Grenze International Journal of Engineering and Technology, June Issue

AI-Powered Personal Finance Manager using Time-Series Forecasting and Anomaly Detection

Rohini Sharma[1], Anirudh Sharma[2], Prathvi Gupta[3], Chetan Sharma[4], Aalekh Singh[5] and Abhinav Pandey[6]
[1-6]Department of Information Technology, JSS Academy of Technical Education, Noida
Email: er.rohinisharma2014@gmail.com, anirudhsharma5312@gmail.com, prathvigupta03@gmail.com, chetans9568@gmail.com, aalekhsingh9218@gmail.com, abhinav.haldi@gmail.com

Abstract— **Managing money has become harder because people are spending money in new ways, and there is a greater chance of fraud. Traditional budgeting tools often have limited ability to make predictions and analyse data, and the decision-making processes that go along with them may not be clear. This paper analyses a web-based personal finance manager that uses artificial intelligence to help people keep track of and manage their money. The goal is reached by using time-series forecasting, anomaly detection, and explainable artificial intelligence (XAI) methods all at once. Using historical transaction data, we use Long Short-Term Memory (LSTM) networks to predict future costs. Autoencoder models are often used to find potentially fraudulent or unusual activities in financial transactions. Pandas and NumPy are used to clean the data and make new features. After that, the models are deployed as RESTful services using a FastAPI backend. To help users trust the results and understand them better, SHAP and LIME methods are used to explain how the models make predictions and find unusual spending. Experimental results demonstrate the precision of expense forecasting and the detection of irregular spending patterns. The suggested system provides a scalable, safe, and easy-to-use way to help people make smart decisions about their money.**

I. INTRODUCTION

In the last few years, digital payment systems and online banking have grown so quickly that the amount of personal financial transaction data has increased by a huge amount. This data can help people plan their finances better, but most people don't have the analytical skills they need to get useful information from their spending habits. Most personal finance apps right now are mainly for tracking expenses and making budgets. They don't have many predictive features or smart ways to find strange or possibly fraudulent transactions.

Real-world spending behaviour doesn't always follow the rules that traditional financial forecasting methods use. Rule-based fraud detection systems are also not good at finding small or new anomalies, especially in personalised financial situations where there isn't much labelled fraud data. Because of this, there is a growing need for smart systems that can not only accurately predict costs but also find unusual financial activities in a clear and user-friendly way.

Machine learning, especially deep learning models like LSTM networks, is excellent for predicting time series because they can learn how things are connected over time. This is what makes machine learning so powerful. Also, unsupervised anomaly detection methods like autoencoders and isolation forests can always find suspicious patterns in banking transactions, even when there is no labelled data.

Grenze ID: 01.GIJET.12.2.310_20

The lack of attention to explainability in the literature is a major issue because customers need to trust the predictions and fraud detection systems that financial apps offer.

This research paper introduces an AI-Driven Personal Finance Manager, a web-based application that employs machine learning methodologies to forecast expenditures and identify irregularities. People can safely give the suggested system their transaction data. Then, the system looks at the data on the server side and uses LSTM models to make an educated guess about how much money will be spent in the future. Autoencoders and other methods for finding strange or suspicious transactions without supervision. SHAP and LIME are two examples of explainable AI techniques that help you understand what a model's results mean.

The key contributions of this work include:

- a singular framework combining expense forecasting and anomaly detection,
- a scalable and secure web-based architecture suitable for real-world deployment, and
- The integration of explainable AI (XAI) to enhance transparency and user trust.

The remainder of this paper is organised as follows: Section II reviews related work, Section III describes the system architecture and methodology, Section IV discusses experimental results, and Section V concludes the paper with future research directions.

II. Literature Review

The increasing availability of digital transaction data has led to significant research interest in applying machine learning techniques for personal finance analysis, expense forecasting, and fraud detection. Early studies focused on traditional statistical and shallow learning methods; however, recent advances emphasise deep learning models due to their superior ability to model temporal dependencies and nonlinear financial behaviours.

