



Advanced Customer Behavior Forecasting for Retail and Financial Decision-Making Using Physics-Based Intelligence

Ebrahim A. Mattar^{1,*} S. K. Towfek^{2,3}

¹ College of Engineering, University of Bahrain, Bahrain

² Computer Science and Intelligent Systems Research Center, Blacksburg 24060, Virginia, USA

³ Jadara Research Center, Jadara University, Irbid 21110, Jordan

Emails: ebmattar@uob.edu.bh · sktowfek@jcsis.org

Received: May 31, 2025 Revised: July 20, 2025 Accepted: September 14, 2025 ★ Corresponding author

ABSTRACT

Accurate prediction of customer behavior remains a core methodological and operational challenge in retail economics and financial decision-making, particularly as institutions increasingly depend on data-driven forecasting systems to improve credit risk assessment, refine customer segmentation, and deliver targeted financial services in competitive and rapidly changing markets. In practice, the economic value of customer behavior prediction lies in its direct connection to profit maximization, loss minimization, and resource allocation efficiency: retailers seek to anticipate spending tendencies and product affinities to reduce marketing waste and optimize inventory, while financial institutions aim to infer creditworthiness and repayment capacity to reduce default exposure and enhance portfolio stability. Despite the demonstrated advantages of data-driven approaches, the predictive performance of advanced learning systems in such contexts is frequently constrained by the dual challenge of high-dimensional, heterogeneous feature spaces and the sensitivity of model outcomes to hyperparameter choices, often resulting in limited generalization, unstable convergence, or performance degradation when applied to unseen customer groups. To address these constraints, this study develops an integrated optimization framework that couples a high-capacity predictive model with a physics-inspired search mechanism, namely the Kirchhoff's Law Algorithm (KLA), and employs it for automated hyperparameter optimization of an End-to-End Attention Long Short-Term Memory model (EALSTM), thereby reducing reliance on manual tuning and improving model reliability under financially meaningful data complexity. In addition to introducing the KLA-driven optimization pipeline, the study conducts a rigorous comparative evaluation against established state-of-the-art metaheuristic optimizers, including Particle Swarm Optimizer (PSO), Biogeography Based Optimizer (BBO), Whale Optimization Algorithm (WOA), Bat Algorithm (BA), Artificial Protozoa Optimizer (APO), Genetic Algorithm (GA), and Stochastic Fractal Search (SFS), enabling a systematic assessment of how different search dynamics influence predictive quality in customer analytics applications. Experimental evaluation is performed using an enhanced customer dataset that integrates demographic descriptors, behavioral spending indicators, and financially meaningful constructs—thereby better reflecting real-world decision environments where customer profiling depends on both consumption behavior and financial capacity—and the results demonstrate that the KLA + EALSTM configuration consistently achieves the strongest predictive performance across the full suite of regression metrics. Specifically, KLA + EALSTM attains an MSE of 3.60×10^{-7} , an RMSE of 0.00728, an MAE of 0.000372, an MBE of 0.000091, a correlation coefficient (r) of 0.972, a coefficient of determination (R^2) of 0.969, an RRMSE of 0.095, an NSE of 0.971, and a Willmott Index of 0.977. These results indicate extremely low error magnitude, minimal systematic bias, strong explanatory power, and high agreement between predicted and observed outcomes, representing a substantial improvement over the unoptimized EALSTM baseline. From an economic and financial viewpoint, these gains are practically consequential because they strengthen the reliability of predictive decision-support systems used for credit scoring, personalized marketing, customer value assessment, and financially efficient resource allocation, where even small prediction errors can translate into measurable cost, risk, or revenue impacts. Overall, the findings provide strong empirical support for physics-inspired, non-parametric optimization as a robust mechanism for improving predictive accuracy, stability, and generalization in customer analytics, and they position KLA-based optimization as a scalable and methodologically efficient solution for next-generation retail and financial analytics systems operating under high-dimensional behavioral and financial data conditions.

Keywords: Customer behavior analytics ■ Financial prediction ■ Deep learning optimization ■ Metaheuristic algorithms ■ Retail and financial decision-making

1. INTRODUCTION

The analysis of customer behavior constitutes a foundational pillar in both retail economics and financial systems, as it directly informs strategic planning, operational efficiency, and risk management decisions [1, 2, 3]. In modern retail environments, understanding how customers interact with products, respond to pricing strategies, and develop long-term purchasing habits enables firms to design data-driven

marketing strategies, optimize inventory allocation, and enhance customer loyalty programs [4, 5, 6]. Such insights extend far beyond short-term sales forecasting and include the identification of high-value customers, the anticipation of customer churn, and the segmentation of heterogeneous customer bases into economically meaningful groups [7, 8]. Effective customer segmentation allows firms to personalize promotional campaigns, tailor product offerings, and improve

overall customer satisfaction, thereby strengthening competitive advantage and long-term profitability.

In parallel, customer behavior analysis plays an equally critical role in the financial sector, where it underpins credit risk assessment, loan approval processes, and the development of personalized financial products [4, 9]. Financial institutions increasingly rely on behavioral indicators—such as spending patterns, income levels, savings capacity, and repayment history—to evaluate customer creditworthiness and manage portfolio risk more effectively [10, 11, 12]. Accurate behavioral insights enable banks and financial service providers to design differentiated lending products, reduce default risk, and comply with regulatory requirements related to responsible lending. In both retail and finance, customer behavior analysis therefore serves as a central mechanism through which organizations align economic objectives with customer-centric strategies.

In recent years, the rapid growth of data availability and computational capacity has fundamentally transformed customer behavior analysis through the adoption of machine learning (ML) techniques. Traditional approaches to customer analytics often rely on rule-based systems, linear statistical models, or heuristic assumptions derived from limited samples [13, 6, 8]. While these methods provide interpretability and simplicity, they are frequently inadequate for capturing the complex, nonlinear, and high-dimensional relationships that characterize modern customer data. Machine learning models, by contrast, are capable of processing large-scale datasets and autonomously learning intricate patterns that may not be evident through conventional analytical techniques. This capability has made ML an indispensable tool for predictive modeling in both retail and financial applications.

In the financial domain, machine learning models have been successfully applied to tasks such as credit score prediction, default risk estimation, and customer lifetime value analysis. These predictive capabilities allow institutions to anticipate adverse outcomes, adjust pricing strategies, and allocate capital more efficiently. In retail contexts, ML-driven models enable the prediction of future spending behavior, product affinity, and customer churn, supporting more effective targeting and personalization strategies. Accurate behavioral forecasting not only enhances customer engagement but also improves operational efficiency by ensuring that marketing budgets and financial resources are allocated to customers with the highest expected return.

The importance of accurate forecasting is particularly pronounced in customer segmentation and credit risk assessment. In marketing analytics, precise predictions enable firms to focus their efforts on customers who are most likely to respond positively to specific campaigns, thereby increasing conversion rates and reducing wasted expenditure. In financial analytics, accurate creditworthiness prediction is essential for minimizing default rates, maintaining portfolio stability, and safeguarding institutional profitability. As financial markets become increasingly competitive and customer expectations continue to rise, institutions must rely on robust predictive models to design personalized products and services that align with individual customer needs. Similarly, in retail operations, accurate demand and behavior forecasts support inventory op-

timization, assortment planning, and promotional scheduling, ultimately contributing to improved supply chain efficiency and reduced operational risk.

Machine learning has fundamentally reshaped the forecasting landscape by enabling automated pattern discovery and adaptive learning from data. Unlike traditional statistical methods, ML algorithms do not require explicit model specifications or predefined functional forms. Instead, they infer relationships directly from data, allowing them to capture complex interactions between demographic, behavioral, and financial variables. Supervised learning techniques leverage labeled historical data to predict future customer outcomes, while unsupervised learning approaches uncover latent structures and groupings within customer populations. These capabilities have made ML models particularly effective in customer analytics, where behavioral heterogeneity and nonlinear dependencies are the norm rather than the exception.

Despite these advances, predicting customer behavior using machine learning presents several methodological and practical challenges. One of the most prominent is **high dimensionality**. As customer datasets expand to include a wide range of demographic attributes, financial indicators, and behavioral features—such as income, spending score, loyalty duration, and product preferences—the dimensionality of the input space increases substantially. High-dimensional data not only imposes greater computational demands but also increases the risk of overfitting, where models learn spurious patterns that fail to generalize to new customers. In economic and financial applications, such overfitting can lead to unreliable predictions and misguided decision-making. Consequently, effective feature selection becomes a critical step in reducing dimensionality, improving model robustness, and enhancing interpretability.

Closely related to high dimensionality is the challenge of **feature redundancy**. Many customer attributes are inherently correlated; for instance, income, estimated savings, and spending behavior often convey overlapping information. Including highly correlated features can inflate model complexity without improving predictive performance, while also exacerbating multicollinearity issues that hinder interpretability and stability. Redundant features increase computational burden and may distort the learned relationships within the model. Addressing feature redundancy through systematic feature selection is therefore essential for constructing efficient and economically meaningful predictive models.

Another significant challenge arises from the **sensitivity of hyperparameters**. Machine learning and deep learning models depend on numerous hyperparameters—such as learning rates, network depth, regularization coefficients, and batch sizes—that strongly influence convergence behavior and predictive accuracy. Identifying optimal hyperparameter configurations is a nontrivial task, particularly in complex models applied to large customer datasets. Conventional tuning approaches, such as grid search or random search, are often computationally expensive and may fail to explore the search space effectively. Moreover, improper hyperparameter selection can lead to overfitting or underfitting, undermining the model's ability to generalize to unseen data.

