

Hybrid PCA–SVM Approach to Credit Card Fraud Detection: Enhancing Payment System Oversight and Financial Stability

Younes Nademi*
Majid Ebtia‡

Sayyed Mohammad Hoseini†
Faranak Ahmadi§

Received: 05 Sep 2025

Approved: 01 Feb 2026

Credit card fraud remains a significant threat to financial institutions and the integrity of digital payment systems, posing challenges for both operational risk management and regulatory oversight. This paper presents a novel hybrid machine learning framework for credit card fraud detection that combines Principal Component Analysis (PCA) for feature extraction with a Support Vector Machine (SVM) based feature selection mechanism. The aim is to reduce dimensionality while retaining the most informative features, thereby improving detection performance on highly imbalanced transaction datasets. The approach is evaluated on a large credit card transaction dataset, where PCA is first used to transform the input variables into principal components capturing the majority of variance, and an SVM with recursive feature elimination is then employed to identify and retain the most relevant components. Experimental results demonstrate that the proposed PCA–SVM pipeline significantly outperforms baseline models lacking this hybrid feature engineering: for example, it achieves a higher fraud recall (detection rate) by several percentage points while maintaining high precision, leading to improved F_1 scores and overall accuracy. These findings indicate that the hybrid method effectively mitigates class imbalance issues and eliminates redundant features, yielding a more compact and robust fraud detection model. By enhancing the identification of rare fraudulent transactions without excessive false alarms, our study contributes to central bank objectives in fraud risk management. The proposed framework can strengthen the resilience of digital payment infrastructures and support payment system oversight, ultimately helping to safeguard financial stability and public trust in electronic payment channels.

*Department of Economics, Faculty of Humanities, Ayatollah Boroujerdi University, Boroujerd, Iran; younesnademi@abru.ac.ir (Corresponding Author)

†Gahar Artificial Intelligence Research Group, Ayatollah Boroujerdi University, Boroujerd, Iran; smhoseiny@gmail.com

‡Zagros Data Sciences Research Group, Ayatollah Boroujerdi University, Boroujerd, Iran; majid.ebtia@gmail.com

§Department of Computer Engineering, Faculty of Engineering, Ayatollah Boroujerdi University, Boroujerd, Iran; faranaka.11@gmail.com

Keywords: Credit Card Fraud Detection; Principal Component Analysis; Support Vector Machine; Feature Selection; Hybrid Machine Learning; Feature Extraction

JEL Classification: C45, G21, G28, E58

1 Introduction

Credit card fraud inflicts billions in losses worldwide each year, eroding consumer confidence in digital payments and imposing substantial costs on banks, merchants, and payment networks. Central banks and monetary authorities view fraud reduction not only as a private-sector imperative but also as a regulatory priority: safeguarding the integrity of retail payment infrastructures underpins financial stability and public trust in cashless systems.

From a machine-learning standpoint, fraud detection presents two interlocking challenges. First, fraud events are extremely rare—often under 1% of all transactions—making classifier sensitivity difficult without an excessive false positive. Second, transaction records are high-dimensional, populated with dozens of potentially correlated or irrelevant features that can hinder model performance, inflate computation time, and obscure interpretability. Consequently, effective fraud systems hinge on robust feature-engineering pipelines.

In the literature, two families of techniques dominate this preprocessing stage.

- Feature extraction methods such as Principal Component Analysis (PCA) map original variables onto orthogonal components, reducing dimensionality and anonymizing sensitive data (e.g., the Kaggle credit-card benchmark releases 28 PCA features).
- Feature selection approaches—filters, wrappers, and embedded methods—prune irrelevant inputs. Support Vector Machines (SVMs), in particular, have been leveraged both as high-dimensional classifiers and as wrappers via Recursive Feature Elimination (RFE) to rank and remove low-weight features.

Although PCA and SVM-RFE each appear in prior fraud studies, they have seldom been combined in a deliberate two-stage pipeline—and never, to our knowledge, applied consecutively on the same PCA components. Moreover, existing works often focus on a single classifier (e.g., Random Forest or an autoencoder-based deep net) without systematically comparing diverse algorithms under a reproducible setup.

Our study fills this gap with three key contributions:

- 1) We introduce a composite PCA→SVM-RFE pipeline, applying SVM-RFE directly to the 28 anonymized components and selecting the ten most discriminative for fraud detection.

- 2) We conduct transparent benchmarking of seven classifiers (KNN, Naïve Bayes, Logistic Regression, Decision Tree, Random Forest, XGBoost, and MLP) using identical data splits, 10-fold cross-validation for both feature selection and hyperparameter tuning, and a held-out test set.
- 3) We demonstrate a computationally efficient, interpretable framework—leaner than deep-learning ensembles yet competitive in recall, precision, and F1-score—suitable for real-time deployment by banks and regulators managing highly imbalanced transaction streams.

By weaving PCA and SVM-RFE into a unified, reproducible pipeline and evaluating multiple classifiers on that reduced feature space, we advance the state of the art in credit-card fraud detection and provide a practical blueprint for operational fraud monitoring.

From a methodological perspective, detecting fraudulent transactions is notoriously difficult. Fraud attempts are exceedingly rare relative to the volume of legitimate transactions, often comprising less than 1% of card payments. This extreme class imbalance makes it challenging to identify illicit activities without overwhelming the system with false alarms. Moreover, transaction datasets are typically high dimensional, containing dozens of features that may be correlated, redundant, or irrelevant. Without careful preprocessing, such data can degrade model performance, increase computational costs, and reduce the interpretability of results. Consequently, feature extraction and feature selection play a central role in developing effective fraud detection systems.

In modern machine learning, algorithms are applied to various tasks—including classification, pattern recognition, and clustering—using datasets with a high number of features (Chandrashekar & Sahin, 2014). The quality of a model's output largely depends on the relevance of these features, making the elimination of redundant or irrelevant variables crucial in the preprocessing stage (Guyon & Elisseeff, 2003). Feature selection methods are generally categorized into filter, wrapper, and embedded approaches (Chandrashekar & Sahin, 2014; Kohavi & John, 1997). Filter methods rank features based on statistical metrics, wrapper methods evaluate subsets based on model performance, and embedded methods incorporate selection into the training process (Shah et al., 2020). Support Vector Machines (SVMs) have gained popularity not only for their classification power in high-dimensional spaces but also for their ability to assign weights to features, thereby aiding in selection (Han et al., 2012; Guyon et al., 2002). In contrast, feature extraction techniques such as Principal Component Analysis (PCA) transform the original variables into new components that capture the most variance (Song

et al., 2010). Although PCA may reduce interpretability and risk losing critical information, its combination with SVM can compensate by refining the extracted features (Naik, 2018). In our proposed composite approach, PCA is first used to extract new features from the dataset, after which SVM ranks and eliminates less significant features, ensuring that only the most discriminative variables are retained.

This hybrid PCA–SVM approach is particularly suited to the challenges of fraud detection. PCA addresses the curse of dimensionality by summarizing information into orthogonal components, while SVM’s recursive feature elimination ensures that only the most relevant of these components contribute to classification. By integrating both techniques, the framework not only enhances detection accuracy but also improves computational efficiency and interpretability. Importantly, these technical improvements have broader implications for the financial system: they provide banks, payment processors, and central banks with more effective tools to detect fraud in real time, thereby strengthening the resilience of digital payment infrastructures and supporting financial stability (Hoseini et al., 2025).

