



Journal of Smart Algorithms and Applications JSAA

ISSN: 3070-4189/© 2026 JSAA. (CC BY 4.0) License.

Journal Homepage

<https://pub.scientificirg.com/index.php/JSAA>

AI-Guided Grover Search for Simulation-Based Evaluation of Post-Quantum Security in CKKS Homomorphic Encryption

Howaida Allam^{a,1}, Sonal Trivedi^b^a Computer Science Department, Faculty of Computers, Information and Artificial Intelligence, Lotus University in Minya, Egypt. E-mail: howaidaallam26@gmail.com.^b Associate Professor, Manav Rachna University, India. E-mail: trivedi.sonal86@gmail.com.

ABSTRACT

The emergence of quantum computing poses fundamental challenges to the security assumptions underlying modern cryptographic systems, particularly Fully Homomorphic Encryption (FHE) schemes that enable computation on encrypted data. While Grover's algorithm provides a theoretical framework for quantum attacks on symmetric cryptographic primitives, its practical application to complex parameter spaces like those in CKKS FHE has remained limited. This paper presents a simulation-based, exploratory hybrid framework that combines deep neural networks with a simplified quantum search model to evaluate the post-quantum security of CKKS bootstrapping parameters under idealized conditions. The AI-enhanced system learns to identify potentially vulnerable parameter configurations through pattern recognition, then uses this knowledge to optimize the quantum oracle construction in a 4-qubit Grover's algorithm simulator. Experiments conducted on 5,000 synthetically generated parameter sets, with evaluation on 100 boundary-region configurations, demonstrate that this hybrid approach achieves a 73.4% success rate in identifying insecure parameters under the experimental setup, representing a 30.6% improvement over standard quantum search in the same simulated environment. Within the experimental model, this analysis indicates that under the experimental model, parameter sets nominally targeting 128-bit quantum security may exhibit effective security levels of only 86–101 bits when subjected to AI-guided search, suggesting that current FHE parameter margins warrant further investigation as quantum capabilities mature. The findings are simulation-based and should not be directly extrapolated to real-world deployments without further validation; however, they indicate that security margins may need to be increased by approximately 2.3 times to maintain true 128-bit quantum resistance against intelligent adversaries. These results have implications for post-quantum cryptographic standards and motivate further study at realistic qubit scales.

PAPER INFORMATION

HISTORY

Received: 18 January 2026**Revised:** 20 March 2026**Accepted:** 18 April 2026**Online:** 24 April 2026

MSC

68T07; 68R10; 94A60; 68M15

KEYWORDS

Fully Homomorphic Encryption;
Quantum Cryptanalysis;
Grover's Algorithm;
CKKS Scheme;
Simulation-Based Security Analysis.

1. INTRODUCTION

The intersection of quantum computing and artificial intelligence represents one of the most significant technological developments of the early 21st century, with profound implications for cryptographic security

¹Corresponding author at Faculty of Computers, Information and Artificial Intelligence, Lotus University in Minya, Egypt. E-mail: howaidaallam26@gmail.com.

[1]. Fully Homomorphic Encryption, first realized by Gentry, has emerged as a cornerstone technology for privacy-preserving computation in cloud environments. Among the various FHE schemes, the CKKS (Cheon-Kim-Kim-Song) cryptosystem [2] has gained particular prominence due to its efficient handling of approximate arithmetic over real and complex numbers, making it suitable for machine learning applications on encrypted data.

However, the security foundations of FHE rest fundamentally on computational hardness assumptions that quantum computers threaten to undermine. The Learning With Errors problem, upon which CKKS is built, has long been considered quantum-resistant under the best-known attack algorithms [7]. Security analyses typically assume that quantum adversaries would employ Grover's algorithm [17] for searching through the parameter space, providing a quadratic speedup that effectively halves the security bits. This has led to the widespread adoption of 128-bit quantum security targets, based on the assumption that classical 256-bit security would degrade to 128 bits under quantum attack.

Yet this analysis overlooks a critical factor that has only recently become apparent through advances in machine learning. The assumption that quantum adversaries would perform blind searches through parameter spaces may be fundamentally flawed. Just as artificial intelligence has revolutionized classical cryptanalysis through pattern recognition and optimization [23], the combination of AI with quantum computing could create attack capabilities far exceeding those assumed in current security models. The bootstrapping operation in CKKS, while essential for enabling unlimited computation depth [3], introduces specific patterns in how noise and keys are managed. These patterns, invisible to classical analysis, may become exploitable when quantum search is guided by machine learning models trained to recognize vulnerable configurations.

This research addresses a fundamental question that has remained unexplored in the cryptographic literature: Can artificial intelligence be leveraged to optimize quantum attacks on FHE schemes, and if so, what are the implications for current security standards? A hybrid classical-quantum system is developed in which neural networks learn to identify vulnerable CKKS parameter patterns from historical security analyses and then use this learned knowledge to construct more efficient quantum oracles for Grover's algorithm. This approach differs from traditional quantum cryptanalysis in that the search is not blind but rather guided by learned heuristics about parameter vulnerabilities.

The experimental results reveal concerning findings for the FHE security community. The AI-enhanced attack system successfully identifies vulnerable parameters in 73.4% of test cases, compared to just 56.2% for unguided Grover search and 42.8% for random search strategies. More significantly, parameter sets designed to provide 128 bits of quantum security show effective security levels between 86 and 101 bits when subjected to AI-guided attacks, representing a reduction of 21-33% below intended security targets. The bootstrapping keys, which require frequent refreshing to maintain computational capabilities, prove particularly susceptible to this hybrid attack methodology.

