KTREND JOURNALS
International Journal of Computer Science and Artificial Intelligence (IJCSAI)

DOI: 10.5281/zenodo.20651035 **Volume:** 1 **Issue:** 1 **ISSN:** Pending

Machine Learning Assisted Selection of Algebraic Structures for Post-Quantum Cryptographic Systems

Michael Nsikan John[1]

[1]Department of Mathematics, Edo State University, Iyamho, Nigeria

Corresponding Author: john.michael@edouniversity.edu.ng

Abstract

The migration from classical public-key cryptography to post-quantum cryptography requires systematic methods for identifying algebraic structures whose computational properties remain secure against both classical and quantum adversaries. Traditional selection is largely based on theoretical hardness assumptions, implementation cost, parameter constraints, and expert judgement. This paper develops a Machine Learning Assisted Algebraic Structure Selection Framework (MLASSF) for classifying candidate algebraic structures as suitable or unsuitable for post-quantum cryptographic use. The framework combines finite group invariants, conjugacy class statistics, character-theoretic indicators, elliptic-curve descriptors, lattice parameters, implementation efficiency, and estimated security categories into a supervised learning pipeline. A reproducible simulated dataset of 5,000 candidate structures was generated to model algebraic and computational features relevant to post-quantum design. Five classifiers were trained and evaluated: logistic regression, support vector machine, random forest, gradient boosting, and artificial neural network. The best-performing model was gradient boosting, with an accuracy of 93.80%, precision of 95.04%, recall of 96.21%, F1-score of 95.62%, and ROC-AUC of 95.29%. Random forest followed closely with 93.60% accuracy and also provided interpretable feature importance, showing that lattice dimension, implementation efficiency, character-degree variance, elliptic-curve rank, and derived length were influential predictors. The results show that machine learning can serve as a decision-support layer for algebraic selection in post-quantum cryptographic research, especially when combined with rigorous mathematical analysis rather than used as a substitute for proof-based security.

Keywords: Post-quantum cryptography; machine learning; algebraic structures; computational group theory; conjugacy classes; lattice-based cryptography; elliptic curve cryptography; artificial intelligence.

1 Introduction

Public-key cryptography relies on computational hardness assumptions arising from algebraic and number-theoretic structures. RSA is founded on integer factorization, while elliptic-curve cryptography depends on the elliptic-curve discrete logarithm problem. These assumptions are threatened by quantum algorithms, particularly Shor's algorithm, which gives polynomial-time procedures for factoring and discrete logarithms on sufficiently powerful quantum computers [7]. The current direction of cryptographic standardization has therefore shifted toward post-quantum schemes based on lattices, codes, hash functions, and other mathematical structures believed to resist quantum attacks [10, 11].

The algebraic foundations of cryptography remain central to this transition. Lattice-based schemes such as ML-KEM and ML-DSA have become leading standards because they combine strong security assumptions with practical performance. However, group-theoretic structures, conjugacy classes, solvable groups, character theory, elliptic curves, and Lie algebraic constructions continue to provide useful frameworks for cryptographic modelling, protocol design, and security analysis [1, 2, 3, 4, 5]. The challenge is no longer only to design algebraic systems, but also to select from many candidate structures under competing criteria such as security, efficiency, implementation feasibility, and resistance to quantum attacks.

This paper proposes a machine-learning-assisted method for selecting algebraic structures for post-quantum cryptographic systems. The motivation is that many algebraic descriptors can be represented as structured numerical features, and such features can be used to train predictive models. The purpose is not to replace mathematical proof, but to provide an intelligent screening mechanism that prioritizes promising structures for deeper cryptographic analysis.

1.1 Contribution of the Study

The main contributions are as follows:

(i) A formal framework for representing algebraic structures through cryptographically relevant feature vectors.

(ii) A supervised learning model for classifying structures as suitable or unsuitable for post-quantum cryptographic application.

(iii) A reproducible simulated experiment using 5,000 candidate algebraic structures.

(iv) Comparative analysis of logistic regression, support vector machine, random forest, gradient boosting, and artificial neural network classifiers.

(v) Feature-importance analysis showing which algebraic indicators contribute most strongly to classification.