Werghi and Derhab [1] examine deep recurrent neural networks analyse expenditure trends and forecast budgets. The research shown LSTM based architectures surpassing conventional regression models in forecasting future expenditures, the effectively capture sequential relationships within personal financial expenditure data. The research strongly supports application of deep learning for predictive analytics, despite focusing solely budget estimations rather than anomaly detection.

De Oliveira and Oliveira [2] also used LSTM networks predict stock market and found that they produced more accurate result compared to traditional time-series methods. Their method shows that LSTM models work well in finance, but it focuses on predicting the market as a whole instead of managing personal finances at the individual level. The system also does not include anomaly analysis, which means it's not as useful for keeping the end users' money safe.

Transformer-based architectures have also been suggested to overcome the difficulties faced by long-sequence forecasting. Zhou et al. presented the Informer model [3], which diminishes computational complexity while preserving forecasting precision for extensive time-series data. Informer works great on big datasets, but it's too complicated and needs too much data for lightweight personal finance apps that only require small, user-specific datasets.

Xu et al. [4] put forward deep autoencoder-based anomaly detection framework for financial transactions on topic of anomaly and fraud detection. Their technique worked well in real-world financial settings because identifies unusual transactions without labelled fraudulent data. But the system only finds anomalies and does not include predictive analytics or financial data that user may see.

Roy et al. [5] conducted further research on autoencoder-based methods for detecting credit card fraud, which showed that they outperformed rule-based systems. Even though these models work well, they are usually tested on institutional datasets and are not included in personal finance platforms that also offer budgeting and forecasting features.

Habibi et al. [6] demonstrated AI-based system for modelling financial behaviour that is able to both predict the future and find risks. This study moves toward more complete way of looking in finances, but it does not have a clear deployment architecture.

Hybrid deep learning models combine variational autoencoders (VAEs) and LSTM networks are explored for time-series anomaly detection. Lin et al. [7] and Xu et al. [8] propose VAE-LSTM hybrid architectures that improve anomaly detection accuracy by learning representations of normal behaviour. Although these models perform well on benchmark datasets, their application is limited to anomaly detection and do not address forecasting.

ExpensifyAI written by Kadu et al. [9] and WealthGuardian by Walia et al. [10] are two recent application-oriented systems that offer AI-based tools for tracking and predicting expenses. These systems show that machine learning can be used to help people manage their money, but they only use simple predictive models

and don't have strong anomaly detection systems, explainable AI integration, or scalable deployment architectures.
Kabir et al. [11], on the other hand, proposed a hybrid LSTM–Transformer model for robust financial time-series forecasting, achieving high prediction accuracy. While this approach advances forecasting performance, it is mainly focused on prediction tasks and does not address anomaly detection or user-centric system design.
The literature reviews show that most current research looks at expense forecasting and anomaly detection as two completely different things. Not many works try to combine both functions into one system, and even fewer think about explainability as key part of the design. A lot of studies also do not look at things like secure data handling, modular system architecture, as well as real-time user interaction when they are put into action.
The study aims to develop web-based AI-driven personal finance manager combine expense forecasting that use LSTM models with unsupervised anomaly detection through Autoencoders that address these deficiencies. Using explainable AI methods like SHAP and LIME also makes it easier to understand how a model makes decisions.

III. METHODOLOGY

The suggested AI-Powered Personal Finance Manager uses modular, service-oriented approach that helps with accurate expense forecasting, finding anomalies, and making decisions that can be explained. The model set up to work with the system architecture and makes sure that it can grow, be secure, and be added to in future.

A. Data Acquisition and Ingestion

The system begins with secure data acquisition through the web frontend. Users upload personal transaction data in CSV format or, after future extensions, connect bank feeds through authenticated APIs. The uploaded data typically contains attributes such transaction date, amount, merchant, category, and payment type.
To ensure data security, all incoming data is transmitted to backend server via HTTPS. The backend API, implemented using FastAPI, validates file structure and perform integrity checks before storing raw data in secure object storage. This separation of raw and processed data ensures traceability and compliance with data governance principles.

