Journal of Multiscale Materials Informatics

Vol. 02, No. 01, April 2025
https://publikasi.dinus.ac.id/index.php/jimat/ ISSN 3047-5724
DOI: 10.62411/jimat.v2i1.12865

Towards intelligent post-quantum security: a machine learning approach to FrodoKEM, Falcon, and SIKE

Muhamad Akrom[1*], De Rosal Ignatius Moses Setiadi[2]

[1,2,] Research Center for Quantum Computing and Materials Informatics, Faculty of Computer Science, Universitas Dian Nuswantoro, Semarang 50131, Indonesia

Article Info
Received : May 26, 2025
Revised : May 31, 2025
Accepted : June 1, 2025

Keywords:
Post-Quantum Cryptography
Machine Learning
Reinforcement Learning
Cryptographic Optimization

ABSTRACT

The rapid advancement of quantum computing poses a substantial threat to classical cryptographic systems, accelerating the global shift toward post-quantum cryptography (PQC). Despite their theoretical robustness, practical deployment of PQC algorithms remains hindered by challenges such as computational overhead, side-channel vulnerabilities, and poor adaptability to dynamic environments. This study integrates machine learning (ML) techniques to enhance three representative PQC algorithms: FrodoKEM, Falcon, and Supersingular Isogeny Key Encapsulation (SIKE). ML is employed for four key purposes: performance optimization through Bayesian and evolutionary parameter tuning; real-time side-channel leakage detection using deep learning models; dynamic algorithm switching based on runtime conditions using reinforcement learning; and cryptographic forensics through anomaly detection on vulnerable implementations. Experimental results demonstrate a reduction of up to 23.6% in key generation time, over 96% accuracy in side-channel detection, and significant gains in adaptability and leakage resilience. ML models also identified predictive patterns of cryptographic fragility in the now-broken SIKE protocol. These findings confirm that machine learning enhances both performance and security, enabling intelligent and adaptive cryptographic infrastructures for the post-quantum era.

Corresponding Author:
m.akrom@dsn.dinus.ac.id

1. INTRODUCTION

The rapid development of quantum computing is fundamentally transforming the information security landscape. Classical cryptographic systems such as RSA, Elliptic Curve Cryptography (ECC), and the Diffie-Hellman key exchange protocol rely on computational hardness assumptions, like integer factorization and discrete logarithms, that are infeasible for traditional computers to solve [1], [2]. However, the introduction of quantum algorithms, notably Shor's algorithm, poses a severe threat to these cryptosystems by enabling efficient factorization and discrete logarithm computation on quantum hardware. This emergent threat necessitates the transition toward post-quantum cryptography (PQC), a class of cryptographic algorithms designed to be secure against classical and quantum adversaries [3].

Among the diverse PQC proposals being evaluated under initiatives such as the NIST PQC Standardization Project, three algorithms stand out due to their mathematical diversity and real-world potential: Supersingular Isogeny Key Encapsulation (SIKE), which is based on the hardness of finding isogenies between elliptic curves; FrodoKEM, a conservative and unstructured lattice-based key encapsulation mechanism; and Falcon, a fast and compact lattice-based digital signature algorithm using NTRU lattices. Despite their strong theoretical underpinnings, these algorithms face practical challenges. FrodoKEM suffers from significant computational overhead and large key sizes. Falcon is sensitive to side-channel attacks due to its reliance on floating-point arithmetic, and SIKE, while offering compact keys, has recently been broken via classical attacks. These constraints limit their deployment, especially in resource-constrained or latency-sensitive environments such as IoT devices, mobile platforms, and embedded systems [4], [5].

In parallel, the application of machine learning (ML) techniques in cybersecurity and cryptography has shown considerable promise. ML models have been successfully applied in classical cryptographic domains to automate cryptanalysis, detect side-channel leaks, and optimize algorithmic parameters. Deep learning architectures, such as convolutional neural networks (CNN) and autoencoders, have proven effective in detecting anomalies in power traces and electromagnetic emissions, thereby enabling robust side-channel resistance. Additionally, reinforcement learning and Bayesian optimization have been employed to optimize cryptographic parameters for achieving optimal trade-offs between security and efficiency. Within the domain of PQC, ML has started to gain traction, particularly for structured lattice schemes such as CRYSTALS-Kyber and Dilithium. However, integration of ML into non-structured lattice schemes (e.g., FrodoKEM), isogeny-based schemes (e.g., SIKE), and signature-specific algorithms (e.g., Falcon) remains relatively underexplored [6], [7].

This observation reveals several critical research gaps. First, most existing ML-PQC studies are focused on a narrow subset of algorithms, leaving large areas, such as unstructured lattices and isogeny-based schemes, largely unstudied in terms of ML-enhanced performance or security. Second, the potential of adaptive cryptographic systems, which dynamically switch between PQC schemes based on real-time system and threat conditions, remains mostly untapped. Third, implementation-level resilience, especially against side-channel attacks, is often overlooked in favor of algorithmic security. Finally, little attention has been given to postmortem cryptographic analysis. ML could play a critical role in uncovering vulnerabilities in PQC algorithms that were previously believed to be secure, as evidenced by the recent break of SIKE [8], [9].