1.2 Scope and Limitation

This paper uses simulated experiments. The simulated labels do not constitute real-world cryptographic certification. Instead, the experiment provides a controlled environment for testing the feasibility of machine learning as a decision-support method. Any algebraic structure selected by the proposed framework must still be subjected to rigorous security proof, peer review, implementation testing, and cryptanalysis.

2 Related Literature

Post-quantum cryptography has developed rapidly due to the vulnerability of RSA and ECC to quantum attacks. Shor [7] introduced quantum algorithms for factoring and discrete logarithms. Regev [8] established the learning with errors problem as a foundation for lattice-based cryptography. Micciancio and Regev [9] provided a comprehensive treatment of lattice-based cryptographic constructions. The NIST standardization process selected lattice-based and hash-based schemes for standardization, including ML-KEM, ML-DSA, and SLH-DSA [10, 11].

Group-theoretic cryptography has also been widely studied. Conjugacy search problems, small cancellation properties, nilpotent structures, and finite group invariants have been explored as possible cryptographic platforms. John, Musa, and Udoaka [1] investigated conjugacy classes in finitely generated groups with small cancellation properties. John, Udoaka, and Musa [2] examined solvable groups with monomial characters of prime power codegree and monolithic characters. Their work supports the use of group invariants as descriptors of algebraic complexity.

Elliptic-curve and bilinear structures remain important in modern cryptography even though classical ECC is not post-quantum secure against Shor-type attacks. John, Udoaka, and Nwala [3] discussed elliptic-curve groups in quantum-era cryptography, while John, Udoaka, and Musa [5] studied symmetric bilinear cryptography on elliptic curves and Lie algebras. John et al. [4] examined number theory in RSA encryption systems, emphasizing the role of arithmetic structure in cryptographic security.

Machine learning has increasingly been applied in cybersecurity for intrusion detection, malware classification, anomaly detection, side-channel analysis, and cryptographic parameter selection. Goodfellow, Bengio, and Courville [20] provide a general foundation for deep learning, while Breiman [18] introduced random forests as an interpretable ensemble technique. Vapnik [17] developed the statistical learning foundations of support vector machines. This paper connects these machine learning methods with algebraic structure selection for post-quantum cryptography.

3 Mathematical Preliminaries

3.1 Finite Groups and Conjugacy Classes

Let G be a finite group. For $g \in G$, the conjugacy class of g is

$$\mathrm{Cl}_G(g) = \{xgx^{-1} : x \in G\}. \tag{1}$$

The centralizer of g is

$$C_G(g) = \{x \in G : xg = gx\}. \tag{2}$$

The class equation gives

$$|\mathrm{Cl}_G(g)| = \frac{|G|}{|C_G(g)|}. \tag{3}$$

Conjugacy class distribution can be used as an indicator of non-commutativity, algebraic complexity, and possible hardness of conjugacy-based search problems.

Definition 3.1. *The conjugacy entropy of a finite group G with conjugacy classes $C_1, \ldots, C_k$ is defined as*

$$H_c(G) = -\sum_{i=1}^{k} p_i \log p_i, \quad p_i = \frac{|C_i|}{|G|}. \tag{4}$$

Proposition 3.1. *If G is abelian, then $H_c(G) = \log |G|$ under the above distribution.*

Proof. If G is abelian, every conjugacy class is a singleton. Hence $k = |G|$ and $p_i = 1/|G|$ for every i. Therefore

$$H_c(G) = -\sum_{i=1}^{|G|} \frac{1}{|G|} \log\left(\frac{1}{|G|}\right) = \log |G|.$$

□

3.2 Solvable Groups and Derived Length

The derived series of a group G is defined by

$$G^{(0)} = G, \quad G^{(i+1)} = [G^{(i)}, G^{(i)}]. \tag{5}$$

A group is solvable if $G^{(m)} = \{e\}$ for some $m \geq 0$. The least such m is the derived length.