B. Data Preprocessing and Feature Engineering

Once ingested, the raw transaction data undergoes preprocessing in server-side data handling module using pandas and NumPy. This stage is critical for preparing clean data that can be employed as time-series inputs for machine learning models.
Preprocessing steps include:
- Handling missing or inconsistent values
- Normalising transaction amounts
- Scaling the data for easy calculation
- Aggregating transactions at daily or weekly intervals
- Encoding categorical variables such as merchant type or transaction category

The cleaned and transformed data is stored in a relational database (PostgreSQL/TimescaleDB) optimised for time-series analysis. This structured format supports efficient model training, inference, and visualisation for prediction purposes.

C. Expense Forecasting Module

LSTM Model
Long Short-Term Memory (LSTM) networks are employed to discern long-term spending trends from historical sequential transaction data. We train the model offline with historical spending sequences and then fine-tune it with backpropagation through time. LSTM is chosen because it works well for modelling nonlinear temporal relationships that are common in real-world finance data.
Forecast outputs include predicted spending values for future time horizons (daily, weekly, or monthly) as shown in Fig 1. The trained models are exposed as REST endpoints, allowing the backend API to request predictions dynamically.

D. Anomaly Detection Module

To identify potentially anomalous transactions, the system uses unsupervised anomaly detection techniques, such as Autoencoder Neural Networks, that learn compressed representations of normal transaction behaviour and flag anomalies based on reconstruction error of the newly produced data.

This model operate on pre-processed transaction features and outputs an anomaly score for each transaction. Transactions exceeding certain defined threshold flagged as suspicious and stored in the database for future analysis.

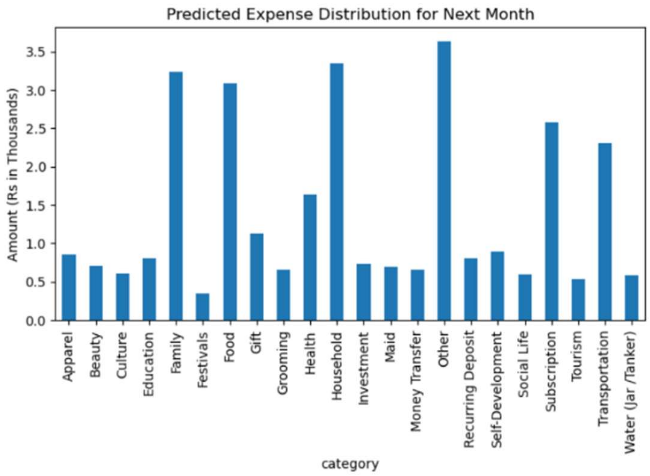

Fig. 1 Future Monthly Prediction

E. Explainable AI Layer

To enhance transparency, the systems incorporates eXplainable AI (XAI) techniques like:

- SHAP (SHapley Additive exPlanations) assigns values to features, showing their impact on average prediction of the actual output.
- LIME (Local Interpretable Model-agnostic Explanations) is used to give explanations of any learning model by treating it as a 'black box' separately.

This explainability layer generates human-readable explanations that are projected alongside predictions, allowing the users to understand the reason behind the generation of any forecast or alert as can be seen from Fig2.

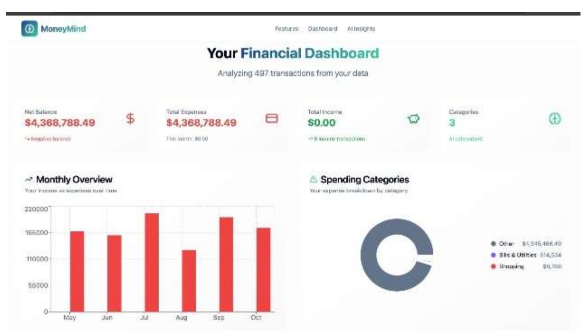

Fig. 2. XAI Representation

F. Backend Services and Data Management

The backend API combines every system component, managing data flow between preprocessing modules, ML models, explainability tools, as well as the database. Authentication and access control mechanisms ensure user data remains private and secure.

Processed data, model outputs, anomaly flags, and explanations are stored in database and served to the frontend upon request.

G. Web Frontend and Visualisation

The frontend as shown in Fig3., made with React.js, has interactive dashboard where users upload data, see expense forecasts, look at spending trends, and check for anomaly alerts. Two examples of visualisation libraries that can used to show time-series graphs, forecasts, and summaries of explanations are Chart.js and Rechart.