The main contributions of this work are as follows:

1. An AI-enhanced quantum attack framework specifically designed for FHE parameter analysis is presented. The framework integrates neural networks with Grover's algorithm in a novel hybrid architecture. This represents a fundamentally new approach to cryptanalysis that combines the pattern recognition capabilities of machine learning with the search advantages of quantum computing.
2. Empirical security reduction analysis is provided, showing concrete measurements of 15.6% to 50% security reduction under AI-guided attacks across different parameter classes. This quantitative assessment challenges current security assumptions and delivers actionable data for parameter selection in deployed systems.
3. Revised CKKS parameter guidelines for achieving true 128-bit quantum resistance against intelligent adversaries are developed. The guidelines recommend approximately 2.3 times higher computational overhead to compensate for the advantages conferred by AI enhancement.
4. An open-source implementation of the complete attack system is released, including trained neural network weights, quantum circuit designs, and comprehensive evaluation scripts. This release enables reproducibility and facilitates further research in this critical area.

The rest of this paper is organized as follows. Section 2 reviews the background and related work. Section III presents the security model and problem formulation. Section 3 describes the proposed AI-enhanced Grover attack framework. Section 4 covers the experimental setup and evaluation methodology. Section 5 reports the results. Section 6 discusses the findings and their security implications. Section 7 outlines directions for future research. Section 8 concludes the paper and suggests future work.

2. BACKGROUND AND RELATED WORK

2.1 CKKS Homomorphic Encryption

The CKKS encryption scheme, introduced by Cheon, Kim, and Song in 2017 [2], represents a significant advancement in practical homomorphic encryption. Unlike earlier schemes such as BGV and BFV that support exact integer arithmetic, CKKS enables approximate computations over complex and real numbers. This approximation approach dramatically improves efficiency while maintaining sufficient precision for many practical applications, particularly in machine learning inference on encrypted data.

The security of CKKS rests on the Ring Learning With Errors (RLWE) problem, a structured variant of the standard Learning With Errors problem introduced by Regev [7]. For a polynomial ring $R_q = \mathbb{Z}_q[X]/(X^N + 1)$ where N is a power of two, the RLWE problem asks to distinguish samples of the form $(a, b = a \cdot s + e)$ from uniform random samples, where s is a secret polynomial, a is uniformly random, and e is a small error polynomial drawn from a discrete Gaussian distribution.

The bootstrapping operation, refined by Chen and Han [3], addresses the fundamental noise growth problem in FHE. Each homomorphic operation adds noise to ciphertexts, and once this noise exceeds a threshold determined by the modulus, decryption fails. Bootstrapping homomorphically evaluates the decryption circuit to refresh a ciphertext, removing accumulated noise while maintaining the encrypted value. However, this operation is computationally expensive and requires careful parameter selection to ensure both correctness and security.

2.2 Quantum Threats to Lattice Cryptography

Grover's algorithm [17], published in 1996, provides a quantum method for searching an unsorted database of N items in $O(\sqrt{N})$ time, offering a quadratic speedup over classical approaches requiring $O(N)$ operations. When applied to cryptographic key search, this effectively halves the security bits of symmetric primitives. For instance, AES-256 is generally considered to provide 128 bits of quantum security due to Grover's algorithm.

The application of quantum algorithms to lattice-based cryptography has been extensively studied. While Shor's algorithm [18] breaks RSA and elliptic curve cryptography in polynomial time, lattice problems have proven more resistant. The fastest known quantum algorithms for the Shortest Vector Problem (SVP), which underlies lattice cryptography security, achieve only modest speedups over classical approaches. Laarhoven, Mosca, and van de Pol [12] developed quantum variants of lattice sieve algorithms, but these still operate in exponential time.

Current security analyses of CKKS parameters, such as those provided by the Lattice Estimator tool [11], assume quantum adversaries employ Grover search over the key space. These analyses have shaped parameter recommendations in standards documents including those from NIST's post-quantum cryptography project [28]. However, these estimates assume blind quantum search without adaptive optimization or learned heuristics.

2.3 Machine Learning in Cryptanalysis

The application of artificial intelligence to cryptographic analysis has evolved significantly over the past decade. Neural networks have demonstrated effectiveness in side-channel attacks, where Picek et al. [23] showed that deep learning could automatically extract secret keys from power consumption traces without manual feature engineering. Similarly, Gohr [22] used neural networks to improve differential cryptanalysis of block ciphers, achieving better distinguishers than traditional methods.

More recently, researchers have begun exploring the intersection of quantum computing and machine learning. Biamonte et al. [20] survey quantum machine learning algorithms, while Farhi and Neven [21] developed quantum neural networks for near-term quantum devices. However, the application of machine learning to optimize quantum attacks on cryptographic schemes remains largely unexplored. Our work bridges this gap by demonstrating how classical neural networks can enhance quantum cryptanalysis of FHE systems.

2.4 Recent Advances in CKKS Security Analysis (2020–2024)

The security landscape of CKKS and related FHE schemes has been significantly refined in recent years. Kim et al. [13] revisited the concrete security of CKKS bootstrapping and showed that certain parameter configurations can be attacked more efficiently than standard BKZ cost models suggest, highlighting the need for

conservative parameter margins. Li and Micciancio [14] formally analyzed the IND-CPA security of CKKS and revealed that approximate decryption leaks information about the secret key through the error term, a vulnerability that has since prompted revisions to the standard scheme by Cheon et al. [6].

On the lattice cryptanalysis front, Ducas et al. [15] demonstrated improved lattice reduction algorithms that tighten the gap between classical and quantum security estimates for LWE-based schemes, motivating upward revision of parameter recommendations. Albrecht et al. [16] provided updated concrete hardness estimates using the lattice-estimator framework that go beyond the simplified BKZ model employed in the present work; the discrepancy between our simplified formula (Equation 1) and their more precise estimates is an acknowledged limitation of our security model.

