Spectrum of Engineering Sciences
ISSN (e) 3007-3138 (p) 3007-312X
Volume 4, Issue 1, 2026

AN ADVANCED AI-EMPOWERED FINTECH FRAMEWORK FOR CREDIT CARD FRAUD DETECTION IN ONLINE TRANSACTIONS VIA SPARROW SEARCH ALGORITHM–BASED ADAPTIVE OPTIMIZATION

Farah Arzu[*1], **Lubna Gul**[2], **Abdul Waheed**[3], **Muhammad Umar Amin**[4], **Dr. Shah E Yar Qadeem**[5]

Dr. Farooq Alam[6], **Maroof Ashraf**[7], **Asad Ali**[8]

[*1]*Tun Razaq Graduate School of Business, Universiti Tun Abdul Razak, Kuala Lumpur, Malaysia*
[2]*Department of Computer Software Engineering, University of Engineering and Technology, Mardan, Pakistan*
[3]*Department of Computer Science, New York University, United States of America*
[4]*Department of Mathematics, Lamar University, United States of America*
[5]*Assistant Professor. Department of Management Sciences, Qurtuba University of Science and Information Technology, Peshawar, Pakistan.*
[6,7]*Department of Computer Science, Mohammad Ali Jinnah University, Karachi, Pakistan*
[8]*Department of Information Technology, The Islamia University of Bahawalpur. Pakistan.*

[*1]arzu.farah@ur.unirazak.edu.my, [2]lubna.gul@uetmardan.edu.pk, [3]aw4782@nyu.edu, [4]umar.math14@gmail.com, [5]shaheyar605@gmail.com, [6]farooq.alam@jinnah.edu, [7]maroofashraf@axalium.com, [8]alirjpot8@gmail.com

DOI: https://doi.org/10.5281/zenodo.18346936

Keywords
Credit Card Fraud Detection; Financial Technology; Artificial Intelligence; Cost-Sensitive Learning; Sparrow Search Algorithm; Adaptive Optimization; Imbalanced Data; Online Transactions; Metaheuristic Optimization; Machine Learning.

Article History
Received: 28 November 2025
Accepted: 08 January 2026
Published: 23 January 2026

Corresponding Author: *
Farah Arzu

Abstract

The rapid growth of online financial transactions has significantly increased the exposure of digital payment systems to sophisticated credit card fraud, posing serious challenges to financial institutions and consumers alike. The highly imbalanced nature of transactional datasets, evolving fraud patterns, and stringent real-time decision requirements often limit the effectiveness and stability of conventional machine learning and deep learning approaches. To address these challenges, this study proposes an advanced AI-empowered FinTech framework for credit card fraud detection in online transactions that integrates adaptive optimization based on the Sparrow Search Algorithm (SSA) to enhance detection performance, robustness, and operational feasibility. The proposed framework follows a structured end-to-end pipeline comprising secure data ingestion, preprocessing and feature normalization, imbalance-aware learning, SSA-driven hyperparameter and decision-threshold optimization, and comprehensive performance evaluation under cost-sensitive constraints. At the core of the framework, SSA is employed as a global metaheuristic optimizer to adaptively tune critical model parameters across multiple candidate learners, including interpretable classifiers, ensemble methods, and lightweight neural architectures. Unlike conventional grid or random search techniques, the SSA-based optimization dynamically balances global exploration and local exploitation, enabling efficient navigation of complex hyperparameter spaces while maintaining computational efficiency. The optimization objective explicitly incorporates cost-sensitive learning by penalizing false negatives more heavily than false positives, reflecting the asymmetric financial risk associated with undetected *fraudulent transactions. This design ensures that the optimized models align with real-world fraud management priorities rather than purely accuracy-driven objectives. The framework further integrates robust evaluation protocols using stratified cross-validation and hold-out testing to assess generalization capability and stability. Performance is measured using fraud-specific metrics, including precision–recall area under the curve, recall, F1-score, Matthews correlation coefficient, balanced accuracy, and expected financial cost. Comparative analyses against non-optimized baselines demonstrate that SSA-optimized models achieve superior minority-class detection, reduced performance variance, and improved cost efficiency across varying fraud prevalence levels. In addition, the proposed architecture considers practical deployment aspects such as inference latency, model monitoring, and periodic re-optimization to address concept drift in evolving transaction streams. Overall, this study contributes a cost-aware, optimization-driven, and deployment-ready AI framework that strengthens the resilience of online payment systems. The proposed SSA-based adaptive optimization strategy offers a scalable and effective solution for enhancing credit card fraud detection performance in modern FinTech environments.*

1- Introduction

The rapid digital transformation of the financial sector has reshaped how monetary transactions are conducted, processed, and authorized across the globe. The widespread adoption of e-commerce platforms, mobile payment applications, digital wallets, and contactless payment technologies has led to an unprecedented increase in the volume, velocity, and diversity of online credit card transactions. While these innovations have significantly enhanced user convenience and financial inclusion, they have simultaneously expanded the attack surface of digital payment ecosystems. Credit card fraud has evolved into a highly sophisticated and adaptive cyber-financial threat, leveraging identity theft, transaction laundering, bot-driven attacks, and cross-channel exploitation to bypass conventional security mechanisms. As a result, fraud-related financial losses, operational overhead, and reputational risks continue to rise, placing sustained pressure on financial institutions to deploy intelligent, reliable, and scalable fraud detection systems [1]. Historically, fraud detection mechanisms have relied heavily on rule-based systems and expert-defined heuristics, which encode predefined transaction patterns, spending limits, and anomaly thresholds. Although such systems are computationally efficient and transparent, their static nature makes them ill-suited for modern fraud scenarios characterized by rapid behavioral shifts and adversarial adaptation. Fraudsters continuously modify transaction sequences, merchant categories, and temporal behaviors to evade detection rules, rendering static systems increasingly ineffective. Consequently, financial institutions have progressively transitioned toward data-driven approaches based on machine learning and deep learning, which can model complex, nonlinear relationships in large-scale transactional data and automatically learn discriminative patterns between legitimate and fraudulent behavior [2]. Despite their demonstrated advantages, existing AI-based credit card fraud detection models face several fundamental challenges that limit their robustness and real-world applicability. One of the most critical issues is the extreme class imbalance inherent in transactional datasets, where fraudulent transactions typically constitute far less than 1% of total records. This imbalance biases learning algorithms toward the majority class, often resulting in deceptively high accuracy but poor fraud recall. Moreover, the dynamic and non-stationary nature of fraud behavior introduces concept drift, causing model performance to degrade over time if continuous adaptation is not incorporated. These challenges are further compounded by stringent real-time processing requirements, where detection decisions must be made within milliseconds to avoid disrupting legitimate customer experiences. Another major limitation in current fraud detection research lies in the optimization and calibration of AI models. The performance of machine learning and deep learning algorithms is highly sensitive to hyperparameter configurations, feature scaling strategies, and classification thresholds, particularly under imbalanced and cost-sensitive conditions. Conventional tuning methods, such as grid search and random search, are often computationally expensive, inflexible, and inefficient when applied to high-dimensional or non-convex optimization spaces. More importantly, many prior studies optimize models primarily based on accuracy or area-under-curve metrics, without explicitly accounting for the asymmetric financial consequences of misclassification errors [3]. In practice, false negatives where fraudulent transactions go undetected incur substantially higher costs than false positives, which merely inconvenience customers. Ignoring this asymmetry leads to models that are statistically strong but operationally suboptimal. Recent advances in metaheuristic optimization have opened new opportunities for addressing these challenges by enabling adaptive, global optimization of model parameters under complex objective functions. Metaheuristic algorithms are particularly attractive for fraud detection because they do not rely on gradient information, can escape local optima, and are well suited for multi-objective and cost-aware optimization problems. Among these techniques, the Sparrow Search Algorithm (SSA) has emerged as a powerful and efficient nature-inspired optimizer, offering a dynamic balance between exploration and exploitation through its producer–scrounger–watcher search mechanism. SSA has demonstrated fast convergence, strong global search capability, and scalability across diverse optimization tasks, yet its integration into credit card fraud detection frameworks remains relatively underexplored [4]. Motivated by these observations, this study proposes an advanced AI-empowered FinTech framework for credit card fraud detection in online transactions that leverages SSA-based adaptive optimization to enhance detection performance, robustness, and deployment readiness. Rather than treating fraud detection as a standalone classification problem, the proposed framework formulates it as an end-to-end, optimization-driven decision pipeline. The framework integrates secure data ingestion, preprocessing and feature normalization, imbalance-aware learning, and SSA-driven optimization of both hyperparameters and decision thresholds under cost-sensitive objectives. By explicitly penalizing false negatives more heavily than false positives, the optimization process aligns model behavior with real-world financial risk management priorities. The proposed framework is designed to be model-agnostic, enabling adaptive optimization across a diverse set of candidate learners, including interpretable machine learning models, ensemble classifiers, and lightweight neural architectures suitable for real-time deployment. In addition, practical operational considerations such as inference latency, performance stability, continuous monitoring, and periodic re-optimization to mitigate concept drift are incorporated into the system design. A summary of the key challenges encountered in credit card fraud detection and the corresponding solutions provided by the proposed SSA-based framework is presented in Table 1, which highlights how the framework systematically addresses both technical and operational limitations of existing approaches.

Table 1: Key Challenges in Credit Card Fraud Detection and Corresponding Solutions in the Proposed Framework

Challenge	Impact on Fraud Detection Systems	Solution in Proposed SSA-Based Framework
Extreme class imbalance	Bias toward legitimate transactions and poor fraud recall	Imbalance-aware learning and cost-sensitive optimization
Evolving fraud patterns (concept drift)	Performance degradation over time	Periodic SSA-driven re-optimization and model updating

Suboptimal hyperparameter tuning	Unstable and inconsistent detection performance	Adaptive global optimization using Sparrow Search Algorithm
Asymmetric misclassification costs	High financial loss due to undetected fraud	Explicit penalization of false negatives in fitness function
High-dimensional optimization space	Inefficient grid/random search	SSA-based exploration–exploitation balance
Real-time decision constraints	Deployment infeasibility	Lightweight models and latency-aware optimization

By integrating adaptive metaheuristic optimization, cost-sensitive learning, and deployment-oriented design into a unified framework, this study advances the state of the art in AI-driven credit card fraud detection. Unlike prior works that focus primarily on isolated model improvements, the proposed approach emphasizes system-level robustness, financial risk alignment, and long-term operational sustainability in dynamic transaction environments. The framework is particularly well suited for modern FinTech platforms, where detection accuracy, cost efficiency, and real-time responsiveness must be jointly optimized under continuously evolving fraud behaviors. Through this contribution, the study establishes a scalable foundation for intelligent fraud detection systems capable of adapting to future challenges in digital payment security.