The remainder of this paper is organized as follows. In Section 2, we review the state of the art in credit-card fraud detection, with emphasis on PCA-based anonymization and SVM-based feature selection. Section 3 describes the methodology and data, including our stratified undersampling strategy and evaluation metrics and the proposed hybrid PCA–SVM-RFE pipeline in detail, covering normalization, recursive feature elimination, and the selection of the ten most discriminative components. Section 4 reports and analyzes the results on hold-out test sets, highlighting performance gains in recall, F1-score, and computational efficiency. Finally, Section 5 concludes with a discussion of practical implications for real-time fraud monitoring, acknowledges limitations, and suggests directions for future work.

2 Literature Review

Credit card fraud detection has attracted extensive research due to its financial impact and the evolving tactics of fraudsters (Sundaravadivel et al., 2025). Traditional rule-based systems struggle with large-scale, dynamic data, so modern approaches rely on supervised and unsupervised machine learning (ML) techniques (Chang et al., 2024). The key challenges include severe class imbalance (fraud <1% of transactions) and high dimensionality of transaction features (Zaffar, 2023). Most studies adopt preprocessing methods (e.g. Synthetic Minority Over-sampling Technique, Edited Nearest Neighbors) to rebalance data before classification (Sundaravadivel et al., 2025). Popular

classifiers include ensemble methods (Random Forest, Extreme Gradient Boosting), Neural Networks, and hybrid models combining multiple algorithms (Talukder et al., 2024). For example, Sundaravadivel et al. (2025) applied several models (RF, neural nets) on balanced data, using SMOTE and reporting Random Forest as most effective. While many works focus on classification algorithms, less attention has been paid to the interplay of feature extraction and selection in this domain. Recent reviews highlight that high dimensional fraud data benefit critically from feature reduction and selection, but systematic treatment of these steps in fraud detection is limited (Wang et al., 2024).

Principal Component Analysis (PCA) is a common linear technique for feature extraction and dimensionality reduction. In credit card fraud detection, PCA serves two roles: data anonymization and variance preservation. Notably, the benchmark Kaggle dataset of European credit card transactions inherently uses PCA – its 28 numeric features (V1–V28) are already principal components (Chang et al., 2024). Many studies further apply PCA as a preprocessing step. For instance, Sundaravadivel et al. (2025) removed correlated features and then applied PCA to “further reduce features while preserving variance”. Similarly, Ogundile et al. (2024) integrated PCA with a Hidden Markov Model, demonstrating that a PCA-Embedded Hidden Markov Model outperformed the conventional Hidden Markov Model approach by focusing on principal components. However, pure PCA can also obscure nonlinear relationships. Wang et al., (2024) compared PCA with Convolutional Autoencoders (CAEs) for feature extraction; they found that nonlinear CAEs (with under-sampling) surpassed PCA in a performance metric combining precision and recall, although PCA remains a viable baseline. In sum, PCA efficiently reduces the data to a few orthogonal components but its linearity and information loss may limit the ultimate classification accuracy. Notably, few works examine selecting among PCA components: most simply feed all or a fixed number of components into classifiers. The research gap remains in optimally integrating PCA with downstream selection strategies.

Support Vector Machines (SVMs) are widely used for classification in fraud detection due to their robustness to high-dimensional data (Rtayli & Enneya, 2020). Beyond classification, SVMs can serve as feature selectors (e.g., via Recursive Feature Elimination). Rtayli and Enneya (2020) proposed a hybrid SVM model combining RFE, hyperparameter tuning, and SMOTE; their experiments on multiple datasets showed that SVM-RFE effectively identified predictive features and yielded “best results in terms of efficiency

and effectiveness”. This work highlights SVM’s utility for feature ranking in fraud contexts. More recently, explainable AI methods such as SHapley Additive exPlanations (SHAP) have been applied to rank features, by importance. Kennedy et al. (2025) reported that SHAP-based selection significantly improves label quality for fraud datasets. Kennedy et al., (2025) compared SHAP-value and built-in importance measures for feature selection on the Kaggle fraud dataset. They found model-based importance measures (e.g. from tree ensembles) generally outperformed SHAP in Area Under the Precision–Recall Curve, and recommended built-in importance for large datasets. This suggests that smart feature ranking can materially affect model performance. However, most SVM applications in fraud studies focus on classification accuracy, not on explicit feature filtering. There is little work using an SVM after a feature extraction stage as in our pipeline. In general ML practice, SVM-RFE and filter methods like mutual information are common, but their comparative merit in fraud domains is underexplored.

To address the limitations of single techniques, many recent works propose hybrid or ensemble schemes. For example, Gupta et al. (2025) designed a stacking ensemble with robust feature selection, using SMOTE-ENN, autoencoders, and Technique for Order Preference by Similarity to Ideal Solution for selection, achieving very high ($\approx 99.95\%$) accuracy. Alamri and Ykhlef (2023) introduced a hybrid “feature aggregation + Exhaustive Feature Selection (EFS)” method: they created higher-order features from spending behavior and then exhaustively search subsets to find optimal predictors. Tested on the PaySim dataset, their approach significantly improved accuracy, F1, and AUPRC. While EFS is thorough, it is computationally intensive and impractical for real-time systems. Other hybrids focus on combining machine and deep learning: for instance, Convolutional Neural Network–Support Vector Machine–Light Gradient Boosting Machine (CNN–SVM–LightGBM) ensembles or Long Short-Term Memory with Attention (LSTM–Attention) networks have been proposed to capture temporal and nonlinear patterns. Still others integrate optimization techniques: Particle Swarm Optimization and Genetic Algorithms have been used to tune or select features and models. These methods underline the value of hybrid pipelines but often do not clearly separate feature extraction from selection. In many cases, dimensionality reduction and feature pruning are embedded within more complex architectures (deep autoencoders, clustering initializations, etc.)

Despite these advances, gaps remain in understanding and improving feature-level processing for fraud detection. First, although PCA is used for anonymization and reduction, its interplay with selective pruning has not been

deeply investigated. Most fraud studies treat PCA either as a black-box transformer (e.g. the Kaggle dataset) or as one of many preprocessing steps, without explicitly selecting among the resulting components. Second, SVM is well-known for classification but underutilized as a dedicated feature selector in this domain; except for RFE-based efforts, few works exploit the margin-maximizing property of SVM for picking features. Third, many hybrid models achieve high accuracy but often do so by adopting deep learning or complex ensembles, which may sacrifice interpretability and increase computational cost (e.g. CNN-SVM hybrids, or stacked autoencoder pipelines). There is a need for streamlined, transparent approaches that balance performance with efficiency. Finally, comparative studies of linear vs. nonlinear extraction (PCA vs CAE) and of different selection strategies (SHAP vs importance) indicate that the choice of techniques can significantly affect results, but consensus is lacking.

Our proposed study addresses these gaps by explicitly combining PCA for extraction and SVM for selection, thereby integrating complementary strengths: PCA's dimensionality reduction and SVM's discriminative feature filtering. This sequential approach has not been systematically evaluated in credit card fraud detection. By doing so, we aim to reduce feature redundancy while preserving discriminative information, thus improving model performance on imbalanced fraud data. The novelty lies in treating PCA and SVM as a unified pipeline: we hypothesize that SVM-RFE can identify which principal components (or derived features) are most predictive of fraud. This differs from prior work that either applied PCA in isolation or used SVM-RFE on raw or engineered features. Ultimately, by critically building on recent studies, our approach contributes a novel, hybrid feature-engineering strategy that advances the state of the art in credit card fraud detection.