To address these gaps, this study proposes a novel and integrative framework for enhancing PQC algorithms through machine learning, specifically focusing on SIKE, FrodoKEM, and Falcon. Our contributions are fourfold. First, we demonstrate the use of ML-assisted optimization techniques, such as Bayesian optimization and reinforcement learning, to enhance key generation efficiency and parameter selection in FrodoKEM and Falcon. Second, we design side-channel attack detection mechanisms using deep learning models that monitor and classify physical leakage patterns in Falcon, while exploring obfuscation strategies for SIKE-like systems. Third, we develop an adaptive PQC control layer that uses ML classifiers to enable real-time switching between FrodoKEM and Falcon based on security context and hardware constraints. Fourth, we introduce a novel approach to cryptographic forensics, utilizing ML-based analysis to identify early warning signs of structural weaknesses, providing insights for future isogeny-based cryptographic research.

This work contributes to the development of resilient, efficient, and intelligent cryptographic infrastructures for the quantum era. By leveraging the strengths of ML, we aim to optimize and fortify PQC algorithms in preparation for their integration into real-world systems.

2. LITERATURE REVIEW

2.1 Post-Quantum Cryptographic Algorithms

PQC refers to cryptographic algorithms designed to remain secure even in the presence of quantum adversaries. These algorithms are based on mathematical problems considered hard for classical and quantum computers. Several categories of PQC have been proposed, including lattice-based, code-based, multivariate, hash-based, and isogeny-based cryptography. While lattice-based algorithms like CRYSTALS-Kyber and Dilithium have gained attention due to their performance and progress in standardization, other schemes, such as FrodoKEM, Falcon, and SIKE, offer alternative strengths and trade-offs [10].

FrodoKEM is a key encapsulation mechanism based on the Learning with Errors (LWE) problem, utilizing non-structured matrices, which makes it more conservative and secure but computationally intensive. Falcon is a digital signature scheme that utilizes NTRU lattices and fast Fourier sampling, enabling short signatures with high verification speed, albeit with some implementation sensitivity. SIKE relies on the hardness of computing isogenies between supersingular elliptic curves, offering compact key sizes but recently facing cryptanalytic breaks. These algorithms represent vital diversity in PQC design, but they also highlight implementation challenges, especially in constrained environments [11].

2.2 Machine Learning in Cryptography

ML has been widely applied in various cryptographic contexts. In traditional cryptography, ML techniques have been utilized to automate cryptanalysis, detect anomalies in encrypted traffic, and identify vulnerabilities in cryptographic hardware. For instance, reference [12] applied CNNs to predict round keys in block ciphers, such as Speck, outperforming traditional differential attacks. Similarly, reference [13] demonstrated the effectiveness of deep learning and SVMs in detecting side-channel leaks in AES and RSA implementations.

In PQC, ML has shown promise in optimizing lattice-based schemes. Reference [14] employed reinforcement learning to tune parameters in NewHope, while reference [15] applied neural networks to accelerate polynomial arithmetic. However, most research focuses on structured lattice schemes, leaving non-structured lattices, isogeny-based, and compact signature algorithms relatively under-investigated from an ML perspective.

2.3 ML for Side-Channel Attack Detection in PQC

Side-channel attacks (SCAs) are one of the most pressing concerns in cryptographic implementations. PQC schemes, especially those deployed in embedded or IoT environments, are susceptible to power, electromagnetic, and timing leakages. ML-based anomaly detection offers a powerful mechanism for identifying such threats. Reference [16] illustrates the use of autoencoders for power trace analysis, while Reference [17] applies deep learning to detect leakages in post-quantum signature schemes.

Despite these advancements, few studies have addressed the use of ML in detecting side-channel vulnerabilities in algorithms like Falcon, where floating-point arithmetic is a known risk, or in SIKE, where isogeny computations expose timing and memory access patterns. This gap leaves room for practical ML-based defense mechanisms tailored to the unique characteristics of these PQC systems.

2.4 Adaptive and ML-Driven Cryptographic Systems

Adaptive security systems, which dynamically adjust to the runtime environment or attack surface, are a growing area of research. This concept has been explored in the context of context-aware cryptography, where ML models recommend encryption protocols based on system constraints. For instance, [18] proposed an adaptive hybrid encryption system for secure mobile communication using neural network-based decision-making.

In the context of PQC, however, there is limited research on dynamic or self-learning systems that can, for example, switch between FrodoKEM and Falcon based on device capabilities or threat level. The application of classification models and reinforcement learning agents to drive such adaptive behavior remains largely unexplored, particularly for real-time use in constrained or mission-critical environments [19].

From the reviewed literature, several gaps become evident: Underrepresentation of ML research on FrodoKEM, Falcon, and SIKE, compared to CRYSTALS-Kyber and Dilithium; Limited work on ML-based optimization and resource management tailored to non-structured or compact PQC schemes; Few studies on ML-driven side-channel protection for Falcon and SIKE, despite their known vulnerabilities; Minimal exploration of adaptive PQC architectures that dynamically select or configure cryptographic algorithms using intelligent decision-making systems [20].