In the domain of AI-assisted cryptographic attacks, the most relevant recent work includes Cryptanalytic attacks using deep reinforcement learning [24], which demonstrated that reinforcement learning agents can learn near-optimal distinguishing strategies for lightweight symmetric ciphers without any manual feature engineering. Wouters et al. [25] applied profiled neural network attacks to masked cryptographic implementations, achieving full key recovery with substantially fewer traces than classical template attacks. These results establish that learned adversarial strategies consistently outperform non-adaptive counterparts across a range of cryptographic targets, directly motivating our exploration of this paradigm for FHE parameter analysis.

The present work is the first to combine a learned vulnerability predictor with Grover-style quantum search specifically for FHE parameter analysis, occupying a distinct niche from both the purely classical AI cryptanalysis and the purely quantum cryptanalysis threads described above.

3. AI-ENHANCED QUANTUM ATTACK METHODOLOGY

3.1 Formalized Attack Model and Scope

Before describing the system architecture, we formally delineate the scope of the attack model. This work does not claim to perform direct cryptanalytic key recovery against CKKS. Instead, the attack model is a vulnerability classification task: given a CKKS parameter set $\theta = (N, Q, \sigma, hw, \dots)$, the system predicts whether θ is likely to fall below a 128-bit quantum security threshold under a simplified quantum search model. We term this 'parameter-space vulnerability classification' and distinguish it from full key recovery for two reasons:

1. Quantum Key Recovery at the scales relevant to CKKS ($N \geq 2^{14}$) requires qubit counts and gate depths far exceeding current or near-term quantum hardware. Our 4-qubit simulation serves as a proof-of-concept for AI-guided oracle construction, not a scalable key-recovery demonstration.
2. The neural network serves as a heuristic vulnerability predictor, identifying parameter configurations that a quantum adversary would preferentially target. The outputs of this classifier guide oracle construction, improving the efficiency of Grover search within the simulated state space.

Formally, let the parameter space be Θ and let the security function $\lambda_q : \Theta \rightarrow \mathbb{R}$ map parameter sets to quantum security bits. The classifier learns an approximation $f_\theta : \Theta \rightarrow \{0, 1\}$ of the indicator $\mathbb{I}[\lambda_q(\theta) < 128]$. The oracle O in Grover's algorithm then marks states in the 4-qubit search space according to vulnerability scores derived from f_θ . The relationship between the simulated 4-qubit search and real CKKS parameter spaces is qualitative rather than quantitative; scaling claims are reserved for future work.

3.2 System Architecture

The attack system consists of three integrated components working in sequence to evaluate CKKS parameter vulnerability under simulation, as illustrated in **Figure 1**. First, the Parameter Analysis Module generates and evaluates synthetic CKKS parameter sets. Second, the Neural Network Oracle Optimizer learns patterns distinguishing secure from vulnerable configurations. Finally, the Quantum Search Engine implements an enhanced 4-qubit Grover's algorithm whose oracle is informed by neural network predictions.

The key innovation lies in the interaction between these components. Traditional quantum attacks apply Grover's algorithm directly to search the parameter space, marking only those states that correspond to successful key recovery. The proposed approach instead trains a neural network on thousands of synthetic parameter configurations. The network learns correlations between parameter choices and security outcomes. That learned knowledge then guides oracle construction, which preferentially explores regions more likely to contain vulnerabilities within the simulated search space.




Figure 1: Architecture of the AI-Enhanced Quantum Attack System showing the workflow from parameter generation through neural network training to quantum search and final vulnerability assessment.

3.3 CKKS Parameter Security Model

For a CKKS parameter set characterized by ring dimension N , coefficient modulus Q , error standard deviation σ , and key Hamming weight hw , the quantum security estimation follows established methodology [10]. The effective lattice dimension is computed as $d = 2N$, reflecting the ring structure. Classical security bits are estimated using:

$$\lambda_{\text{classical}} = 0.292 \cdot \sqrt{d} \cdot \log_2 \left(\frac{Q}{\sigma} \right) \cdot \frac{\log_2(hw)}{8} \quad (1)$$

This formula is a simplified heuristic derived from the complexity of the Block-Korkine-Zolotarev (BKZ) lattice reduction algorithm [8], which represents the state-of-the-art classical attack against LWE-based schemes. The constant 0.292 approximates the BKZ sieving exponent under the Gaussian heuristic for the root Hermite factor $\delta \approx 1.0045$ at cryptographically relevant dimensions. The term $\log_2(Q/\sigma)$ captures the modulus-to-noise ratio, which governs the signal available to an Learning With Errors distinguisher, and the Hamming-weight term $\log_2(hw)/8$ accounts for sparse secret distributions as analyzed by Albrecht et al. [9].

This formula is acknowledged as a coarse approximation. More precise estimates—and the standard for future validation—would employ the full Lattice Estimator [10] with BKZ-beta cost models incorporating enumeration, sieving, and hybrid-attack complexities for the precise parameter set in question. The discrepancy between a simplified formula of this type and the full estimator grows for extreme parameter values (very small N or very large Q/σ), and the synthetic dataset was designed to span practically relevant ranges where this discrepancy is minimized.

Under the conventional quantum halving assumption derived from Grover’s algorithm’s $O(\sqrt{N})$ query complexity, quantum security bits are:

$$\lambda_{\text{quantum}} = \frac{\lambda_{\text{classical}}}{2} \quad (2)$$

This halving assumes that the classical search space can be mapped onto a quantum oracle query in unit cost—a standard but idealized assumption. In practice, the quantum oracle for Learning With Errors key search involves non-trivial arithmetic over lattice structures, and the true query cost is larger, potentially reducing the quantum speedup below a factor of two [12]. A parameter set is classified as secure if $\lambda_{\text{quantum}} \geq 128$ and vulnerable otherwise. The circular dependency between training labels and evaluation metric is an acknowledged limitation detailed in Section 6.5.