Ensuring robust **generalization** remains a persistent concern in customer behavior prediction. Models that perform well

on historical data may degrade when applied to new customer populations or evolving economic conditions. This issue is especially critical in retail and financial environments, where customer preferences, market dynamics, and macroeconomic factors change over time. Techniques such as cross-validation, regularization, and robust optimization are therefore required to balance model flexibility with stability and to ensure reliable performance in real-world deployment scenarios.

Against this backdrop, the objective of this study is to systematically evaluate and compare multiple machine learning models for customer behavior prediction, with particular emphasis on forecasting **credit scores**, **preferred shopping categories**, and **savings**. The study examines the predictive performance of several advanced models, including EALSTM, FTLTM, BiLSTM, and GNN, each of which embodies distinct modeling paradigms suited to complex customer data. In addition, the study investigates the role of metaheuristic algorithms, such as PSO, BBO, WOA, and GA, in enhancing model performance through systematic feature selection and hyperparameter tuning. These optimization strategies aim to improve predictive accuracy, reduce computational overhead, and address the challenges associated with high-dimensional and redundant feature spaces.

Furthermore, this study explores the synergistic effect of combining feature selection with hyperparameter optimization. By jointly optimizing input representations and model configurations, the proposed framework seeks to achieve superior generalization performance, mitigate overfitting, and enhance computational efficiency. This integrated approach is particularly relevant for economically sensitive applications, where model reliability and scalability are paramount.

The contributions of this research extend beyond empirical performance evaluation. By integrating advanced optimization techniques with machine learning models, the study provides actionable insights into the design of predictive systems for retail and financial decision-making. The findings offer guidance on selecting appropriate modeling and optimization strategies for practical deployment and highlight the economic value of optimized predictive analytics in customer-centric environments.

The remainder of this paper is organized as follows. Section 2 presents a comprehensive literature review on customer behavior prediction, customer segmentation, and the application of machine learning and deep learning techniques in retail and financial analytics. Section 3 describes the materials and methods, including dataset characteristics, preprocessing procedures, and model architectures. Section 4 reports the empirical results, detailing model performance before and after optimization. Section 5 discusses the findings in relation to existing studies and outlines key limitations. Finally, Section 6 concludes the paper and identifies directions for future research.

2. LITERATURE REVIEW

The rapid digitalization of business processes and the exponential growth of customer-related data have positioned machine learning (ML) and deep learning (DL) as foundational technologies for customer analytics. Across domains such as customer churn prediction, customer lifetime value

(CLV) estimation, customer experience (CX) management, demand forecasting, and customer relationship management (CRM), these techniques enable organizations to move beyond descriptive analytics toward predictive and prescriptive decision-making. The following literature review synthesizes recent and foundational contributions, emphasizing methodological trends, empirical findings, and persistent research gaps.

Customer churn prediction has emerged as one of the most extensively studied applications of ML and DL due to its direct impact on revenue and long-term business sustainability. Recent review research synthesizes the principal modeling stages, feature categories, algorithms, and evaluation practices used in customer churn prediction [11]. The findings show that ensemble-based ML models, particularly boosting and tree-based techniques, are highly competitive on structured churn datasets. Concurrently, DL models—including Long Short-Term Memory (LSTM) networks and Convolutional Neural Networks (CNNs)—are increasingly applied to sequential and high-dimensional data. Despite performance gains, the literature identifies unresolved challenges, notably class imbalance, changing customer behavior, limited interpretability, and the insufficient adoption of profit-oriented evaluation metrics. Explainable and adaptive learning frameworks therefore remain important directions for practical deployment.

Closely related to churn prediction is the estimation of customer lifetime value, which plays a central role in strategic marketing and resource allocation. One study focuses on predictive modeling techniques for forecasting consumer behavior and CLV in e-commerce contexts, highlighting ML and DL as essential tools for competitive advantage [14]. By integrating historical transaction data with customer segmentation and predictive analytics, the proposed framework enables businesses to target high-value customers more effectively. The work emphasizes that dynamic adaptation of advertising strategies based on predicted CLV can significantly enhance customer retention and optimize marketing expenditures.

While deep learning has gained popularity across analytics applications, its effectiveness in business contexts has been critically examined. An empirical study investigating multiple industry use cases demonstrates that DL does not consistently outperform traditional ML models when applied to structured datasets with fixed-length feature vectors [13]. Factors such as computational complexity, lack of transparency inherent to black-box models, insufficient big data infrastructure, and organizational skill shortages are identified as key barriers to adoption. The study concludes that gradient boosting models remain the most reliable choice for structured business data, positioning DL as a complementary rather than dominant solution in business analytics.

The telecom sector continues to serve as a benchmark domain for churn prediction research due to data availability and well-defined churn behaviors. One empirical investigation combines extensive data preprocessing, hybrid imputation strategies, and exploratory data analysis to uncover churn drivers such as service quality and competitive offerings [15]. Multiple ML classifiers—including Decision Trees, Random Forests, K-Nearest Neighbors, and XGBoost—are evaluated,

with Random Forest achieving the highest accuracy. This study reinforces the role of ensemble learning in churn prediction and demonstrates how predictive insights can support targeted retention strategies.

Beyond churn, customer experience analytics has gained prominence as organizations seek to improve satisfaction and advocacy metrics. Traditional CX indicators such as Net Promoter Score (NPS) often fail to capture complex relationships between experience attributes and customer loyalty. Addressing this limitation, one study proposes a novel classification framework combining logistic regression with advanced ML algorithms to model CX more accurately [16]. Applied to telecom survey data, the approach yields significant improvements across multiple evaluation metrics, offering a more robust analytical foundation for CX-driven optimization.

In e-commerce environments, CRM systems increasingly rely on advanced analytics to manage complex customer interactions. A DL-based CRM framework leverages artificial neural networks and deep neural networks for customer segmentation, churn prediction, and recommendation tasks [17]. Using a rich dataset of transactional and behavioral attributes, the study demonstrates that DL models can capture nuanced customer patterns that traditional methods may overlook. The reported improvements in predictive accuracy highlight DL's potential to bridge the gap between academic research and real-world CRM applications.

Several large-scale review studies provide broader context for these developments. A comprehensive review of 212 churn prediction studies published between 2015 and 2023 examines the full modeling pipeline, from feature engineering to evaluation practices [11]. The review identifies a strong focus on demographic, behavioral, and usage-related features, while highlighting the underrepresentation of profitability-based evaluation metrics. The authors advocate the integration of ensemble methods, deep learning, and explainable models to enhance both predictive performance and business relevance. Similarly, a survey of ML and DL techniques in e-commerce from 2018 to 2023 categorizes key application areas such as recommendation systems, sentiment analysis, fraud detection, and churn prediction [6]. Persistent challenges—including imbalanced data, overfitting, interpretability, and personalization—are identified as critical directions for future research.

Customer insight generation has also expanded beyond prediction toward optimization and decision support. Recent studies combine big-data analytics, transactional indicators, RFM analysis, clustering, and machine learning to transform heterogeneous customer data into actionable segments [5, 7]. These approaches enable automated analysis, more precise targeting, and more efficient allocation of marketing resources in data-intensive digital business environments.

Advancements in CLV forecasting further illustrate the convergence of AI and strategic business planning. A hierarchical ensemble framework for business-to-business software services treats CLV estimation as a supervised regression problem and combines temporal modeling with customer-segment ensembling [18]. The approach accommodates heterogeneous customer populations and limited historical data while supporting marketing-spend allocation and long-term profitability analysis.

Demand forecasting represents another domain where customer behavior modeling intersects with ML and DL. A recent supply-chain demand prediction study applies an LSTM architecture to learn temporal demand patterns and demonstrates the suitability of recurrent modeling for operational forecasting [19]. Complementary research evaluates ML and DL approaches for demand forecasting based on advertising expenditures, finding that LSTM models outperform traditional regressors and tree-based methods in capturing temporal dependencies [10]. These findings underscore the importance of sequential modeling for marketing-driven demand prediction.

Broader customer analysis studies further compare ML classifiers to assess trade-offs between accuracy and computational efficiency. One study evaluates multiple classifiers—including K-NN, C4.5, Random Forest, Logistic Regression, and Naive Bayes—demonstrating that while C4.5 achieves superior accuracy, Naive Bayes offers notable advantages in execution time [8]. Finally, innovative applications of deep learning extend to qualitative analytics, where autoencoder-based anomaly detection is shown to effectively identify innovation-relevant customer reviews, outperforming traditional ML techniques in stability and relevance [20].

In summary, the literature indicates that ML and DL have become integral to modern customer analytics across churn prediction, CLV estimation, CRM, CX analysis, and demand forecasting. Nevertheless, challenges related to interpretability, real-world deployment, ethical considerations, and alignment with business-oriented metrics persist. Addressing these gaps represents a critical avenue for future research and practical advancement.

3. MATERIALS AND METHODS

3.1 Dataset Description

The dataset used in this study is an enhanced version of the widely recognized *Mall Customers* dataset. This dataset, originally designed to explore customer segmentation in a retail context, includes basic demographic data such as age, gender, annual income, and spending scores. While the original dataset provides useful insights for simple clustering tasks, it lacks the depth required for more sophisticated analysis of customer behavior in dynamic environments such as retail and finance. To overcome this limitation, the dataset has been augmented with additional synthetic but logically consistent features that provide a richer and more realistic representation of customer behavior. These enhanced features have been designed to better reflect the complexity of customer interactions in both retail and financial contexts, where understanding financial behavior, creditworthiness, and customer loyalty is crucial for decision-making.