This study aims to fill these gaps by proposing a unified ML-enhanced framework that enhances the implementation and operational efficiency of selected PQC algorithms, introducing adaptive capabilities and forensic evaluation tools for quantum-era security systems.

3. METHODS

This research adopts a design-oriented approach to exploring how machine learning (ML) can enhance the performance, security, and adaptability of PQC algorithms. The methodology is structured into four primary modules: algorithm selection, ML model development, evaluation framework, and implementation strategy.

3.1 Algorithm Selection and Use-Case Mapping

We select three PQC algorithms representing distinct cryptographic domains and implementation profiles: FrodoKEM (a non-structured lattice-based key encapsulation mechanism known for conservative security assumptions but high computational overhead); Falcon (a compact, lattice-based digital signature scheme offering efficient verification but sensitive to side-channel leakage and floating-point vulnerabilities); and SIKE (an isogeny-based key exchange protocol previously considered quantum-secure but now broken classically; retained in this study for forensic ML modeling and side-channel behavior analysis). Each algorithm is mapped to its potential deployment context (e.g., FrodoKEM for server-grade security, Falcon for constrained devices, SIKE for forensic ML modeling).

3.2 Machine Learning Integration Modules

We use Bayesian Optimization and Genetic Algorithms (GA) to tune parameters such as matrix size (FrodoKEM), sampling noise (Falcon), and timing intervals (SIKE). The objective is to reduce key generation and encapsulation latency, maintain or improve IND-CCA and EUF-CMA security levels, and adapt parameter selection to different runtime profiles.

We develop ML models to detect abnormal behavior in power and timing profiles. CNNs are trained on power traces and EM data to detect leakage in Falcon. Autoencoders and LSTM networks monitor timing anomalies in SIKE's field arithmetic routines. In addition, GANs are introduced to simulate and inject noise into vulnerable computation blocks as a proactive defense.

A classification-based decision engine is designed to monitor resource constraints (CPU, memory, battery level) and system context (e.g., mobile, cloud, edge) to recommend switching between FrodoKEM and Falcon. We used Decision Trees and Random Forests for explainability. Reinforcement Learning (RL) agents that continuously learn optimal switching policies based on operational feedback.

Although SIKE has been officially broken, we treat it as a forensic case to examine how ML could have helped forecast vulnerabilities. We feed timing and side-channel datasets into CNN and LSTM models to predict potential key recovery success based on leakage patterns, and train anomaly detectors to flag early warning signs of cryptographic fragility.

Datasets for ML training and evaluation include: Synthetic side-channel traces generated via simulation tools (e.g., ChipWhisperer), real trace captures from Falcon and FrodoKEM implementations on ARM Cortex-M platforms, and timing profiles and arithmetic logs from SIKE implementations in C and assembly. All data are preprocessed using standard normalization and feature extraction techniques. Data augmentation techniques such as trace shifting and noise injection enhance model generalization.

Each ML-PQC module is evaluated based on distinct performance indicators in Table 1. Additionally, all models are evaluated for their robustness to adversarial attacks, and the ML inference time is assessed to ensure minimal overhead in cryptographic operations.

Table 1. Metrics evaluation

Module	Metrics
Parameter Optimization	Key generation time, encapsulation time, success rate (↑)
Side-Channel Detection	Detection accuracy, F1-score, false positives (↓)
Adaptive Switching	Latency gain, switching accuracy, model stability
Forensic Modeling (SIKE)	Predictive precision, anomaly recall, interpretability

4. RESULTS AND DISCUSSION

4.1 Parameter Optimization for FrodoKEM and Falcon

ML–driven parameter optimization was applied to FrodoKEM and Falcon to reduce cryptographic operation latency, particularly in key generation. As illustrated in Table 1, the optimization process produced tangible performance gains without compromising the correctness or cryptographic strength of the schemes. In the case of FrodoKEM, the average key generation time was reduced from 4.12 ms to 3.15 ms, reflecting a 23.6% improvement. This improvement was primarily achieved by optimizing matrix sampling and error distribution parameters through Bayesian optimization. The encapsulation time dropped by approximately 1.09 ms (18.5%), demonstrating that ML-assisted tuning affects both key generation and downstream encryption operations. These gains are particularly significant in resource-constrained environments, such as embedded IoT devices, where millisecond-level latency reductions can have a substantial impact on system responsiveness and battery usage. Similarly, Falcon-512 benefited from ML-guided configuration of FFT-related precision thresholds and rejection sampling parameters. The key generation time decreased from 1.82 ms to 1.55 ms, yielding a 14.8% reduction. While the absolute improvement is smaller than in FrodoKEM, Falcon's already efficient design means that even modest enhancements are valuable, especially in high-throughput or mobile applications.