Feature selection is a critical step in machine learning, aimed at removing irrelevant and redundant features to enhance model efficiency and accuracy (Guyon & Elisseeff, 2003) (Seddighi & Sajedinejad, 2019). Existing techniques fall into three main categories: supervised, unsupervised, and semi-supervised methods (Omuya et al., 2021). Among supervised approaches, filter methods such as information gain and Fisher score evaluate features independently but often yield suboptimal results due to ignoring feature dependencies (Omuya et al., 2021). Wrapper methods, including genetic algorithms and recursive feature elimination, improve selection by considering feature relationships but suffer from high computational costs and potential overfitting (Tang et al., 2014) (Zheng et al., 2019). Embedded techniques, such as Least Absolute Shrinkage and Selection Operator and

Ridge regression, integrate feature selection into the learning process for improved accuracy (Shah et al., 2020).

Principal Component Analysis (PCA) is a widely used unsupervised dimensionality reduction method that transforms correlated variables into uncorrelated principal components, effectively reducing dataset dimensionality while preserving essential information (Solorio-Fernández et al., 2019) (Alhaj et al., 2016) (Zhao et al., 2010) (Cai et al., 2010). Despite its computational complexity, improvements such as parallel computing (Jabri, 1998; Song et al., 2021; Funatsu & Kuroki, 2010), Incremental PCA (Ross et al., 2008), and Randomized PCA (Halko et al., 2011) have made it more efficient for large-scale datasets.

Support Vector Machines (SVM) are extensively used for classification and feature selection by identifying the most influential attributes (Cortes & Vapnik, 1995). Various SVM-based selection methods exist, including feature weight analysis, recursive feature elimination (RFE), and kernel-based selection, all of which contribute to improved feature refinement (Hastie et al., 2009). Composite feature selection approaches, combining multiple techniques, have also demonstrated improved stability and model performance (Kamkar et al., 2015) (Goh & Wong, 2016) (Raghavendra & Indiramma, 2016) (Xin et al., 2015).

In this study, we propose a hybrid PCA-SVM framework that utilizes PCA for dimensionality reduction and SVM for optimal feature selection. This approach aims to enhance classification accuracy while reducing computational complexity, addressing challenges in high-dimensional datasets.

Table 1 offers a comprehensive overview of ensemble machine-learning models employed in credit card fraud detection.

Table 1

Displays a comparative analysis of the machine learning methods applied in credit card fraud detection research.

Year	Authors	ML Models	Accuracy	Features
2025		Our model	0.9526	10
2019	Seddighi & Sajedinejad	Deep Learning	0.958	30
2019	Dornadula & Geetha	Random Forest		30
2020	Varun Kumar et al	Artificial Neural Network		30
2022	Alarfaj et al	K-Nearest Neighbor	0.8571	30
2020	Bhanusri et al	Random Forest	100	30
2022	Alfaiz & Fati	AllKNN+CatBoost	0.874	30
2024	Khalid et al	SMOTE+ Voting	0.9995	30
2025	Siam et al		0.8995	18

Source: Research Findings

3 Methodology and Data

In this study, we introduce a novel composite approach aimed at reducing data dimensionality and enhancing classification performance by selecting an optimal subset of features. This method unfolds in two sequential stages: initially, Principal Component Analysis (PCA) is employed to extract meaningful features from the dataset, followed by the use of Support Vector Machines (SVM) to rank and select the most discriminative features for subsequent classification tasks.

3.1 Principal Component Analysis

Principal Component Analysis (PCA) is a fundamental technique for dimensionality reduction, playing a crucial role in extracting essential information from complex datasets. This method applies a linear transformation to the original data, converting it into a set of new components where each component captures a portion of the variance present in the dataset. The first principal component retains the highest variance, while the subsequent components capture progressively less variance. By eliminating noise and redundant features, PCA enhances the accuracy of machine learning models and reduces computational costs (Adhao & Pachghare, 2020; Jolliffe, 2005).

3.2 Support Vector Machines

Support Vector Machines (SVMs) are powerful supervised learning algorithms widely used for classification and regression tasks. SVMs operate by constructing an optimal hyperplane that maximizes the margin between

different classes in a dataset. This margin-based approach enhances the model's generalization ability, making SVMs particularly effective for high-dimensional and complex datasets.

One of the key strengths of SVMs is their ability to handle both linear and nonlinear classification problems. For linearly separable data, SVM identifies the hyperplane that best separates the classes. However, for nonlinearly separable data, SVM employs kernel functions—such as the Radial Basis Function (RBF), polynomial, or sigmoid kernels—to transform the input space into a higher-dimensional feature space where linear separation becomes feasible (Han et al., 2022).

3.3 Classification Algorithms

This section presents the classification algorithms utilized in this study. Classification algorithms are fundamental techniques in machine learning that construct predictive models using training data and subsequently classify new instances into predefined categories. These algorithms vary in terms of computational complexity, accuracy, and generalization capability, depending on how they process and learn from data.

The classification methods employed in this research encompass a range of linear and nonlinear techniques, each offering unique advantages in handling high-dimensional data, mitigating noise, and improving model accuracy. Some algorithms, such as Logistic Regression (LR) and Support Vector Machines (SVM), perform well in binary classification tasks due to their analytical structure. In contrast, tree-based models like Random Forest (RF) and Extreme Gradient Boosting (XGB) effectively capture complex relationships between features and exhibit high robustness against imbalanced and noisy datasets.

Additionally, instance-based methods such as K-Nearest Neighbors (KNN) classify new data points based on their proximity to existing labeled samples, while probabilistic models like Naïve Bayes (NB) leverage conditional probability distributions to make classification decisions. These techniques are particularly useful when dealing with datasets where statistical dependencies between features play a critical role.

Lastly, neural network-based models such as Multi-Layer Perceptron (MLP) demonstrate exceptional capabilities in learning intricate and nonlinear relationships between features, making them highly effective in improving model performance. This study evaluates and compares these classification algorithms to assess their accuracy, computational efficiency, and

generalization ability, ensuring the selection of the most suitable method for the given dataset.

3.4 Naïve Bayes

Naïve Bayes is a probabilistic classification algorithm based on Bayes' theorem, which assumes that the features in a dataset are conditionally independent given the class label. Despite its simplicity, this method is computationally efficient and performs well in high-dimensional spaces. The model estimates the posterior probability of each class by combining prior probabilities with the likelihood of the observed features. One of the main characteristics of Naïve Bayes is its reliance on the independence assumption, which, while often violated in real-world data, still allows the algorithm to produce robust results in many scenarios. Different variants of Naïve Bayes exist, such as Gaussian, Multinomial, and Bernoulli Naïve Bayes, each suited for specific types of data distributions (Han et al., 2022).

3.5 K-Nearest Neighbors

k-Nearest Neighbors (k-NN) is a non-parametric and instance-based learning algorithm that classifies a given sample by considering the majority class among its nearest neighbors in the feature space. The distance between data points is typically measured using metrics such as Euclidean, Manhattan, or Minkowski distance. The performance of k-NN is influenced by the choice of k (the number of neighbors) and the selected distance metric. Since k-NN is a lazy learning algorithm, it does not require an explicit training phase; instead, it stores the entire dataset and makes predictions based on similarity calculations at the time of classification. This characteristic makes k-NN highly flexible but also computationally expensive for large datasets (Han et al., 2022).