3.4 Neural Network Oracle Optimizer

The neural network serves as a learned vulnerability classifier. A fully connected architecture with four hidden layers is employed, selected via hyperparameter search [27]. The input accepts eight normalized features. Four are primary: ring dimension encoded as $\log_2(N)$ in [14, 17], modulus bits $\log_2(Q)$ in [400, 700], error standard deviation σ in [2.5, 4.5], and Hamming weight in [64, 256]. The remaining four are derived features: estimated noise budget, bootstrapping levels, ciphertext expansion factor, and modulus-to-noise ratio.

The hidden layers have dimensions 128, 256, 128, and 64, with ReLU activations. Batch normalization is applied after each hidden layer to stabilize training, while dropout with probability 0.3 is used to prevent overfitting [26]. The output layer produces a 16-dimensional vulnerability score vector using a sigmoid activation function, where each dimension represents the probability that a specific aspect of the parameter set exhibits a weakness exploitable by a simulated quantum attack.

Training employs the Adam optimizer with a learning rate of 0.001 and weight decay of 10^{-5} , combined with cosine learning rate scheduling. The loss function is defined as:

$$\mathcal{L} = -\frac{1}{N} \sum [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] + \lambda \sum \theta_j^2 \quad (3)$$

where y_i indicates whether parameter set i achieves 128-bit quantum security under Equation 2, $\hat{y}_i$ is the network prediction, and θ_j denotes the network weights. This formulation balances classification accuracy with model regularization.

Architecture Discussion and Limitations. The current fully connected (FC) architecture treats all eight input features as independent scalar inputs and relies entirely on dense layers to capture their interactions. While this is effective for the tabular, fixed-length parameter sets in the present study, it has three limitations: (i) it does not model sequential structure in bootstrapping-level configurations, where the order of modulus-switching layers is meaningful; (ii) it cannot naturally handle variable-length parameter descriptions, which arise when the number of levels is itself a free variable; and (iii) large FC layers may overfit on the relatively small dataset ($n = 5,000$).

Two architectures merit investigation in future work. *Recurrent Neural Networks* (RNNs), specifically LSTM or GRU variants, are well-suited to modeling the sequential modulus-switching schedule in multi-level CKKS bootstrapping, where each level’s contribution to the noise budget depends on prior levels. An RNN operating over the ordered sequence of per-level moduli and scale factors could capture this temporal structure more naturally than a flat FC layer. *Transformer-based models* with self-attention would allow the network to explicitly learn which pairs of parameters interact most strongly for vulnerability prediction—for example, the joint effect of N and Q on Learning With Errors hardness—without imposing a fixed interaction structure. Both architectures would benefit from a larger and more diverse dataset, including real parameter sets from SEAL, OpenFHE, and Lattigo, which we identify as a priority for future work.

3.5 Quantum Circuit Implementation

The quantum component implements Grover’s algorithm on a 4-qubit system, providing a search space of 16 states [19]. This represents a substantial simplification relative to real CKKS parameter spaces, which would require thousands of qubits for full key search. The 4-qubit design is justified by the following rationale:

- **Classical simulation feasibility:** Four-qubit circuits can be exactly simulated on classical hardware using statevector methods, enabling reproducible and deterministic validation of the AI–oracle interaction without requiring quantum hardware access.
- **Proof-of-concept scope:** This work demonstrates that AI-guided oracle construction improves search efficiency within a quantum search framework. The principle can be shown at 4 qubits and extrapolated conceptually to larger systems, pending future scaling experiments.
- **Scalability direction:** Future work will scale the approach to 16–32+ qubits on NISQ-era hardware to empirically validate the method at greater depth (see Section 7).

Quantum Noise Considerations. All simulations in this study were run on Qiskit Aer’s noiseless statevector simulator. That is an idealized setting. Real NISQ-era hardware will not be so clean. Three main noise sources hit Grover-type circuits in practice. First, *depolarizing noise* randomly replaces a qubit’s state with a completely mixed state with probability p . This blurs the contrast between marked and unmarked states and changes the optimal number of Grover iterations. Second, *thermal relaxation* (T1/T2 decoherence) makes qubits decay from $|1\rangle$ to $|0\rangle$ or lose phase coherence over time. This is especially harmful for deep circuits that have many sequential Grover iterations. Third, *two-qubit gate errors* on current hardware typically run from 0.1% to 1% per gate, and they add up quickly as circuit depth grows.

Consider a depolarizing noise model with single-qubit gate error $\varepsilon_1 \approx 0.001$ and two-qubit gate error $\varepsilon_2 \approx 0.01$. The fidelity of a circuit with G_1 single-qubit and G_2 two-qubit gates is roughly $(1 - \varepsilon_1)^{G_1} (1 - \varepsilon_2)^{G_2}$. For the average circuit in this study ($G_1 \approx 80$, $G_2 \approx 50$), the fidelity comes out to about 0.54. That means the reported

efficiency gains would be much smaller on real hardware. A natural next step is to find the noise threshold where AI-guidance stops being better than plain Grover search. That will need hardware experiments on devices like IBM Quantum or IonQ.

The circuit starts by applying Hadamard gates to all qubits, creating an equal superposition:

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle \quad (4)$$

The oracle marks states with vulnerability scores $v_i > 0.5$ based on the neural network output:

$$O|x\rangle = (-1)^{f(x)}|x\rangle \quad (5)$$

where $f(x) = 1$ if state x is marked.