The enhanced dataset includes a variety of features that offer insight into different aspects of customer behavior. These features include **CustomerID**, a unique identifier for each customer, allowing the dataset to maintain privacy and traceability of customer data. **Gender** provides information on whether the customer is male or female, which can be useful for segmenting customers based on gender-specific behaviors. **Age** is a continuous feature representing the customer's age in years, which is essential for understanding demographic pat-

terns and tailoring marketing strategies to different age groups. **Annual Income (k\$)** is another key feature that indicates the customer's annual income, measured in thousands of dollars. This feature is critical for assessing the customer's purchasing power and overall financial status, which are important factors for both retail marketing and credit risk assessment in the financial sector.

The dataset also includes a **Spending Score (1–100)**, a measure of how much a customer is likely to spend, which can be useful for predicting future purchases or customer lifetime value. **Age Group** categorizes customers into predefined age brackets, providing a way to group customers for demographic analysis. **Estimated Savings (k\$)** is an important feature that reflects a customer's savings, derived from income and spending behavior. This feature helps to assess a customer's financial stability and potential for making future purchases or investments. **Credit Score**, a synthetic feature, represents the customer's creditworthiness and is crucial in the financial industry for assessing the risk of lending to that customer. The **Loyalty Years** feature indicates the length of the customer's relationship with the business, providing insights into customer retention and loyalty trends. Finally, the **Preferred Category** feature reveals the customer's most preferred shopping category, such as Luxury, Budget, Fashion, or Electronics, and can be used to tailor marketing campaigns and product offerings.

The enriched dataset now includes a comprehensive range of features that capture both demographic and behavioral factors, providing a well-rounded resource for tasks such as customer segmentation, classification, and regression. The inclusion of features such as credit score, estimated savings, and preferred category adds a layer of financial and behavioral insight, making the dataset more suitable for applications in both retail marketing and financial services.

The prediction goals vary depending on the task at hand. Specifically, the study aims to predict key customer outcomes such as **credit score**, **category preference**, and **savings**. Predicting credit score is particularly valuable in the financial sector, where such predictions can be used to assess creditworthiness and make informed lending decisions. Understanding category preference is essential for retail businesses, as it allows more targeted marketing and personalized product offerings, ultimately improving customer engagement and increasing sales. Predicting customer savings is important for assessing financial stability, which can influence decisions related to product offerings, loan approvals, and other financial products.

For this study, no external datasets were incorporated. While external data such as retail transactions or broader financial market data could further enhance the analysis, the dataset used in this study is sufficiently comprehensive for the current research objectives. The features provided within the enhanced *Mall Customers* dataset are rich enough to capture a variety of customer behaviors and provide valuable insights into customer segmentation and prediction tasks. However, future research could consider integrating such external data sources to create an even more comprehensive dataset.

The dataset was split into three distinct subsets: training, validation, and testing. Specifically, 70% of the data was used for training the machine learning models, 15% was

reserved for validation during model development to fine-tune hyperparameters, and the remaining 15% was used for final testing to assess performance on unseen data. This split ensures that the model is trained on a sufficient amount of data while also being evaluated on data it has not seen before, helping to prevent overfitting and ensuring that performance is generalizable to new customers.

Understanding the joint behavior of customer income distribution and spending tendencies across demographic segments is a fundamental step in customer analytics, as it provides essential context for subsequent modeling and segmentation tasks. In this regard, Fig. 1 presents a combined visual exploration of annual income levels and spending scores across different age groups and genders. The left panel illustrates the kernel density estimation (KDE) of annual income, which offers a smooth, non-parametric representation of income distribution and highlights central tendencies, dispersion, and potential skewness. Such a representation is useful for identifying dominant income ranges and assessing the heterogeneity of purchasing power. The right panel complements this analysis with a swarm plot of spending scores segmented by age group and gender, enabling a granular comparison of consumption behavior across demographic cohorts. By visualizing individual observations rather than aggregated statistics, this plot reveals intra-group variability, overlap between genders, and age-related patterns in spending intensity. Together, these visualizations provide an integrated descriptive overview of how income and demographic factors jointly influence customer spending behavior.

A comprehensive understanding of the interrelationships among customer attributes is essential for informed feature selection, multicollinearity assessment, and the interpretation of behavioral patterns. In this context, Fig. 2 presents the correlation matrix of all key features considered in the analysis, offering a concise quantitative summary of the strength and direction of linear associations between demographic, financial, and behavioral variables. By visualizing pairwise correlation coefficients in a heatmap, the figure enables rapid identification of strongly correlated features, such as the pronounced associations between annual income, estimated savings, and credit score, as well as the relatively weaker or inverse relationships involving spending score and age. This representation also facilitates the detection of potentially redundant variables and highlights features that may carry complementary information for predictive modeling. Overall, the correlation patterns provide critical insights into the underlying structure of the dataset and inform subsequent analytical and modeling decisions.

Analyzing the distribution of financial resources across demographic groups provides valuable insights into underlying economic behavior and potential heterogeneity within a customer population. Figure 3 illustrates the distribution of estimated savings for male and female customers using violin plots, which combine features of boxplots and kernel density estimations to convey both central tendencies and distributional shapes. This visualization enables a detailed comparison of savings behavior between genders by highlighting differences in median values, interquartile ranges, and the overall spread of observations. The width of each violin reflects the relative concentration of data points at different savings levels,




Figure 1. Kernel density estimation of annual income distribution (left) and swarm plot of spending score by age group and gender (right).




Figure 2. Heatmap of pairwise correlation coefficients among customer demographic, financial, and behavioral features.



Figure 3. Violin plots comparing the distribution of estimated savings between male and female customers.

thereby revealing potential skewness and multimodality that may not be apparent from summary statistics alone. By presenting the full distribution rather than aggregated measures, the figure supports a nuanced interpretation of gender-based financial patterns and provides an informative foundation for subsequent behavioral or predictive analyses.

3.2 Data Preprocessing

Data preprocessing represents a fundamental stage in the development of reliable and robust machine learning models, particularly in economically and financially oriented customer analytics where data quality directly affects predictive accuracy and decision validity. Given the heterogeneous nature of the dataset—comprising demographic attributes, behavioral indicators, and financially meaningful variables—careful preprocessing was required to ensure consistency, reduce noise, and enhance model generalization. Each preprocessing step was designed to address challenges commonly encountered in real-world retail and financial data, such as missing information, mixed data types, scale disparities, and multicollinearity.

The first and most critical preprocessing task involved handling missing values. Incomplete observations are common in customer datasets due to reporting errors, data integration issues, or privacy-related omissions and, if left untreated, can introduce systematic bias and distort learned relationships. For numerical features such as annual income and estimated savings, missing values were imputed using either mean or median imputation, depending on the underlying statistical distribution of each feature. When the distribution was approximately symmetric, mean imputation was employed to preserve central tendency. For skewed distributions—often observed in income- or savings-related variables—the median was preferred because it is more robust to extreme values. For categorical variables, including gender and preferred category, missing values were imputed using mode imputation. In instances where a feature exhibited an excessive proportion of missing values that could compromise model reliability and distort economic interpretation, the corresponding records were removed to maintain overall data integrity.

Following missing-value treatment, categorical feature encoding was performed to transform non-numeric variables into a format suitable for machine learning algorithms. Since most predictive models operate on numerical inputs, categorical attributes required systematic conversion without introducing artificial ordinal relationships. One-hot encoding was applied to nominal variables such as gender and preferred category. This scheme generates a binary indicator for each category, allowing the model to independently assess the influence of each group. For example, gender was transformed into two binary variables representing Male and Female, while pre-

ferred category was expanded into separate indicators for Luxury, Budget, Fashion, and Electronics. This approach ensures that categorical information is preserved in a mathematically neutral form, which is especially important in economic modeling where imposing unintended rankings could lead to misleading interpretations.

Subsequently, feature scaling was applied to address disparities in the numerical ranges of the input variables. Customer datasets typically include features measured on vastly different scales, such as income expressed in thousands of monetary units and spending scores defined within a limited numeric range. Without appropriate scaling, models—particularly those relying on distance-based or gradient-based optimization—may disproportionately weight features with larger magnitudes, leading to biased learning outcomes. To mitigate this issue, standardization was employed, whereby each numerical feature was transformed by subtracting its mean and dividing by its standard deviation. This process produces features with zero mean and unit variance, ensuring balanced contribution across all variables. Standardization is particularly important in financial prediction tasks, where both small-scale behavioral indicators and large-scale monetary values carry economically meaningful information.

The final preprocessing step involved feature correlation analysis to identify and mitigate redundancy among input variables. In customer behavior and financial datasets, certain features are inherently correlated due to underlying economic relationships. For instance, annual income and estimated savings often convey overlapping information about a customer's financial capacity. Including highly correlated features can lead to multicollinearity, which inflates model variance, reduces interpretability, and degrades predictive stability. A correlation matrix was constructed to quantify pairwise relationships among numerical features. Features exhibiting correlation coefficients above a predefined threshold, typically 0.9, were flagged as redundant. In addition, variance inflation factor (VIF) analysis was conducted to quantify how much the variance of a regression coefficient is inflated due to feature correlation. Features with excessively high VIF values were considered for removal or consolidation.