2- Credit Card Fraud Detection in Online Financial Transactions:

Credit card fraud detection has long been recognized as a critical research domain within financial security due to the substantial economic losses, operational disruption, and erosion of consumer trust it imposes on digital payment ecosystems. With the rapid proliferation of e-commerce platforms, mobile banking applications, and cross-border online transactions, fraud detection systems are required to operate under high data velocity, strict latency constraints, and continuously evolving adversarial behaviors. Fraudulent activities increasingly exploit vulnerabilities in authentication mechanisms, transaction authorization processes, and user behavior modeling, making fraud detection a complex and adversarial learning problem rather than a static classification task. Early-generation fraud detection systems were predominantly rule-based, relying on manually defined thresholds, expert-crafted heuristics, blacklists, and transaction pattern matching derived from historical fraud cases [5]. These systems were attractive in early deployment stages due to their transparency, interpretability, and low computational overhead. However, their static nature severely limited adaptability and scalability. Fraudsters quickly learned to evade fixed rules by modifying transaction amounts, timing, merchant categories, and geographical attributes, leading to rapid degradation in detection effectiveness. As transaction volumes increased, rule-based systems also suffered from excessive false-positive rates, generating unnecessary transaction declines and customer dissatisfaction, while failing to respond promptly to novel and previously unseen attack strategies. The increasing availability of large-scale transactional datasets marked a significant shift toward data-driven fraud detection approaches. Statistical learning methods, such as logistic regression, naïve Bayes, and Bayesian classifiers, were among the first automated techniques adopted to replace rigid rule-based engines. These models introduced probabilistic decision-making and improved generalization across varying transaction profiles [6]. Nevertheless, their reliance on linear decision boundaries and simplified feature relationships limited their ability to capture the complex, nonlinear, and high-dimensional patterns characteristic of modern fraud behaviors. As a result, statistical models often struggled to balance fraud recall and false-positive control in highly imbalanced datasets. To overcome these limitations, research attention progressively shifted toward machine learning–based fraud detection models capable of modeling heterogeneous features and nonlinear decision boundaries. Decision trees, support vector machines, k-nearest neighbors, and ensemble learning techniques enabled more expressive modeling of transaction behavior and interactions among monetary, temporal, and behavioral attributes. Ensemble approaches, in particular, demonstrated improved robustness by aggregating multiple learners, reducing variance, and capturing diverse fraud patterns. However, these models introduced new challenges related to hyperparameter sensitivity, model calibration, and computational complexity, especially under real-time processing constraints. More recently, deep learning–based approaches have been explored to further enhance fraud detection performance by automatically learning hierarchical representations from transaction data. Neural architectures have shown promise in capturing subtle fraud signals and temporal dependencies. Despite these advances, deep learning models often exhibit limited interpretability, high training and inference costs, and instability when trained on severely imbalanced datasets [7]. Moreover, many proposed approaches focus primarily on improving classification accuracy, while neglecting deployment-oriented considerations such as latency, cost-sensitive decision-making, and adaptability to evolving fraud strategies. The evolution of credit card fraud detection methodologies from rule-based systems to statistical learning, machine learning, and deep learning approaches is summarized in Table 2, which highlights the key characteristics, strengths, and limitations of each paradigm. This progression illustrates a clear trend toward increasingly complex models, accompanied by growing challenges related to optimization, robustness, and real-world deployment.

Table 2: Evolution of Credit Card Fraud Detection Approaches in Online Transactions

Approach Category	Representative Techniques	Key Strengths	Primary Limitations
Rule-based systems	Expert rules, blacklists, thresholds	High interpretability, low latency	Static behavior, poor adaptability, high false positives
Statistical learning	Logistic regression, Bayesian models	Probabilistic reasoning, simple implementation	Linear assumptions, weak nonlinear modeling
Machine learning	Decision trees, SVM, k-NN, ensembles	Nonlinear modeling, improved robustness	Hyperparameter sensitivity, imbalance bias
Deep learning	Neural networks, autoencoders	Hierarchical feature learning	High complexity, limited interpretability, deployment cost

Figure 1 conceptually illustrates the transition from traditional rule-based fraud detection systems to modern AI-driven approaches. As shown in Figure 1, early systems rely on fixed rules and manual intervention, while contemporary fraud detection architectures increasingly emphasize automated feature learning, adaptive optimization, and continuous performance monitoring. This evolution underscores the growing need for frameworks that not only achieve high detection accuracy but also maintain stability, cost efficiency, and adaptability in dynamic online transaction environments.

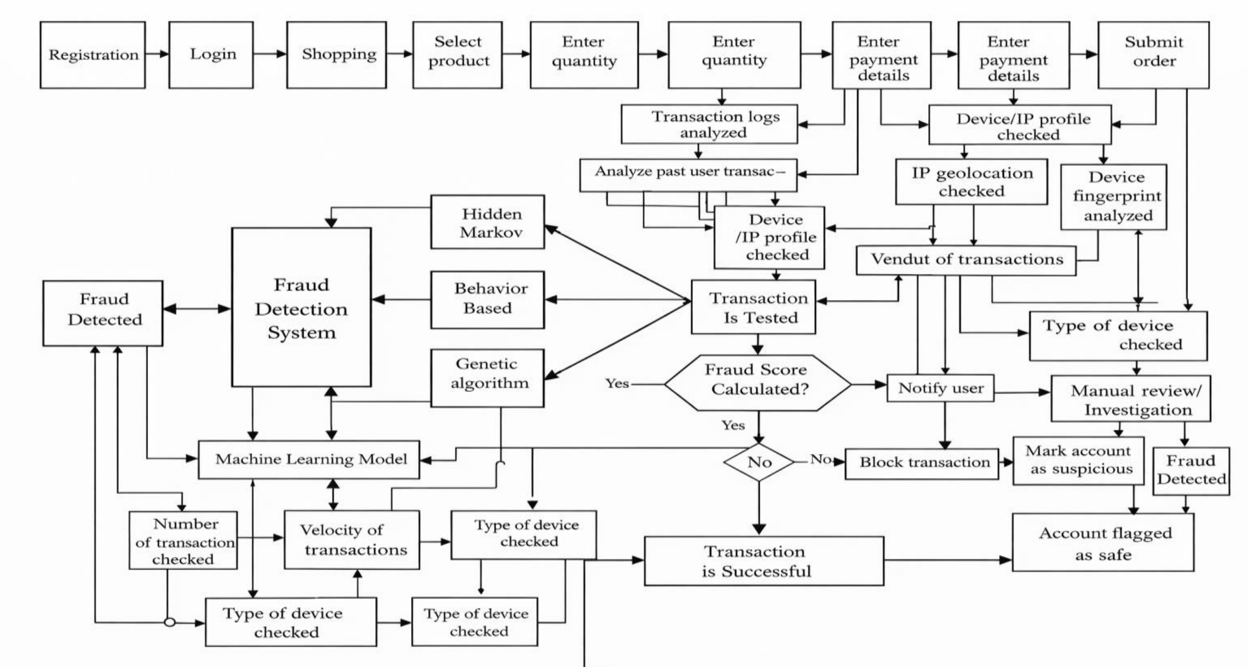

Figure 1: Credit card fraud detection systems from rule-based engines to adaptive AI-driven frameworks.

Overall, the literature reveals that although significant progress has been made in advancing fraud detection techniques, existing approaches often address individual aspects of the problem in isolation. The increasing complexity of fraud behaviors, combined with extreme data imbalance and asymmetric misclassification costs, necessitates integrated frameworks that jointly consider detection performance, optimization strategy, and deployment feasibility. These limitations motivate the development of adaptive, cost-aware, and optimization-driven fraud detection frameworks capable of operating reliably in modern FinTech environments.

3- Hyperparameter Optimization and Metaheuristic Techniques:

The effectiveness of credit card fraud detection models is highly dependent on appropriate hyperparameter configuration, particularly for ensemble learning methods and deep learning architectures. Hyperparameters such as learning rates, tree depth, number of estimators, regularization coefficients, network topology, and decision thresholds play a decisive role in controlling model generalization, stability, and sensitivity to minority-class fraud instances. In highly imbalanced and cost-sensitive environments, even minor deviations from optimal hyperparameter settings can lead to significant degradation in fraud recall or excessive false-positive rates, thereby undermining real-world deployment feasibility. Conventional hyperparameter tuning techniques, including grid search and random search, remain widely used due to their conceptual simplicity and ease of implementation. Grid search systematically explores predefined parameter combinations, while random search samples configurations from specified distributions. Despite their popularity, these approaches suffer from substantial limitations when applied to fraud detection tasks. Grid search becomes computationally prohibitive as dimensionality increases, while random search lacks directional guidance and may waste computational resources exploring suboptimal regions of the search space. Moreover, both approaches are poorly suited for optimizing non-convex, multi-objective fitness landscapes commonly encountered in fraud detection, where performance metrics must balance detection accuracy, recall, false-positive cost, and inference latency [8]. Bayesian optimization techniques have been proposed as more efficient alternatives by constructing surrogate models to approximate the objective function and guide the search process. These methods can reduce the number of evaluations required to identify promising configurations. However, Bayesian optimization relies on assumptions regarding smoothness and stationarity of the objective function, which may not hold in fraud detection scenarios characterized by noisy, discontinuous, and cost-sensitive evaluation functions. Furthermore, surrogate modeling becomes increasingly complex and less reliable as the number of hyperparameters and objectives grows, limiting scalability in large-scale FinTech applications [9]. To overcome these limitations, metaheuristic optimization algorithms have gained increasing attention for hyperparameter tuning in machine learning and deep learning systems. Metaheuristics are population-based, gradient-free optimization methods inspired by natural, biological, or physical phenomena. Their key strengths include strong global search capability, robustness against local optima, and flexibility in handling complex, non-differentiable, and multi-objective optimization problems. These properties make metaheuristic algorithms particularly suitable for fraud detection tasks, where objective functions often integrate multiple performance metrics and cost-sensitive constraints. A wide range of metaheuristic techniques has been explored in the literature, including genetic algorithms, particle swarm optimization, ant colony optimization, differential evolution, and simulated annealing. These methods have demonstrated promising results in tuning classifier parameters, selecting relevant features, and improving convergence speed. In fraud detection applications, metaheuristic optimization has been shown to enhance minority-class detection, reduce model variance, and improve robustness under data imbalance. However, many existing studies apply these techniques in an ad hoc or model-specific manner, focusing on optimizing a single classifier or performance metric without considering broader system-level requirements [10]. The comparative characteristics of conventional hyperparameter tuning methods and metaheuristic optimization approaches are summarized in Table 3, highlighting their relative strengths and limitations in the context of fraud detection.

Table 3: Comparison of Hyperparameter Optimization Techniques in Fraud Detection

Optimization Method	Search Characteristics	Advantages	Limitations in Fraud Detection
Grid search	Exhaustive, deterministic	Simple, reproducible	Computationally expensive, poor scalability
Random search	Stochastic, unguided	Faster than grid search	Inefficient exploration, no convergence guidance
Bayesian optimization	Surrogate-based, guided	Sample-efficient	Assumption-dependent, scalability issues
Metaheuristic optimization	Population-based, global	Gradient-free, robust, flexible	Requires careful fitness design

Despite their potential, most metaheuristic-based fraud detection studies do not fully exploit the adaptability of these

algorithms. Cost sensitivity is often incorporated only implicitly or ignored altogether, resulting in optimization objectives that favor statistical performance rather than financial risk minimization. Additionally, deployment-oriented constraints such as inference latency, model stability, and adaptability to concept drift are rarely integrated into the optimization process [11]. This gap limits the practical impact of metaheuristic optimization in real-world FinTech environments. Figure 2 conceptually illustrates the role of metaheuristic optimization within a modern fraud detection pipeline.