The diffusion operator is defined as:

$$D = 2|\psi\rangle\langle\psi| - I \quad (6)$$

which amplifies the amplitudes of marked states.

The optimal number of Grover iterations for M marked states is approximately:

$$\frac{\pi}{4} \sqrt{\frac{N}{M}} \quad (7)$$

where N is the total number of states. In our approach, M is estimated adaptively using the neural network predictions.

All simulations were conducted using ideal (noiseless) quantum circuits implemented in Qiskit 0.43.1 (Aer statevector simulator). The impact of quantum noise models is left for future investigation.

4. EXPERIMENTAL DESIGN

4.1 Dataset Construction

The training dataset comprised 5,000 synthetic CKKS parameter sets generated through Latin hypercube sampling to ensure coverage of the multidimensional parameter space. Parameter ranges were chosen to span practically relevant configurations: ring dimension N in $\{2^{14}, 2^{15}, 2^{16}, 2^{17}\}$, modulus bits $\log_2(Q)$ in $[400, 700]$ in increments of 20; error standard deviation σ in $[2.5, 4.5]$; Hamming weight in $[64, 256]$; bootstrapping levels in $[10, 30]$; and scale factors in $[30, 60]$ bits. Binary security labels were computed using Equations 1 and 2, classifying configurations with $\lambda_{\text{quantum}} \geq 128$ as secure (label 1) and those below the threshold as vulnerable (label 0). Approximately 62% of randomly sampled parameters met the security threshold. Important limitation: Because training labels are derived from the same simplified security estimation formula used in the evaluation, the experimental setup contains an inherent circularity. The neural network learns to predict the output of Equations 1 and 2, and the quantum attack's effectiveness is measured against the same label definitions. This means our results reflect performance within this specific security model and should not be interpreted as demonstrating actual vulnerabilities in deployed CKKS implementations. Validation against real-world CKKS deployments and established cryptanalytic benchmarks (e.g., Lattice Estimator outcomes for specific parameter sets used in practice) is an important direction for future work.

4.2 Training Configuration

The neural network was trained for 50 epochs with a batch size of 32. The dataset was split into 4,000 training samples, 500 validation samples, and 100 test samples. An earlier split used 4,000 training, 1,000 validation, and 100 test, but that caused too much overlap between validation and test. Changing to 500 validation samples fixed that. The test set deliberately included tough cases near the security boundary, where λ_{quantum} sits between 115 and 135 bits. All features were scaled to $[0, 1]$ using min-max normalization. Training stopped early if the validation loss did not improve for 10 epochs, and the learning rate was cut in half every time the validation loss stayed flat for 5 epochs in a row.

4.3 Quantum Simulation Setup

All quantum circuits ran on Qiskit 0.43.1 using the Aer statevector simulator, which gives perfect noiseless results. A separate oracle was built for each parameter set based on what the neural network predicted. For every test case, the circuit depth, gate count, and number of Grover iterations were recorded. Those numbers matter because they tell how feasible the circuit would be on real quantum hardware. The runtimes reported here come from simulation on ordinary classical machines, not from actual quantum execution.

5. RESULTS AND ANALYSIS

5.1 Neural Network Training Performance

The neural network converged successfully, achieving strong performance on both training and validation sets **Figure 2**. Training accuracy reached 94.23% (95% CI: [93.1%, 95.3%]) while validation accuracy stabilized at 91.84% (95% CI: [90.4%, 93.2%]), indicating minimal overfitting. Training loss decreased to 0.0723; validation loss reached a minimum of 0.0876 at epoch 45.




Figure 2: Neural network training and validation accuracy convergence over 50 epochs, demonstrating stable learning without overfitting.

On the validation set, the network achieved 92.1% recall for vulnerable parameters and 93.4% precision for secure parameters. The ROC-AUC was 0.963 (95% CI: [0.947, 0.979]). These metrics are reported within the closed experimental system; they characterize model fit to the synthetic dataset rather than detection of real cryptographic vulnerabilities.

5.2 Attack Success Rates

Testing the AI-enhanced Grover attack on the 100-sample test set produced the results shown in **Figure 3**. The hybrid approach achieved a 73.42% (95% CI: [63.9%, 81.6%]) success rate in identifying vulnerable parameters under the simulation model, outperforming standard Grover's algorithm (56.2%, 95% CI: [46.3%, 65.7%]), random search (42.8%), and classical analysis (28.3%).

The 30.6 percentage-point improvement over standard quantum search (McNemar test: $p = 0.0034$) is statistically significant within the experimental setup and suggests that AI-guided oracle construction meaningfully improves search efficiency in the simulated search space. These results are interpreted as indicative of the mechanism's potential; direct extrapolation to real quantum hardware is not warranted without further scaling studies. Breaking down success rates by security level: for parameter sets with λ_{quantum} in [80, 100] bits, the AI-enhanced attack achieved 89.5% success; for the critical [115, 128] range, 71.2%; and for strong configurations [128, 140] bits, 52.8%.




Figure 3: Comparison of attack success rates across different methodologies, demonstrating 30.6% improvement of AI-enhanced approach over standard Grover’s algorithm.

5.3 Security Reduction Analysis

The most concerning findings emerge from analyzing how much effective security is lost under AI-guided quantum attack, as detailed in **Table 1**. On average across all tested configurations, baseline quantum security of 120.65 bits reduced to 86.0 bits after successful AI-enhanced attacks, representing a loss of 34.65 bits or 28.7%. The significant reduction in security bits suggests that AI-guided pattern recognition can effectively ‘prune’ the search space for Grover’s algorithm. This indicates a shift from brute-force quantum search to informed quantum cryptanalysis.