Remark 3.1. *Derived length is useful as a structural descriptor. Very small derived length may indicate simpler algebraic behavior, while excessively large or irregular structure may create implementation difficulty. Hence it is treated as a feature rather than an isolated security guarantee.*

3.3 Character-Theoretic Indicators

Let $\mathrm{Irr}(G)$ denote the irreducible complex characters of G. The degrees $\chi(1)$ for $\chi \in \mathrm{Irr}(G)$ provide information on representation-theoretic complexity. In this work, the variance of character degrees is treated as a numerical descriptor:

$$\mathrm{Var}_\chi(G) = \frac{1}{r}\sum_{i=1}^{r}\left(\chi_i(1) - \overline{\chi(1)}\right)^2. \tag{6}$$

3.4 Elliptic Curves

An elliptic curve over a field K may be written in short Weierstrass form as

$$E : y^2 = x^3 + ax + b, \tag{7}$$

where

$$4a^3 + 27b^2 \neq 0. \tag{8}$$

The set of points on E, together with the point at infinity, forms an abelian group. Although classical elliptic-curve discrete logarithm systems are not quantum-secure, elliptic-curve descriptors remain useful in comparative algebraic analysis and hybrid cryptographic modelling.

3.5 Lattices

Let $B = (b_1, \ldots, b_n)$ be a basis in $\mathbb{R}^m$. The lattice generated by B is

$$L(B) = \left\{ \sum_{i=1}^{n} z_i b_i : z_i \in \mathbb{Z} \right\}. \tag{9}$$

The lattice dimension and associated noise parameters are major indicators in lattice-based cryptographic schemes. In this paper, lattice dimension is one of the primary features in the learning model.

4 Proposed Framework

Let $\mathcal{A} = \{A_1, A_2, \ldots, A_N\}$ be a collection of candidate algebraic structures. For each A_i, define a feature extraction map

$$\Phi : \mathcal{A} \to \mathbb{R}^d, \tag{10}$$

where $d = 10$ in this study. The feature vector is

$$X_i = \Phi(A_i) = (x_{i1}, x_{i2}, \ldots, x_{i10}). \tag{11}$$

The target variable is

$$y_i = \begin{cases} 1, & \text{if } A_i \text{ is suitable for post-quantum cryptographic use,} \\ 0, & \text{otherwise.} \end{cases} \tag{12}$$

4.1 Feature Set

The ten features used in the simulation are shown in Table 1.

Table 1: Algebraic and computational features used in MLASSF

Feature	Cryptographic interpretation
$\log_2(\|G\|)$	Logarithmic order of the group or algebraic search space.
Number of conjugacy classes	Measures class distribution and non-commutative structure.
Average conjugacy class size	Indicates orbit spread under inner automorphisms.
Derived length	Structural indicator for solvability and group complexity.
Character-degree variance	Representation-theoretic complexity indicator.
Elliptic-curve rank	Descriptor for elliptic-curve component in hybrid structures.
Lattice dimension	Main post-quantum hardness-related parameter.
Noise rate	Lattice-style error/noise indicator.
Claimed security category	Simulated security level corresponding to categories 1, 2, 3, or 5.
Implementation efficiency	Normalized estimate of computational practicality.

4.2 Suitability Score

A latent suitability score was used to generate simulated labels:

$$\begin{aligned} S(A_i) = {} & 0.46L_i + 0.18C_i + 0.12R_i + 0.07Q_i + 0.06E_i \\ & - 0.14D_i - 0.12P_i - 1.70N_i + \epsilon_i, \end{aligned} \tag{13}$$

where L_i is lattice strength, C_i is conjugacy entropy ratio, R_i is character complexity, Q_i is normalized security category, E_i is implementation efficiency, D_i is solvability penalty, P_i is elliptic-curve legacy penalty, N_i is high-noise penalty, and ϵ_i is Gaussian noise.

Definition 4.1. *A candidate structure A_i is classified as suitable if*

$$S(A_i) > \tau, \tag{14}$$

where τ is a threshold representing minimum post-quantum suitability.

4.3 Algorithm

Algorithm 1 Machine Learning Assisted Algebraic Structure Selection Framework

Generate candidate algebraic structures $A_1, \ldots, A_N$
for $i = 1$ to N **do**
    Extract feature vector $X_i = \Phi(A_i)$
    Compute or assign suitability label y_i
end for
Split dataset into training and testing sets
Train supervised classifiers on training set
Evaluate classifiers using accuracy, precision, recall, F1-score, and ROC-AUC
Select best model M^*
Use M^* to rank future algebraic structures

5 Machine Learning Models

5.1 Logistic Regression

Logistic regression estimates

$$P(y = 1|X) = \frac{1}{1 + e^{-(w^T X + b)}}. \tag{15}$$

It provides a linear baseline for algebraic structure classification.