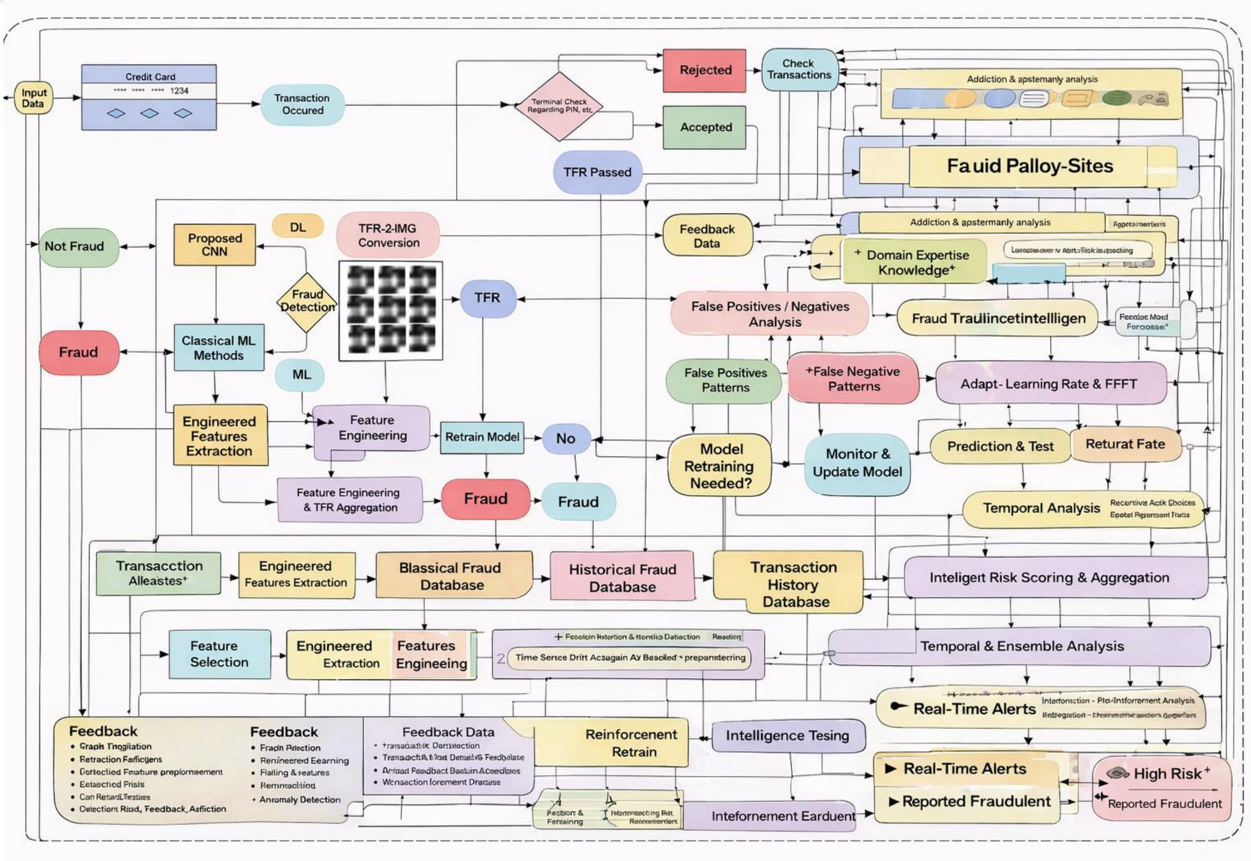

Figure 2: Conceptual role of metaheuristic optimization in AI-based credit card fraud detection systems.

While hyperparameter optimization is widely acknowledged as a critical component of fraud detection model performance, existing tuning strategies remain inadequate for addressing the complexity, imbalance, and cost sensitivity inherent in online transaction environments. Metaheuristic optimization offers a promising alternative, but its full potential can only be realized through integration into a holistic, cost-aware, and deployment-ready framework [12]. These observations motivate the adoption of adaptive metaheuristic strategies such as the Sparrow Search Algorithm to systematically optimize fraud detection models under realistic FinTech constraints, forming the foundation for the proposed approach in this study.

4- Methodology:

This section establishes the methodological foundation of the proposed AI-empowered FinTech framework for credit card fraud detection in online transaction environments. The methodology is explicitly designed to ensure robustness, adaptability, and real-world deployability by formulating fraud detection as an integrated, end-to-end, optimization-driven decision process rather than a conventional standalone classification problem. Unlike traditional approaches that focus primarily on improving predictive accuracy in isolation, the proposed methodology emphasizes system-level performance by jointly optimizing data handling, learning behavior, and decision-making under realistic operational constraints [13]. The framework systematically integrates secure data acquisition and handling mechanisms, comprehensive preprocessing and feature normalization, imbalance-aware learning strategies, and adaptive optimization of both model hyperparameters and classification decision thresholds using the Sparrow Search Algorithm. By embedding SSA within the learning pipeline, the methodology enables dynamic exploration and exploitation of complex parameter spaces, allowing the detection models to adapt effectively to evolving fraud patterns and non-stationary data distributions. Furthermore, cost-sensitive evaluation criteria are explicitly incorporated into the optimization process to reflect the asymmetric financial risks associated with misclassification errors, particularly the substantially higher cost of false negatives in fraud detection scenarios. Each methodological component is carefully structured to address the key challenges inherent in online credit card fraud detection, including extreme class imbalance, concept drift, and stringent real-time processing requirements. At the same time, computational efficiency and inference latency are considered throughout the design to ensure compatibility with high-throughput transaction streams typical of modern FinTech systems [14]. By adopting this holistic and optimization-driven design philosophy, the proposed methodology bridges the gap between high-performing AI models and practical deployment needs in dynamic financial environments. The following subsections provide a detailed description of each stage of the proposed framework, beginning with data acquisition and secure ingestion, and progressing through preprocessing, learning, adaptive optimization, evaluation, and deployment considerations. Figure 3 shows the overall research design and methodology framework.

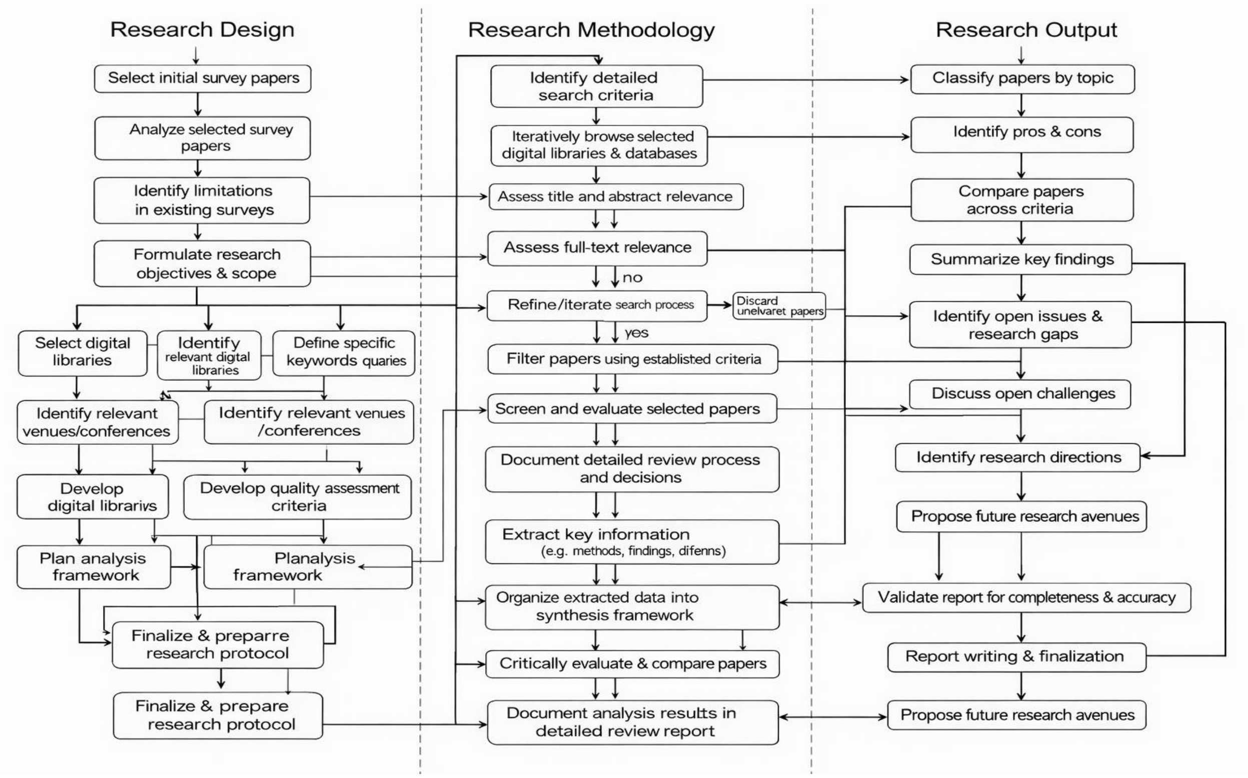
Research Design
Research Methodology
Research Output
Select initial survey papers
Analyze selected survey papers
Identify limitations in existing surveys
Formulate research objectives & scope
Select digital libraries
Identify relevant digital libraries
Define specific keywords quaries
Identify relevant venues/conferences
Identify relevant venues /conferences
Develop digital librarivs
Develop quality assessment criteria
Plan analysis framework
Planalysis framework
Finalize & preparre research protocol
Finalize & prepare research protocol
Identify detailed search criteria
Iteratively browse selected digital libraries & databases
Assess title and abstract relevance
Assess full-text relevance
no
Refine/iterate search process
Discard unelvaret papers
yes
Filter papers using established criteria
Screen and evaluate selected papers
Document detailed review process and decisions
Extract key information (e.g. methods, findings, difenns)
Organize extracted data into synthesis framework
Critically evaluate & compare papers
Document analysis results in detailed review report
Classify papers by topic
Identify pros & cons
Compare papers across criteria
Summarize key findings
Identify open issues & research gaps
Discuss open challenges
Identify research directions
Propose future research avenues
Validate report for completeness & accuracy
Report writing & finalization
Propose future research avenues

Figure 3: Research design of proposed framework

4.2- Data Acquisition and Secure Ingestion:

Accurate and reliable credit card fraud detection critically depends on the quality, integrity, and security of transactional data. In the proposed AI-empowered FinTech framework, transactional data are obtained from online credit card payment systems and structured into standardized records representing individual transaction events. Each record comprises a diverse set of attributes capturing monetary characteristics (e.g., transaction amount and currency), temporal indicators (e.g., transaction time and frequency), merchant-related information (e.g., merchant category and transaction channel), and customer behavioral features (e.g., spending patterns and transaction velocity). This heterogeneous feature composition enables comprehensive modeling of both legitimate and fraudulent transaction behaviors. Given the sensitive nature of financial transaction data, secure data ingestion is treated as a foundational component of the methodology. Data are collected through protected channels using encryption and secure communication protocols to prevent interception or tampering during transmission [15]. To ensure compliance with financial regulations and data protection standards, all personally identifiable information is anonymized or tokenized prior to storage and analysis. Sensitive identifiers such as card numbers, customer IDs, and account details are replaced with non-reversible tokens, ensuring that individual users cannot be re-identified from the dataset. In addition, role-based access control mechanisms are enforced to restrict data access exclusively to authorized processes and analytical modules within the fraud detection system [16]. Following secure ingestion, transactions are validated for completeness and consistency before entering the analytical pipeline. Incomplete or corrupted records are flagged for further inspection or removed to prevent contamination of the learning process. Transactions are then chronologically ordered to preserve temporal dependencies and sequential transaction behavior, which are essential for accurately modeling fraud patterns such as rapid transaction bursts or abnormal spending sequences. This temporal ordering also minimizes information leakage during model development by ensuring that training data precede validation and testing data in time, thereby simulating realistic deployment conditions [17]. To further enhance methodological rigor, the ingested data are partitioned into training, validation, and testing subsets using time-aware splitting strategies. This approach prevents future information from inadvertently influencing model training and enables a more realistic assessment of generalization performance under evolving fraud behaviors. A summary of the key data components and security measures employed during the acquisition and ingestion process is provided in Table 4, highlighting how data integrity and confidentiality are maintained throughout the pipeline.