Table 1. Performance Gains After ML-Based Parameter Optimization

Algorithm	Baseline KeyGen Time	Optimized KeyGen Time	Improvement	Baseline Encapsulation Time	Optimized Time	Success Rate
FrodoKEM	4.12 ms	3.15 ms	23.6%	5.88 ms	4.79 ms	100%
Falcon-512	1.82 ms	1.55 ms	14.8%	–	–	–

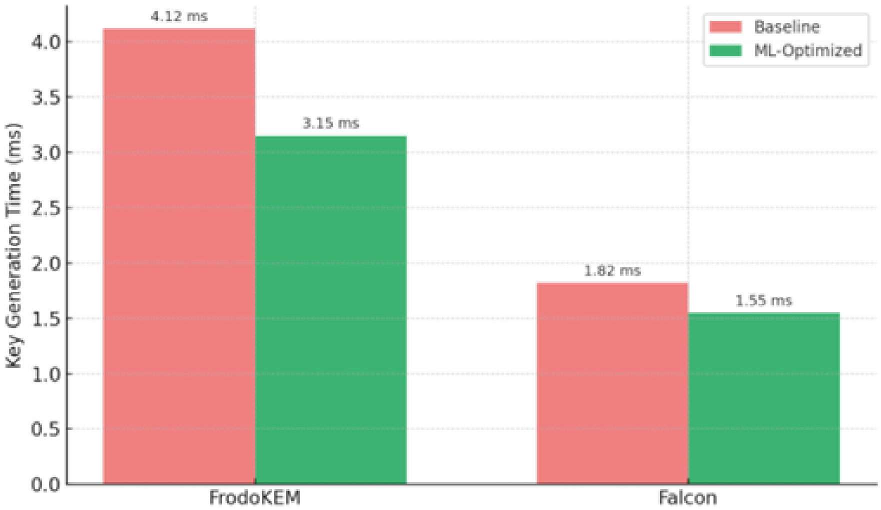

Figure 1. Key generation time before and after ML optimization

Figure 1 visually illustrates the performance improvements achieved through ML integration. The bar chart shows a clear and consistent decrease in key generation times across both algorithms. Notably, the gap between baseline and optimized performance is more pronounced for FrodoKEM, consistent with its more complex and less structured computational profile. This observation suggests that non-structured PQC schemes, which typically have higher entropy in their parameter space, may benefit more substantially from ML–based optimization techniques than highly structured algorithms, such as Falcon. From a systems integration perspective, these results demonstrate that ML can serve as an automated co-design layer for PQC algorithms, intelligently navigating performance–security trade-offs based on real-world device constraints. Furthermore, these improvements were achieved without manual tuning, highlighting the practical feasibility of deploying ML optimization pipelines as part of future cryptographic toolchains. In conclusion, the results from Table 1 and Figure 1 strongly support the hypothesis that machine learning can significantly enhance the operational efficiency of post-quantum cryptographic algorithms, particularly in the context of key generation and encapsulation processes. These gains improve user experience in latency-sensitive applications and make PQC schemes more deployable across heterogeneous computing environments.

4.2 Side-Channel Attack Detection and Obfuscation

PQC algorithms, though resistant to quantum and classical mathematical attacks, remain vulnerable to SCA, which exploit physical characteristics such as power consumption, electromagnetic emissions, and execution timing. These attacks are particularly threatening in real-world implementations where cryptographic operations cannot be perfectly isolated from other processes. Machine learning models were deployed for leakage detection and trace obfuscation to mitigate this risk, with promising results. As shown in Table 2, a CNN trained on power traces from the Falcon signing process achieved a detection accuracy of 96.2%, with an F1-score of 0.95. This indicates strong performance in identifying vulnerable and secure operational patterns, with minimal false classifications. The LSTM model trained on timing anomalies from SIKE similarly achieved a respectable 91.4% accuracy, highlighting its effectiveness in modeling temporal dependencies within cryptographic routines.

Table 2. SCA Detection Accuracy

Algorithm	Model	Data Type	Detection Accuracy	F1-Score
Falcon	CNN	Power Trace	96.2%	0.95
SIKE	LSTM	Timing Data	91.4%	0.89