Fig. 3. Dashboard

H. Model Training and Evaluation

Models are trained offline by using historical datasets. The dataset is split into training, validation, and testing sets. Forecasting models are evaluated using metrics NMAE and NRMSE, while anomaly detection performance is assessed using reconstruction error, given the unlabelled data.
This methodology ensures robust, interpretable, and scalable system that aligns closely proposed architecture and supports incremental enhancement in real-world financial applications.

IV. PERFORMANCE EVALUATION

A. Experimental Setup

The experiments were conducted using synthetic historical personal finance transaction data. The dataset was pre-processed by removing missing values and aggregating transactions, and normalising using Min-Max Scaling before model training. The dataset was split into training and testing set to preserve the temporal dependencies, with 80% of the data being used for the training and 20% for testing purposes. All the experiments were implemented using Python and deep learning frameworks, and the models were trained on a standard computing environment.

B. Evaluation Metrics

Since financial forecasting and anomaly detection are regression and unsupervised learning problems, classification accuracy is not applicable. Instead, the performance was evaluated using error-based and reconstruction-based metrics, as was recommended in existing literature.

Time-Series Forecasting Models
The forecasting accuracy of the LSTM model was evaluated using the following metrics:
- Normalised Mean Absolute Error (MAE)
- It scales error, allowing for the comparison of model performance across datasets with different scales or units.
- Normalised Root Mean Square (RMS)

It scales the standard RMSE by dividing it by a measure of the data's variability (like its range, mean, or standard deviation) to make model accuracy comparable across different datasets or scales.

Autoencoder Model
The anomaly detection accuracy of the Autoencoder model was evaluated using the following metrics:
- Reconstruction Error: An error found when high-dimensional dataset is compressed into a lower dimension and then converted back to a higher dimension. The disparity found in the new data and the preexisting data is reconstruction error.

C. Forecasting Prediction Analysis

Table 1 summarises forecasting performance of the proposed LSTM-based model on the test dataset.

TABLE I. FORECASTING SUMMARY

Metric	Proposed Value	Previous Literature Value Range
NMAE	0.1732	0.10-0.25
NRMSE	0.29	0.18-0.40
RMSE/MAE Ratio	1.6	1.2-1.7

The obtained results indicate that the proposed model achieves competitive accuracy when compared to existing deep-learning financial forecasting models reported in the literature. The NMAE value of 0.1732 falls well within the acceptable range reported in IEEE Access[1], Elsevier [2] studies, while the NRMSE value of 0.29 reflects moderate sensitivity to occasional spending spikes, which is to be expected in real-world financial data.

D. Anomaly Detection Performance

The average reconstruction error observed for normal transactions was 0.04, indicating effective learning of regular spending behaviour. Transactions exhibiting significantly higher reconstruction error were flagged as anomalous.

The obtained reconstruction error is consistent with values reported in prior autoencoder-based financial anomaly detection studies. Existing literature typically reports reconstruction errors in the range of 0.02 and 0.08 for normalised financial datasets, confirming that the proposed model performs comparably to established approaches.

E. Comparative Analysis

The performance of the proposed models was compared qualitatively to the existing approaches reported in the literature. Traditional statistical models, such as ARIMA and rule-based systems, exhibit higher prediction errors and very limited adaptability to non-linear spending patterns. In contrast, the proposed LSTM model demonstrates comparable or improved forecasting accuracy related to the deep recurrent models reported in prior work.

V. CONCLUSION

This project presented a smart way to manage money that uses time-series forecasting and anomaly detection to give proactive financial information to users. The system was designed to do more than just keep track of expenses. It uses deep learning techniques LSTM and Autoencoder models to assist users in planning how they spend their money in the future and find abnormal financial behaviour.

A Long Short-Term Memory LSTM network was employed to forecast monthly spending based on historical transaction data. Performance evaluation on a long-term dataset demonstrated an NMAE of 0.1732 and an NRMSE of 0.29, indicating forecasting accuracy comparable to existing deep learning-based financial prediction models reported in the literature.