By systematically addressing missing data, categorical representation, feature scaling, and redundancy, the preprocessing pipeline significantly reduced noise and complexity within the dataset. This not only improved computational efficiency but also enhanced the interpretability and generalization capability of the subsequent machine learning models. From an economic and financial perspective, such rigorous preprocessing is essential to ensure that predictive insights are both statistically sound and economically meaningful, thereby supporting reliable decision-making in retail and financial applications.

3.3 Deep Learning Models

In this study, five deep learning models were selected for evaluation based on their suitability for modeling complex customer behavior in retail and financial analytics. Customer-related datasets, even when represented in tabular form, often embody intricate nonlinear relationships among demographic variables, financial indicators, and behavioral attributes. Variables such as annual income, spending score, estimated sav-

ings, credit score, and loyalty years are not independent; rather, they interact in ways that reflect underlying economic behavior, financial decision-making processes, and consumption patterns. Accurately modeling these interactions requires learning frameworks that can capture high-dimensional nonlinear mappings while remaining robust to noise, feature heterogeneity, and potential redundancy. Deep learning models are particularly appropriate in this context because of their strong representational capacity and ability to learn hierarchical abstractions directly from data.

The selected models cover a diverse range of modeling paradigms relevant to customer segmentation and prediction. Specifically, the set includes architectures capable of capturing sequential dependencies, learning expressive latent representations, and exploiting relational structure among customers. This diversity allows the study to assess how different deep learning perspectives perform when applied to economically meaningful prediction tasks, such as estimating customer savings, predicting credit scores, or identifying preferred shopping categories. The models considered are EALSTM, FTLM, BiLSTM, GNN, and VAE.

- **EALSTM:** This model belongs to the family of Long Short-Term Memory architectures, which are designed to handle complex dependency structures through gated memory mechanisms. LSTM-based models are well established in financial forecasting and customer analytics because of their ability to retain relevant information over extended input sequences and mitigate vanishing-gradient issues during training. In customer behavior modeling, EALSTM-type architectures are particularly useful when attributes can be organized into meaningful ordered representations, such as customer lifecycle stages, loyalty progression, or age-group transitions. These models can learn nonlinear interactions between financial indicators and behavioral features, making them suitable for prediction tasks that require sensitivity to both short- and long-range dependencies.
- **FTLM:** FTLM was included because of its ability to model complex feature interactions within structured customer data. Models of this type are generally designed to enhance representation learning in settings where inputs consist of mixed data types, including continuous financial variables and categorical behavioral indicators. In retail and financial analytics, such architectures are valuable for capturing latent relationships that are not easily modeled by shallow learners, particularly when customer behavior is influenced by subtle interactions between income levels, spending tendencies, and preference categories.
- **BiLSTM:** Bidirectional Long Short-Term Memory networks extend conventional LSTM architectures by processing the input sequence in both forward and backward directions. This bidirectional processing allows the model to incorporate contextual information from both earlier and later positions within an ordered representation. In customer behavior analysis, this property is advantageous when records are indexed according to economically meaningful orderings rather than strict time, such as loyalty-duration classes, age brackets, or derived behavioral stages. By exploiting information from both directions, BiLSTM models can generate richer feature representations, improving predictive

accuracy and segmentation quality.

- **GNN:** Graph Neural Networks were selected to account for the relational nature of customer behavior. Although the dataset is primarily tabular, economically meaningful graphs can be constructed by connecting customers based on similarity in income, spending score, estimated savings, or preferred category. GNNs operate through message-passing mechanisms, enabling each customer representation to be influenced by neighboring customers. This capability is relevant in marketing economics and financial analytics, where customers often exhibit clustered behavior patterns and shared financial characteristics. By leveraging relational structure, GNNs can capture peer effects and segment-level dependencies that are difficult to model using purely independent observations.
- **VAE:** Variational Autoencoders are probabilistic deep learning models designed for representation learning and dimensionality reduction. VAEs learn compact latent representations that summarize high-dimensional input data while preserving essential structural information. In customer analytics, VAEs are useful for compressing correlated financial features, such as annual income and estimated savings, into lower-dimensional latent factors that can support more stable segmentation and prediction. These latent representations can also facilitate the identification of atypical customer profiles, which is important for risk awareness and for identifying high-value or unusual spending behaviors.

The selection of these models reflects the need to address multiple analytical perspectives within customer behavior modeling. Recurrent architectures emphasize dependency learning and nonlinear temporal or ordered relationships, graph-based models capture relational and peer-driven effects, and generative models support robust representation learning in high-dimensional financial data. Together, they provide a comprehensive experimental framework for evaluating deep learning approaches to customer segmentation and prediction in retail and financial contexts.

All deep learning models are trained and evaluated under a consistent experimental setup using the same preprocessed dataset to ensure fairness and comparability. Model performance is assessed using standard regression and goodness-of-fit metrics appropriate for financial and economic prediction tasks, allowing differences in predictive capability to be attributed to architectural characteristics rather than data inconsistencies or preprocessing effects.

3.4 Metaheuristic Optimization Algorithms

Metaheuristic optimization algorithms have become an essential component in modern machine learning workflows, particularly in economic and financial applications where model performance is highly sensitive to parameter configuration and data heterogeneity. In customer behavior analytics, deep learning models such as EALSTM, FTLN, BiLSTM, GNN, and VAE involve a large number of tunable hyperparameters, including learning rates, network depth, hidden-unit sizes, regularization coefficients, and batch sizes. The selection of these parameters has a direct impact on predictive accuracy, convergence stability, and generalization capability.

Manual or grid-based tuning strategies are often computationally infeasible and prone to suboptimal solutions, especially when the search space is high-dimensional. Consequently, metaheuristic optimization algorithms provide an effective alternative by enabling automated, adaptive, and population-based exploration of complex parameter spaces.

3.4.1 Role of Metaheuristics in Hyperparameter Optimization

In this study, metaheuristic algorithms are employed to automate the hyperparameter optimization process for the deep learning models under investigation. The primary objective is to identify parameter configurations that enhance predictive performance while maintaining strong generalization ability. In financial and retail customer modeling, overfitting is a critical concern, as models that perform well on historical customer data may fail to generalize to new customers or evolving market conditions. Metaheuristic algorithms address this challenge by balancing exploration and exploitation during the search process, thereby reducing the risk of convergence to local optima and improving robustness across different data partitions.

The optimization framework operates by iteratively evaluating candidate solutions, where each candidate represents a specific hyperparameter configuration for a given deep learning model. Performance feedback is obtained using validation-based objective functions that reflect both accuracy and stability considerations. Through population evolution, stochastic search operators, and adaptive learning mechanisms, metaheuristics efficiently navigate the hyperparameter space without requiring gradient information. This property is particularly advantageous for deep learning models applied to customer behavior prediction, where objective functions may be non-convex, noisy, or computationally expensive to evaluate.

From an economic and financial perspective, metaheuristic hyperparameter optimization supports the development of reliable decision-support models. By improving generalization performance, optimized models are better suited for tasks such as credit scoring, savings estimation, and customer segmentation, where inaccurate predictions can lead to financial risk, misallocation of marketing resources, or suboptimal customer targeting. Therefore, integrating metaheuristic optimization into the modeling pipeline is a key methodological contribution of this study.

3.5 Kirchhoff's Law Algorithm

The Kirchhoff's Law Algorithm (KLA) used in this study is formulated as a physics-inspired, population-based metaheuristic derived from Kirchhoff's Current Law (KCL), which states that the algebraic sum of currents entering and leaving a node in an electrical circuit is zero. In the optimization context, each candidate solution is modeled as a node in an equivalent resistive electrical network, while the objective-function value is mapped to an electrical resistance. This analogy enables the optimization process to be governed by current-flow dynamics rather than manually tuned heuristic control parameters. The design follows the broader class of physics-based population metaheuristics, which construct search dynamics from physical laws and seek an effective balance between exploration and exploitation [21].

Let $\mathbf{X}_i^k \in \mathbb{R}^D$ denote the position (solution vector) of the i th individual at iteration k , and let $f(\mathbf{X}_i^k)$ be its corresponding objective-function value. In KLA, the resistance between two nodes i and j is defined as a function of their objective values and stochastic components:

$$R_{ij}^k = \frac{1}{(\text{rand}_1 + \text{rand}_2) \left[f(\mathbf{X}_i^k) - f(\mathbf{X}_j^k) \right]^2 \text{rand}_3}. \quad (1)$$

Here, rand_1 , rand_2 , and rand_3 are random numbers uniformly distributed in $[0, 1]$. This formulation ensures that solution pairs with smaller objective differences yield lower resistance, thereby attracting higher current flow in accordance with Ohm's law.

Based on Kirchhoff's Current Law, the net current at each node must satisfy

$$\sum_{j \neq i} I_{ij}^k = 0, \quad (2)$$

which enforces equilibrium among all population members. The current-driven position update for each candidate solution is computed as

$$\mathbf{X}_i^{k+1} = \mathbf{X}_i^k + \Delta \mathbf{X}_i^k, \quad (3)$$

where the displacement term $\Delta \mathbf{X}_i^k$ is obtained from the aggregated current influence of neighboring nodes:

$$\Delta \mathbf{X}_i^k = \sum_{j \neq i} \frac{f(\mathbf{X}_j^k) - f(\mathbf{X}_i^k)}{|f(\mathbf{X}_j^k) - f(\mathbf{X}_i^k)| + \varepsilon} (\mathbf{X}_j^k - \mathbf{X}_i^k), \quad (4)$$

where ε is a small positive constant introduced to avoid division by zero. This coefficient determines the direction and magnitude of current flow, ensuring that candidate solutions are guided from regions of higher resistance (poorer fitness) toward regions of lower resistance (better fitness).