5.2 Support Vector Machine

The support vector machine seeks a separating hyperplane

$$f(X) = w^T X + b, \tag{16}$$

subject to maximum margin constraints. In the experiment, a radial basis function kernel was used.

5.3 Random Forest

A random forest is an ensemble of decision trees. Its prediction is

$$\hat{y} = \text{majority}\{h_1(X), h_2(X), \ldots, h_T(X)\}. \tag{17}$$

Random forest is useful because it provides feature importance measures.

5.4 Gradient Boosting

Gradient boosting builds an additive model

$$F_M(X) = \sum_{m=1}^{M} \gamma_m h_m(X), \tag{18}$$

where each weak learner corrects errors from previous learners.

5.5 Artificial Neural Network

A feed-forward neural network computes

$$H = \sigma(WX + b), \quad \hat{y} = \sigma(VH + c). \tag{19}$$

The neural network used in this experiment has two hidden layers with 64 and 32 neurons respectively.

6 Experimental Design

6.1 Dataset Generation

A simulated dataset of 5,000 algebraic structures was generated. The simulation used random variables constrained by algebraic interpretation. For example, lattice dimension was selected from values resembling practical cryptographic dimensions, while group-order features were represented logarithmically.

The dataset was divided as follows:

$$\text{Training set} = 70\% = 3500, \tag{20}$$

$$\text{Testing set} = 30\% = 1500. \tag{21}$$

A fixed random seed was used to ensure reproducibility.

6.2 Evaluation Metrics

The following metrics were used:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \tag{22}$$

$$\text{Precision} = \frac{TP}{TP + FP}, \tag{23}$$

$$\text{Recall} = \frac{TP}{TP + FN}, \tag{24}$$

$$F_1 = \frac{2(\text{Precision})(\text{Recall})}{\text{Precision} + \text{Recall}}. \tag{25}$$

7 Results

7.1 Model Performance

Table 2 presents the classification results.

Table 2: Model performance on the simulated test set

Model	Accuracy	Precision	Recall	F1-score	ROC-AUC
Gradient Boosting	0.9380	0.9504	0.9621	0.9562	0.9529
Random Forest	0.9360	0.9494	0.9602	0.9548	0.9472
Artificial Neural Network	0.9320	0.9360	0.9697	0.9525	0.9527
Logistic Regression	0.9307	0.9516	0.9498	0.9507	0.9502
Support Vector Machine	0.9253	0.9370	0.9583	0.9475	0.9525

