands is the result of non-linear operation, i.e., 0 or 1. So that the TFHE circuit would be much less than linear operation between two random nature integers. On the other hand, Scheme Switching presents the lowest efficiency across all input bits and array sizes because of the frequent switching since every non-linear operation will follow a linear operation.

6.6 Private Database Aggregation

Private Database Aggregation (PDBA) is a protocol that protects client privacy during aggregate queries on cloud-stored databases, such as Microsoft Azure SQL Server [2] and AWS Aurora [1]. Recently, Leidos has partnered with AWS to explore the solution of using FHE to protect database [12]. It allows clients to securely store encrypted data on a cloud server and perform operations, like summation, without disclosing data or the details of the queries. The server processes these functions on encrypted data and returns only the encrypted aggregate results, ensuring the server does not access actual data or query specifics.

Recently, researchers have developed encrypted databases to ensure data privacy while enabling encrypted data processing. Cryptography-based solutions [43, 82, 83, 91], utilize cryptographic primitives such as DET [14], SE [22, 32], partially PHE [81, 86] and order OPE/ORE [17, 25]. Recent advancements have explored the application of FHE in HEDA [85], HE3DB [16], ArcEDB [106] and Engorgio [15].

An aggregation query typically involves a combination of linear operations and non-linear operations. Consider the following SQL query as an example:

```
SELECT ID FROM emp
WHERE salary * work_hours BETWEEN 5000 AND 6000
AND salary + bonus BETWEEN 700 AND 800;
```

Listing 1: SQL Query for Employee Salaries.

In this query, the filtering predicates "salary * work_hours BETWEEN 5000 AND 6000" and "salary + bonus BETWEEN 700 AND 800" involve a non-linear operation following a linear operation. This requirement necessitates universal FHE methods and aligns with workload-1 defined in Section 5.

While word-wise FHE methods work well for linear operations, they are less effective for non-linear operations like comparisons and logic filtering. To address this limitation and leverage the high efficiency of SIMD, Ren et al. introduce a novel framework HEDA [85]. The core idea is to utilize scheme switching between RLWE and LWE ciphertexts. This approach conducts non-linear operations on LWE ciphertexts, where such operations are more naturally supported. Then, it converts the results back to RLWE ciphertexts to perform linear operations in a SIMD fashion.

Experimental Analysis. As described in Section 6.6, SQL queries involve combinations of non-linear and linear operations. To evaluate the performance of each universal FHE method on the private database query, we conducted experiments on Query 1 across databases with different numbers of rows under 8-bit precision, as illustrated in Figure 8 (b).

The results indicate that the word-wise universal FHE methods generally outperform the bit-wise TFHE, largely due to their use of SIMD capabilities, a feature that the bit-wise TFHE lacks. Furthermore, as the number of rows increases, Encoding Switching demonstrates better efficiency than Scheme Switching. For example, Encoding Switching takes 115 seconds for a 512-row database, compared to 119 seconds for Scheme Switching. This advantage is attributed to Encoding Switching's ability to perform non-linear operations in a SIMD-style, while Scheme Switching processes non-linear operations on bit-wise ciphertexts and cannot leverage SIMD.




Figure 8: (a) compares the performance of various FHE DNN works on CIFAR-10 using VGG-9. (b) Private database implemented by different schemes.

6.7 Private Neural Network

One of the most valuable applications of FHE is in facilitating privacy-preserving deep learning, i.e., Deep Learning as a Service (DLaaS) [78, 89], which is commonly used in information sensitive areas such as financial [45] and healthcare [51]. This scenario involves a server that possesses a deep learning model, i.e., neural network, and a client who encrypts their input data before sharing with the untrusted server. The server processes encrypted input to provide predictions without accessing any sensitive information, thereby ensuring a secure inference service.

Neural networks consist of linear layers such as convolutions and non-linear layers like ReLU and MaxPool [54, 56]. The combination of linear and non-linear operations required by these layers necessitates universal FHE methods; indeed, the iterative sequence

of linear and non-linear layers forms the basis of workload-3 as defined in Section 5.