The exploration–exploitation balance in KLA emerges naturally from the resistance–current relationship. During early iterations, large variations in objective values result in diverse resistance levels, allowing currents to traverse multiple branches of the network and encouraging extensive exploration of the search space. As iterations progress, resistance values increasingly concentrate around high-quality solutions, strengthening current flow toward these regions and promoting exploitation. Unlike conventional metaheuristic algorithms, this transition does not rely on predefined control parameters; instead, it is intrinsically governed by physical laws that minimize electrical power loss:

$$P = I^2 R, \quad (5)$$

which directly corresponds to minimizing the optimization cost function.

Through this formulation, KLA achieves a self-adaptive and parameter-free optimization process that preserves population diversity, avoids premature convergence, and steadily drives the search toward globally competitive solutions. These characteristics make KLA particularly suitable for hyperparameter optimization in deep learning models, where complex, high-dimensional search spaces and strong generalization requirements are common in retail and financial customer analytics.

3.5.1 State-of-the-Art Metaheuristic Optimization Algorithms

To comprehensively evaluate the effectiveness of metaheuristic-driven hyperparameter optimization, several state-of-the-art algorithms were considered. These algorithms were selected based on their demonstrated success in complex optimization tasks and their relevance to machine learning applications in economics and finance.

- **PSO:** Particle Swarm Optimizer is inspired by the social behavior of swarms. PSO updates candidate solutions by combining individual experience with collective knowledge, enabling efficient convergence in continuous search spaces. In hyperparameter optimization, PSO is valued for its fast convergence and relatively low computational overhead, making it suitable for tuning deep learning models on customer datasets with moderate to large dimensionality.
- **BBO:** Biogeography Based Optimizer is inspired by the migration and distribution of species across habitats. Candidate solutions are treated as habitats whose suitability is determined by fitness values. Through migration and mutation operations, BBO promotes the sharing of high-quality parameter configurations across the population. This mechanism is particularly useful in preserving diversity during optimization, which is essential for avoiding premature convergence in financial prediction models.
- **WOA:** Whale Optimization Algorithm is inspired by the bubble-net hunting strategy of humpback whales. WOA alternates between exploitation and exploration phases using encircling and spiral updating mechanisms. Its adaptive search behavior makes it effective for navigating complex, nonlinear hyperparameter spaces commonly encountered in deep learning-based customer analytics.
- **BA:** Bat Algorithm is motivated by the echolocation behavior of bats. It dynamically adjusts search parameters such as loudness and pulse rate to balance global exploration and local exploitation. BA is particularly effective in optimization scenarios where fine-grained local search is required after identifying promising regions of the parameter space, which is often the case when tuning deep neural architectures.
- **APO:** Artificial Protozoa Optimizer was included to broaden the comparative analysis. APO contributes an alternative search dynamic that complements the exploration–exploitation balance of the other algorithms, allowing a more comprehensive evaluation of metaheuristic performance in hyperparameter tuning for customer behavior prediction models.
- **GA:** Genetic Algorithm is one of the most established evolutionary optimization techniques, inspired by natural selection and genetic inheritance. GA employs crossover, mutation, and selection operators to evolve candidate solutions over generations. Its robustness and flexibility make it well suited for optimizing both continuous and discrete hyperparameters in deep learning models applied to financial and retail datasets.
- **SFS:** Stochastic Fractal Search is a population-based metaheuristic that combines diffusion processes with stochastic sampling. SFS is designed to maintain diversity while

progressively refining candidate solutions, which is particularly beneficial for avoiding stagnation in high-dimensional optimization problems. In this study, SFS provides a competitive alternative for hyperparameter optimization in deep learning-based customer analytics.

The inclusion of these metaheuristic algorithms enables a systematic comparison of different optimization strategies under a unified experimental framework. By evaluating their effectiveness in tuning deep learning models for customer behavior prediction, this study aims to identify robust optimization techniques that enhance predictive accuracy, improve generalization, and support economically meaningful decision-making in retail and financial applications.

3.6 Evaluation Metrics

To rigorously assess the predictive performance and generalization capability of the proposed deep learning models, a comprehensive set of regression-based evaluation metrics was employed. These metrics were selected to capture different aspects of model accuracy, bias, correlation strength, and goodness-of-fit, which are particularly important in retail and financial customer analytics. In such applications, predictive models must not only minimize average error but also demonstrate stability, unbiased estimation, and strong agreement with observed values, as inaccuracies may lead to suboptimal financial decisions, ineffective customer targeting, or increased economic risk.

Let y_i denote the observed value, $\hat{y}_i$ the predicted value, $\bar{y}$ the mean of observed values, $\bar{\hat{y}}$ the mean of predicted values, and n the total number of samples. The evaluation metrics used in this study are summarized in Table 1, along with their corresponding mathematical formulations.

The combined use of these metrics ensures a balanced evaluation of predictive accuracy and reliability. Error-based metrics such as MSE, RMSE, MAE, and RRMSE quantify the magnitude of prediction errors, while MBE provides insight into systematic bias. Correlation-based measures (r and R^2) evaluate the strength of linear association and explanatory power, whereas NSE and WI assess model efficiency and agreement relative to observed variability. Together, these metrics provide a robust and multidimensional framework for evaluating deep learning models in economically and financially driven customer behavior prediction tasks.

4. EXPERIMENTAL RESULTS

4.1 Baseline Model Performance (Before Optimization)

This subsection presents a comprehensive evaluation of the baseline predictive performance of the deep learning models prior to the application of any optimization strategies. Establishing baseline performance is a critical step, as it reflects the intrinsic modeling capability of each architecture under identical data preprocessing, training, and testing conditions. In retail and financial customer analytics, such baseline assessment is essential to understand how well different model structures capture the underlying relationships among demographic attributes, financial indicators, and behavioral features before enhancement through metaheuristic optimization.

The evaluated models include EALSTM, FTLM, BiLSTM, GNN, and VAE. Their predictive performance was assessed using the regression metrics defined in Section 3.6. These metrics jointly quantify prediction-error magnitude, systematic bias, correlation strength, explanatory power, and agreement between predicted and observed values. The numerical results for all baseline models are summarized in Table 2.

As shown in Table 2, EALSTM demonstrates the strongest baseline performance across nearly all evaluation metrics. It achieves the lowest MSE (0.0093) and RMSE (0.0965), indicating minimal overall prediction error. The MAE value of 0.0602 confirms that, on average, EALSTM predictions deviate only slightly from observed values. The MBE of 0.0409 suggests limited systematic bias. Furthermore, EALSTM exhibits strong correlation and explanatory capability, with an r value of 0.875 and an R^2 of 0.853. High agreement with observed data is further supported by NSE and WI values of 0.892 and 0.891, respectively.

FTLM ranks second among the baseline models, with an MSE of 0.0134 and RMSE of 0.1158. Although these values are higher than those of EALSTM, they still indicate relatively accurate baseline predictions. FTLM records an MAE of 0.0681 and an MBE of 0.0520, reflecting moderate absolute error and bias. Its correlation coefficient (0.867) and R^2 value (0.845) demonstrate strong explanatory power, while NSE (0.878) and WI (0.875) indicate reliable agreement with observed customer behavior.

BiLSTM exhibits moderate baseline performance, yielding an MSE of 0.0191 and RMSE of 0.1382. The MAE (0.0773) and MBE (0.0594) values indicate increased prediction error and bias relative to EALSTM and FTLM. Nevertheless, BiLSTM maintains a correlation coefficient of 0.855 and an R^2 of 0.835, suggesting that bidirectional dependency modeling captures a substantial portion of variability in the data. NSE (0.865) and WI (0.864) values further reflect reasonable predictive consistency.

The GNN model shows a further reduction in baseline accuracy, with an MSE of 0.0261 and RMSE of 0.1616. Its MAE (0.0889) and MBE (0.0743) indicate higher error dispersion and bias. Despite this, GNN maintains an r value of 0.844 and an R^2 of 0.823, highlighting its ability to capture relational patterns within the customer dataset. NSE (0.848) and WI (0.852) confirm moderate agreement between predictions and observations.

Among all evaluated models, VAE exhibits the weakest baseline performance. It records the highest MSE (0.0347) and RMSE (0.1863), along with an MAE of 0.1026 and an MBE of 0.0869. Correlation and explanatory metrics are correspondingly lower, with $r = 0.834$ and $R^2 = 0.813$. NSE (0.833) and WI (0.838) further indicate reduced agreement with observed values. This behavior is consistent with the primary design objective of VAEs, which emphasizes representation learning rather than direct regression accuracy in baseline configurations.

Overall, the numerical results reveal a clear performance hierarchy among the deep learning models. EALSTM and FTLM demonstrate superior baseline accuracy and stability, BiLSTM and GNN show moderate performance, and VAE exhibits limited predictive capability in its unoptimized form.