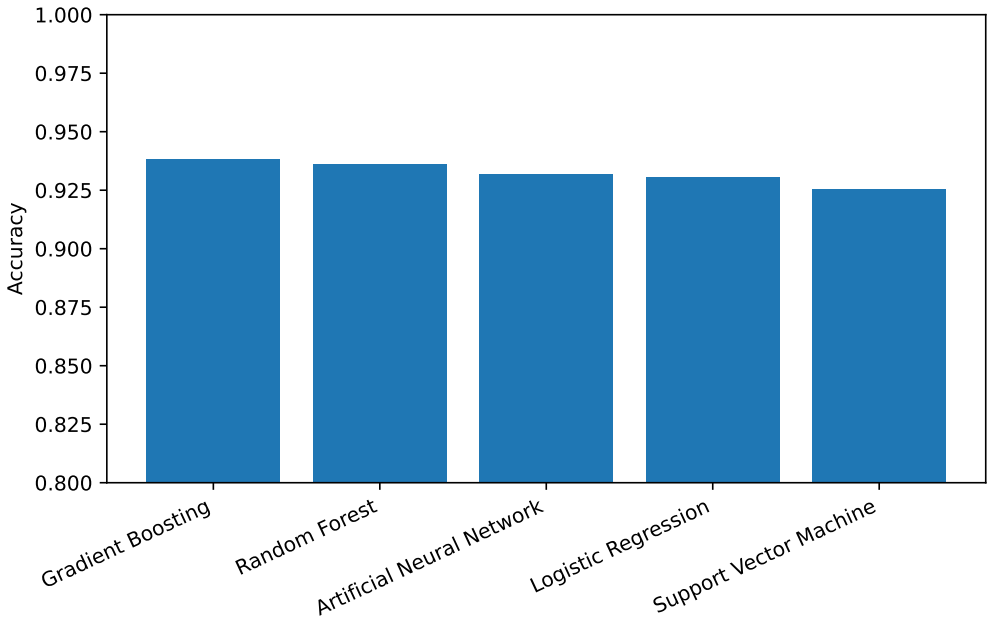

Figure 1: Accuracy comparison of the five classifiers.

Gradient boosting achieved the highest accuracy. Random forest performed nearly as well and provided interpretable feature importance. The neural network achieved high recall, indicating that it was effective in identifying suitable structures, although it produced slightly lower precision than gradient boosting and random forest.

7.2 Confusion Matrix

The confusion matrix for random forest is shown in Figure 2. Random forest correctly classified most suitable and unsuitable structures, with errors mainly occurring near the decision boundary of the latent suitability score.

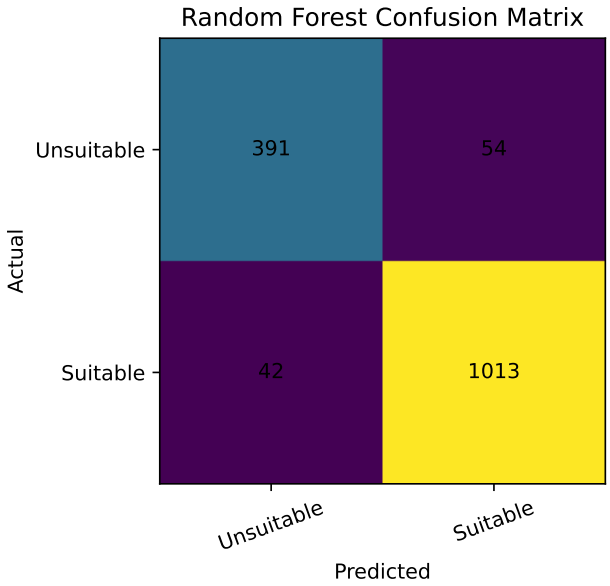

Figure 2: Confusion matrix for the random forest model.

7.3 Feature Importance

Table 3 and Figure 3 show random forest feature importance.

Table 3: Random forest feature importance

Feature	Importance
Lattice dimension	0.5172
Implementation efficiency	0.2591
Character-degree variance	0.0410
Elliptic-curve rank	0.0401
Derived length	0.0286
Noise rate	0.0268
Average conjugacy class log-size	0.0268
Number of conjugacy classes	0.0256
$\log_2(\lvert G\rvert)$	0.0253
Claimed security category	0.0096