Despite the use of universal FHE methods, some research leveraged the specific feature of neural network to improve efficiency but sacrificing some accuracy. Specifically, polynomial approximation-based methods [58, 61, 62] leverage the noise resilience of neural networks to enhance efficiency while reducing multiplicative depth and complexity. Although these approximations are typically accurate only within a narrow range (often between -1 and 1), activation layer inputs are usually constrained (e.g., within $[-32, 32]$ or $[-64, 64]$). By scaling these inputs to fit within $[-1, 1]$ using a predetermined scalar, servers can effectively implement these approximation techniques throughout the network, as demonstrated in approaches like DaCapo [30], OPT-CNN [50], and Hetal [63] among others [33, 59]. However, determining the appropriate scaling factor is critical: if it is too large, inputs may not be sufficiently constrained within $[-1, 1]$; if too small, the inputs become overly diminished, resulting in significant approximation errors near 0.

Another approximation strategy simulates non-linear activation functions using linear operations. For example, CryptoNets [41] replaces ReLU with the square function, a method later adopted by frameworks like LoLa [21] and Falcon [71]. This approach avoids the overhead of homomorphic non-linear operations by training the neural network with a substitute activation function.

Experimental Analysis. To evaluate the performance of various private neural network inference implementations described in Section 6.7, we conducted experiments using the CIFAR-10 dataset [53] on a VGG-9 model. VGG-9 is a streamlined version of the VGG network family, comprising 9 layers including convolutional and fully connected layers. Typically, it consists of 6 convolutional layers followed by 3 fully connected layers, with ReLU activations and max pooling operations interleaved between the convolutional layers. Encryption parameters were set as specified in each referenced paper to ensure at least 128-bit security.

Figure 8 (a) shows that implementations using the square function as an activation substitute, such as Falcon [71] and LoLa [21], yield lower performance. This is primarily because training with degree-2 polynomial activations can lead to instability issues like exploding or vanishing gradients [10, 93]. Although the square function enables faster inference, it is generally only suitable for smaller, less complex neural networks.

In contrast, private neural networks that employ approximation-based activation functions, such as OPP-CNN [50] and EVA [33], achieve higher accuracy levels. Meanwhile, TFHE’s performance is significantly slower compared to implementations using interpolation approximations. This discrepancy arises because neural network architectures predominantly consist of linear layers, which benefit from the SIMD capabilities of word-wise CKKS schemes. Bit-wise TFHE, lacking SIMD support and optimized linear operations, exhibits lower performance. Additionally, the Scheme Switching method (SHE [69]) achieves the highest accuracy but at slower inference times than the approximation-based approaches.

7 RECOMMENDATION AND DISCUSSION

Considering the capabilities of each FHE method, we offer recommendations tailored to the types of operations required by a

given privacy-preserving application and whether parallelism can be exploited. Figure 9 summarizes these guidelines.

First, if the application involves only linear operations, word-wise FHE methods (BGV, BFV, CKKS) are the most efficient choice. Conversely, if the application requires only non-linear operations, the next consideration should be whether the application can benefit from the word-wise SIMD technique. If SIMD does not offer an advantage, a bit-wise FHE method, such as TFHE, is the most efficient option. However, if SIMD is useful, several techniques enable non-linear functionalities in a word-wise context, including direct polynomial interpolation [80], decomposition-based polynomial interpolation [48], general functional bootstrapping [39], or special encoding methods like XCMP [6]. These word-wise methods enable parallel computation to improve efficiency.

When an application demands both linear and non-linear computations (i.e., universality-required applications), the next question is whether the workload can benefit from SIMD parallelism. If SIMD does not offer advantages, e.g., private sorting for an array, a bit-wise FHE method, such as TFHE, remains the most practical choice. If SIMD is beneficial, the final consideration is whether exact, error-free results are necessary. If approximate results are acceptable, as is common in tasks like neural network inference, where noise tolerance is inherent, CKKS with approximation [27, 60] offers the most efficiency. If exact computation is required, the Encoding Switching [102] between beFV and FV, and the Scheme Switching [74] between word-wise and bit-wise FHE can deliver the best performance.