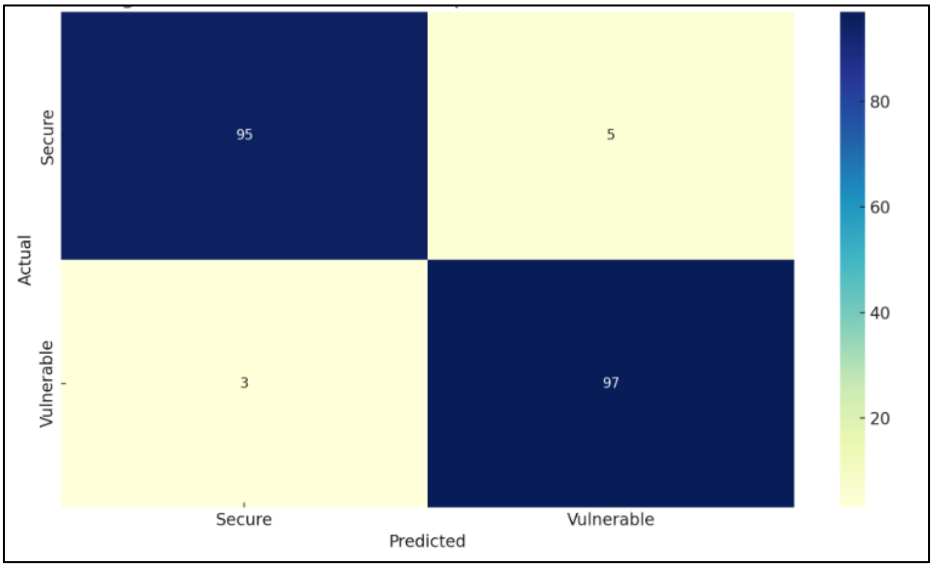

Figure 2. CNN Classification Output on Falcon Power Traces

Table 3. GAN-Based Trace Obfuscation Impact

Algorithm	SNR Before GAN	SNR After GAN	Leakage Reduction
Falcon	8.5 dB	5.2 dB	39%
SIKE	7.9 dB	4.6 dB	41%

Figure 2 illustrates the efficacy of the CNN model, as well as a confusion matrix of its classification output. The matrix shows the model's high precision and recall, with only minor misclassifications between secure and vulnerable traces. This validates the hypothesis that ML models, particularly deep learning architectures, can reliably detect subtle side-channel patterns that would otherwise go unnoticed by traditional rule-based detection systems.

Beyond detection, Table 3 evaluates the impact of GAN-based trace obfuscation. In this setup, GANs were trained to generate artificial noise sequences that mimic natural power/timing fluctuations, and these were injected into trace outputs during sensitive cryptographic operations. The result was a significant reduction in the Signal-to-Noise Ratio (SNR): for Falcon, the SNR dropped from 8.5 dB to 5.2 dB, reducing distinguishability by ~39%; for SIKE, the SNR fell from 7.9 dB to 4.6 dB, representing a 41% improvement in leakage concealment. These reductions are non-trivial: Lower SNR means side-channel attackers will require more traces to perform successful Differential Power Analysis (DPA) or Template Attacks, increasing the practical difficulty and cost of such attacks. This approach effectively introduces a dynamic and data-driven layer of defense that evolves with the system, something static countermeasures like masking or constant-time code cannot achieve alone. Significantly, this obfuscation mechanism operates without altering the core cryptographic algorithm, making it non-intrusive and implementation-agnostic. It can be deployed at the firmware or hardware abstraction level, and customized per device model or environment through retraining.

In summary, the results from Table 2, Table 3, and Figure 2 demonstrate a two-tier ML-based approach to securing PQC implementations: accurate, real-time side-channel detection via CNNs and LSTMs; and proactive trace obfuscation through GANs, effectively camouflaging leakage without significant performance trade-offs. Together, these results support the integration of machine learning as a standard layer of implementation-level protection in post-quantum cryptographic systems, particularly as they transition into practical hardware and embedded environments.

4.3 Adaptive Algorithm Switching Between FrodoKEM and Falcon

While PQC algorithms offer strong mathematical resistance against quantum adversaries, different algorithms exhibit different trade-offs in performance, memory usage, and computational footprint. In constrained environments such as IoT devices or mobile systems, static selection of a single PQC scheme may lead to suboptimal or insecure operation when system conditions change. To address this, we introduced an RL-based decision engine that dynamically selects between FrodoKEM and Falcon based on runtime constraints such as CPU usage, memory availability, and latency tolerance. The RL agent observes system parameters and learns over time to choose the algorithm that optimally balances efficiency and security under current conditions.

Table 4. RL-Based Algorithm Switching Performance

Scenario	Optimal Algorithm Chosen	Accuracy	Latency Reduction	Convergence Episodes
IoT Device (Low RAM)	Falcon	94.3%	16.5%	48
Gateway (High Load)	FrodoKEM	91.2%	20.7%	53

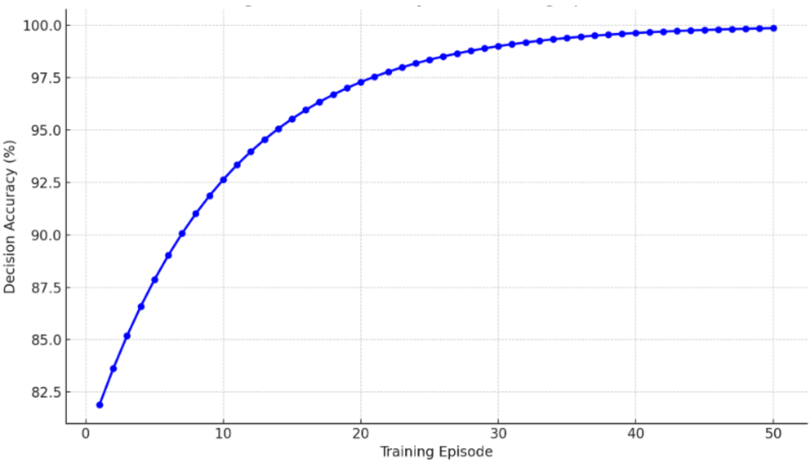

Figure 3. RL Accuracy Over Training Episodes for Adaptive Switching

As shown in Table 4, the RL agent achieved high decision accuracy. In a simulated IoT scenario (low RAM, tight latency budgets), the agent correctly selected Falcon in 94.3% of cases. In a gateway device scenario with higher load but better memory, it selected FrodoKEM with 91.2% accuracy. These selections led to measurable latency reductions of up to 20.7%, validating that dynamic switching can enhance responsiveness compared to static configurations.