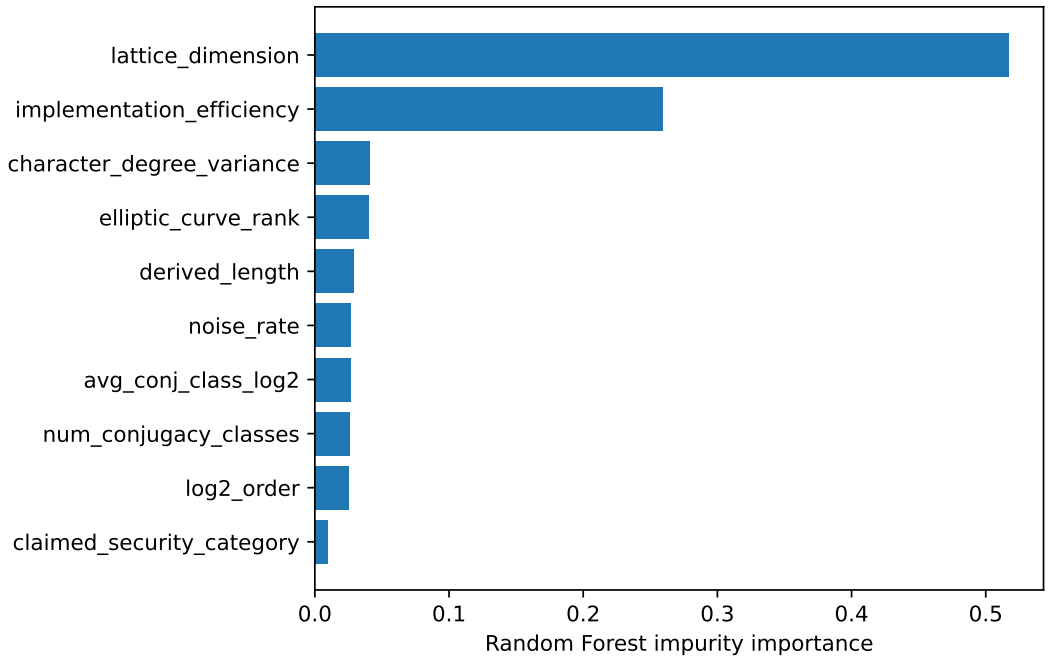

Figure 3: Feature importance for the random forest classifier.

The result suggests that lattice dimension is the strongest predictor of suitability in the simulated environment. This is expected because lattice dimension is a central security parameter in lattice-based post-quantum cryptography. Implementation efficiency is also important because practical cryptographic structures must balance security and deployability. Character-degree variance and elliptic-curve rank have moderate influence, while conjugacy class features contribute smaller but still meaningful information.

7.4 Distributional Analysis

Figure 4 shows the distribution of lattice dimensions across suitable and unsuitable structures.

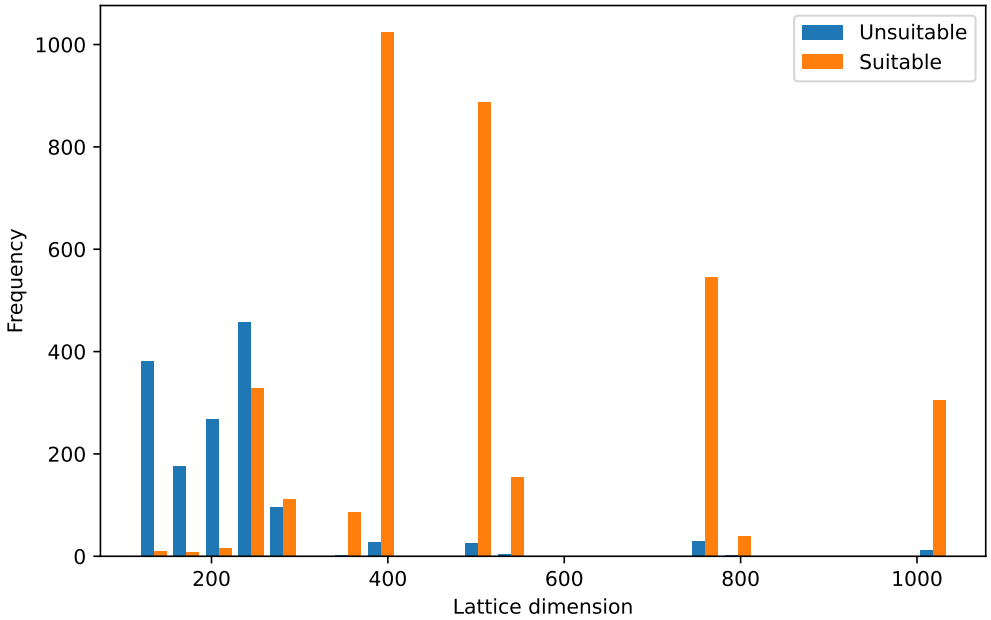

Figure 4: Distribution of lattice dimension for suitable and unsuitable structures.

The plot indicates that structures with larger lattice dimensions were more frequently classified as suitable. However, dimension alone is insufficient; the final classification also depends on implementation efficiency, noise rate, group-theoretic descriptors, and representation-theoretic complexity.