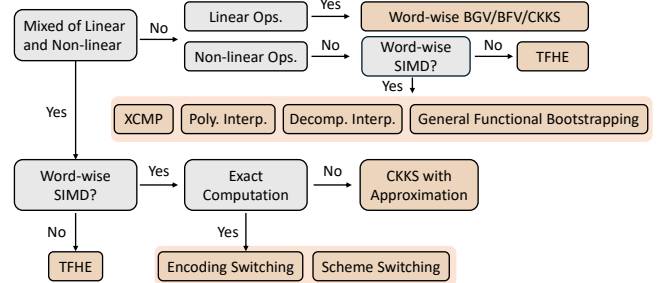


Figure 9: Flow chart to guide users to the FHE method that best fits the requirements of their privacy-sensitive applications and available resources.

8 CONCLUSION

In this paper, we systematically evaluate the universality of Fully Homomorphic Encryption (FHE)—its ability to perform non-interactive mixed linear and non-linear operations—a critical requirement for real-world privacy-sensitive applications like neural networks, graph analytics, and secure database queries. Through a three-stage methodology combining theoretical complexity analysis, micro-benchmarks on mixed-operation workloads, and empirical evaluations across five applications, we reveal significant overheads in existing universal FHE solutions, emphasizing the urgent need for optimizations that bridge efficiency gaps while preserving security. To guide practitioners, we provide recommendations for selecting

optimal FHE methods based on operation types, parallelism opportunities, and error tolerance, enabling developers to balance efficiency, precision, and security in privacy-preserving systems.