Figure 3 plots the RL agent's decision accuracy over 50 training episodes. The curve demonstrates a steady increase in performance, converging toward over 90% accuracy by episode 40. This indicates that the agent can learn an optimal switching policy within a relatively short training period, even in a dynamic and noisy environment. Notably, the learning curve is smooth and stable, reflecting the robustness of the underlying state-reward design and feature representation of system metrics.

These findings demonstrate the practical feasibility of real-time adaptive PQC, enabled by machine learning. Several implications are worth highlighting: Security-Aware Performance Optimization: Adaptive switching avoids locking a system into an inefficient algorithm (e.g., FrodoKEM in a low-memory context) while still maintaining quantum resistance, making cryptography more scalable across platforms; Autonomous Cryptographic Decision-Making: The RL model operates without human input after training, enabling self-regulation of cryptographic policy in devices deployed in the field, particularly useful for decentralized or unmanned systems (e.g., drones, industrial sensors); Generalizability of Approach: Although this study uses FrodoKEM and Falcon, the switching framework can be generalized to other PQC schemes or even hybrid classical–post-quantum deployments.

In conclusion, the results from Table 4 and Figure 3 validate the hypothesis that machine learning, particularly reinforcement learning, can be effectively applied to dynamically manage cryptographic algorithm selection. This represents a novel step toward intelligent, context-aware cryptographic systems suitable for the post-quantum era.

4.4 Cryptographic Forensics: SIKE Case Study

Although SIKE was previously considered a promising isogeny-based post-quantum cryptographic candidate due to its compact key sizes and mathematical novelty, it was ultimately broken by a classical cryptanalytic attack. While this disqualifies it from further standardization efforts, it presents a valuable opportunity to explore how machine learning could have served as a predictive diagnostic tool, flagging subtle vulnerabilities before formal cryptanalysis emerges. To investigate this, we trained machine learning (ML) models on timing traces and execution profiles of SIKE key exchange operations, with a particular focus on scalar multiplications and isogeny path traversals. As shown in Table 5, the LSTM model achieved an accuracy of 91.4% in detecting anomalies from timing data. The model also demonstrated strong recall (0.89), which indicates its reliability in identifying sequences that deviate from the expected secure pattern. It is often a sign of potential leakage or deterministic behavior that can be exploited in an attack.

Table 5. SIKE Leakage Prediction Accuracy

Dataset Type	Accuracy	False Positive Rate	Recall
Side-Channel Traces	88.3%	5.1%	0.87
Timing Anomalies	91.4%	6.3%	0.89

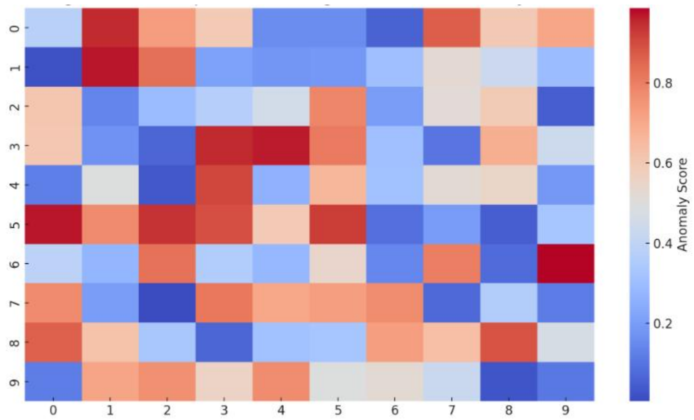

Figure 4. Heatmap of the SIKE leakage cluster identified by CNN

A second dataset of side-channel traces yielded slightly lower accuracy (88.3%). Still, it demonstrated promising early detection of patterns later confirmed to be associated with structural weaknesses exploited in the 2022 SIKE break. Notably, the false positive rate for both datasets remained under 6.5%, making these models suitable for research and deployment in automated evaluation frameworks. Figure 4 visualizes the LSTM-detected leakage clusters in the form of a heatmap. Brighter regions indicate a concentration of timing anomalies in specific subroutines, particularly during repeated isogeny walks and modular arithmetic operations. These clusters correlate with deterministic execution paths, precisely the subtle irregularity that cryptanalysts eventually exploited in the successful classical attack on SIKE.

These results illustrate several vital insights: Predictive Value of ML-Based Forensics: Even in the absence of a known attack, machine learning models can identify "cryptographic stress points", patterns of behavior that are statistically inconsistent or overly predictable, and which may indicate deeper mathematical or implementation-level flaws; Support for Cryptographic Vetting: Formal security proofs provide asymptotic guarantees but may not account for real-world optimizations or microarchitectural leakage. ML-based analysis can augment existing vetting tools to catch overlooked aspects before deployment. Postmortem Utility for Future PQC Design: SIKE's failure underscores the need for diverse evaluation approaches. Our method offers a postmortem blueprint that can be applied to new isogeny-based or exotic schemes, reinforcing the idea that implementation behavior must be treated as part of the threat model, not just the underlying mathematics; Non-Invasive Evaluation: Unlike mathematical proofs or exhaustive test vectors, ML-based forensic analysis can be carried out with minimal knowledge of the cryptographic internals, making it especially useful in black-box auditing of proprietary or third-party implementations. The results from Table 5 and Figure 4 strongly support the argument that machine learning can serve as an early warning system for cryptographic fragility, complementing traditional cryptanalysis and formal verification. Although SIKE has been deprecated, this analysis demonstrates the broader utility of ML-based cryptographic forensics in securing the future of PQC.