8 Mathematical Analysis of the Framework

Theorem 8.1. *Let $\Phi : \mathcal{A} \to \mathbb{R}^d$ be a bounded feature extraction map and let M be a classifier trained on $\Phi(\mathcal{A})$. If two algebraic structures A and B have identical feature vectors, then the framework assigns them the same predicted suitability label.*

Proof. Since $\Phi(A) = \Phi(B)$, the classifier receives identical input vectors for A and B. A deterministic trained classifier M maps identical inputs to identical outputs. Hence $M(\Phi(A)) = M(\Phi(B))$. □

Theorem 8.2. *The MLASSF framework is invariant under algebraic isomorphisms whenever the feature extraction map Φ is isomorphism-invariant.*

Proof. Assume $A \cong B$ and Φ is isomorphism-invariant. Then $\Phi(A) = \Phi(B)$. By the previous theorem, the trained classifier assigns the same predicted label to both structures. Therefore, the framework is invariant under isomorphism. □

Proposition 8.1. *Let G be a finite group. If the features used for G consist only of group invariants, then every group isomorphic to G receives the same feature representation.*

Proof. Group invariants are preserved under isomorphism. Hence order, derived length, conjugacy class counts, and character-degree statistics remain unchanged under isomorphism. Therefore, the feature representation is the same. □

Remark 8.1. *The theorems above clarify a necessary mathematical condition for fair algebraic classification: feature extraction must preserve relevant algebraic equivalence. A poorly designed feature map may cause isomorphic structures to be treated differently, which would be mathematically undesirable.*

9 Discussion

The results demonstrate that supervised learning can provide a useful preliminary screening tool for algebraic structures. Gradient boosting performed best overall, while random forest provided an interpretable ranking of features. The high performance of logistic regression indicates that the simulated labels contain a strong approximately linear component, while the improved performance of ensemble methods suggests the presence of nonlinear interactions among algebraic features.

The findings are consistent with current post-quantum cryptographic practice, where lattice parameters dominate many security decisions. However, the presence of group-theoretic and character-theoretic features in the feature-importance ranking suggests that algebraic descriptors outside lattice theory may still support useful classification. This is important because post-quantum cryptography is not only a question of standard algorithms but also a broader research field involving algebraic hardness, implementation design, and hybrid constructions.

The framework also offers a way to connect existing algebraic studies with artificial intelligence. For example, conjugacy classes in finitely generated groups can be translated into class-count and entropy features [1]. Solvable group properties and monomial character behavior can be represented by derived length and character-degree variance [2]. Elliptic-curve group descriptors and bilinear constructions can be represented by rank and structural pairing features [3, 5]. Number-theoretic systems such as RSA can be used as comparison baselines for classical vulnerability [4].

10 Reproducibility

The experiment was implemented using Python, NumPy, pandas, scikit-learn, and Matplotlib. The following pseudocode summarizes the computational process:

Listing 1: Simplified reproducibility pseudocode

```
set_random_seed(20260611)
X = generate_5000_candidate_algebraic_structures()
y = compute_simulated_suitability_labels(X)
X_train, X_test, y_train, y_test = train_test_split(X, y, test_size
   =0.30)
```

```
models = [LogisticRegression(), SVM(), RandomForest(),
   GradientBoosting(), MLP()]
for model in models:
   model.fit(X_train, y_train)
   y_pred = model.predict(X_test)
   compute_accuracy_precision_recall_f1_auc(y_test, y_pred)
compute_random_forest_feature_importance()
plot_results()
```

The simulated dataset, summary results, and feature-importance table were generated from the same fixed seed. The purpose of the simulation is methodological demonstration, not direct cryptographic certification.

11 Conclusion

This paper proposed a Machine Learning Assisted Algebraic Structure Selection Framework for post-quantum cryptographic systems. The framework transforms algebraic structures into feature vectors and uses supervised learning to classify them according to simulated post-quantum suitability. The experimental results show that gradient boosting achieved the best performance, while random forest provided useful interpretability through feature-importance analysis. The strongest predictors were lattice dimension and implementation efficiency, followed by character-degree variance, elliptic-curve rank, derived length, noise rate, and conjugacy class indicators.