REFERENCES

- [1] 2022. AWS Aurora. <https://aws.amazon.com/rds/aurora/>.
- [2] 2022. Microsoft Azure SQL Server. <https://azure.microsoft.com/en-us/products/azure-sql/database>.
- [3] 2022. Microsoft SEAL. <https://github.com/microsoft/SEAL>.
- [4] Adi Akavia, et al. 2022. Privacy-preserving decision trees training and prediction. *ACM Transactions on Privacy and Security* 25, 3 (2022), 1–30.
- [5] Cuneyt Akcora, et al. 2012. Privacy in social networks: How risky is your social graph?. In *2012 IEEE 28th International Conference on Data Engineering*. IEEE, 9–19.
- [6] Rasoul Akhavan Mahdavi, et al. 2023. Level up: Private non-interactive decision tree evaluation using levelled homomorphic encryption. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 2945–2958.
- [7] Ahmad Al Badawi, et al. 2022. Openfhe: Open-source fully homomorphic encryption library. In *proceedings of the 10th workshop on encrypted computing & applied homomorphic cryptography*. 53–63.
- [8] Martin R Albrecht, et al. 2015. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology* 9, 3 (2015), 169–203.
- [9] Andreea Alexandru, et al. 2024. General functional bootstrapping using CKKS. *Cryptology ePrint Archive* (2024).
- [10] Ramy E Ali, et al. 2020. On polynomial approximations for privacy-preserving and verifiable relu networks. *arXiv preprint arXiv:2011.05530* (2020).
- [11] Frederik Armknecht, et al. 2015. A guide to fully homomorphic encryption. *Cryptology ePrint Archive* (2015).
- [12] AWS. 2023. Enable fully homomorphic encryption with Amazon SageMaker. <https://aws.amazon.com/cn/blogs/machine-learning/enable-fully-homomorphic-encryption-with-amazon/>.
- [13] Sofiane Azogagh, et al. 2022. Probonite: Private one-branch-only non-interactive decision tree evaluation. In *Proceedings of the 10th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*. 23–33.
- [14] Mihir Bellare, et al. 2007. Deterministic and efficiently searchable encryption. In *Advances in Cryptology-CRYPTO 2007: 27th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2007. Proceedings* 27. Springer, 535–552.
- [15] Song Bian, et al. 2025. Engorgio: An Arbitrary-Precision Unbounded-Size Hybrid Encrypted Database via Quantized Fully Homomorphic Encryption. *Cryptology ePrint Archive* (2025).
- [16] Song Bian, et al. 2023. He3db: An efficient and elastic encrypted database via arithmetic-and-logic fully homomorphic encryption. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 2930–2944.
- [17] Alexandra Boldyreva, et al. 2009. Order-preserving symmetric encryption. In *Advances in Cryptology-EUROCRYPT 2009: 28th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Cologne, Germany, April 26-30, 2009. Proceedings* 28. Springer, 224–241.
- [18] Christina Boura, et al. 2020. Chimera: Combining ring-lwe-based fully homomorphic encryption schemes. *Journal of Mathematical Cryptology* 14, 1 (2020), 316–338.
- [19] Zvika Brakerski. 2012. Fully homomorphic encryption without modulus switching from classical GapSVP. In *Annual cryptography conference*. Springer, 868–886.
- [20] Zvika Brakerski, et al. 2014. (Leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)* 6 (2014), 1–36.
- [21] Alon Brutzkus, et al. 2019. Low latency privacy preserving inference. In *International Conference on Machine Learning*. PMLR, 812–821.
- [22] David Cash, et al. 2013. Highly-scalable searchable symmetric encryption with support for boolean queries. In *Advances in Cryptology-CRYPTO 2013: 33rd Annual Cryptology Conference, Santa Barbara, CA, USA, August 18-22, 2013. Proceedings, Part I*. Springer, 353–373.
- [23] Gizem S Çetin, et al. 2015. Depth optimized efficient homomorphic sorting. In *Progress in Cryptology-LATINCRYPT 2015: 4th International Conference on Cryptology and Information Security in Latin America, Guadalajara, Mexico, August 23-26, 2015. Proceedings* 4. Springer, 61–80.
- [24] Hao Chen, et al. 2018. Homomorphic lower digits removal and improved FHE bootstrapping. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 315–337.
- [25] Nathan Chenette, et al. 2016. Practical order-revealing encryption with limited leakage. In *Fast Software Encryption: 23rd International Conference, FSE 2016, Bochum, Germany, March 20-23, 2016, Revised Selected Papers* 23. Springer, 474–493.
- [26] Jung Hee Cheon, et al. 2017. Homomorphic encryption for arithmetic of approximate numbers. In *Advances in Cryptology-ASIACRYPT 2017: 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I* 23. Springer, 409–437.