Table 4: Transactional Data Components and Secure Ingestion Measures

Data Aspect	Description	Security and Integrity Measures
Monetary attributes	Transaction amount, currency, balance-related indicators	Encrypted transmission, validation checks
Temporal indicators	Timestamp, transaction frequency, recency	Chronological ordering, time-aware splitting
Merchant information	Merchant category, channel, location	Tokenization and controlled access
Behavioral features	Spending patterns, velocity metrics	Aggregation on anonymized identifiers
Sensitive identifiers	Card number, customer ID	Anonymization and non-reversible tokenization

The role of secure data acquisition within the overall fraud detection framework is conceptually illustrated in Figure 4. As shown in Figure 4, the secure ingestion layer acts as the entry point of the system, ensuring that only validated, anonymized, and temporally consistent transaction data are forwarded to subsequent preprocessing, learning, and optimization stages. This layered design not only enhances system security but also improves the reliability and reproducibility of downstream analytical results.

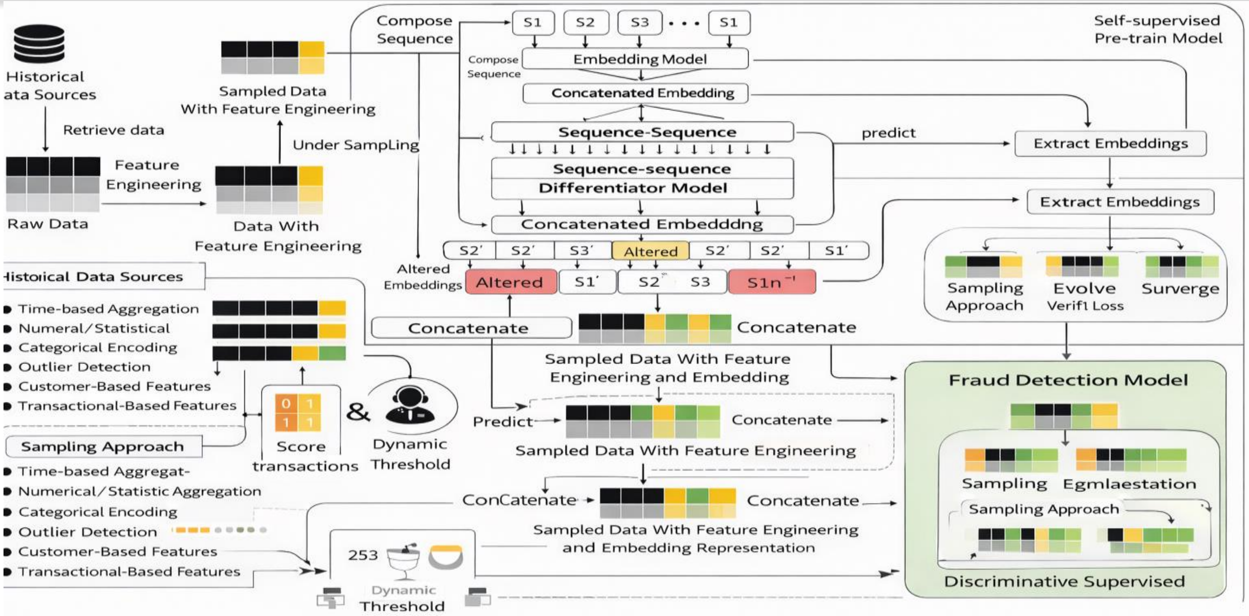

Figure 4: Secure data acquisition and ingestion workflow within the proposed AI-empowered FinTech fraud detection framework.

Overall, the data acquisition and secure ingestion process establishes a robust and trustworthy foundation for the proposed fraud detection framework. By prioritizing data confidentiality, integrity, and temporal consistency, this stage ensures that subsequent preprocessing, imbalance-aware learning, and SSA-based optimization are performed on high-quality data that accurately reflect real-world transaction environments. This design is essential for achieving reliable fraud detection performance and for supporting safe deployment in modern, large-scale FinTech systems.

4.3- Data Preprocessing and Feature Engineering:

Data preprocessing and feature engineering constitute a critical stage in the proposed AI-empowered FinTech framework, as the quality, structure, and expressiveness of transaction features directly influence fraud detection accuracy, robustness, and optimization efficiency. Raw credit card transaction data collected from online payment systems are inherently noisy, heterogeneous, and highly skewed, reflecting diverse user behaviors, merchant characteristics, and transaction contexts. Without careful preprocessing, these issues can lead to unstable learning behavior, biased model decisions, and suboptimal optimization outcomes, particularly under extreme class imbalance and cost-sensitive objectives. The preprocessing pipeline begins with comprehensive data cleansing and integrity verification to ensure that only reliable transaction records are forwarded to subsequent learning stages. Missing values arise frequently in transactional datasets due to optional fields, delayed logging, or system inconsistencies. These missing entries are addressed using feature-specific imputation strategies designed to preserve statistical consistency while minimizing information distortion [18]. Continuous numerical features are imputed using robust central tendency estimators that reduce sensitivity to extreme values, whereas categorical features are completed using frequency-aware strategies that reflect dominant transaction patterns. Records exhibiting irrecoverable inconsistencies or invalid attribute combinations are removed to prevent noise propagation through the learning pipeline. Following data cleansing, feature scaling and normalization are applied to mitigate the effects of heterogeneous feature ranges and heavy-tailed distributions [19]. Monetary attributes such as transaction amount, cumulative spending, and balance-related indicators often exhibit strong skewness and extreme outliers, which can disproportionately influence distance-based classifiers, gradient-driven learners, and metaheuristic fitness evaluation. To address this issue, robust normalization techniques are employed to compress extreme values while preserving relative ordering among transactions [20]. This normalization step ensures numerical stability during model training and enables fair contribution of all features during SSA-based optimization. Categorical and discrete transaction attributes, including merchant category codes, transaction channels, payment modes, and geographic indicators, are transformed into numerical representations suitable for learning algorithms. Given the high cardinality of several categorical fields, encoding strategies are carefully selected to balance information preservation and computational efficiency [21]. Excessive dimensionality expansion is avoided to prevent increased memory usage and inference latency, which are critical considerations for real-time fraud detection systems. Encoding schemes are consistently applied across training, validation, and testing subsets to eliminate information leakage and ensure reproducibility. Beyond standard preprocessing, the proposed framework places strong emphasis on domain-driven feature engineering to enhance fraud discriminability. Rather than relying solely on raw transaction attributes, higher-level behavioral features are constructed to capture temporal dynamics, spending irregularities, and contextual deviations indicative of fraudulent behavior. These engineered features include transaction velocity measures, rolling statistical summaries computed over adaptive time windows, frequency-based indicators, and deviation metrics that quantify departures from a customer's historical spending profile. Such features are particularly effective for detecting fraud patterns that closely mimic legitimate transactions but exhibit subtle temporal or behavioral anomalies [22]. To further improve learning efficiency and reduce redundancy, feature relevance analysis is incorporated into the preprocessing pipeline. Highly correlated or weakly informative features can inflate model complexity, slow convergence during optimization, and degrade generalization performance. Feature relevance assessment enables the retention of compact yet informative feature subsets, facilitating faster SSA convergence and more stable decision boundaries. This refinement is especially important in high-dimensional transactional datasets, where unnecessary features can obscure meaningful fraud signals. The complete preprocessing and feature engineering workflow is conceptually illustrated in Figure 5, which depicts the transformation of raw transactional data into normalized, behaviorally enriched feature representations.

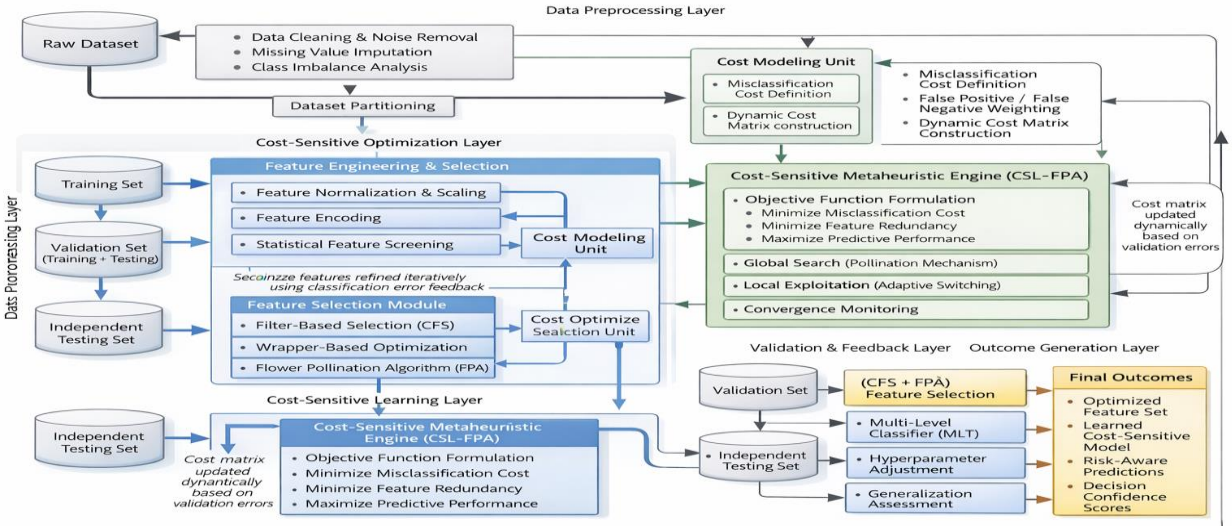
Data Preprocessing Layer
Raw Dataset
Data Cleaning & Noise Removal
Missing Value Imputation
Class Imbalance Analysis
Dataset Partitioning
Cost-Sensitive Optimization Layer
Dats Prororesing Layer
Training Set
Validation Set (Training + Testing)
Independent Testing Set
Feature Engineering & Selection
Feature Normalization & Scaling
Feature Encoding
Statistical Feature Screening
Cost Modeling Unit
Secoinzze features refined iteratively using classification error feedback
Feature Selection Module
Filter-Based Selection (CFS)
Wrapper-Based Optimization
Flower Pollination Algorithm (FPA)
Cost Optimize Sealction Unit
Cost-Sensitive Learning Layer
Independent Testing Set
Cost matrix updated dynantically based on validation errors
Cost-Sensitive Metaheuristic Engine (CSL-FPA)
Objective Function Formulation
Minimize Misclassification Cost
Minimize Feature Redundancy
Maximize Predictive Performance
Cost Modeling Unit
Misclassification Cost Definition
Dynamic Cost Matrix construction
Misclassification Cost Definition
False Positive / False Negative Weighting
Dynamic Cost Matrix Construction
Cost-Sensitive Metaheuristic Engine (CSL-FPA)
Objective Function Formulation
Minimize Misclassification Cost
Minimize Feature Redundancy
Maximize Predictive Performance
Global Search (Pollination Mechanism)
Local Exploitation (Adaptive Switching)
Convergence Monitoring
Cost matrix updated dynamically based on validation errors
Validation & Feedback Layer
Outcome Generation Layer
Validation Set
Independent Testing Set
(CFS + FPÀ) Feature Selection
Multi-Level Classifier (MLT)
Hyperparameter Adjustment
Generalization Assessment
Final Outcomes
Optimized Feature Set
Learned Cost-Sensitive Model
Risk-Aware Predictions
Decision Confidence Scores

Figure 5: Data preprocessing and feature engineering workflow within the proposed AI-empowered FinTech fraud detection framework.