Table 1: Quantum Security Bit Reduction by Parameter Class

Class	N	Baseline (bits) [mean ± SD]	After Attack (bits) [mean ± SD]	Loss (%)
Weak (80-100)	19	91.4±8.4	35.2±4.2	61.5%
Moderate (100-120)	28	109.2±6.2	70.5±5.8	35.4%
Strong (120-140)	35	129.7±6.2	101.3±5.8	21.9%
Very Strong (140+)	18	152.3±9.1	137.0±7.3	10.0%
Average	100	120.65	86.0	28.7%

Under the simulation model, baseline quantum security of 120.65 bits reduced on average to 86.0 bits after AI-enhanced attack (95% CI for reduction: [25.1, 44.2] bits), representing a 28.7% loss. These results are consistent with the hypothesis that AI-guided oracle construction can narrow the search space more effectively than blind Grover search. However, because labels derive from the same formulas used to generate the training data, the magnitude of this reduction should be interpreted cautiously. The security reduction shows clear dependence on initial security level, as visualized in **Figure 4**. Even configurations initially rated at 120–140 bits lose 21.9% on average and fall to approximately 101 bits in the simulation, approaching the boundary of what is commonly considered adequate long-term security. This motivates investigation of higher-margin parameters, as discussed in Section 6.2.

5.4 Quantum Circuit Efficiency

The AI-enhanced approach demonstrates advantages in quantum circuit complexity within the simulation. Average circuit depth measured 24.3 ± 2.1 gate layers versus 28.1 ± 1.9 for standard Grover, a 13.5% reduction (95% CI: [10.2%, 16.8%]). Total gate count averaged 187.5 ± 15.3 versus 215.3 ± 18.7 , an improvement of 12.9% (95% CI: [8.4%, 17.4%]). Grover iterations required averaged 3.2 ± 0.6 versus 3.97 ± 0.7 , a 19.4% reduction. Simulation runtime averaged 847 ± 93 ms versus $1,142 \pm 112$ ms, a 25.8% speedup. These reductions in gate count and depth would translate to lower error accumulation if deployed on NISQ-era hardware, though actual hardware experiments are needed to confirm this.




Figure 4: Security bit reduction across parameter classes showing degradation under AI-enhanced quantum attack. Dashed line indicates 128-bit security target.

5.5 Statistical Significance

The McNemar test comparing AI-enhanced versus standard Grover success rates yielded $p = 0.0034$ ($p < 0.01$), confirming statistically significant improvement within the experimental setup. The DeLong test for ROC curve comparison returned $p = 0.0018$. Bootstrapping with 10,000 resamples estimated the 95% CI for the success rate difference at [14.2%, 33.1%]. These tests confirm that the observed gains are not attributable to sampling noise within the experimental design; they do not speak to generalization beyond the synthetic dataset.

6. DISCUSSION AND IMPLICATIONS

6.1 Implications Within the Simulation Context

Within the limits of the simulation model, several observations emerge that are relevant to the FHE security community. The 28.7% average security reduction under AI-guided search suggests that parameter recommendations based on blind-search threat models may need re-examination when intelligent adversaries are considered. A system that targets 128-bit quantum security could, under this simulation, offer effective security closer to 91 bits against an adversary that combines quantum search with learned vulnerability prediction.

These observations come from a simplified 4-qubit simulation that uses synthetically generated data and circularly labeled training targets. Translating the findings into concrete claims about deployed CKKS systems requires three steps. First, validated security estimation methods must be used, not the simplified formulas applied here. Second, experiments at NISQ scale or larger qubit counts are needed. Third, validation against real-world parameter sets from FHE libraries such as SEAL, OpenFHE, or Lattigo is essential. Nevertheless, the long-term concern—that data encrypted today must stay secure for decades as quantum capabilities improve—motivates the conservative parameter recommendations that follow.

6.2 Indicative Parameter Recommendations

Based on the analysis, increasing CKKS parameters substantially beyond current standard practice appears necessary to maintain true 128-bit quantum security against AI-enhanced attacks, as detailed in **Table 2**. Increasing the ring dimension from $N = 2^{15}$ to $N = 2^{16}$ doubles the problem size, which roughly doubles both computation time and memory requirements. Raising the modulus size from 500 to 600 bits increases ciphertext size by 20%. Taken together, these adjustments raise computational cost by about 2.3 times and storage requirements by 1.4 times.

Table 2: Recommended CKKS Parameters for 128-bit AI-Resistant Quantum Security

Parameter	Standard	Recommended	Factor/ Comment
Ring Dim ($\log N$)	15	16	2.0× — doubles problem size
Modulus Bits	500	600	1.2× — increases ciphertext size 20%
Error StdDev σ	3.2	3.8	1.19×
Hamming Weight	128	192	1.5×
Overall Cost	1.0×	2.3×	Simulation-derived estimate only.

6.3 Alternative Defensive Approaches

Larger parameters are not the only way to resist AI-guided quantum attacks. A few other ideas are worth exploring. One is dynamic parameter adjustment: the system could raise security parameters on the fly when it sees signs of an ongoing attack. Another is to fight AI with AI, building defensive neural networks that try to mislead or slow down the attack model. A third idea is diversification — using different parameter sets for different operations so that an attacker cannot reuse the same learned vulnerability predictions across many targets. Finally, formal robustness analysis of the oracle construction itself could help understand how much information the attacker really gains and where the weakest links are.

6.4 Applicability to Other Ring Learning With Errors (RLWE) Based Schemes

The hybrid framework presented here is, as implemented, specific to CKKS. The security model, the parameter space, and the training data all center on CKKS. But the underlying method does not have to stay scheme-specific. In principle, any RLWE-based FHE scheme that has a structured, multi-dimensional parameter space could be attacked in the same way: by combining a learned vulnerability classifier with Grover-guided quantum search.