- [27] Jung Hee Cheon, et al. 2020. Efficient homomorphic comparison methods with optimal complexity. In *Advances in Cryptology-ASIACRYPT 2020: 26th International Conference on the Theory and Application of Cryptology and Information Security, Daejeon, South Korea, December 7–11, 2020, Proceedings, Part II* 26. Springer, 221–256.
- [28] Jung Hee Cheon, et al. 2019. Numerical method for comparison on homomorphically encrypted numbers. In *International conference on the theory and application of cryptology and information security*. Springer, 415–445.
- [29] Jung Hee Cheon, et al. 2022. Efficient homomorphic evaluation on large intervals. *IEEE Transactions on Information Forensics and Security* 17 (2022), 2553–2568.
- [30] Seonyoung Cheon, et al. [n. d.]. DaCapo: Automatic Bootstrapping Management for Efficient Fully Homomorphic Encryption. ([n. d.]).
- [31] Ilaria Chillotti, et al. 2020. TFHE: fast fully homomorphic encryption over the torus. *Journal of Cryptology* 33, 1 (2020), 34–91.
- [32] Reza Curtmola, et al. 2006. Searchable symmetric encryption: improved definitions and efficient constructions. In *Proceedings of the 13th ACM conference on Computer and communications security*. 79–88.
- [33] Roshan Dathathri, et al. 2020. EVA: An encrypted vector arithmetic language and compiler for efficient homomorphic computation. In *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*. 546–561.
- [34] Xianglong Deng, et al. 2024. Trinity: A General Purpose FHE Accelerator. *Accepted by MICRO'24* (2024).
- [35] Mark Dackendorf, et al. 2022. Graph Algorithms over Homomorphic Encryption for Data Cooperatives. In *SECRYPT*. 205–214.
- [36] Léo Ducas, et al. 2015. FHEW: bootstrapping homomorphic encryption in less than a second. In *Annual international conference on the theory and applications of cryptographic techniques*. Springer, 617–640.
- [37] Junfeng Fan, et al. 2012. Somewhat practical fully homomorphic encryption. *Cryptology ePrint Archive* (2012).
- [38] Robert W Floyd. 1962. Algorithm 97: shortest path. *Commun. ACM* 5, 6 (1962), 345–345.
- [39] Robin Geelen, et al. 2023. Bootstrapping for BGV and BFV Revisited. *Journal of Cryptology* 36, 2 (2023), 12.
- [40] Craig Gentry, et al. 2012. Fully homomorphic encryption with polylog overhead. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 465–482.
- [41] Ran Gilad-Bachrach, et al. 2016. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International conference on machine learning*. PMLR, 201–210.
- [42] Charles Gouert, et al. 2023. Sok: New insights into fully homomorphic encryption libraries via standardized benchmarks. *Proceedings on privacy enhancing technologies* (2023).
- [43] Timon Hackenjos, et al. 2020. SAGMA: secure aggregation grouped by multiple attributes. In *Proceedings of the 2020 ACM SIGMOD international conference on management of data*. 587–601.
- [44] Shai Halevi, et al. 2021. Bootstrapping for helib. *Journal of Cryptology* 34, 1 (2021), 7.
- [45] Kyoohyung Han, et al. 2019. Logistic regression on homomorphic encrypted data at scale. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 9466–9471.
- [46] Mingqin Han, et al. 2022. coxHE: A software-hardware co-design framework for FPGA acceleration of homomorphic computation. In *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 1353–1358.
- [47] IBM. 2024. Intesa Sanpaolo and IBM secure digital transactions with fully homomorphic encryption. <https://www.ibm.com/case-studies/blog/intesa-sanpaolo-ibm-secure-digital-transactions-fhe>.
- [48] Ilia Iliashenko, et al. 2021. Faster homomorphic comparison operations for BGV and BFV. *Proceedings on Privacy Enhancing Technologies* 2021, 3 (2021), 246–264.
- [49] Lei Jiang, et al. 2022. MATCHA: A Fast and Energy-Efficient Accelerator for Fully Homomorphic Encryption over the Torus. In *The Design Automation Conference (DAC 2022)*.
- [50] Dongwoo Kim, et al. 2023. Optimized privacy-preserving cnn inference with fully homomorphic encryption. *IEEE Transactions on Information Forensics and Security* 18 (2023), 2175–2187.
- [51] Miran Kim, et al. 2015. Private genome analysis through homomorphic encryption. In *BMC medical informatics and decision making*, Vol. 15. Springer, 1–12.
- [52] Ágnes Kiss, et al. 2019. SoK: Modular and efficient private decision tree evaluation. *Proceedings on Privacy Enhancing Technologies* (2019).
- [53] Alex Krizhevsky, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [54] Alex Krizhevsky, et al. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012).