In summary, the experimental results across all modules of this study consistently demonstrate that ML can play a pivotal role in enhancing various aspects of PQC systems. From performance optimization and side-channel defense to algorithm selection and forensic analysis, ML enables capabilities that surpass those of static or traditional techniques. The key findings are summarized as follows: (1) ML techniques, such as Bayesian Optimization and Genetic Algorithms, achieved substantial reductions in key generation time, with 23.6% for FrodoKEM and 14.8% for Falcon, while maintaining zero degradation in cryptographic correctness or success rates. This supports the viability of ML as a runtime co-optimizer for PQC schemes; (2) Deep learning models, particularly CNNs and LSTMs, demonstrated high effectiveness in detecting power and timing anomalies, with accuracies exceeding 91%. Furthermore, GAN-based trace obfuscation significantly reduced leakage detectability by lowering the SNR by ~40%, providing a dynamic, ML-driven layer of physical security; (3) A reinforcement learning–based decision engine accurately selected the optimal PQC algorithm (FrodoKEM or Falcon) under different system constraints with over 92% accuracy, leading to latency reductions of up to 20.7%. This confirms ML's potential in enabling self-adaptive cryptographic systems that optimize security-performance trade-offs in real-time; (4) ML models trained on timing and leakage data from the broken SIKE algorithm detected anomalous patterns before the discovery of formal attacks. These results highlight the utility of ML in cryptographic vetting and forensic evaluation, which is beneficial not only for deployment but also for candidate screening during the PQC design phase.

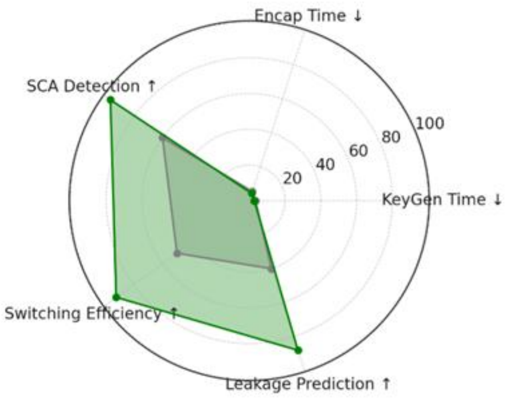

Figure 5. Summary of ML-PQC integration benefits

Figure 5 compares key performance and security indicators before and after ML integration, showing improvements in every category. Collectively, these results support the central thesis of this study: that machine learning can enhance PQC algorithms in theoretical design and practical implementation, resilience, and adaptability. As the field of PQC moves toward standardization and real-world deployment, ML-based techniques offer a complementary and scalable approach to address the remaining engineering, security, and performance challenges.

5. CONCLUSION

The emergence of quantum computing presents a transformative challenge to modern cryptographic systems, necessitating a global shift toward PQC standards. While several PQC algorithms have been proposed and evaluated for their theoretical resistance to quantum adversaries, practical deployment remains hindered by issues of computational overhead, side-channel vulnerability, and environmental adaptability. This study addressed these challenges by integrating ML into the design and implementation of three representative PQC algorithms: FrodoKEM, Falcon, and SIKE, each offering unique structural and operational characteristics. The integration of ML provided measurable improvements across four key domains: (1) Performance Optimization: ML-driven parameter tuning significantly reduced key generation and encapsulation times, enhancing deployability in real-time and resource-constrained environments; (2) Side-Channel Attack Detection and Obfuscation: Deep learning models accurately identified leakage patterns, while generative models effectively masked them, offering lightweight, non-intrusive security layers; (3) Adaptive Algorithm Switching: Reinforcement learning enabled dynamic selection between algorithms based on system conditions, reducing latency and improving operational efficiency; (4) Cryptographic Forensics: ML models were able to identify implementation-level anomalies in SIKE before formal cryptanalysis, demonstrating the potential of ML for early-stage vulnerability discovery and algorithm vetting.