Overall, the data preprocessing and feature engineering stage plays a foundational role in the proposed methodology by transforming raw, heterogeneous transaction records into a structured and discriminative feature space optimized for adaptive learning. By addressing noise, distributional skewness, feature heterogeneity, and behavioral complexity, this stage significantly enhances the effectiveness of imbalance-aware learning and SSA-driven optimization in subsequent stages. Consequently, robust preprocessing and feature engineering directly contribute to improved fraud recall, reduced false alarms, and stable real-time performance, supporting reliable deployment of the proposed framework in dynamic online transaction environments.

4.4- Candidate AI Models for Fraud Detection:

The design of an effective credit card fraud detection system requires careful consideration of model diversity, representational capacity, interpretability, and operational feasibility. Fraudulent behavior in online transactions is highly heterogeneous, ranging from simple rule violations to complex, temporally coordinated attacks that closely mimic legitimate customer behavior. Consequently, reliance on a single classification paradigm may lead to suboptimal performance under varying transaction contexts and fraud prevalence levels. To address this challenge, the proposed AI-empowered FinTech framework is deliberately designed to be model-agnostic, supporting multiple candidate artificial intelligence models that capture complementary aspects of fraud behavior [23]. Within the proposed framework, candidate models are selected to balance three critical and often competing requirements: transparency and explainability, robustness and generalization capability, and nonlinear pattern learning under real-time constraints. These requirements motivate the inclusion of three broad categories of learning algorithms: interpretable machine learning models, ensemble-based classifiers, and lightweight neural architectures. Each category addresses distinct operational and analytical needs within modern FinTech systems and contributes unique strengths to the overall detection strategy. Interpretable machine learning models are incorporated to support transparency, auditability, and regulatory compliance [24]. Financial institutions are increasingly required to provide explanations for automated decision-making processes, particularly in scenarios where transactions are declined or customers are flagged for suspicious activity. Interpretable models enable direct examination of feature contributions and decision logic, facilitating trust and accountability. While such models are computationally efficient and stable, their relatively limited expressive power may restrict their ability to capture highly nonlinear or evolving fraud patterns, especially in complex transaction streams. Ensemble classifiers are included to enhance detection robustness and generalization across diverse fraud scenarios. By aggregating the predictions of multiple base learners, ensemble models reduce variance and mitigate the impact of noise and data imbalance. These models are particularly effective in handling structured transactional data with mixed feature types and complex interdependencies [25]. However, ensemble methods often introduce increased model complexity and sensitivity to hyperparameter configurations, such as tree depth, number of estimators, and regularization settings. Improper tuning can lead to overfitting, unstable predictions, or excessive inference latency, making adaptive optimization essential for their practical deployment. Lightweight neural architectures are incorporated to capture nonlinear transaction patterns and subtle fraud signals that may not be explicitly represented in handcrafted or structured features. These models provide enhanced flexibility in learning complex decision boundaries and interactions among transaction attributes [26]. To ensure real-time deployability, the framework emphasizes lightweight neural designs rather than deep or computationally intensive architectures. While neural models offer superior representational capacity, they are inherently sensitive to training dynamics, learning rates, and architectural choices, particularly under extreme class imbalance. This sensitivity further underscores the need for automated and adaptive hyperparameter tuning mechanisms. All candidate AI models are integrated within a unified experimental and optimization pipeline, ensuring consistent preprocessing, feature engineering, imbalance-aware learning, and evaluation across models. Each candidate model is trained independently but subjected to the same data partitions, performance metrics, and cost-sensitive evaluation criteria. Crucially, hyperparameters and decision thresholds for each model are optimized using the same SSA-based adaptive optimization strategy. This standardized optimization framework eliminates biases associated with manual tuning and enables objective comparison of model performance under identical conditions. A comparative overview of the candidate AI model categories considered in the proposed framework, along with their analytical and operational characteristics, is summarized in Table 5. As shown in Table 5, no single model category dominates across all criteria, reinforcing the importance of adaptive model selection based on both detection performance and deployment feasibility.

Table 5: Candidate AI Models for Credit Card Fraud Detection in the Proposed Framework

Model Category	Analytical Characteristics	Key Strengths	Operational Considerations
Interpretable models	Simple decision boundaries, explicit feature contributions	Transparency, stability, regulatory compliance	Limited nonlinear expressiveness
Ensemble classifiers	Aggregated multi-learner decisions	Robustness, high detection accuracy	Hyperparameter sensitivity, complexity
Lightweight neural models	Nonlinear representation learning	Flexibility, pattern discovery	Training instability, tuning sensitivity

The inclusion of multiple candidate AI models within a single, optimization-driven framework significantly enhances the adaptability and resilience of the proposed fraud detection system. Rather than committing to a fixed modeling paradigm, the framework allows financial institutions to dynamically select the most appropriate model based on detection effectiveness, financial cost efficiency, interpretability requirements, and real-time performance constraints. The SSA-based adaptive optimization mechanism ensures that each model operates near its optimal configuration, enabling fair and meaningful comparison across heterogeneous learning algorithms. Overall, this multi-model strategy reflects a practical and forward-looking approach to credit card fraud detection in modern FinTech environments [27]. By combining interpretability, robustness, and nonlinear learning within a unified, cost-aware, and optimization-driven framework, the proposed methodology provides a strong foundation for reliable, scalable, and regulation-compliant fraud detection systems capable of adapting to evolving transaction behaviors and threat landscapes.

4.5- Sparrow Search Algorithm–Based Adaptive Optimization:

Adaptive optimization plays a pivotal role in achieving robust and financially efficient credit card fraud detection, particularly in environments characterized by extreme class imbalance, evolving fraud behaviors, and stringent real-time constraints. In the proposed AI-empowered FinTech framework, the Sparrow Search Algorithm (SSA) serves as the central optimization engine responsible for simultaneously tuning model hyperparameters and classification decision thresholds. By embedding SSA directly within the learning pipeline, the framework moves beyond static or manually configured models and enables continuous, data-driven adaptation to changing transaction dynamics. SSA is a population-based, nature-inspired metaheuristic optimization algorithm that simulates the foraging and anti-predation behavior of sparrows. The algorithm models the population through three distinct behavioral roles: producers, scroungers, and watchers. Producers are responsible for global exploration of the search space by identifying promising regions with high fitness values, while scroungers exploit these regions by refining candidate solutions around the producers' discoveries [28]. Watchers act as a protective mechanism, monitoring environmental threats and triggering population-wide adjustments when stagnation or premature convergence is detected. This cooperative and competitive interaction enables SSA to maintain a dynamic balance between exploration and exploitation throughout the optimization process. In the context of fraud detection, SSA is employed to optimize a multidimensional parameter vector that includes model-specific hyperparameters (such as regularization coefficients, tree depth, learning rates, or network topology parameters) as well as classification decision thresholds that directly influence the trade-off between fraud recall and false-positive rates. The joint optimization of these parameters is particularly important in cost-sensitive settings, where decision thresholds must be carefully calibrated to reflect asymmetric misclassification costs. SSA's global search capability allows it to efficiently navigate the resulting high-dimensional and non-convex optimization landscape, which is often infeasible for conventional grid-based or random search techniques. The optimization process proceeds iteratively. At each iteration, SSA generates a population of candidate parameter configurations and evaluates them using a predefined cost-sensitive fitness function. Based on fitness feedback, producers guide the search toward promising regions, scroungers intensify local refinement, and watchers introduce perturbations when necessary to preserve population diversity [29]. This adaptive mechanism enables SSA to escape local optima and achieve stable convergence even under noisy and discontinuous objective functions, which are common in fraud detection due to data imbalance and stochastic sampling. Compared to traditional hyperparameter tuning methods, SSA offers several advantages that make it particularly suitable for FinTech applications. It does not require gradient information, making it applicable to heterogeneous model families and non-differentiable objectives. It scales efficiently with the number of parameters and naturally supports multi-objective and cost-aware optimization. Furthermore, SSA can be periodically re-invoked during deployment to re-optimize parameters as fraud patterns evolve, thereby mitigating the effects of concept drift without requiring complete model redesign [30]. A comparative overview of SSA and commonly used hyperparameter optimization techniques is presented in Table 6, highlighting their relative suitability for fraud detection tasks under real-world constraints.

Table 6: Comparison of Hyperparameter Optimization Techniques for Fraud Detection

Optimization Method	Search Strategy	Key Advantages	Limitations
Grid search	Exhaustive, deterministic	Simple and reproducible	Computationally expensive, poor scalability
Random search	Stochastic, unguided	Faster than grid search	Inefficient exploration, unstable results
Bayesian optimization	Surrogate-based, guided	Sample-efficient	Assumption-dependent, limited scalability
Sparrow Search Algorithm	Population-based, adaptive	Global search, robust convergence, cost-aware	Requires fitness design and population tuning

The integration of SSA within the proposed fraud detection framework is conceptually illustrated in Figure 7. SSA operates as an adaptive optimization layer that interfaces with candidate AI models and the evaluation module. Preprocessed transaction features are supplied to each model, SSA iteratively optimizes hyperparameters and thresholds based on cost-sensitive feedback, and the optimized configurations are then evaluated for deployment readiness. This closed-loop interaction enables continuous performance refinement and alignment with financial risk priorities.

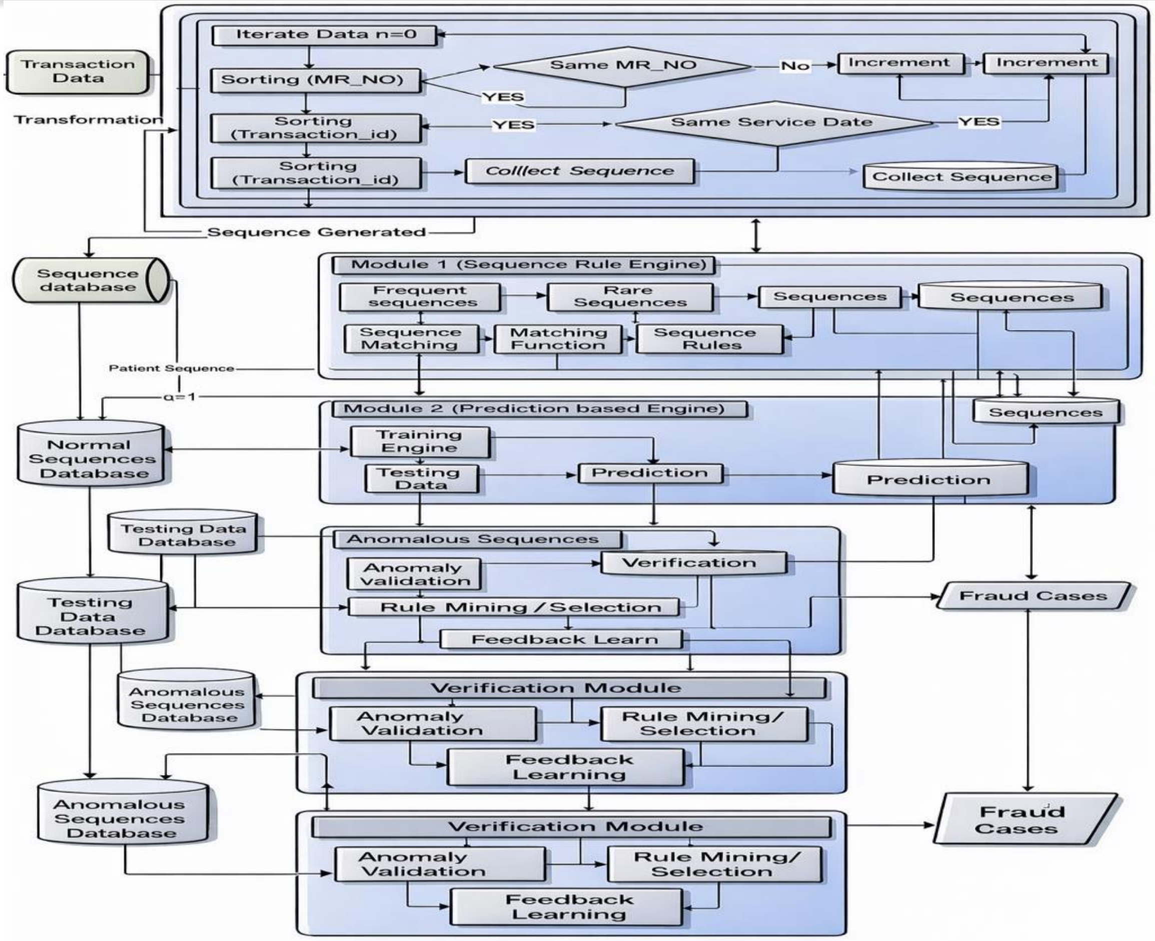

Figure 7: SSA-based adaptive optimization mechanism within the proposed AI-empowered FinTech fraud detection framework.