- [55] Mayank Kumar, et al. 2025. TFHE-Coder: Evaluating LLM-agentic Fully Homomorphic Encryption Code Generation. *arXiv preprint arXiv:2503.12217* (2025).
- [56] Yann LeCun, et al. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [57] Dongwon Lee, et al. 2024. Functional bootstrapping for packed ciphertexts via homomorphic LUT evaluation. *Cryptology ePrint Archive* (2024).
- [58] Eunsang Lee, et al. 2022. Optimization of homomorphic comparison algorithm on rns-ckks scheme. *IEEE Access* 10 (2022), 26163–26176.
- [59] Eunsang Lee, et al. 2022. Low-complexity deep convolutional neural networks on fully homomorphic encryption using multiplexed parallel convolutions. In *International Conference on Machine Learning*. PMLR, 12403–12422.
- [60] Eunsang Lee, et al. 2021. Minimax approximation of sign function by composite polynomial for homomorphic comparison. *IEEE Transactions on Dependable and Secure Computing* 19, 6 (2021), 3711–3727.
- [61] Junghyun Lee, et al. 2023. Precise approximation of convolutional neural networks for homomorphically encrypted data. *IEEE Access* 11 (2023), 62062–62076.
- [62] Joon-Woo Lee, et al. 2022. Privacy-preserving machine learning with fully homomorphic encryption for deep neural network. *IEEE Access* 10 (2022), 30039–30054.
- [63] Seewoo Lee, et al. 2023. HETAL: efficient privacy-preserving transfer learning with homomorphic encryption. In *International Conference on Machine Learning*. PMLR, 19010–19035.
- [64] Yang Li. 2021. Checking chordality on homomorphically encrypted graphs. *arXiv preprint arXiv:2106.13560* (2021).
- [65] Zeyu Liu, et al. 2023. Amortized functional bootstrapping in less than 7 ms, with $O(1)$ polynomial multiplications. In *International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 101–132.
- [66] Junzhen Lou. 2024. Homomorphic Encryption for Healthcare Data Privacy in Industry Use Cases. (2024).
- [67] Qian Lou, et al. 2019. Glyph: Fast and accurately training deep neural networks on encrypted data. *NeurIPS 2020 (Advances in Neural Information Processing Systems)* (2019).
- [68] Qian Lou, et al. 2019. AutoQ: Automated Kernel-Wise Neural Network Quantization. In *International Conference on Learning Representations (ICLR) 2020*.
- [69] Qian Lou, et al. 2019. She: A fast and accurate deep neural network for encrypted data. *Advances in neural information processing systems* 32 (2019).
- [70] Qian Lou, et al. 2021. HEMET: A Homomorphic-Encryption-Friendly Privacy-Preserving Mobile Neural Network Architecture. *ICML 2021* (2021).
- [71] Qian Lou, et al. 2020. Falcon: Fast spectral inference on encrypted data. *Advances in Neural Information Processing Systems* 33 (2020), 2364–2374.
- [72] Qian Lou, et al. 2021. SAFENet: A Secure, Accurate and Fast Neural Network Inference. (2021).
- [73] Qian Lou, et al. 2020. AutoPrivacy: Automated Layer-wise Parameter Selection for Secure Neural Network Inference. *NeurIPS 2020* (2020).
- [74] Wen-jie Lu, et al. 2021. PEGASUS: bridging polynomial and non-polynomial evaluations in homomorphic encryption. In *2021 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1057–1073.
- [75] Wen-jie Lu, et al. 2018. Non-interactive and output expressive private comparison from homomorphic encryption. In *Proceedings of the 2018 on Asia Conference on Computer and Communications Security*. 67–74.
- [76] Xianrui Meng, et al. 2015. Grecs: Graph encryption for approximate shortest distance queries. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. 504–517.
- [77] Prateek Mittal, et al. 2012. Preserving link privacy in social network based systems. *arXiv preprint arXiv:1208.6189* (2012).
- [78] Payman Mohassel, et al. 2017. SecureML: A system for scalable privacy-preserving machine learning. In *2017 IEEE symposium on security and privacy (SP)*. IEEE, 19–38.
- [79] Koki Morimura, et al. 2023. Accelerating Polynomial Evaluation for Integer-wise Homomorphic Comparison and Division. *Journal of Information Processing* 31 (2023), 288–298.
- [80] Harika Narumanchi, et al. 2017. Performance analysis of sorting of FHE data: integer-wise comparison vs bit-wise comparison. In *2017 IEEE 31st International Conference on Advanced Information Networking and Applications (AINA)*. IEEE, 902–908.
- [81] Pascal Paillier. 1999. Public-key cryptosystems based on composite degree residuosity classes. In *International conference on the theory and applications of cryptographic techniques*. Springer, 223–238.
- [82] Antonis Papadimitriou, et al. 2016. Big data analytics over encrypted datasets with seabed. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*. 587–602.
- [83] Raluca Ada Popa, et al. 2011. CryptDB: Protecting confidentiality with encrypted query processing. In *Proceedings of the twenty-third ACM symposium on operating systems principles*. 85–100.