Table 1. Regression evaluation metrics used for model performance assessment

Metric	Mathematical Definition
Mean Squared Error (MSE)	$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$
Root Mean Squared Error (RMSE)	$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}$
Mean Absolute Error (MAE)	$\text{MAE} = \frac{1}{n} \sum_{i=1}^n y_i - \hat{y}_i $
Mean Bias Error (MBE)	$\text{MBE} = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)$
Correlation Coefficient (r)	$r = \frac{\sum_{i=1}^n (y_i - \bar{y})(\hat{y}_i - \bar{\hat{y}})}{\sqrt{\sum_{i=1}^n (y_i - \bar{y})^2 \sum_{i=1}^n (\hat{y}_i - \bar{\hat{y}})^2}}$
Coefficient of Determination (R^2)	$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$
Relative Root Mean Squared Error (RRMSE)	$\text{RRMSE} = \frac{\sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}}{\bar{y}} \times 100$
Nash–Sutcliffe Efficiency (NSE)	$\text{NSE} = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$
Willmott Index (WI)	$\text{WI} = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (\hat{y}_i - \bar{y} + y_i - \bar{y})^2}$


Table 2. Baseline performance of deep learning models before optimization

Model	MSE	RMSE	MAE	MBE	r	R^2	RRMSE	NSE	WI
EALSTM	0.0093	0.0965	0.0602	0.0409	0.875	0.853	1.69	0.892	0.891
FTLM	0.0134	0.1158	0.0681	0.0520	0.867	0.845	1.97	0.878	0.875
BiLSTM	0.0191	0.1382	0.0773	0.0594	0.855	0.835	2.29	0.865	0.864
GNN	0.0261	0.1616	0.0889	0.0743	0.844	0.823	2.63	0.848	0.852
VAE	0.0347	0.1863	0.1026	0.0869	0.834	0.813	2.98	0.833	0.838

These findings provide a rigorous quantitative benchmark and strongly motivate the application of advanced metaheuristic optimization techniques to enhance predictive accuracy, reduce bias, and improve generalization in economically and financially sensitive customer behavior prediction tasks.

A clear statistical characterization of model performance metrics is essential for evaluating both accuracy and robustness across different evaluation criteria. Figure 4 illustrates a comparative line plot of the mean values and corresponding standard deviations for commonly used performance metrics, including error-based, correlation-based, and efficiency-based indicators. The mean curves provide an overview of the central tendency of each metric, while the accompanying standard-deviation curves quantify variability and stability across multiple experiments or runs. This dual representation facilitates the identification of metrics that not only exhibit strong average performance but also demonstrate consistent behavior with limited dispersion. The figure also highlights contrasts between low-magnitude error metrics and higher-magnitude efficiency metrics, offering a unified perspective on how different criteria respond to model uncertainty.

A robust evaluation of predictive models requires not only comparison of average performance metrics but also explicit consideration of the associated uncertainty across experimen-

**Figure 4.** Line plot showing the mean and standard deviation of selected performance metrics.

tal runs. Figure 5 presents a bar plot of key performance metrics augmented with error bars that represent the variability of each metric. The height of each bar denotes the mean value, while the error bars provide insight into the dispersion and stability of model performance. This visualization enables direct comparison between error-based measures and efficiency- or correlation-based indicators, facilitating the




Figure 5. Bar plot of model performance metrics with error bars




Figure 6. Pairplot of model performance metrics with kernel density estimates along the diagonal.

identification of metrics that demonstrate both strong central performance and low variability. It also highlights differences in scale among the metrics, emphasizing their relative sensitivity to model uncertainty.

Exploring the joint behavior of performance metrics is essential for understanding potential trade-offs, redundancies, and complementarities in model evaluation. Figure 6 presents a comprehensive pairwise visualization of key model performance metrics, augmented with kernel density estimates along the diagonal. The diagonal panels summarize the marginal distributions of individual metrics, providing insight into central tendencies and variability, while the off-diagonal panels depict pairwise relationships that reveal the strength and direction of inter-metric associations. This combined representation enables identification of strong linear dependencies between error-based metrics, as well as inverse relationships between error measures and efficiency or correlation indicators. It also facilitates detection of collinearity among metrics, which is particularly relevant when selecting evaluation criteria for comparative analysis or multi-objective optimization.

A comparative evaluation of baseline deep learning architectures provides an essential reference point for assessing

more advanced or hybrid modeling approaches. Figure 7 illustrates the MSE obtained by the EALSTM, FTLM, BiLSTM, GNN, and VAE models. Each subplot presents the MSE value corresponding to a specific model, enabling a clear and isolated comparison of predictive accuracy across architectural paradigms. This visualization highlights relative performance differences among sequence-based, graph-based, and generative learning models, revealing how architectural choices influence error magnitude. It also facilitates an intuitive assessment of which baseline models provide stronger predictive capability and which may require further optimization or hybridization.




Figure 7. Mean squared error comparison of baseline deep learning models.

4.2 Optimized Model Analysis

This subsection presents a detailed comparative analysis of the optimized deep learning models, with particular emphasis on the impact of metaheuristic-driven hyperparameter optimization on predictive accuracy, bias reduction, and generalization capability. All optimized configurations are based on the EALSTM architecture, selected due to its superior baseline performance, and are enhanced using different metaheuristic optimization algorithms. The objective is to quantify how various optimizers influence model behavior under identical experimental conditions and to identify the most effective optimization strategy for economically and financially oriented customer behavior prediction.

Table 3 reports the numerical results obtained after optimization using KLA, PSO, BBO, WOA, BA, APO, GA, and SFS. The evaluation uses the same regression metrics defined in Section 3.6, ensuring direct comparability across optimization strategies.

As shown in Table 3, all metaheuristic optimization strategies lead to substantial improvements over the baseline EALSTM performance reported in Section 4.1, confirming the effectiveness of automated hyperparameter tuning in enhancing predictive accuracy and stability. Among all evaluated optimizers, the KLA-optimized EALSTM consistently achieves the best overall performance across all metrics. Specifically, it records the lowest MSE (3.60×10^{-7}), RMSE (0.00728), MAE (0.000372), and MBE (0.000091), indicating an exceptional reduction in both absolute error magnitude and systematic bias. In addition, the KLA-based model achieves the highest correlation coefficient ($r = 0.972$) and coefficient of determination ($R^2 = 0.969$), reflecting superior explana-

Table 3. Performance comparison of optimized EALSTM models using different metaheuristic algorithms

Model	MSE	RMSE	MAE	MBE	r	R^2	RRMSE	NSE	WI
KLA + EALSTM	3.60×10^{-7}	0.00728	0.000372	0.000091	0.972	0.969	0.095	0.971	0.977
PSO + EALSTM	5.15×10^{-7}	0.00905	0.000418	0.000123	0.964	0.960	0.135	0.963	0.971
BBO + EALSTM	6.55×10^{-7}	0.01010	0.000458	0.000138	0.956	0.953	0.175	0.956	0.966
WOA + EALSTM	7.75×10^{-7}	0.01092	0.000482	0.000159	0.950	0.948	0.215	0.949	0.961
BA + EALSTM	9.30×10^{-7}	0.01184	0.000501	0.000176	0.943	0.941	0.270	0.942	0.956
APO + EALSTM	1.09×10^{-6}	0.01285	0.000517	0.000191	0.938	0.935	0.340	0.936	0.951
GA + EALSTM	1.27×10^{-6}	0.01379	0.000534	0.000209	0.931	0.928	0.440	0.928	0.945
SFS + EALSTM	1.44×10^{-6}	0.01471	0.000549	0.000224	0.925	0.921	0.580	0.920	0.939

tory power. High agreement with observed values is further confirmed by NSE (0.971) and WI (0.977), which are the largest among all optimized configurations.

The PSO-optimized EALSTM ranks second in performance, with an MSE of 5.15×10^{-7} and RMSE of 0.00905. Although its error values are slightly higher than those of KLA, PSO still achieves strong correlation ($r = 0.964$) and explanatory capability ($R^2 = 0.960$). Similarly, BBO and WOA demonstrate competitive optimization behavior, progressively increasing error metrics while maintaining relatively high agreement and efficiency scores. These results indicate that population-based evolutionary and swarm-inspired algorithms are effective in navigating the hyperparameter search space, though their convergence behavior differs in terms of precision and stability.

BA, APO, GA, and SFS exhibit a gradual decline in optimization effectiveness, as reflected by increasing MSE, RMSE, MAE, and RRMSE values, along with decreasing r , R^2 , NSE, and WI. In particular, the SFS-optimized EALSTM yields the highest error metrics and the lowest agreement indicators among the optimized models, suggesting comparatively weaker exploitation of high-quality regions in the hyperparameter space. Nevertheless, even the least effective optimized configuration substantially outperforms the unoptimized baseline, underscoring the overall benefit of metaheuristic optimization.

From an economic and financial perspective, the numerical hierarchy observed in Table 3 highlights the critical role of optimizer selection in developing reliable customer behavior prediction models. The superior performance of KLA can be attributed to its physics-inspired, parameter-free design and intrinsic balance between exploration and exploitation, which enables robust convergence toward high-quality hyperparameter configurations. These characteristics make KLA particularly suitable for financially sensitive prediction tasks, where minimizing prediction error, bias, and uncertainty is essential for effective decision-making and risk management. Examining the distributional characteristics of performance metrics provides deeper insight into model stability, variability, and overall reliability beyond point estimates alone. Figure 8 presents a mixed visualization that combines empirical density representations with KDE curves for a comprehensive set of model performance metrics. Each subpanel corresponds to an individual metric and depicts both the smoothed probability density and the underlying distributional shape,

enabling assessment of central tendency, spread, and potential skewness. This dual representation is useful for identifying metrics that exhibit concentrated distributions with limited dispersion, which may indicate consistent model behavior across multiple evaluations. It also facilitates comparison across error-based, correlation-based, and efficiency-based measures by placing all distributions within a unified visual framework.