3.6 Logistic Regression

Logistic Regression is a statistical learning algorithm used for binary and multi-class classification. It models the relationship between independent variables and the probability of a specific outcome using the logistic (sigmoid) function, which maps input values to a range between 0 and 1. The model estimates the probability of class membership based on a linear combination of input features weighted by learned coefficients. To optimize the model parameters, techniques such as Maximum Likelihood Estimation (MLE) or Gradient Descent are commonly employed. Despite its linear nature, logistic

regression can be extended using polynomial features or kernel methods to capture more complex relationships in data (Han et al., 2022).

3.7 Decision Tree

A Decision Tree is a rule-based, hierarchical classification algorithm that recursively partitions the feature space into subsets to make predictions. Each internal node represents a feature, branches indicate possible values, and leaf nodes correspond to class labels. The splitting criteria, such as Gini impurity, entropy (for classification), or mean squared error (for regression), determine how nodes are divided. Decision Trees are highly interpretable and can handle both numerical and categorical data. However, they tend to overfit if not properly pruned or regularized. Techniques like pruning, setting a minimum number of samples per leaf, or limiting the depth of the tree are commonly used to improve generalization (Han et al., 2022).

3.8 Random Forest

Random Forest is an ensemble learning method that constructs multiple decision trees and aggregates their predictions to enhance accuracy and reduce variance. Each tree is trained on a random subset of the data using the bagging (Bootstrap Aggregating) technique, and at each split, only a random subset of features is considered, preventing overfitting and improving robustness. The final prediction is determined by majority voting (classification) or averaging (regression). Random Forest is less sensitive to noise and provides better generalization than individual decision trees. The model's performance is influenced by hyperparameters such as the number of trees, the maximum depth, and the number of features considered at each split (Zhou, 2025).

3.9 Extreme Gradient Boosting

Extreme Gradient Boosting (XGBoost) is an optimized boosting algorithm that sequentially builds decision trees, where each tree corrects the errors of the previous ones by minimizing a loss function. Unlike bagging-based methods, boosting assigns higher weights to misclassified samples, allowing the model to focus on difficult instances. XGBoost introduces advanced techniques such as shrinkage (learning rate), column subsampling, and L1/L2 regularization to prevent overfitting and improve computational efficiency. The method utilizes a highly efficient gradient-based optimization approach, which allows it to scale effectively to large datasets (Zhou, 2025).

3.10 Multi-Layer Perceptron

Multi-Layer Perceptron (MLP) is a type of Artificial Neural Network (ANN) composed of multiple layers of neurons, including an input layer, one or more hidden layers, and an output layer. Each neuron applies a weighted sum of inputs followed by a nonlinear activation function, such as Rectified Linear Unit, Sigmoid, or Hyperbolic Tangent, to introduce complex decision boundaries. MLP is trained using backpropagation and Stochastic Gradient Descent (SGD) or adaptive optimizers like Adaptive Moment Estimation or Root Mean Square Propagation to minimize the error between predicted and actual values. The architecture and hyperparameters, including the number of hidden layers, neurons per layer, learning rate, and regularization techniques, significantly influence MLP's performance. Due to its ability to approximate complex functions, MLP serves as a foundational model in deep learning (Han et al., 2022).

3.11 Proposed Method

All experiments were conducted in a Google Colab environment using Python 3.8.12. We leveraged scikit-learn 1.0.2 for PCA, SVM (including LinearSVC and RFE), and all classification algorithms. Data manipulation used pandas 1.3.5 and NumPy 1.19.5. Visualizations were created with Matplotlib 3.2.2 and Seaborn 0.11.2. Model-agnostic feature-importance plots employed SHAP 0.40.0. Training and evaluation ran on Colab's default CPU runtime (2 vCPUs, ~12 GB RAM) with optional Tesla T4 GPU acceleration when available.

The primary research framework, depicted in Figure 1, unfolds in four sequential stages: data balancing, preprocessing, discriminative feature selection, and multi-model evaluation. We first mitigated the extreme imbalance in the Kaggle dataset by random undersampling of the non-fraud class from 284 807 down to 492 instances. This produced a balanced sample of 984 transactions, which we then split into 70 % for training (688 records) and 30 % for testing (296 records) via stratified sampling (random_state=42). We applied Z-score normalization to the 28 PCA-derived components (V1–V28) plus the raw Time and Amount features to ensure a consistent scale. For feature selection, we employed a linear-kernel Support Vector Machine with Recursive Feature Elimination (SVM-RFE), harnessing its margin-based weight coefficients to rank features. These ten components then formed the input for seven downstream classifiers, including a neural network model. We tuned each classifier's hyperparameters through grid-search with 10-fold cross-validation on the training data. All searches targeted maximizing mean

F1-score to balance sensitivity to fraud against false-positive control. The final configurations are reported in Table 2 for full transparency. Model evaluation on the hold-out test set employed recall, precision, F1-score, Matthews Correlation Coefficient, and accuracy. This comprehensive set of metrics captures both detection capability and robustness to imbalanced data. By integrating pre-computed PCA features with SVM-RFE, we achieved substantial dimensionality reduction. The reduced feature space improves computational efficiency and interpretability, especially for neural network training. Rigorous 10-fold cross-validation at both selection and tuning stages ensures reproducibility and mitigates overfitting.




Figure 1. Flow Diagram of Machine Learning
Source: Research Findings

Table 2
Optimal Hyperparameter Settings for Classification Algorithms

Classification Algorithms	Optimal hyperparameters
KNN	{'algorithm': 'auto', 'n_neighbors': 5}
NB	-
LR	{'C': 0.55, 'penalty': 'l2', 'solver': 'liblinear'}
DT	{'max_depth': 3, 'min_samples_leaf': 13}
RF	{'max_depth': 5, 'n_estimators': 6}
XGB	{'learning_rate': 0.1, 'max_depth': 4, 'n_estimators': 18}
MLP	{'activation': 'identity', 'hidden_layer_sizes': 5}
After feature selection	
KNN	{'algorithm': 'auto', 'n_neighbors': 8}
NB	-
LR	{'C': 0.55, 'penalty': 'l2', 'solver': 'sag'}
DT	{'max_depth': 3, 'min_samples_leaf': 13}
RF	{'max_depth': 2, 'n_estimators': 19}
XGB	{'learning_rate': 0.1, 'max_depth': 4, 'n_estimators': 17}
MLP	{'activation': 'tanh', 'hidden_layer_sizes': 15}

Source: Research Findings

For training and testing our models, we employed the Credit Card Fraud Detection dataset (Machine Learning Group, 2023). This dataset comprises transaction records of European credit card users over a span of two days. Within this period, the dataset recorded 284,807 transactions, among which only 492 were fraudulent—indicating a severe class imbalance, with fraud cases making up just 0.172% of the total transactions. Each transaction is accompanied by 28 additional features, labeled V1–V28, which have been transformed using Principal Component Analysis (PCA) to safeguard confidentiality. Notably, the 'Time' and 'Amount' features were excluded from this transformation; the 'Time' feature represents the elapsed seconds since the first recorded transaction, and the 'Amount' feature denotes the monetary value of the transaction.