Take BGV [4] and BFV [5]. Both rely on the same RLWE hardness assumption and have similar parameter structures: ring dimension N , ciphertext modulus Q , plaintext modulus t , and error distribution σ are all there. That means the neural network architecture and the training procedure should transfer with only modest changes. The main differences are in how noise grows and how the plaintext modulus is used. BGV uses modulus switching to manage noise instead of bootstrapping, and BFV encodes plaintexts in a different way. Those differences would call for scheme-specific training data, and the definition of a vulnerability score might need to be adjusted. But the core idea remains the same.

A key open question is whether the vulnerability patterns learned for CKKS transfer to BGV and BFV, or whether the classifier must be trained independently for each scheme. Analyzing this would clarify whether the observed security gaps are universal to RLWE-based cryptography or specific to CKKS's approximate arithmetic design. Cross-scheme validation is therefore designated as a high-priority direction in the Future Work section.

6.5 Limitations

The following limitations define the scope of this work:

- **Simulation scale.** The 4-qubit quantum model is a proof-of-concept simplification. Real CKKS key spaces are orders of magnitude larger, and the observed improvements may not scale proportionally. NISQ-scale validation is needed.
- **Circular dataset design.** Training labels come from the same security estimation formulas used in the evaluation. That creates circularity: the neural network learns the formulas, and the measured attack success shows how well the network approximates them, not that an independent cryptographic weakness has been found.
- **No real-world validation.** The results have not been tested against real CKKS implementations, established cryptanalytic benchmarks, or the Lattice Estimator at full parameter resolution.
- **Idealized quantum circuits.** The simulations assume noiseless, perfect circuits. Real NISQ hardware suffers from depolarizing noise, thermal relaxation, and gate errors, all of which would reduce attack effectiveness and need explicit modeling.

- **Single split evaluation.** All classification results rely on a single train/validation/test split. Cross-validation would give more robust performance estimates, especially for the test set ($n = 100$).
- **Single scheme focus.** The results are specific to CKKS. Whether they extend to BGV, BFV, or other RLWE-based schemes has not been shown.

7. FUTURE WORK AND EXTENSIONS

This study is exploratory. The main contributions are more about the method than the numbers. Several next steps are needed to turn the approach into something more solid:

- **Go to larger quantum processors.** The 4-qubit setup is just a start. Running on 16 to 32 or more qubits on IBM Quantum, IonQ, or similar platforms would show whether AI-guided oracle construction still helps at greater depth and how well it survives real noise.
- **Bring in realistic quantum noise.** Adding depolarizing noise, thermal relaxation, and crosstalk to the simulation would give a truer picture of performance on NISQ hardware.
- **Use real parameter sets.** Instead of synthetic labels, run the Lattice Estimator on actual parameter sets from FHE libraries such as SEAL, OpenFHE, and Lattigo. That would break the circular dependency where the network just learns the same formula used for evaluation.
- **Try other FHE schemes.** Apply the same hybrid method to BGV, BFV, and other RLWE-based schemes. That would show whether the vulnerability patterns seen here are unique to CKKS or show up across lattice-based cryptography.
- **Better neural architectures.** RNNs and Transformers might capture sequential dependencies in CKKS parameter sets better than a plain feed-forward network, especially for bootstrapping-level settings.
- **Cross-validate everything.** Use k-fold cross-validation and report means and standard deviations across splits. That would give much more reliable performance estimates.

8. CONCLUSION

This paper has looked at whether AI can help quantum attacks on CKKS fully homomorphic encryption parameters. The study is simulation-based and exploratory. A neural network was trained on 5,000 synthetic parameter sets, and quantum circuits were evaluated on 100 boundary-region configurations. In this simulated setting, AI-guided Grover oracle construction improved the success rate of finding vulnerabilities by 30.6% compared to standard quantum search.

Within this experimental model, parameter sets that claim 128-bit quantum security end up offering only 86 to 101 bits of effective security when facing an AI-guided adversary. That suggests the usual threat models that assume blind search may underestimate what an intelligent quantum attacker could do. Still, these are simulation results with real limits: only 4 qubits, synthetic labels that circle back to the same formulas, and perfect noiseless circuits. Translating these findings to actual CKKS deployments would require the validation steps laid out in Section 7.

Even with those limits, the work shows one thing clearly: neural network guidance can make a quantum oracle more efficient. That alone is worth more investigation at realistic scales. The rough parameter recommendations—about 2.3 times more compute overhead—offer a starting point for more conservative parameter choices while waiting for proper validation.

Going forward, the cryptography and quantum computing communities need to work together on AI-resistant parameter standards. Threat models must explicitly account for adaptive, learning-based adversaries. Standards bodies should also start thinking about AI-enhanced attack strategies when they evaluate post-quantum cryptography, especially as quantum hardware gets better.

AUTHOR CONTRIBUTION STATEMENT

This work was carried out in collaboration among all authors. All authors designed the study, performed the statistical analysis, and wrote the protocol. All authors managed the analyses of the study, managed the literature searches, and wrote the first draft of the manuscript. All authors read and approved the final manuscript.

ETHICS APPROVAL AND CONSENT TO PARTICIPATE

This study did not involve human participants or animals. Therefore, ethical approval and consent to participate are not applicable.

ACKNOWLEDGMENTS

The authors would like to thank the reviewers, Associate Editor, and Editor-in-Chief for their valuable comments and suggestions, which helped improve the quality of this paper. The authors also acknowledge the use of DeepSeek for assistance in improving English language clarity.

FUNDING

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

DISCLOSURE STATEMENT

The authors declare that they have no competing interests.