- [84] Jean Louis Raisaro, et al. 2018. Protecting privacy and security of genomic data in i2b2 with homomorphic encryption and differential privacy. *IEEE/ACM transactions on computational biology and bioinformatics* 15, 5 (2018), 1413–1426.
- [85] Xuanle Ren, et al. 2022. HEDA: multi-attribute unbounded aggregation over homomorphically encrypted database. *Proceedings of the VLDB Endowment* 16, 4 (2022), 601–614.
- [86] Ronald L Rivest, et al. 1978. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 21, 2 (1978), 120–126.
- [87] Nigel P Smart, et al. 2014. Fully homomorphic SIMD operations. *Designs, codes and cryptography* 71 (2014), 57–81.
- [88] Yan-Yan Song, et al. 2015. Decision tree methods: applications for classification and prediction. *Shanghai archives of psychiatry* 27, 2 (2015), 130.
- [89] Wenting Zheng Srinivasan, et al. 2019. DELPHI: A cryptographic inference service for neural networks. In *Proc. 29th USENIX secur. symp.*, Vol. 3.
- [90] Benjamin Hong Meng Tan, et al. 2020. Efficient private comparison queries over encrypted databases using fully homomorphic encryption with finite fields. *IEEE Transactions on Dependable and Secure Computing* 18, 6 (2020), 2861–2874.
- [91] Stephen Lyle Tu, et al. 2013. Processing analytical queries over encrypted data. (2013).
- [92] Pengtao Xie, et al. 2014. Cryptgraph: Privacy preserving graph analytics on encrypted graph. *arXiv preprint arXiv:1409.5021* (2014).
- [93] Jiaqi Xue, et al. 2022. Audit and improve robustness of private neural networks on encrypted data. *arXiv preprint arXiv:2209.09996* (2022).
- [94] Jiaqi Xue, et al. 2023. CryptoTrain: Fast Secure Training on Encrypted Dataset. In *Proceedings of the 1st ACM Workshop on Large AI Systems and Models with Privacy and Safety Analysis*. 97–104.
- [95] Jiaqi Xue, et al. 2024. CryptoTrain: Fast Secure Training on Encrypted Dataset. In *Proceedings of the 1st ACM Workshop on Large AI Systems and Models with Privacy and Safety Analysis*. 97–104.
- [96] Ardhi Wiratama Baskara Yudha, et al. 2024. BoostCom: Towards Efficient Universal Fully Homomorphic Encryption by Boosting the Word-wise Comparisons. In *(PACT’24) The International Conference on Parallel Architectures and Compilation Techniques (PACT)*.
- [97] Zama. 2022. Build Apps with Fully Homomorphic Encryption (FHE). <https://www.zama.ai/>.
- [98] Zama. 2022. Concrete ML: a Privacy-Preserving Machine Learning Library using Fully Homomorphic Encryption for Data Scientists. <https://github.com/zama-ai/concrete-ml>.
- [99] Zama. 2022. Concrete: TFHE Compiler that converts python programs into FHE equivalent. <https://github.com/zama-ai/concrete>.
- [100] Zama. 2022. Health Insurance Portability and Accountability Act. <https://www.hhs.gov/hipaa/index.html>.
- [101] Zama. 2022. TFHE-rs: A Pure Rust Implementation of the TFHE Scheme for Boolean and Integer Arithmetics Over Encrypted Data. <https://github.com/zama-ai/tfhe-rs>.
- [102] Yancheng Zhang, et al. 2023. HEBridge: Connecting Arithmetic and Logic Operations in FV-style HE Schemes. In *Proceedings of the 12th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*. 23–35.
- [103] Yuchen Zhang, et al. 2015. Foresee: Fully outsourced secure genome study based on homomorphic encryption. In *BMC medical informatics and decision making*, Vol. 15. Springer, 1–11.
- [104] Yancheng Zhang, et al. 2025. CipherPrune: Efficient and Scalable Private Transformer Inference. *arXiv preprint arXiv:2502.16782* (2025).
- [105] Yancheng Zhang, et al. 2024. Heprune: Fast private training of deep neural networks with encrypted data pruning. *Advances in Neural Information Processing Systems* 37 (2024), 51063–51084.
- [106] Zhou Zhang, et al. 2024. ArcEDB: An Arbitrary-Precision Encrypted Database via (Amortized) Modular Homomorphic Encryption. *Cryptology ePrint Archive* (2024).
- [107] Mengxin Zheng, et al. 2023. PriML: An Electro-Optical Accelerator for Private Machine Learning on Encrypted Data. In *2023 24th International Symposium on Quality Electronic Design (ISQED)*. IEEE, 1–7.
- [108] Mengxin Zheng, et al. 2024. OFHE: An Electro-Optical Accelerator for Discretized TFHE. In *Proceedings of the 29th ACM/IEEE International Symposium on Low Power Electronics and Design*. 1–6.
- [109] Mengxin Zheng, et al. 2023. Cofhe: Software and hardware co-design for fhe-based machine learning as a service. *Frontiers in Electronics* 3 (2023), 1091369.
- [110] Mengxin Zheng, et al. 2022. CryptoLight: An Electro-Optical Accelerator for Fully Homomorphic Encryption. In *Proceedings of the 17th ACM International Symposium on Nanoscale Architectures*. 1–2.
- [111] Mengxin Zheng, et al. 2023. Primer: Fast Private Transformer Inference on Encrypted Data. *DAC 2023* (2023).