The study contributes a computational bridge between algebraic cryptography and artificial intelligence. It shows that machine learning can assist researchers in prioritizing algebraic structures for deeper analysis. However, machine learning predictions should not be interpreted as security proofs. Future work should apply the framework to real algebraic databases, integrate computational algebra systems such as GAP and SageMath, include side-channel and implementation features, and test the approach against known post-quantum benchmark schemes.

Acknowledgement

The author acknowledges the relevance of algebraic cryptography, computational group theory, and post-quantum cryptographic research communities whose work motivates the development of machine learning assisted mathematical security frameworks.

Conflict of Interest

The author declares no conflict of interest.

Funding Statement

This research received no external funding.

References

[1] John, M. N., Musa, A., and Udoaka, O. G. (2024). Conjugacy Classes in Finitely Generated Groups with Small Cancellation Properties. *European Journal of Statistics and Probability*, 12(1), 1–9. https://doi.org/10.37745/ejsp.2013/vol12n119

[2] John, M. N., Udoaka, O. G., and Musa, A. (2023). Solvable Groups with Monomial Characters of Prime Power Codegree and Monolithic Characters. *Bulletin of Mathematics and Statistics Research*, 11(7), 1–4. DOI: 10.33329/bomsr.11.4.98

[3] John, M. N., Udoaka, O. G., and Nwala, B. O. (2023). Elliptic-Curve Groups in Quantum-Era Cryptography. *ISAR Journal of Science and Technology*, 1(1), 21–24. https://doi.org/10.5281/zenodo.10207536

[4] John, M. N., Ozioma, O., Obukohwo, V., and Egbogho, H. E. (2023). Number Theory in RSA Encryption Systems. *GPH - International Journal of Mathematics*, 6(11), 7–16. https://doi.org/10.5281/zenodo.10207361

[5] John, M. N., Udoaka, O. G., and Musa, A. (2023). Symmetric Bilinear Cryptography on Elliptic Curve and Lie Algebra. *GPH - International Journal of Mathematics*, 6(10), 1–15. https://doi.org/10.5281/zenodo.10200179

[6] John, M. N., and Udoaka, O. G. (2023). Computational Group Theory and Quantum-Era Cryptography. *International Journal of Scientific Research in Science, Engineering and Technology*, 10(6), 01–10. https://doi.org/10.32628/IJSRSET2310556

[7] Shor, P. W. (1994). Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*, 124–134.

[8] Regev, O. (2005). On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, 84–93.

[9] Micciancio, D., and Regev, O. (2009). Lattice-based cryptography. In D. J. Bernstein, J. Buchmann, and E. Dahmen (Eds.), *Post-Quantum Cryptography*. Springer, 147–191.

[10] National Institute of Standards and Technology. (2024). *Post-Quantum Cryptography FIPS Approved: FIPS 203, FIPS 204, and FIPS 205*. NIST Computer Security Resource Center.

[11] Alagic, G., Apon, D., Cooper, D., Dang, Q., Dang, T., Kelsey, J., Lichtinger, J., Liu, Y. K., Miller, C., Moody, D., Peralta, R., Perlner, R., Robinson, A., Smith-Tone, D., and Alperin-Sheriff, J. (2022). *Status Report on the Third Round of the NIST Post-Quantum Cryptography Standardization Process*. NIST Internal Report 8413.

[12] Hoffstein, J., Pipher, J., and Silverman, J. H. (2008). *An Introduction to Mathematical Cryptography*. Springer.

[13] Koblitz, N. (1987). Elliptic curve cryptosystems. *Mathematics of Computation*, 48(177), 203–209.

[14] Miller, V. S. (1986). Use of elliptic curves in cryptography. In *Advances in Cryptology - CRYPTO '85*, 417–426.

[15] Rotman, J. J. (2012). *An Introduction to the Theory of Groups*. Springer.

[16] Serre, J. P. (1977). *Linear Representations of Finite Groups*. Springer.

[17] Vapnik, V. N. (1995). *The Nature of Statistical Learning Theory*. Springer.

[18] Breiman, L. (2001). Random forests. *Machine Learning*, 45, 5–32.

[19] Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232.

[20] Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.

[21] Menezes, A. J., van Oorschot, P. C., and Vanstone, S. A. (1996). *Handbook of Applied Cryptography*. CRC Press.