Assessing the distributional assumptions of performance metrics is a critical step in statistical analysis, particularly when interpreting summary statistics or applying parametric evaluation techniques. Figure 9 presents a comprehensive set of quantile–quantile (Q–Q) plots for all considered model performance metrics. Each subplot compares the empirical quantiles of a given metric with the corresponding theoretical quantiles of a normal distribution, thereby providing a visual diagnostic of normality. The degree to which the plotted points align with the reference line indicates how closely the empirical distribution conforms to normality, while systematic deviations may suggest skewness, heavy tails, or other distributional irregularities. Examining all metrics within a unified framework enables a consistent assessment of distributional behavior across error-based, correlation-based, and efficiency-based indicators.

A systematic comparison of multiple hybrid models across a diverse set of performance metrics is essential for identifying the most accurate and reliable modeling framework. Figure 10 presents a heatmap-based comparison of hybrid optimization–deep learning models evaluated using error-based, correlation-based, and efficiency-based performance indicators. Each row corresponds to a distinct hybrid configuration, while each column represents a specific evaluation metric, with color intensity reflecting the relative magnitude of performance. This visual structure enables rapid identification of high-performing models that simultaneously achieve low error values and high efficiency scores. It also highlights consistent performance trends across metrics, facilitating an integrated assessment of model robustness and generalization capability.

5. DISCUSSION

This section provides an expanded and integrative interpretation of the empirical findings, situating the results within the broader context of retail economics and financial analytics. Beyond reporting comparative performance, the discus-

Mixed Plot: Density + KDE for Metrics

**Figure 8.** Combined density and kernel density estimation plots for selected model performance metrics.**Figure 9.** Quantile–quantile plots for all model performance metrics against a theoretical normal distribution.

sion aims to explain why the observed performance patterns emerge, how these patterns connect to established knowledge in customer segmentation and financial prediction, and what methodological constraints should be considered when interpreting the outcomes. The discussion is structured around three themes: interpretation of observed performance behavior across baseline and optimized configurations, comparison with prior studies in customer segmentation, marketing analytics, and financially oriented predictive modeling, and a critical assessment of limitations and future methodological considerations.

A central observation emerging from the empirical evaluation is that predictive performance is strongly dependent on hyperparameter configuration and optimizer selection, even when advanced model architectures are employed. This is a particularly important insight in economically sensitive applications, where predictive systems are frequently deployed as decision-support tools and performance variability can directly affect costs, revenues, and risk exposure. The consistent improvements observed from baseline to optimized configurations indicate that the learning capacity embedded in model architecture alone is not sufficient to guarantee reliable perfor-



Figure 10. Heatmap comparison of hybrid model performance across multiple evaluation metrics.

mance; rather, that capacity must be operationalized through appropriate optimization of learning dynamics, regularization strength, and convergence behavior. In statistical terms, the optimization process alters the effective bias–variance trade-off by steering models toward parameter regions that reduce systematic error while maintaining adequate flexibility. In economic terms, optimized models are better positioned to learn stable relationships between customer financial capacity, behavioral intensity, and preference structure without overfitting to idiosyncratic noise in the training sample.

The metric-level evidence supports this interpretation. Error-based measures such as MSE, RMSE, and MAE decrease markedly after optimization, indicating stronger pointwise predictive accuracy and reduced dispersion of residuals. In financial prediction tasks, RMSE reduction is particularly meaningful because it penalizes larger errors more strongly and is therefore closely aligned with risk management objectives, where rare but large deviations may create disproportionate costs. The systematic reduction in MBE also indicates that optimization not only reduces noise-related error but mitigates directional error, which is critical in finance because biased estimates can systematically misprice credit, misallocate marketing budgets, or misclassify customer value tiers. Simultaneously, goodness-of-fit indicators—including R^2 , NSE, and WI—increase, suggesting that optimized models improve not only absolute error but the overall agreement structure between observed and predicted behaviors. In customer economics, stronger agreement metrics imply that models better capture relative ranking and proportional patterns in customer behavior, which is essential for segmentation, targeting, and portfolio prioritization.

Among the optimized configurations, the KLA-based framework exhibits the most substantial and consistent performance gains. This outcome can be interpreted through the algorithmic structure of KLA itself. Unlike many optimization schemes that rely on multiple externally tuned parameters, KLA employs a physics-inspired search process grounded in circuit-equilibrium principles, thereby enabling a self-regulating exploration–exploitation balance. In practical opti-

mization terms, this balance is essential because hyperparameter landscapes for high-capacity models are typically non-convex, multimodal, and sensitive to stochasticity in training. The superior performance of KLA suggests that it provides a more reliable mechanism for escaping poor local optima while converging toward stable, high-quality regions in the hyperparameter space. In economic modeling terms, such stability is valuable because it reduces the risk that predictive performance is an artifact of particular training realizations rather than a robust reflection of underlying customer structure.

An additional noteworthy observation is that the improvement achieved through optimization is broad-based across multiple categories of metrics. Gains are not limited to MSE-type measures but extend consistently to correlation (r), explanatory power (R^2), efficiency (NSE), and agreement (WI). This coherence is important because it implies that the optimization process improves both error magnitude and structural predictive alignment, rather than merely compressing errors in a narrow numerical sense. Such coherence strengthens confidence that the optimized model is capturing economically meaningful relationships rather than exploiting metric-specific artifacts. Furthermore, no unexpected metric deterioration is observed in the optimized configurations, suggesting that improvements in accuracy are not achieved at the expense of increased bias or reduced agreement, which can occur in some optimization settings when models become overly aggressive in fitting training patterns.

When positioned against prior studies in customer segmentation and marketing analytics, the present findings align with a growing body of evidence that advanced learning models benefit substantially from enriched feature spaces that include behavioral and financially relevant attributes. Previous research in customer analytics has repeatedly shown that segmentation quality improves when demographic variables are complemented with spending-intensity measures and financial-capacity indicators, because customer heterogeneity is multidimensional and cannot be adequately captured through demographics alone. The present study extends this conclusion by illustrating that, even with a richer feature space, predictive benefit depends strongly on optimizing model configuration. This supports the broader methodological argument that modern customer analytics is no longer driven solely by model choice, but increasingly by the integration of model architecture with systematic optimization mechanisms that govern how the model learns from economically structured data.

In financial prediction literature, especially in credit scoring and customer value modeling, prior work has emphasized that hyperparameter tuning is often decisive in determining whether a model generalizes reliably across customer populations. The findings here reinforce that perspective while also highlighting the potential of physics-inspired optimization frameworks to deliver strong performance without extensive algorithm-specific tuning. Many existing studies adopt evolutionary or swarm-based optimizers and report improvements, but these approaches frequently introduce additional tuning burdens related to optimizer parameters. In contrast, the behavior observed here suggests that KLA may reduce such burdens while still achieving high-quality solutions, which

is valuable in operational financial analytics where reproducibility, automation, and computational efficiency are often prioritized.

Compared with optimization-driven approaches in the literature, the observed performance hierarchy among optimizers is broadly consistent with the expectation that population-based algorithms tend to outperform manual tuning or simplistic search in complex hyperparameter landscapes. However, the dominance of KLA relative to other established optimizers highlights a key contribution of the present work. This outcome suggests that optimization designs rooted in physical equilibrium processes may provide an effective pathway to robust convergence, particularly when applied to customer behavior prediction tasks characterized by heterogeneous features and nonlinear relationships. For financial applications, this is especially important because stable optimization reduces the likelihood of deployment instability, supports consistent model retraining cycles, and can improve governance outcomes in regulated environments by reducing variance across training runs.

Despite the contributions of this study, several limitations should be acknowledged. First, although the dataset is enhanced to better reflect retail and financial customer behavior, some financially meaningful features—including estimated savings and credit score—are synthetically generated. While their construction is logically consistent, synthetic features may not capture the full distributional complexity, reporting noise, and structural biases present in real financial records. In practice, real-world credit and savings indicators are shaped by institutional rules, market frictions, behavioral biases, and macroeconomic shocks, which may not be fully represented in synthetic variables. Second, the analysis is conducted on a single dataset with a fixed experimental setup, which may limit generalizability across countries, income distributions, retail formats, or economic regimes. Third, the present evaluation is centered on predictive performance and goodness-of-fit metrics; interpretability and explainability, which are critical in financial decision-making and regulatory settings, are not explicitly examined. Fourth, although multiple optimizers are evaluated, the study focuses on hyperparameter optimization and does not fully explore feature-selection dynamics or joint optimization regimes that could further enhance efficiency.

Future research can address these limitations in several directions. Incorporating real transaction histories, credit bureau records, or longitudinal customer profiles would enhance realism and support stronger external validity. Integrating macroeconomic indicators, such as inflation, unemployment rates, and interest-rate regimes, could further improve financial realism by accounting for environment-level drivers of customer spending and credit behavior. Methodologically, extending the framework to joint feature selection and hyperparameter optimization would be particularly valuable in high-dimensional financial datasets, where redundancy and multicollinearity may increase over time due to derived features and behavioral aggregations. Moreover, incorporating explainable artificial intelligence techniques would improve transparency and interpretability, enabling institutions to understand the economic logic underlying predictions and supporting compliance and audit requirements. Finally, evaluating the proposed framework under streaming or real-time

deployment settings, such as adaptive marketing platforms or continuously updated credit-scoring engines, would provide critical insights into computational scalability and robustness under changing customer distributions.