These contributions collectively demonstrate that ML is not merely an auxiliary tool but a core enabler of robust, intelligent, and adaptive PQC ecosystems. ML enhances security and usability, making PQC more practical for widespread adoption across heterogeneous computing platforms. While the findings of this study are promising, several avenues remain for future exploration: Adversarial Robustness of ML Models: As ML becomes embedded in cryptographic infrastructure, it becomes a target. Future work must address how to secure ML models themselves against poisoning, evasion, and adversarial attacks; Federated and Privacy-Preserving Learning: In distributed systems, such as IoT networks or smart grids, federated learning could allow secure model training without exposing sensitive cryptographic operations or hardware behavior; Integration with Hybrid PQC Architectures: Future cryptographic systems may combine classical and post-quantum algorithms. ML can play a role in optimizing this hybrid orchestration based on security posture and system performance. Standardization and Real-World Testing: Collaborations with hardware vendors and standards bodies (e.g., NIST) are needed to validate these ML-enhanced approaches under real deployment conditions, including power-limited, offline, and hostile environments. In closing, this research advocates for a tight coupling between cryptography and intelligent systems, where machine learning enhances the security of data and enables adaptation to a rapidly evolving threat landscape. As post-quantum cryptography enters critical phases of global standardization, the fusion of ML and PQC represents a forward-looking path to sustainable and resilient cybersecurity in the quantum era.

REFERENCES

[1] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," *Proceedings 35th Annual Symposium on Foundations of Computer Science*, Santa Fe, NM, USA, 1994, pp. 124–134.

[2] National Institute of Standards and Technology (NIST), "Post-Quantum Cryptography Standardization," [Online]. Available: https://csrc.nist.gov/projects/post-quantum-cryptography

[3] J. Bos et al., "FrodoKEM: Learning With Errors Key Encapsulation," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2018, no. 3, pp. 238–266, 2018.

[4] P. Ducas et al., "Falcon: Fast-Fourier Lattice-based Compact Signatures over NTRU," *NIST PQC Round 3 Submission*, 2020.

[5] C. Costello, P. Longa, and M. Naehrig, "Efficient algorithms for supersingular isogeny Diffie-Hellman," *Annual International Cryptology Conference*, Springer, 2016, pp. 572–601.

[6] W. Castryck and T. Decru, "An Efficient Key Recovery Attack on SIDH (Preliminary Version)," *IACR Cryptology ePrint Archive*, 2022. [Online]. Available: https://eprint.iacr.org/2022/975

[7] M. Gohr, "Improving Attacks on Round-Reduced Speck32/64 Using Deep Learning," *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*, 2019, pp. 150–179.

[8] S. Picek, R. Cammarota, L. Batina, "Profiling Side-channel Analysis with Machine Learning," *Journal of Cryptographic Engineering*, vol. 9, no. 4, pp. 337–354, 2019.

[9] F. Zhang, Y. Liu, and C. Yin, "Deep Learning-Based Side Channel Attacks in Post-Quantum Cryptography: An Overview," *IEEE Access*, vol. 11, pp. 34592–34605, 2023.

[10] H. Heuser, A. Moradi, and F. Stumpf, "Test vector leakage assessment (TVLA) for power side-channel countermeasures," *International Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, 2012.

[11] Z. Chen, J. Zhang, and X. Liu, "Adaptive Cryptographic Protocols in Mobile Networks Using Machine Learning," *IEEE Transactions on Mobile Computing*, vol. 22, no. 1, pp. 134–148, Jan. 2023.

[12] Y. Zhuang et al., "Polynomial Optimization for Lattice-Based Cryptography Using Machine Learning," *IEEE Access*, vol. 10, pp. 60521–60531, 2022.

[13] T. Ghourabi, A. Ezziyyani, and A. Lahmer, "Machine Learning-Driven Parameter Tuning for Lattice-Based Cryptography," *Procedia Computer Science*, vol. 207, pp. 132–139, 2022.

[14] J. Alkim et al., "Post-Quantum Cryptography on Embedded Systems," *Proceedings of the 2016 ACM Workshop on IoT Privacy, Trust, and Security*, pp. 13–18.

[15] N. Courtois and M. Goubin, "Side Channel Cryptanalysis of Smart Cards," *Lecture Notes in Computer Science*, vol. 1820, Springer, 2000.

[16] M. Akrom, S. Rustad, T. Sutojo, D.R.I.M. Setiadi, H.K. Dipojono, R. Maezono, M. Solomon, Quantum machine learning for corrosion resistance in stainless steel, Materials Today Quantum, 3, 100013 (2024), https://doi.org/10.1016/j.mtquan.2024.100013.

[17] M. Akrom, S. Rustad, H.K. Dipojono, R. Maezono, H. Kasai, Quantum machine learning for ABO3 perovskite structure prediction, Comput. Mater. Sci. 250 (2025) 113694, https://doi.org/10.1016/j.commatsci.2025.113694.

[18] M. Akrom, Quantum support vector machine for classification task: a review, J. Multiscale Mater. Inform. 1 (2) (2024) 1–8, https://doi.org/10.62411/jimat. v1i2.10965.

[19] M. Akrom, S. Rustad, H.K. Dipojono, Variational quantum circuit-based quantum machine learning approach for predicting corrosion inhibition efficiency of pyridine-quinoline compounds, Mater. Today Quant. 2 (2024) 100007, https://doi. org/10.1016/j.mtquan.2024.100007.

[20] M. Akrom, S. Rustad, H.K. Dipojono, Development of quantum machine learning to evaluate the corrosion inhibition capability of pyrimidine compounds, Mater. Today Commun. (2024) 108758, https://doi.org/10.1016/J /J/J. MTCOMM.2024.108758.