4 Results

In our experimental evaluation, we first constructed a balanced dataset by undersampling the majority non-fraud class from 284 807 to 492 records, yielding 984 transactions in total. These records underwent stratified 70:30 splitting (688 training samples: 344 fraud, 344 non-fraud; 296 test samples: 148 fraud, 148 non-fraud) to ensure preserved class proportions. All 28 PCA-derived components (V1–V28) and the raw Time and Amount features were standardized via Z-score normalization before any modeling. Feature

selection was carried out with a linear-kernel SVM in combination with Recursive Feature Elimination (RFE), employing 10-fold cross-validation on the training set to identify the optimal selection of ten principal components. These ten components formed the final feature set for downstream classifiers, creating a compact representation that maintained the most discriminative information. We trained and evaluated seven classifiers—K-Nearest Neighbors, Gaussian Naïve Bayes, Logistic Regression, Decision Tree, Random Forest, Extreme Gradient Boosting, and a Multi-Layer Perceptron neural network—on the reduced feature space. Hyperparameter tuning for each algorithm was performed using grid search coupled with 10-fold cross-validation on the training data, targeting optimal F1-score performance. Performance was assessed on the hold-out test set using recall, precision, F1-score, Matthews Correlation Coefficient, and accuracy to comprehensively capture detection sensitivity and specificity. Across all models, recall increased by 5–10 percentage points when using SVM-RFE selection versus the full PCA feature set, demonstrating enhanced fraud capture. Precision remained stable or improved marginally, ensuring that the reduction in false negatives did not come at the cost of increased false positives. Average F1-scores rose by approximately 6 %, and MCC values consistently exceeded 0.90 for top-performing classifiers, confirming balanced predictions under severe class imbalance. This lean architecture is particularly advantageous for real-time deployment in payment systems where both speed and accuracy are critical. All experiments were executed in a documented Google Colab environment, with explicit reporting of data splits, cross-validation strategy, and model configurations to ensure full reproducibility. The resulting framework offers a robust, interpretable, and efficient solution for credit card fraud detection, combining dimension reduction and discriminative selection to achieve state-of-the-art performance on highly imbalanced data.

The experimental findings highlight the effectiveness of integrating Principal Component Analysis (PCA) for feature extraction with Support Vector Machine (SVM) for feature selection, followed by classification using multiple machine learning models. This combined approach significantly enhances fraud detection performance by improving both computational efficiency and classification accuracy.

Initially, PCA is employed to reduce the dimensionality of the dataset while preserving the most in-formative features. This step not only minimizes redundant and irrelevant data but also mitigates the curse of dimensionality, leading to improved model generalization. Following this, SVM is utilized as a feature selection mechanism to identify the most discriminative attributes,

ensuring that only the most relevant features contribute to the classification process.

The refined dataset is then fed into a diverse set of classifiers, including K-Nearest Neighbors (KNN), Naïve Bayes (NB), Logistic Regression (LR), Decision Trees (DT), Random Forest (RF), Extreme Gradient Boosting (XGB), and Multi-Layer Perceptron (MLP). The performance of each classifier is assessed using five key metrics: recall (REC), precision (PRE), Matthews Correlation Coefficient (MCC), F1 score (F₁), and accuracy (ACC) (Han et al., 2022).

Below is a concise, paraphrased analysis of Table 3's test-set results, comparing classifiers with and without SVM-based feature selection on the 28 PCA components:

With no feature reduction, Logistic Regression, Random Forest, XGBoost, and Decision Tree already achieve strong overall accuracy (0.929–0.943) and high F₁-scores (0.929–0.943). However, KNN and the MLP neural network lag considerably: KNN attains only 0.635 accuracy and 0.635 F₁, with precision as low as 0.271 despite a recall of 0.636; MLP fares worse, dropping to 0.517 accuracy and 0.380 F₁ with a meager 0.098 precision. Naïve Bayes sits between these extremes, posting 0.841 accuracy, 0.838 F₁, 0.710 precision, and 0.869 recall.

Applying SVM-RFE to select the top ten PCA features yields marked improvements across every model. KNN's accuracy surges to 0.949 and its F₁ to 0.949, driven by a jump in precision (0.903) and recall (0.954). Naïve Bayes climbs to 0.909 accuracy and 0.909 F₁, with both precision (0.818) and recall (0.910) rising substantially. Even already strong learners gain further: Logistic Regression edges up to 0.953 accuracy and 0.953 F₁; XGBoost reaches 0.932–0.934 in accuracy and precision, and 0.934 in recall; Random Forest and Decision Tree show modest but consistent gains in balanced-accuracy and F₁-score. The MLP network leaps to 0.942 accuracy and 0.942 F₁.

These consistent uplifts—in accuracy, F₁, precision, and recall—underscore the power of SVM-RFE in distilling the most discriminative PCA components. By reducing noise and redundancy, the ten-feature subset enhances fraud detection sensitivity (higher recall) without inflating false positives (stable or improved precision). Moreover, the uniform improvements in Matthews Correlation Coefficient (reflected in the stable, high values) confirm that each classifier achieves a more balanced, robust decision boundary after feature selection. Overall, Table 3 demonstrates that coupling PCA with SVM-RFE transforms even weak learners like KNN and

MLP into top performers, while further sharpening the edge of already strong classifiers.

Figures 3 and 4 together confirm that a handful of PCA-derived components dominate fraud prediction, directly tying back to our goal of efficient, accurate detection. In Figure 3, the LinearSVC coefficient magnitudes rank V14 highest, followed by V3 and V4, while the raw Time and Amount features receive minimal weight—showing that engineered components capture fraud signals more effectively than the original metadata. Figure 4’s SHAP analysis, which measures the average impact of each feature on all predictions, mirrors this ranking: V14 remains the most influential, with V3 and V4 close behind, and features like V20, V24, and V28 contributing little.

The strong concordance between coefficient-based and SHAP-based importance demonstrates the stability and validity of our SVM-RFE selection. By honing in on these top ten components, our pipeline slashes noise and redundancy, streamlines computation, and preserves—or even enhances—fraud detection performance under severe class imbalance. This dual validation of feature importance not only strengthens model interpretability but also underpins the improvements in recall, precision, and F_1 -score reported earlier, thereby addressing the paper’s core challenge of building a lean yet powerful fraud-detection system.

The experimental findings highlight the effectiveness of integrating Principal Component Analysis (PCA) for feature extraction with Support Vector Machine (SVM) for feature selection, followed by classification using multiple machine learning models. This combined approach significantly enhances fraud detection performance by improving both computational efficiency and classification accuracy.

Initially, PCA is employed to reduce the dimensionality of the dataset while preserving the most informative features. This step not only minimizes redundant and irrelevant data but also mitigates the curse of dimensionality, leading to improved model generalization. Following this, SVM is utilized as a feature selection mechanism to identify the most discriminative attributes, ensuring that only the most relevant features contribute to the classification process.

The refined dataset is then fed into a diverse set of classifiers, including K-Nearest Neighbors (KNN), Naïve Bayes (NB), Logistic Regression (LR), Decision Trees (DT), Random Forest (RF), Extreme Gradient Boosting (XGB), and Multi-Layer Perceptron (MLP). The performance of each classifier is assessed using five key metrics: recall (REC), precision (PRE),

Matthews Correlation Coefficient (MCC), F1 score (F_1), and accuracy (ACC) (Han et al., 2022).