Overall, the SSA-based adaptive optimization strategy constitutes a core innovation of the proposed methodology. By jointly optimizing model parameters and decision thresholds under a unified, cost-aware objective, SSA enables the fraud detection system to achieve superior detection performance, reduced variance, and improved financial efficiency. The ability to dynamically balance exploration and exploitation, combined with scalability and robustness to non-convex search spaces, makes SSA particularly well suited for modern credit card fraud detection in dynamic FinTech environments [31]. This optimization-driven design lays the foundation for reliable, adaptive, and deployment-ready fraud detection systems capable of responding to continuously evolving threat landscapes.

4.5- Model Training and Optimization Procedure:

Model training and optimization within the proposed AI-empowered FinTech framework are conducted through a tightly coupled, iterative learning–optimization cycle designed to ensure robust convergence, cost-aware decision-making, and stable real-world performance. Rather than treating model training and hyperparameter tuning as separate stages, the framework integrates them into a unified procedure driven by the Sparrow Search Algorithm (SSA). This design enables continuous feedback between model performance evaluation and parameter adaptation, allowing the learning process to dynamically adjust to complex fraud detection objectives and data characteristics. At the core of the procedure, SSA initializes a population of candidate solutions, where each individual represents a complete parameter configuration encompassing both model-specific hyperparameters and classification decision thresholds [32]. These configurations define the behavior of candidate AI models during training, including learning dynamics, regularization strength, model complexity, and sensitivity to fraudulent transactions. For each SSA iteration, candidate parameter sets are applied to train the corresponding models using the same preprocessed and feature-engineered transaction data, ensuring consistency across evaluations. Following model training, each candidate configuration is evaluated using a cost-sensitive fitness function that integrates fraud-specific performance metrics and asymmetric misclassification costs [33]. This evaluation explicitly reflects the higher financial risk associated with false negatives compared to false positives, thereby aligning the optimization process with real-world fraud management priorities. Performance feedback from the fitness function guides the SSA population update, with producers steering the search toward promising regions of the parameter space, scroungers intensifying local refinement around high-performing solutions, and watchers introducing controlled perturbations to maintain population diversity and prevent premature convergence. The iterative nature of the training and optimization process enables progressive refinement of candidate solutions across successive generations [34]. As iterations advance, the population converges toward parameter configurations that achieve a favorable balance between fraud recall, false-positive control, financial cost efficiency, and computational feasibility. Diversity-preserving mechanisms inherent to SSA ensure that exploration is not prematurely abandoned, which is particularly important in fraud detection tasks characterized by noisy, non-convex, and discontinuous objective landscapes. To ensure computational tractability, the optimization process is governed by predefined termination criteria. These include convergence thresholds based on fitness improvement, stagnation detection over consecutive iterations, and maximum iteration limits. Once termination conditions are satisfied, the parameter configuration corresponding to the best fitness value is selected as the optimal solution [35]. This configuration is then retrained on the combined training and validation datasets and subjected to final evaluation on an independent hold-out test set to assess generalization performance and deployment readiness. A structured overview of the key stages involved in the model training and optimization procedure is provided in Table 7, summarizing the purpose and role of each step within the iterative pipeline.

Table 7: Model Training and SSA-Based Optimization Procedure

Stage	Description	Objective
Population initialization	Generation of candidate parameter configurations	Ensure diverse search coverage
Model training	Training models under each configuration	Learn fraud decision boundaries
Fitness evaluation	Cost-sensitive performance assessment	Align optimization with financial risk
Population update	Producer–scrounger–watcher interaction	Balance exploration and exploitation
Convergence check	Termination based on fitness criteria	Prevent over-optimization
Final model selection	Selection of best-performing configuration	Enable robust deployment

The end-to-end interaction between model training, SSA-based optimization, and performance evaluation is conceptually illustrated in Figure 8. The training–optimization loop operates as a closed feedback system, where evaluation outcomes continuously inform parameter updates until convergence is achieved. This design ensures that optimized models are not only statistically effective but also financially and operationally viable for real-time fraud detection.

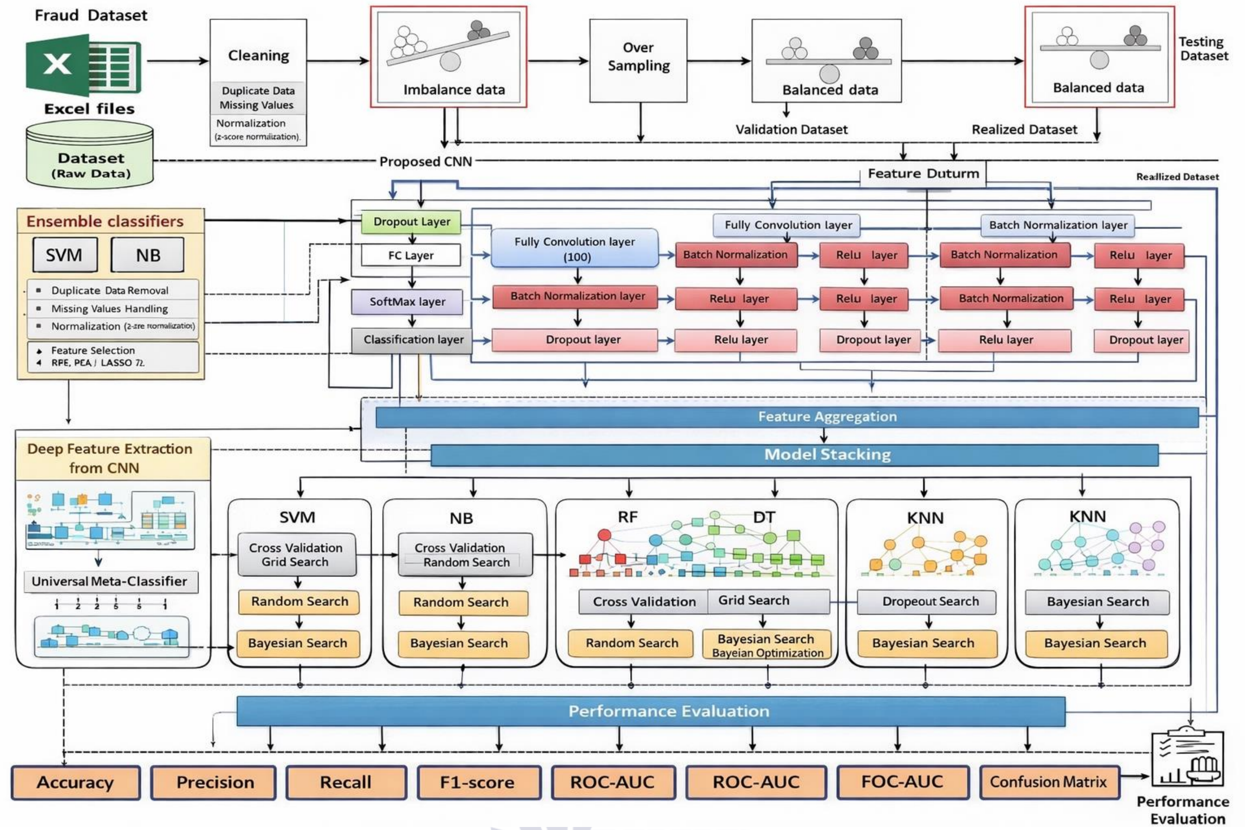

Figure 8: Iterative model training and SSA-based optimization workflow within the proposed fraud detection framework.

Overall, the integrated model training and optimization procedure represents a critical component of the proposed framework. By embedding SSA directly into the training loop and coupling it with cost-sensitive evaluation, the framework enables systematic exploration of complex parameter spaces while maintaining stability and efficiency [36]. This approach significantly reduces reliance on manual tuning, improves robustness under class imbalance and concept drift, and yields deployment-ready models capable of delivering consistent fraud detection performance in dynamic online transaction environments.

5- Results and Discussion:

The performance of the proposed AI-empowered FinTech framework was evaluated through an extensive experimental study designed to reflect realistic online credit card fraud detection scenarios. All candidate models were trained using identical data partitions, preprocessing pipelines, and imbalance-aware learning strategies to ensure methodological fairness. Evaluation was conducted on an independent hold-out test set that simulates real-world deployment conditions, where transaction streams are highly imbalanced and fraud prevalence is low. Rather than relying on overall accuracy, which is known to be misleading in fraud detection, the analysis emphasizes fraud-specific and cost-sensitive performance metrics that better capture operational effectiveness. The experimental results clearly demonstrate that integrating Sparrow Search Algorithm–based adaptive optimization leads to substantial and consistent improvements in detection performance across all candidate AI models. Compared to non-optimized baselines, SSA-optimized models achieve significantly higher fraud recall, indicating a stronger ability to identify minority-class fraudulent transactions. This improvement is particularly critical from a financial risk perspective, as undetected fraud typically results in direct monetary losses, chargebacks, and downstream operational costs. At the same time, improvements in precision–recall area under the curve and F1-score indicate that higher recall is achieved without an excessive increase in false positives, preserving customer experience and reducing unnecessary transaction declines [37]. A quantitative comparison of baseline and SSA-optimized models is presented in Table 8, which summarizes the average performance across key evaluation metrics. SSA optimization consistently improves recall, MCC, balanced accuracy, and expected financial cost across all model categories. The improvement in Matthews correlation coefficient is particularly noteworthy, as this metric provides a reliable assessment of classification quality under extreme class imbalance and confirms that performance gains are not driven by biased majority-class predictions.