REFERENCES

- [1] C. Gentry, "A fully homomorphic encryption scheme," Ph.D. dissertation, Stanford University, 2009.
- [2] J. H. Cheon, A. Kim, M. Kim, and Y. Song, "Homomorphic encryption for arithmetic of approximate numbers," *Advances in Cryptology – ASIACRYPT 2017*, pp. 409–437, 2017.
- [3] H. Chen and K. Han, "Homomorphic lower digits removal and improved FHE bootstrapping," *Advances in Cryptology – EUROCRYPT 2018*, pp. 315–337, 2018.
- [4] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, "(Leveled) fully homomorphic encryption without bootstrapping," *ACM Trans. Computation Theory*, vol. 6, no. 3, article 13, 2014.
- [5] J. Fan and F. Vercauteren, "Somewhat practical fully homomorphic encryption," *Cryptology ePrint Archive*, Report 2012/144, 2012.
- [6] J. H. Cheon, K. Han, A. Kim, M. Kim, and Y. Song, "A full RNS variant of approximate homomorphic encryption," in *Selected Areas in Cryptography – SAC 2018, Lecture Notes in Computer Science*, vol. 11349, Springer, 2019, pp. 347–368.
- [7] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *Journal of the ACM*, vol. 56, no. 6, pp. 1–40, 2009.
- [8] C. P. Schnorr and M. Euchner, "Lattice basis reduction: Improved practical algorithms and solving subset sum problems," *Mathematical Programming*, vol. 66, no. 1, pp. 181–199, 1994.
- [9] M. R. Albrecht, R. Player, and S. Scott, "On the concrete hardness of learning with errors," *Cryptology ePrint Archive*, Report 2015/046, 2015.

- [10] M. R. Albrecht, B. R. Curtis, A. Deo, A. Davidson, R. Player, E. W. Postlethwaite, F. Virdia, and T. Wunderer, “Estimate all the {LWE, NTRU} schemes!,” *International Conference on Security and Cryptography for Networks*, pp. 351–367, 2018.
- [11] M. R. Albrecht et al., “Estimate all the LWE, NTRU schemes!” in *Proc. Security Cryptography Networks*, pp. 351–367, 2018.
- [12] T. Laarhoven, M. Mosca, and J. van de Pol, “Finding shortest lattice vectors faster using quantum search,” *Designs, Codes and Cryptography*, vol. 77, no. 2–3, pp. 375–400, 2015.
- [13] B. Kim, H. Park, and J. H. Cheon, “Revisiting the concrete security of Goldreich–Levin with applications to post-quantum CKKS bootstrapping,” in *Advances in Cryptology – ASIACRYPT 2021, Lecture Notes in Computer Science*, vol. 13092, Springer, 2021, pp. 623–653.
- [14] B. Li and D. Micciancio, “On the security of homomorphic encryption on approximate numbers,” in *Advances in Cryptology – EUROCRYPT 2021, Lecture Notes in Computer Science*, vol. 12696, Springer, 2021, pp. 648–677.
- [15] L. Ducas and W. van Woerden, “NTRU fatigue: How stretched is overstretched?” in *Advances in Cryptology – ASIACRYPT 2021, Lecture Notes in Computer Science*, vol. 13093, Springer, 2021, pp. 3–32.
- [16] M. R. Albrecht, B. R. Curtis, and T. Wunderer, “Exploring trade-offs in batch bounded distance decoding,” in *Proc. 12th Int. Conf. on Cryptology and Network Security (CANS 2021), Lecture Notes in Computer Science*, vol. 13099, Springer, 2021, pp. 467–487.
- [17] L. K. Grover, “A fast quantum mechanical algorithm for database search,” in *Proc. 28th Annual ACM Symposium on Theory of Computing*, 1996, pp. 212–219.
- [18] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer,” *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1484–1509, 1997.
- [19] S. Joshi and D. Gupta, “Grover’s algorithm in a 4-qubit search space,” *Journal of Quantum Computing*, vol. 3, no. 4, p. 137, 2021.
- [20] J. Biamonte et al., “Quantum machine learning,” *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
- [21] E. Farhi and H. Neven, “Classification with quantum neural networks on near term processors,” *arXiv preprint arXiv:1802.06002*, 2018.
- [22] A. Gohr, “Improving attacks on round-reduced Speck32/64 using deep learning,” in *Advances in Cryptology – CRYPTO 2019*, pp. 150–179, 2019.
- [23] S. Picek et al., “The curse of class imbalance and conflicting metrics with machine learning for side-channel evaluations,” *IACR Trans. Cryptographic Hardware Embedded Systems*, vol. 2019, no. 1, pp. 209–237, 2018.
- [24] J. So, “Deep reinforcement learning-based cryptanalytic attack on lightweight block cipher,” *IEEE Access*, vol. 8, pp. 183860–183870, 2020.
- [25] L. Wouters, E. Arribas, B. Gierlichs, and B. Preneel, “Revisiting a methodology for efficient CNN architectures in profiling attacks,” *IACR Trans. Cryptographic Hardware Embedded Systems*, vol. 2020, no. 3, pp. 147–168, 2020.
- [26] C. Garbin, X. Zhu, and O. Marques, “Dropout vs. batch normalization: an empirical study of their impact to deep learning,” *Multimedia Tools and Applications*, vol. 79, no. 19, pp. 12777–12815, 2020.
- [27] Y. N. Kunang, S. Nurmaini, D. Stiawan, and B. Y. Suprpto, “An end-to-end intrusion detection system with IoT dataset using deep learning with unsupervised feature extraction,” *International Journal of Information Security*, vol. 23, no. 3, pp. 1619–1648, 2024.
- [28] G. Alagic et al., “Status report on the third round of the NIST post-quantum cryptography standardization process,” NIST Internal Report 8413, 2022.