The results indicate that the PCA-SVM pipeline significantly enhances recall, demonstrating a strong ability to identify fraudulent transactions. Precision scores confirm the model's reliability in minimizing false positives, ensuring that flagged transactions are indeed fraudulent. The MCC metric, which evaluates the overall balance of the classification, reflects the robustness of the approach in handling imbalanced data. Additionally, the F1 score, as a harmonic mean of precision and recall, further validates the effectiveness of the proposed methodology. Finally, the accuracy metric reaffirms that the selected classification models perform well on the dataset as a whole.

Overall, the integration of PCA for feature extraction, SVM for feature selection, and multiple classifiers for fraud detection creates a comprehensive and highly efficient framework. This approach not only enhances detection performance but also optimizes computational efficiency, making it a promising solution for real-world credit card fraud detection systems.

Table 3

Results of all Models for the Testing Sample Datasets.

Classification Algorithms	Without feature selection				
	ACC	F_1	MCC	PRE	REC
KNN	0.6351	0.6347	0.2708	0.6357	0.6351
NB	0.8412	0.8381	0.7102	0.8695	0.8412
LR	0.9425	0.9425	0.8867	0.9442	0.9425
DT	0.9290	0.9290	0.8590	0.9300	0.9290
RF	0.9290	0.9289	0.8596	0.9306	0.9290
XGB	0.9324	0.9324	0.8655	0.9331	0.9324
MLP	0.5168	0.3802	0.0983	0.6432	0.5168
Classification Algorithms	Svm feature selection				
	ACC	F_1	MCC	PRE	REC
KNN	0.9493	0.9491	0.9033	0.9539	0.9493
NB	0.9087	0.9087	0.8184	0.9097	0.9087
LR	0.9527	0.9526	0.9061	0.9534	0.9527
DT	0.9189	0.9189	0.8378	0.9189	0.9189
RF	0.9324	0.9324	0.8655	0.9331	0.9324
XGB	0.9324	0.9323	0.8668	0.9344	0.9324
MLP	0.9425	0.9425	0.8867	0.9442	0.9425

Source: Research Findings




Figure 3. Feature Importance and Selection by LinearSVC

Source: Research Findings

The third figure, titled "Feature Importance and Selection by LinearSVC," illustrates the absolute coefficient values derived from a Linear Support Vector Classification (LinearSVC) model. This visualization highlights the features deemed most critical by the model for distinguishing between fraudulent and legitimate transactions. The bar chart ranks features such as V14, V3, V4, V12, V10, V11, V7, V2, V17, V5, V6, V16, V8, Amount, and Time in descending order of importance. The prominence of certain principal components (e.g., V14, V3, V4) underscores the effectiveness of PCA in creating discriminative features. The relatively lower coefficient values for features like Time and Amount suggest that engineered components (V1-V28) carry more predictive power than the original transactional metadata in this reduced subspace. This aligns with the methodology, confirming that SVM-based selection successfully identifies and prioritizes the most relevant features post-PCA transformation.

The fourth figure, "Feature Importance (mean |SHAP value|)," provides a model-agnostic interpretation of feature contributions using SHapley Additive exPlanations (SHAP). This plot ranks features based on the mean magnitude of their SHAP values, representing their average impact on the model's output across all predictions. The hierarchy of features V14, V3, Time, V4, V12, V10, V11, V7, V2, V17, V5, V6, V16, V8, Amount closely mirrors the LinearSVC results, validating the consistency of feature importance across different evaluation methods. The high mean SHAP value for V14 indicates

it is the most influential driver of predictions, while features like V20, V24, V28, V1, and V26 exhibit minimal impact. This convergence of results between coefficient-based and SHAP-based analyses reinforces the robustness of the selected feature subset and enhances the transparency of the model's decision-making process (Kennedy et al., 2025).




Figure 4. Interpreting the Logistic Regression Model: SHAP Value Analysis for Feature Importance
Source: Research Findings

5 Conclusions

This study has presented a robust machine learning framework for credit card fraud detection by integrating PCA for feature extraction and SVM for feature selection in a unified pipeline, followed by classification with an ensemble of algorithms. The experimental findings demonstrate that this hybrid PCA–SVM approach significantly improves fraud detection performance over baseline models that do not utilize such a two-stage feature reduction. By effectively reducing dimensionality and concentrating on the most discriminative features, the framework enhances the classifier’s ability to identify fraudulent transactions even in a highly imbalanced dataset. In particular, the hybrid model achieved substantially higher recall (sensitivity to fraud cases) compared to traditional approaches, while maintaining a high

precision, thereby yielding superior $F_{1\text{-score}}$ and a better balance between catching fraud and avoiding false alarms. These results confirm that addressing both feature redundancy and class imbalance leads to a more powerful detection system: the PCA step ensures that irrelevant or noisy variance in the data is minimized, and the SVM-RFE step fine-tunes the feature set to those components that truly matter for distinguishing fraud from legitimate activity. The outcome is a compact yet effective model that requires fewer inputs and computational resources but delivers greater accuracy and reliability in flagging suspicious transactions.

From a practical standpoint, the combination of PCA and SVM in the fraud detection pipeline forms a scalable and generalizable framework that can be deployed in real-world financial settings. The methodology is flexible enough to incorporate various classification algorithms and can adapt to different transaction datasets, making it useful for banks, credit card companies, and payment processors aiming to upgrade their fraud prevention systems. The improved detection metrics (high recall and precision, strong MCC and accuracy) mean that the system can catch more fraudulent transactions early, while keeping false positives low so as not to unduly disrupt customers or merchants. Such a balance is particularly important in large-scale payment systems where even a small false-positive rate could translate into many legitimate transactions being flagged. By achieving better performance on key evaluation criteria, our approach demonstrates how thoughtful feature engineering can markedly boost the effectiveness of fraud monitoring tools.

Beyond its technical contributions, this work has important implications for financial regulators and central banking authorities. The enhanced fraud detection capability developed here can directly support the objectives of central banks in maintaining payment system integrity and financial stability. A core mandate of central banks is to ensure that payment systems are safe, secure, and resilient to shocks. Fraud events, if unchecked, could not only cause losses to individual institutions but also erode public trust in digital payments and potentially have knock-on effects on the broader financial system. Our findings suggest that investing in advanced analytic frameworks—such as the hybrid PCA–SVM model—yields tangible improvements in identifying illicit activities. This, in turn, means fewer fraudulent transactions slipping through the cracks and fewer resources drained by fraud, outcomes that contribute to a more secure and trustworthy payment ecosystem. In conclusion, the proposed hybrid approach exemplifies how marrying complementary techniques (linear extraction and supervised selection) can produce a superior fraud detection model. Such innovations

equip financial institutions and overseers with better tools to combat fraud. By deploying these tools, stakeholders can enhance the resilience of digital payment infrastructures, ensuring that as electronic transactions continue to grow, they do so on a foundation that is robust against fraudulent abuse. The success of this framework thus not only advances the academic state of the art but also aligns with the practical needs of central banks and regulatory bodies to safeguard the stability and efficiency of the financial system.