Overall, the expanded discussion emphasizes that the proposed predictive framework contributes not only to methodological understanding but also to practical economic and financial analytics, where predictive reliability, robustness, and generalization are indispensable for effective and responsible decision-making.

6. CONCLUSION AND FUTURE WORK

This study has presented an extensive investigation into customer behavior prediction within retail and financial contexts, emphasizing that predictive value emerges from the interaction between model capacity and optimization quality. By leveraging an enriched customer dataset that integrates demographic descriptors, spending-behavior indicators, and financially meaningful variables such as estimated savings and credit score, the experimental design was structured to reflect economically relevant decision environments. The baseline evaluation demonstrated that even advanced models exhibit meaningful performance variation prior to optimization, underscoring that architecture alone does not guarantee reliable forecasting in high-dimensional customer datasets.

The empirical findings provide strong evidence that metaheuristic-driven hyperparameter optimization is a decisive determinant of predictive performance, stability, and generalization. In particular, the use of KLA as the primary optimization mechanism produced the most consistent improvements across all evaluation metrics, outperforming both unoptimized baselines and alternative optimization strategies. The performance behavior of KLA can be attributed to its physics-inspired formulation, which promotes a self-regulating balance between exploration and exploitation without requiring extensive parameter calibration. This property is especially valuable in financial analytics, where reproducibility, stable convergence, and robustness across retraining cycles are central requirements for operational deployment.

From an economic and financial standpoint, the implications of these findings are substantial. Accurate prediction of customer behavior—including spending intensity, savings capacity, and creditworthiness—supports a wide range of high-impact applications such as credit-risk management, customer segmentation, targeted marketing, portfolio prioritization, and personalized financial product design. Improvements in predictive accuracy and reductions in bias enhance institutional decision quality and can reduce financial losses arising from misclassification or inefficient resource allocation. In retail economics, such improvements translate into better promotion design, more effective loyalty programs, and improved inventory strategies. In financial services, they support more reliable credit scoring, improved customer profitability modeling, and better alignment between product offerings and customer risk profiles.

Beyond the immediate empirical contributions, this study highlights the broader methodological relevance of combining high-capacity learning models with advanced, physics-inspired optimization in economic analytics. As customer

datasets become larger, more heterogeneous, and more financially complex, traditional tuning approaches become increasingly impractical. Metaheuristic optimization provides a systematic pathway for navigating complex hyperparameter landscapes, while physics-inspired mechanisms such as KLA offer the additional benefit of reduced tuning overhead and potentially stronger convergence robustness. Accordingly, the framework proposed in this work contributes to the development of scalable and reliable predictive pipelines for data-driven economic decision systems.

Several future research directions remain relevant and necessary. First, extending KLA to jointly optimize feature selection and hyperparameters in a unified scheme would improve efficiency and may further enhance generalization, particularly in high-dimensional financial datasets characterized by redundancy and multicollinearity. Second, future work should incorporate real-world transactional histories, longitudinal customer profiles, and verified credit data to strengthen external validity and ensure that predictive performance holds under realistic institutional conditions. Third, incorporating macroeconomic context variables would enable dynamic forecasting under changing economic regimes and allow the framework to reflect how external shocks influence customer behavior. Fourth, explainability should be treated as a central design objective in future extensions, particularly for regulated financial applications where transparency, fairness, and accountability are mandatory. Finally, evaluating the proposed framework in real-time deployment scenarios would provide insights into computational scalability, retraining stability, and robustness under non-stationary customer behavior patterns.

In summary, this study demonstrates that a Kirchhoff's Law Algorithm-driven optimization framework can substantially enhance customer behavior prediction in retail and financial analytics by improving accuracy, reducing bias, and strengthening generalization. The outcomes contribute to both methodological advancements in optimization-driven learning and the practical development of high-precision decision-support tools for economically and financially sensitive environments.

DATA AVAILABILITY

The dataset used in this study is publicly available on Kaggle at <https://www.kaggle.com/datasets/vikasjigupta786/customer-analytics-practice-dataset>.

REFERENCES

- [1] P. Singh, L. Arora, and A. Choudhry, "Consumer behavior in the service industry: An integrative literature review and research agenda," *Sustainability*, vol. 15, no. 1, p. 250, 2023.
- [2] İ. H. Efendioğlu, "Digital consumer behavior: A systematic literature review," *Prizren Social Science Journal*, vol. 8, no. 1, pp. 67–80, 2024.
- [3] A. N. Korkmaz and M. U. Altan, "A systematic literature review of sustainable consumer behaviours in the context of industry 4.0 (I4.0)," *Sustainability*, vol. 16, no. 1, p. 126, 2024.
- [4] R. K. Behera, P. K. Bala, N. P. Rana, R. S. Algharabat, and K. Kumar, "Transforming customer engagement with artificial intelligence e-marketing: An e-retailer perspective in the era of retail 4.0," *Marketing Intelligence & Planning*, vol. 42, no. 7, pp. 1141–1168, 2024.
- [5] L. Theodorakopoulos and A. Theodoropoulou, "Leveraging big data analytics for understanding consumer behavior in digital marketing: A systematic review," *Human Behavior and Emerging Technologies*, vol. 2024, p. 3641502, 2024.
- [6] X. Zhang, F. Guo, T. Chen, L. Pan, G. Beliaikov, and J. Wu, "A brief survey of machine learning and deep learning techniques for e-commerce research," *Journal of Theoretical and Applied Electronic Commerce Research*, vol. 18, no. 4, pp. 2188–2216, 2023.
- [7] I. Lewaelhamd, "Customer segmentation using machine learning model: An application of RFM analysis," *Journal of Data Science and Intelligent Systems*, vol. 2, no. 1, pp. 29–36, 2024.
- [8] D. S. K. T. Anish, C. Balakrishnan, D. S. Lakumarapu, D. P. B. Pajila, and R. Siva Subramanian, "Leveraging machine learning for customer intelligence: An experimental analysis learning classifiers," *Procedia Computer Science*, vol. 230, pp. 128–137, 2023.
- [9] G. H. Lee, K. J. Lee, B. Jeong, and T. Kim, "Developing personalized marketing service using generative AI," *IEEE Access*, vol. 12, pp. 22 394–22 402, 2024.
- [10] S. Birim, I. Kazancoglu, S. K. Mangla, A. Kahraman, and Y. Kazancoglu, "The derived demand for advertising expenses and implications on sustainability: A comparative study using deep learning and traditional machine learning methods," *Annals of Operations Research*, vol. 339, no. 1, pp. 131–161, 2024.
- [11] A. Manzoor, M. A. Qureshi, E. Kidney, and L. Longo, "A review on machine learning methods for customer churn prediction and recommendations for business practitioners," *IEEE Access*, vol. 12, pp. 70 434–70 463, 2024.
- [12] N. Kshetri, Y. K. Dwivedi, T. H. Davenport, and N. Panteli, "Generative artificial intelligence in marketing: Applications, opportunities, challenges, and research agenda," *International Journal of Information Management*, vol. 75, p. 102716, 2024.
- [13] M. Schmitt, "Deep learning in business analytics: A clash of expectations and reality," *International Journal of Information Management Data Insights*, vol. 3, no. 1, p. 100146, 2023.
- [14] B. Jamalpur, D. Singh, B. S. Kumar, A. Nagpal, S. Pawar, and D. Banerjee, "Applications of deep learning in marketing analytics: Predictive modeling and segmenting customers," *Proceedings of the 2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE)*, pp. 1473–1477, 2024.

-
- [15] M. Afzal, S. Rahman, D. Singh, and A. Imran, "Cross-sector application of machine learning in telecommunications: Enhancing customer retention through comparative analysis of ensemble methods," *IEEE Access*, vol. 12, pp. 115 256–115 267, 2024.
 - [16] I. Markoulidakis, I. Rallis, I. Georgoulas, G. Kopsiaftis, A. Doulamis, and N. Doulamis, "A machine learning based classification method for customer experience survey analysis," *Technologies*, vol. 8, no. 4, p. 76, 2020.
 - [17] S. R. Ahmed *et al.*, "Deep learning for customer relationship management in e-commerce," *Proceedings of the 2024 International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA)*, pp. 1–7, 2024.
 - [18] S. Curiskis, X. Dong, F. Jiang, and M. Scarr, "A novel approach to predicting customer lifetime value in B2B SaaS companies," *Journal of Marketing Analytics*, vol. 11, pp. 587–601, 2023.
 - [19] N. Na, "Research on supply chain demand prediction model based on LSTM," *Procedia Computer Science*, vol. 243, pp. 313–322, 2024.
 - [20] X. Cui, Z. Zhu, L. Liu, Q. Zhou, and Q. Liu, "Anomaly detection in consumer review analytics for idea generation in product innovation: Comparing machine learning and deep learning techniques," *Technovation*, vol. 134, p. 103028, 2024.
 - [21] E. H. Houssein, M. K. Saeed, G. Hu *et al.*, "Metaheuristics for solving global and engineering optimization problems: Review, applications, open issues and challenges," *Archives of Computational Methods in Engineering*, vol. 31, pp. 4485–4519, 2024.