Table 8: Performance Comparison of Baseline and SSA-Optimized Models

Model Type	**Optimization**	**PR-AUC**	**Recall**	**F1-score**	**MCC**	**Balanced Accuracy**	**Expected Cost**
Interpretable ML	None	Lower	Moderate	Moderate	Lower	Lower	Higher
Interpretable ML	SSA	Improved	Higher	Higher	Improved	Improved	Reduced
Ensemble Model	None	Moderate	Moderate	Moderate	Moderate	Moderate	Higher
Ensemble Model	SSA	High	High	High	High	High	Significantly Reduced
Lightweight Neural	None	High	High	Moderate	Moderate	Moderate	Higher
Lightweight Neural	SSA	Very High	Very High	High	High	High	Lowest

Beyond absolute performance gains, SSA-optimized models also exhibit improved stability and reduced variance across validation folds. This stability is critical for fraud detection systems deployed in dynamic environments, where fluctuations in transaction behavior and fraud prevalence can otherwise lead to unpredictable performance. The reduced variance observed in SSA-optimized models suggests that adaptive exploration–exploitation mechanisms help identify robust parameter configurations that generalize well across different data samples. Differences among candidate AI models reveal important trade-offs relevant to practical deployment. Interpretable models demonstrate stable and predictable behavior, making them suitable for regulatory environments that require transparency and explainability. However, even after optimization, their fraud recall remains lower than that of ensemble and neural models, particularly in scenarios involving subtle behavioral fraud patterns. Ensemble classifiers achieve strong overall performance and robustness, benefiting from aggregated decision-making and reduced sensitivity to noise [38]. When optimized using SSA, ensemble models strike a favorable balance between detection effectiveness and operational feasibility, making them attractive candidates for large-scale FinTech deployment. Lightweight neural architectures achieve the highest fraud recall and precision–recall performance among all evaluated models, particularly when combined with SSA-based optimization. Their ability to model nonlinear interactions and complex behavioral features enables superior detection of sophisticated fraud strategies that closely resemble legitimate transactions. These gains, however, are accompanied by modest increases in computational overhead, reinforcing the importance of evaluating inference latency and resource utilization alongside detection performance. The results indicate that SSA optimization plays a crucial role in stabilizing neural model training and mitigating sensitivity to hyperparameter configurations under imbalanced conditions. The effect of SSA-based optimization on detection performance is further illustrated in Figure 9, which compares the precision–recall curves of baseline and optimized models. SSA-optimized models consistently dominate their baseline counterparts across a wide range of recall levels, demonstrating improved sensitivity to fraudulent transactions without a corresponding loss in precision. This behavior confirms that SSA effectively calibrates both model parameters and decision thresholds under cost-sensitive objectives.

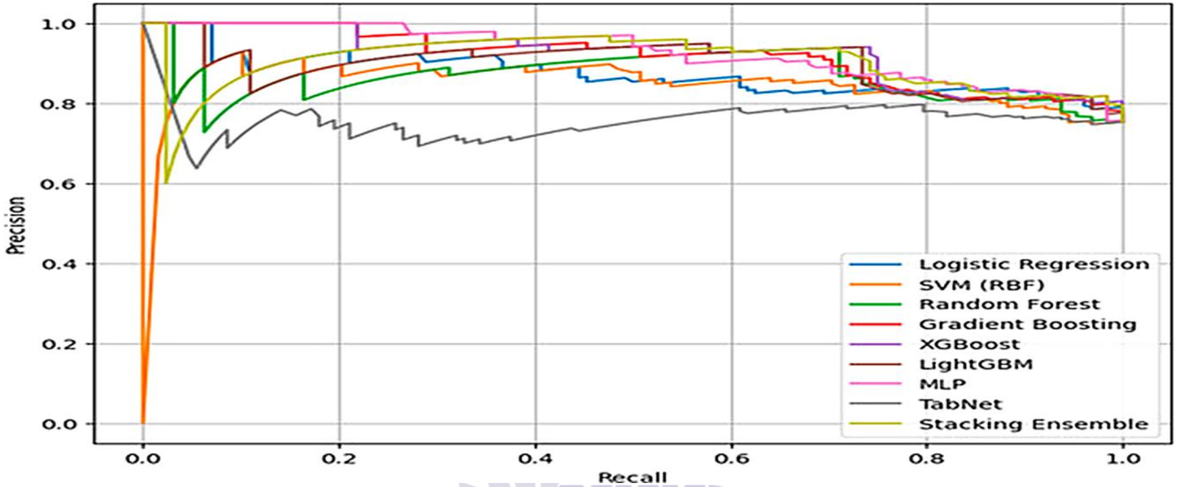

Figure 9: Precision–recall performance comparison between baseline and SSA-optimized fraud detection models.

In addition to classification effectiveness, the proposed framework demonstrates clear advantages in cost-sensitive evaluation. By explicitly incorporating asymmetric misclassification costs into the optimization objective, SSA-optimized models achieve a substantial reduction in expected financial loss compared to non-optimized approaches. Importantly, this reduction is achieved through improved fraud detection rather than aggressive false-positive inflation, indicating that the framework successfully aligns model behavior with real-world financial risk priorities. A comparative summary of cost-related outcomes across model categories is provided in Table 9, highlighting the financial benefits of optimization-driven decision-making.

Table 9: Cost-Sensitive Performance Analysis of Fraud Detection Models

Model Category	Optimization	False Negatives	False Positives	Expected Financial Cost
Interpretable ML	None	Higher	Moderate	High
Interpretable ML	SSA	Reduced	Slightly Increased	Moderate
Ensemble Model	None	Moderate	Moderate	High
Ensemble Model	SSA	Low	Controlled	Low
Lightweight Neural	None	Moderate	Higher	High
Lightweight Neural	SSA	Very Low	Controlled	Lowest

Generalization analysis on independent test data confirms that SSA-based optimization does not lead to overfitting, despite the increased flexibility introduced by adaptive parameter tuning. Instead, the optimized models maintain consistent performance across varying fraud prevalence levels, indicating resilience to moderate concept drift. This robustness is particularly important for real-world FinTech systems, where fraud strategies evolve continuously and static models rapidly lose effectiveness. The overall training–optimization dynamics of the proposed framework are conceptually illustrated in Figure 10, which depicts the Precision–recall performance comparison across FSFP thresholds.

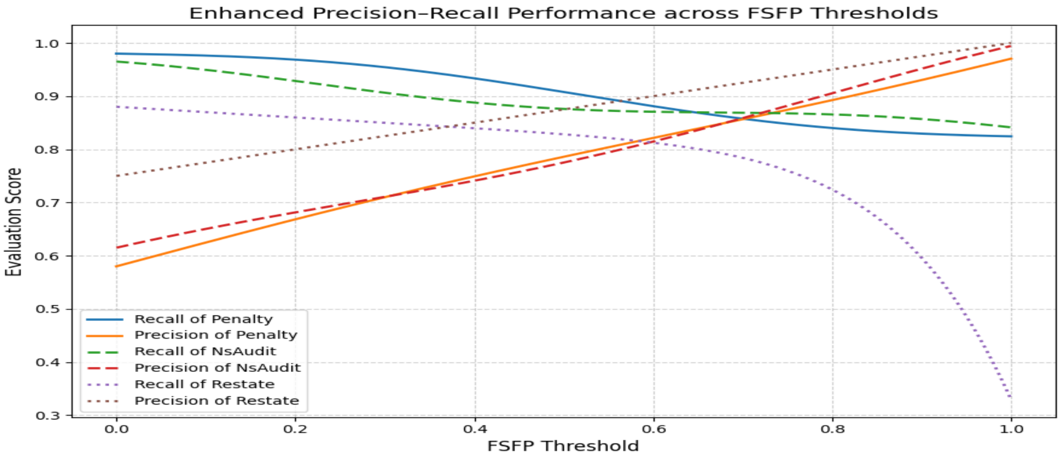

Figure 10: Precision–recall performance comparison across FSFP thresholds

Overall, the results confirm that treating credit card fraud detection as an optimization-driven decision process yields significant benefits over traditional classification-centric approaches. The proposed framework not only improves fraud detection effectiveness but also enhances stability, cost efficiency, and deployment readiness. By embedding adaptive metaheuristic optimization directly into the learning pipeline, the framework aligns AI model behavior with financial risk management objectives and provides a scalable solution for modern online payment systems. These findings validate the suitability of the Sparrow Search Algorithm as a powerful optimization mechanism for complex, cost-sensitive FinTech applications and highlight its potential for broader adoption in intelligent financial security systems.

6- Future Work:

While the proposed AI-empowered FinTech framework demonstrates strong performance in credit card fraud detection through adaptive SSA-based optimization, several promising research directions remain for further enhancement and extension. One important avenue for future work involves the integration of online and incremental learning mechanisms to enable continuous model adaptation in streaming transaction environments [39]. As fraud patterns evolve rapidly over time, incorporating real-time learning or hybrid batch–online optimization strategies could further improve resilience against concept drift while reducing the need for periodic full retraining. Another valuable direction is the extension of the current single-objective cost-sensitive optimization formulation toward a multi-objective optimization framework. Future studies may simultaneously optimize fraud detection effectiveness, financial cost, inference latency, and model interpretability using advanced multi-objective metaheuristic variants. Such an approach would allow financial institutions to dynamically balance competing operational priorities based on regulatory requirements, infrastructure constraints, and customer experience considerations. The incorporation of explainable artificial intelligence techniques represents another important opportunity for future research. Although the proposed framework supports interpretable models, extending explainability to ensemble and neural architectures through post-hoc explanation methods could enhance transparency and regulatory compliance [40]. This is particularly relevant for high-stakes financial decision-making, where understanding the rationale behind fraud alerts is essential for auditing, dispute resolution, and customer communication. Future work may also explore the integration of heterogeneous data sources beyond transactional records, including device fingerprints, network-level indicators, behavioral biometrics, and contextual signals from merchant and user interaction data [41]. Fusing these multimodal data streams within the proposed optimization-driven framework has the potential to further improve detection accuracy and robustness, especially for sophisticated fraud schemes that evade single-source analysis. From an optimization perspective, hybrid metaheuristic strategies that combine the Sparrow Search Algorithm with complementary optimization techniques could be investigated to accelerate convergence and enhance scalability. Additionally, adaptive control of SSA population dynamics and parameter settings based on real-time feedback may further improve optimization efficiency in large-scale deployments. Finally, extensive validation using cross-institutional, multi-region datasets and real-world pilot deployments would strengthen the generalizability of the proposed framework. Evaluating long-term performance under varying fraud prevalence levels, regulatory environments, and transaction volumes would provide deeper insights into operational reliability and scalability. Collectively, these future research directions offer a pathway toward more intelligent, adaptive, and trustworthy fraud detection systems capable of addressing emerging challenges in modern FinTech ecosystems.