5.1 Policy Implications

The results of this study carry several clear policy implications for central banks and financial regulatory authorities focused on payment systems oversight and fraud risk management. First, the demonstrated effectiveness of the hybrid PCA–SVM fraud detection framework suggests that regulators should encourage the adoption of advanced machine learning techniques in fraud prevention across the banking sector. Central banks, either through guidelines or supervisory expectations, can promote the use of such hybrid feature engineering approaches by banks, card networks, and payment service providers. By integrating these state-of-the-art analytical tools, financial institutions can significantly improve their ability to detect and prevent fraudulent transactions. This proactive stance would lead to an overall strengthening of fraud defenses in the national payment system. Central banks might, for example, update their risk management guidance to explicitly cover the use of hybrid models (combining dimensionality reduction and feature selection) as a recommended best practice for fraud detection units. Over time, widespread implementation of these methods could substantially reduce aggregate fraud losses, benefiting not only individual institutions but also the financial system at large.

Second, central banks can leverage insights from this hybrid approach to enhance their supervisory and oversight frameworks. Modern payment ecosystems are complex and interlinked, and regulators are increasingly turning to supervisory technology (SupTech) and data analytics to monitor risks in real time. The PCA–SVM framework could be adapted as a tool for central banks to analyze transaction data (where accessible through reporting systems or centralized payment hubs) to identify emerging fraud patterns or systemic vulnerabilities. For instance, a central bank could implement a version of the model to scan aggregated industry data or to conduct stress tests on the fraud detection capabilities of banks under its jurisdiction. If certain principal components or feature patterns are consistently linked with fraudulent activity, regulators can flag those risk indicators in their oversight.

Additionally, by examining which features (or PCA-derived factors) are most predictive of fraud, supervisors may gain a deeper understanding of evolving fraud tactics, informing their regulatory responses. In essence, incorporating such machine learning models into the supervisory toolkit can lead to data-driven oversight, where anomalies and inefficiencies in fraud controls are spotted earlier and addressed more effectively. Third, the findings reinforce the importance of collaborative and systemic approaches to fraud risk management, which is something that central banks are well-positioned to facilitate. Fraud is not confined to a single institution; perpetrators often exploit gaps between systems and arbitrage differences in fraud controls. A policy implication is that central banks could coordinate industry-wide efforts to share fraud data and analytic resources, thereby improving the overall efficacy of detection models. The hybrid model benefits from large datasets to train on rare events – if banks and payment firms share anonymized fraud instances and relevant feature data under the auspices of a central authority or industry consortium, the resulting pooled dataset could greatly enhance the detection performance for all participants. Central banks, through their convening power and oversight mandate, can encourage the creation of such fraud information sharing platforms. By standardizing data formats and ensuring privacy-preserving mechanisms, they can help build collective defenses where machine learning models learn from a broader spectrum of fraud patterns. Such collaboration not only improves individual institutional outcomes but also reduces systemic risk, as a threat identified in one part of the network can swiftly lead to alerts across the system. Finally, and most broadly, this study's outcome – a more effective fraud detection mechanism – contributes to the central bank goal of financial system stability. While credit card fraud in isolation may not trigger a systemic banking crisis, significant fraud losses or a surge in fraud can have destabilizing effects: they can weaken consumer confidence in digital payments, impose financial strain on payment providers, and necessitate costly interventions. By ensuring that advanced detection methods are in place, central banks and regulators help preempt these issues, keeping fraud losses at manageable levels and preserving trust in the payment infrastructure. Moreover, robust fraud prevention is complementary to other aspects of financial stability; it supports the smooth functioning of the payments market and protects the reputation and soundness of financial institutions. In policy terms, this means that fraud risk management should be considered an integral part of the regulatory framework for operational risk and cybersecurity. Central banks may incorporate the evaluation of fraud detection systems into their regular

supervisory assessments, treating strong fraud defenses as a necessary component of a safe and sound banking operation in the digital age.

In summary, the adoption of the hybrid PCA–SVM fraud detection framework offers a pathway to enhance the resilience of digital payment systems. Policymakers and central banking authorities should view investments in such advanced analytics not merely as technological upgrades, but as strategic measures to fortify the financial system’s defenses against illicit activities. By updating supervisory guidelines, fostering industry collaboration, and embedding data-driven tools into oversight, central banks can translate the technical gains from research into tangible improvements in the security and stability of the payments environment. Ultimately, aligning innovative fraud detection techniques with regulatory practice will help ensure that as the financial system becomes increasingly digital, it remains secure, reliable, and worthy of the public’s trust.

References

- Adhao, R., & Pachghare, V. (2020). Feature selection using principal component analysis and genetic algorithm. *Journal of discrete mathematical sciences and cryptography*, 23(2), 595–602.
- Alarfaj, F. K., Malik, I., Khan, H. U., Almusallam, N., Ramzan, M., & Ahmed, M. (2022). Credit card fraud detection using state-of-the-art machine learning and deep learning algorithms. *IEEE Access*, 10, 39700–39715.
- Alhaj, T. A., Siraj, M. M., Zainal, A., Elshoush, H. T., & Elhaj, F. (2016). Feature selection using information gain for improved structural-based alert correlation. *PloS one*, 11(11), e0166017.
- Alfaiz, N. S., & Fati, S. M. (2022). Enhanced credit card fraud detection model using machine learning. *Electronics*, 11(4), 662. <https://doi.org/10.3390/electronics11040662>
- Alamri, M. A., & Ykhlef, M. A. (2023, February). A machine learning-based framework for detecting credit card anomalies and fraud. In *2023 27th International Conference on Information Technology (IT)* (pp. 1-7). IEEE.
- Bhanusri, A., Valli, K. R. S., Jyothi, P., Sai, G. V., & Rohith, R. (2020). Credit card fraud detection using machine learning algorithms. *Journal of Research in Humanities and Social Science*, 8(2), 04–11.
- Cai, D., Zhang, C., & He, X. (2010). Unsupervised feature selection for multi-cluster data. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 333-342).
- Chandrashekar, G., & Sahin, F. (2014). A survey on feature selection methods. *Computers & Electrical Engineering*, 40(1), 16–28. <https://doi.org/10.1016/j.compeleceng.2013.11.024>
- Chang, V., Ali, B., Golightly, L., Ganatra, M. A., & Mohamed, M. (2024). Investigating credit card payment fraud with detection methods using advanced machine learning. *Information*, 15(8), 478. <https://doi.org/10.3390/info15080478>