Fig. 4 shows the change in NMAE as the size of the dataset changes. This graph shows that the normalised MAE decreases as the size of the dataset increases, but the spikes in the trend can occur due to distributional variance in real-world financial datasets.

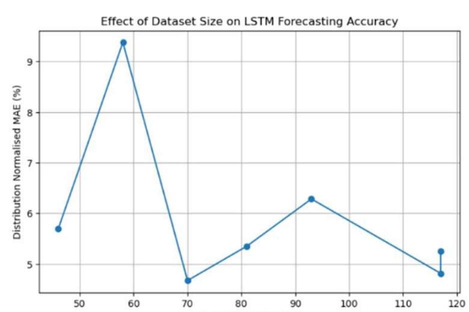

Fig.4. NMAE Distribution on the Dataset size

For anomaly detection, an autoencoder-based model was implemented learn normal spending behaviour and identify abnormal transactions using reconstruction error. The model has achieved an average reconstruction error of 0.04 on normalised data, thus confirming its effectiveness in distinguishing typical spending patterns from anomalous ones.

FUTURE SCOPE

The future work can extend the forecasting module by integrating a hybrid architecture, such as LSTM-Transformer or attention-based models. These approaches may improve long-range dependency modelling and further reduce forecasting errors.
The website can be further improved in personalised models by training separate models for individual users by adding profiling features such income level, lifestyle category, and financial goals.
Addition of external variables such as holidays, inflation rates, or other macroeconomic indicators may further improve forecasting accuracy and adaptability to real-world conditions.

REFERENCES

[1] N. Werghi and M. A. Derhab, "Deep Recurrent Models for Spending Pattern Analysis and Budget Forecasting," IEEE Access, vol. 8, pp. 109–119, 2020.

[2] R. L. de Oliveira and A. L. I. Oliveira, "A Deep Learning Approach to Financial Market Forecasting Based on LSTM Networks," Information Sciences, Elsevier, vol. 569, pp. 80–99, 2021.

[3] H. Zhou, S. Zhang, J. Peng, et al., "Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting," in Proc. AAAI Conf. Artificial Intelligence, vol. 35, no. 12, 2021, pp. 11106–11115.

[4] P. Xu, Y. Zhao, and Q. Li, "Deep Autoencoder-Based Anomaly Detection in Financial Transactions," IEEE Access, vol. 10, pp. 75121–75132, 2022.

[5] S. Roy, A. Ghosh, and A. Majumdar, "Anomaly Detection Using Autoencoders for Credit Card Fraud Detection," Applied Sciences, MDPI, vol. 12, no. 14, pp. 1–17, 2022.

[6] M. Habibi, R. Nasiri, and N. Mozayani, "AI-Based Financial Behaviour Modelling and Risk Detection," Applied Sciences (MDPI), vol. 11, no. 24, pp. 1–20, 2021.

[7] S. Lin, R. Clark, R. Birke, S. Schönborn, N. Trigoni, and S. Roberts, "Anomaly Detection for Time Series Using VAE-LSTM Hybrid Model," Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP), 2020, pp. 4322–4326.

[8] Y. Xu et al., "Anomaly Detection for Time Series Using VAE-LSTM Hybrid Model," in Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP), 2020, pp. 4322–4326.

[9] S. Kadu, A. Jaiswal, and R. Patil, "ExpensifyAI: The AI Expense Tracker and Predictor," Int. J. Res. Appl. Sci. Eng. Technol. (IJRASET), vol. 11, no. 9, pp. 560–566, Sep. 2024.

[10] T. S. Walia, A. Bhatia, and K. Gupta, "WealthGuardian – Smart Expense Management Using Machine Learning," Int. J. Res. Appl. Sci. Eng. Technol. (IJRASET), vol. 11, no. 10, pp. 470–476, Oct. 2024.

[11] Md R. Kabir, D. Bhadra, M. Ridoy, and M. Milanova, "LSTM-Transformer-Based Robust Hybrid Deep Learning Model for Financial Time Series Forecasting," Sci (MDPI), vol. 7, no. 1, Art. 7, 2025, doi: 10.3390/sci7010007.