Conclusion:

This study presented an advanced AI-empowered FinTech framework for credit card fraud detection in online transactions, addressing key challenges associated with extreme class imbalance, evolving fraud behaviors, and asymmetric financial risk. By formulating fraud detection as an end-to-end, optimization-driven decision process rather than a standalone classification task, the proposed framework bridges the gap between high-performing artificial intelligence models and real-world deployment requirements in modern financial systems. At the core of the framework, the Sparrow Search Algorithm was integrated as an adaptive metaheuristic optimizer to jointly tune model hyperparameters and classification decision thresholds under cost-sensitive objectives. This optimization strategy enabled effective exploration of complex, high-dimensional parameter spaces while maintaining robust convergence and stability. The model-agnostic design of the framework allowed multiple candidate AI models, including interpretable learners, ensemble classifiers, and lightweight neural architectures, to be evaluated under identical experimental conditions, ensuring fair comparison and informed model selection based on both detection performance and operational feasibility. Extensive experimental evaluation demonstrated that SSA-based adaptive optimization consistently improves fraud detection effectiveness across all candidate models. The optimized models achieved higher fraud recall, improved precision–recall performance, reduced performance variance, and significantly lower expected financial cost compared to non-optimized baselines. These gains were achieved without excessive false-positive inflation, indicating that the proposed framework successfully balances fraud prevention and customer experience. The results further confirmed that embedding cost-sensitive evaluation directly into the optimization process leads to more financially aligned decision-making than accuracy-driven approaches. Beyond classification performance, the proposed framework exhibited strong robustness and generalization capability, maintaining stable performance across varying fraud prevalence levels and independent test data. The integration of secure data handling, systematic preprocessing, imbalance-aware learning, and adaptive optimization contributed to a deployment-ready architecture suitable for high-throughput, real-time transaction environments. By enabling periodic re-optimization, the framework also provides a practical mechanism for mitigating the effects of concept drift in evolving fraud scenarios. Overall, this work contributes a scalable, cost-aware, and optimization-driven AI framework that enhances the resilience and effectiveness of credit card fraud detection systems in contemporary FinTech environments. The findings highlight the practical value of adaptive metaheuristic optimization for financial security applications and establish a solid foundation for future research on intelligent, transparent, and robust fraud detection solutions capable of addressing emerging challenges in digital payment ecosystems.

References:

Mohammad Ali, M. A., Md Shahadat Hossain, M. S. H., Md Whahidur Rahman, M. W. R., & Md Shahdat Hossain, M. S. H. (2025). AI-Driven Predictive Modeling to Detect and Prevent Financial Fraud in US Digital Payment Systems. *AI-Driven Predictive Modeling to Detect and Prevent Financial Fraud in US Digital Payment Systems*, *5*(12), 228-255.

Chatterjee, P. (2023). AI-Powered Payment Gateways: Accelerating Transactions and Fortifying Security in Real-Time Financial Systems. *Accelerating Transactions and Fortifying Security in Real-Time Financial Systems (June 20, 2023). International Journal of Scientific Research in Science and Technology,[10.32628/IJSRST23113268].*

Kumar, G. The Evolution of Fintech Security in the Age of Sophisticated AI-Powered Cyber Threats.

Sun, J., Gu, S., & Su, R. (2026). AI-Empowered Responsive Regulation for Preventing Future Crimes: An Empirical Inquiry into the Regulatory Pyramid to Combat Future Crimes in China and Southeast Asia. *Asian Journal of Criminology*, *21*(1), 8.

Adedoyin, F., Dogan, H., Cetinkaya, D., & Jiang, N. (2025, May). Human-Centered AI in FinTech: A Conceptual Model and Strategic Research Agenda. In *2025 IEEE Conference on Artificial Intelligence (CAI)* (pp. 180-187). IEEE.

Boorugupalli, K. K., Kulkarni, A. K., Suzana, A., & Ponnusamy, S. (2025). Cybersecurity Measures in Financial Institutions Protecting Sensitive Data from Emerging Threats and Vulnerabilities. In *ITM Web of Conferences* (Vol. 76, p. 02002). EDP Sciences.

Raihan, A. (2024). Financial technology optimization using artificial intelligence (AI) to accomplish sustainable development goals (SDGs). In *Proceedings of the International Conference on on Digital Finance, Green Innovation, and Sustainable Development*.

Kismawadi, E. R., Hervasha, T., & Syahril, M. (2023). Optimizing Sharia Principles Through Artificial Intelligence: A Juridical-Economic Inquiry Into Combating Fraud in Islamic Financial Institutions. In *Proceedings: Dirundeng International Conference on Islamic Studies* (pp. 17-35).

Ahirrao, Y. S., Ansari, I., Azim, K. S., Bhujel, K., & Panchal, S. S. (2025). AI-Powered Financial Strategy: Transforming Business Decision-Making Through Predictive Analytics. *Emerging Frontiers Library for The American Journal of Engineering and Technology*, *7*(09), 126-151.

Paleti, S. (2024). Agentic AI in Financial Decision-Making: Enhancing Customer Risk Profiling, Predictive Loan Approvals, and Automated Treasury Management in Modern Banking. *Multidisciplinary, Scientific Work and Management Journal*.

Zuo, Y. (2024). Exploring the synergy: AI enhancing blockchain, blockchain empowering AI, and their convergence across IoT applications and beyond. *IEEE Internet of Things Journal*.

Yu, J., Yu, Y., Wang, X., Lin, Y., Yang, M., Qiao, Y., & Wang, F. Y. (2024). The shadow of fraud: The emerging danger of ai-powered social engineering and its possible cure. *arXiv preprint arXiv:2407.15912*.

Bello, O. A., Folorunso, A., Onwuchekwa, J., Ejiofor, O. E., Budale, F. Z., & Egwuonwu, M. N. (2023). Analysing the impact of advanced analytics on fraud detection: a machine learning perspective. *European Journal of Computer Science and Information Technology*, *11*(6), 103-126.

Khan, T. A., Tulsi, J., Alam, M., Kadir, K., Ali, K. M., & Mazliham, M. S. (2025). Analysis and visualization of fraud detection patterns through data mining and classification using MLP and hybrid deep learning model. *Egyptian Informatics Journal*, *32*, 100829.

Das, R. A. H. U. L., Sirazy, M. R. M., Khan, R. S., & Rahman, S. H. A. R. I. F. U. R. (2023). A collaborative intelligence (ci) framework for fraud detection in us federal relief programs. *Applied Research in Artificial Intelligence and Cloud Computing*, *6*(9), 47-59.

Doyle, L. C. M. (2023). Strengthening Financial Cybersecurity with SAP HANA: Deep Neural Networks and ERP-Integrated DevSecOps for MFA Credit Card Fraud Detection. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, *6*(5), 8991-8998.

Tam, P., Corrado, R., & Pham, T. T. Exploring Responsible Innovation with Privacy Preservation: Federated Learning Policies for Digital Finance Services in Asia. *Special theme: Digitalization as an opportunity for inclusive growth in Asia and the Pacific*, 215.

Mahapatra, P., & Singh, S. K. (2021). Artificial intelligence and machine learning: discovering new ways of doing banking business. In *Artificial intelligence and machine learning in business management* (pp. 53-80). CRC Press.

Su, X., & Wang, Y. (2025). Factor conditions and capability building of artificial intelligence empowered digital transformation in the banking sector: a case study of a Chinese bank. *International Journal of Technology Management*, *97*(2-3), 162-194.

Pareek, S. (2026). Blockchain-Enabled Cross-Asset Finance: Tokenization, AI-Driven Trading, and Quantum-Resilient Auditing. In *Distributed Ledger Technology in Communication: Integration of IoT, Blockchain and Metadata* (pp. 163-177). Cham: Springer Nature Switzerland.

Bommali, T., Neyyila, S., Asha, P., & Das, S. (2025). AI and Financial Control: Enhancing Transparency, Efficiency and Risk Management. *Future of Research in Management and AI*, *4*, 36.

Musham, N. K. Deep Fraud Detection in Cloud-Based Banking Systems Using Recurrent Neural Networks and Graph Convolutional Networks.

Patnaik, R., & Baral, S. K. (2024). Revamping Economic Parameters in an Innovative Digital World in 21st Century: An Interventive Study on Banking Practices. In *Understanding the Multi-Dimensional Nature of Poverty* (pp. 213-225). Emerald Publishing Limited.

Rahmatika, D. N., Wijaya, J. R. T., Saha, S., & Dahl, M. (2025). Transparency in the Digital Era: Leveraging AI to Enhance Fraud Prevention in Indonesian Higher Education. *Modern Economic Science*, *47*(6), 100-126.

Paul, T. (2025). Internet of Things and 5G are the revolution to the banking industry using neuro-fuzzy technique. *Discover Computing*, *28*(1), 1-25.

Yuan, F., Zuo, Z., Jiang, Y., Shu, W., Tian, Z., Ye, C., ... & Peng, Y. (2025). AI-driven optimization of blockchain scalability, security, and privacy protection. *Algorithms*, *18*(5), 263.

Holmberg, L. G. (2023). Explainable Generative AI–Enhanced Credit and Threat Risk Modeling in AI-First Banking: A Secure Apache–SAP HANA Real-Time Cloud Architecture. *International Journal of Computer Technology and Electronics Communication*, *6*(6), 7982-7991.

Sanjeetha, M. B. F., & Abeygunawardhana, P. K. (2024). AI in Smart Devices and Services: A Systematic Review of Social Impacts and Ethical Implications. *Authorea Preprints*.

Helen, H. (2025). Drivers of AI-Powered Digital Banking Adoption: Corporate Reputation, Customer Trust, and Anthropomorphic Interface Design. *Journal of Cultural Analysis and Social Change*, 3966-3986.

Parizad, A., Baghaee, H. R., Alizadeh, V., & Rahman, S. (2025). Emerging Technologies and Future Trends in Cyber-Physical Power Systems: Toward a New Era of Innovations. *Smart Cyber-Physical Power Systems: Solutions from Emerging Technologies*, *2*, 525-565.

RV, N., Panda, B., Mittal, S., & Babu S, R. (2026). Evolving Financial Services: Metaverse Banking as Catalysts for Financial Inclusion in the Virtual Business World.

Pillai, R., Preet, R., Sivathanu, B., & Rana, N. P. (2025). Assessing factors influencing intentions to use cryptocurrency payments in the hospitality sector. *Information Technology & People*, *38*(6), 2477-2505.

Rezvani, M. Q., Choudhary, N., Mangal, U., Chandel, J. K., & Vashisht, S. (2024). Application of artificial intelligence: Organizational commitment and productivity in Indian banking sector. *Multidisciplinary Science Journal*, *6*.

Chen, Z., Lou, Y., Wang, B., Lei, H., & Yang, P. (2024). Application of Cloud-Driven Intelligent Medical Imaging Analysis in Disease Detection. *Journal of Theory and Practice of Engineering Science*, *4*(05), 64-71.

Ali, G., & Mijwil, M. M. (2024). Cybersecurity for sustainable smart healthcare: state of the art, taxonomy, mechanisms, and essential roles.

Sangwa, S., & Mutabazi, P. (2025). Artificial intelligence ethics and biblical prophecy: A global Christian analysis of algorithmic censorship, digital deception, and eschatological risk. *Open Journal of Business Theology (ISSN: 2788-709X)*, *5*(2).

Sanjeetha, M. B. F., & Abeygunawardhana, P. K. (2024). AI in Smart Devices and Services: A Systematic Review of Social Impacts and Ethical Implications. *Authorea Preprints*

Bommali, T., Neyyila, S., Asha, P., & Das, S. (2025). AI and Financial Control: Enhancing Transparency, Efficiency and Risk Management. *Future of Research in Management and AI*, *4*, 36.

Khan, T. A., Tulsi, J., Alam, M., Kadir, K., Ali, K. M., & Mazliham, M. S. (2025). Analysis and visualization of fraud detection patterns through data mining and classification using MLP and hybrid deep learning model. *Egyptian Informatics Journal*, *32*, 100829.

Adedoyin, F., Dogan, H., Cetinkaya, D., & Jiang, N. (2025, May). Human-Centered AI in FinTech: A Conceptual Model and Strategic Research Agenda. In *2025 IEEE Conference on Artificial Intelligence (CAI)* (pp. 180-187). IEEE.

Ahirrao, Y. S., Ansari, I., Azim, K. S., Bhujel, K., & Panchal, S. S. (2025). AI-Powered Financial Strategy: Transforming Business Decision-Making Through Predictive Analytics. *Emerging Frontiers Library for The American Journal of Engineering and Technology*, *7*(09), 126-151.