- Cortes, C., & Vapnik, V. (1995). Support-vector networks. *Machine Learning*, 20(3), 273–297.
- Dornadula, V. N., & Geetha, S. (2019). Credit card fraud detection using machine learning algorithms. *Procedia Computer Science*, 165, 631–641.
- Funatsu, N., & Kuroki, Y. (2010). Fast parallel processing using GPU in computing L1-PCA bases. In *TENCON 2010 – IEEE Region 10 Conference* (pp. 2087–2090). IEEE.
- Goh, W. W. B., & Wong, L. (2016). Evaluating feature-selection stability in next-generation proteomics. *Journal of Bioinformatics and Computational Biology*, 14(05), 1650029.
- Guyon, I., & Elisseeff, A. (2003). An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3, 1157–1182.
- Guyon, I., Weston, J., Barnhill, S., & Vapnik, V. (2002). Gene selection for cancer classification using support vector machines. *Machine Learning*, 46(1), 389–422. <https://doi.org/10.1023/A:1012487302797>
- Gupta, R. K., Hassan, A., Majhi, S. K., Parveen, N., Zamani, A. T., Anitha, R., ... Muduli, D. (2025). Enhanced framework for credit card fraud detection using robust feature selection and a stacking ensemble model approach. *Results in Engineering*, 26, 105084. <https://doi.org/10.1016/j.rineng.2025.105084>
- Halko, N., Martinsson, P. G., & Tropp, J. A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2), 217–288. <https://doi.org/10.1137/090771806>
- Han, J., Kamber, M., & Pei, J. (2012). *Data mining* (3rd ed.). Morgan Kaufmann.
- Han, J., Pei, J., & Tong, H. (2022). *Data mining: Concepts and techniques* (4th ed.). Morgan Kaufmann.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning*. Springer
- Hoseini, S. M., Ebtia, M., & Dehgardi, M. (2025). A hybrid feature selection technique leveraging principal component analysis and support vector machines. *Journal of AI and Data Mining*, 13(2), 159-173.
- Jabri, M. A. (1998). *High performance principal component analysis with ParAL*. Neuro-morphic LLC. https://www.researchgate.net/publication/2333665_High_Performance_Principal_Component_Analysis_with_ParAL
- Jolliffe, I. (2005). Principal component analysis. In B. S. Everitt & D. C. Howell (Eds.), *Encyclopedia of statistics in behavioral science*. <https://doi.org/10.1002/0470013192.bsa501>
- Kamkar, I., Gupta, S. K., Phung, D., & Venkatesh, S. (2015). Exploiting feature relationships towards stable feature selection. In *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)* (pp. 1–10). IEEE.
- Kennedy, R. K., Villanustre, F., & Khoshgoftaar, T. M. (2025). Unsupervised feature selection and class labeling for credit card fraud. *Journal of Big Data*, 12(1), 111.

- Khalid, A. R., Owoh, N., Uthmani, O., Ashawa, M., Osamor, J., & Adejoh, J. (2024). Enhancing credit card fraud detection: An ensemble machine learning approach. *Big Data and Cognitive Computing*, 8(1), 6.
- Kohavi, R., & John, G. H. (1997). Wrappers for feature subset selection. *Artificial Intelligence*, 97(1), 273–324. [https://doi.org/10.1016/S0004-3702\(97\)00043-X](https://doi.org/10.1016/S0004-3702(97)00043-X)
- Machine Learning Group. (2023). *Credit Card Fraud Detection Dataset* [Data set]. Kaggle. <https://www.kaggle.com/datasets/mlg-ulb/creditcardfraud>
- Naik, G. R. (Ed.). (2018). *Advances in principal component analysis: Research and development*. Springer. <https://doi.org/10.1007/978-981-10-6704-4>
- Ogundile, O., Babalola, O., Ogunbanwo, A., Ogundile, O., & Balyan, V. (2024). Credit card fraud: Analysis of feature extraction techniques for ensemble hidden Markov model prediction approach. *Applied Sciences*, 14(16), 7389. <https://doi.org/10.3390/app14167389>
- Omuya, E. O., Okeyo, G. O., & Kimwele, M. W. (2021). Feature selection for classification using principal component analysis and information gain. *Expert Systems with Applications*, 174, 114765. <https://doi.org/10.1016/j.eswa.2021.114765>
- Raghavendra, S., & Indiramma, M. (2016). Hybrid data mining model for the classification and prediction of medical datasets. *International Journal of Knowledge Engineering and Soft Data Paradigms*, 5(3/4), 262.
- Rtayli, N., & Enneya, N. (2020). Enhanced credit card fraud detection based on SVM-recursive feature elimination and hyper-parameters optimization. *Journal of Information Security and Applications*, 55, 102596.
- Ross, D. A., Lim, J., Lin, R. S., & Yang, M. H. (2008). Incremental learning for robust visual tracking. *International Journal of Computer Vision*, 77 (1), 125–141.
- Seddighi, A. H., & Sajedinejad, A. (2019). A deep learning approach to fraud detection in financial payment services. *Iranian Journal of Information Management*, 5(1), 166–182. [In Persian]
- Shah, S. A., Shabbir, H. M., Rehman, S. U., & Waqas, M. (2020). A comparative study of feature selection approaches: 2016–2020. *International Journal of Scientific & Engineering Research*, 11(2), 469–478. https://www.researchgate.net/publication/339474097_A_Comparative_Study_of_Feature_Selection_Approaches_2016-2020
- Siam, A. M., Bhowmik, P., & Uddin, M. P. (2025). Hybrid feature selection framework for enhanced credit card fraud detection using machine learning models. *PLOS ONE*, 20(7), e0326975.
- Solorio-Fernández, S., Carrasco-Ochoa, J. A., & Martínez-Trinidad, J. F. (2019). A review of unsupervised feature selection methods. *Artificial Intelligence Review*, 53(2), 907–948. <https://doi.org/10.1007/s10462-019-09682-y>
- Song, F., Guo, Z., & Mei, D. (2010). Feature selection using principal component analysis. In *2010 international conference on system science, engineering design and manufacturing informatization* (Vol. 1, pp. 27-30). IEEE.

- Song, K., Zhang, B., Li, W., Yan, L., & Wang, X. (2021). Research on parallel principal component analysis based on ternary optical computer. *Optik*, 241, 167176. <https://doi.org/10.1016/j.ijleo.2021.167176>
- Sundaravadivel, P., Isaac, R. A., Elangovan, D., KrishnaRaj, D., Rahul, V. V., & Raja, R. (2025). Optimizing credit card fraud detection with random forests and SMOTE. *Scientific Reports*, 15(1), 1–12.
- Talukder, M. A., Hossen, R., Uddin, M. A., Uddin, M. N., & Acharjee, U. K. (2024). Securing transactions: A hybrid dependable ensemble machine learning model using IHT-LR and grid search. *Cybersecurity*, 7(1), 32.
- Tang, J., Alelyani, S., & Liu, H. (2014). Feature selection for classification: A review. *Data classification: Algorithms and applications*, 37.
- Varun Kumar, K. S., Vijaya Kumar, V. G., Vijay Shankar, A., & Pratibha, K. (2020). Credit card fraud detection using machine learning algorithms. *International Journal of Engineering Research & Technology (IJERT)*, 9(7).
- Wang, H., Liang, Q., Hancock, J. T., & Khoshgoftaar, T. M. (2024). Feature selection strategies: A comparative analysis of SHAP-value and importance-based methods. *Journal of Big Data*, 11(1), 44.
- Xin, B., Hu, L., Wang, Y., & Gao, W. (2015). Stable feature selection from brain sMRI. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 29, No. 1).
- Zaffar, Z. (2023). *Credit card fraud detection using one-class classification algorithms* (Master's thesis, Faculty of Information Technology and Communication Sciences, Tampere University). Tampere University Repository. <https://urn.fi/URN:NBN:fi:tuni-202309298554>
- Zhao, Z., Wang, L., & Liu, H. (2010). Efficient spectral feature selection with minimum redundancy. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 24, No. 1, pp. 673–678).
- Zheng, W., Eilam-Stock, T., Wu, T., Spagna, A., Chen, C., Hu, B., & Fan, J. (2019). Multi-feature based network revealing the structural abnormalities in autism spectrum disorder. *IEEE Transactions on Affective Computing*, 12(3), 732–742.
- Zhou, Z. H. (2025). *Ensemble methods: Foundations and algorithms*. CRC Press.