Tharaa: An Integrated Mobile System for Personal Finance Management with OCR, Predictive Analytics, and Intelligent Guidance

Ghaida Alotaibi
Department of Computer Information Systems
Imam Abdulrahman bin Faisal University
Dammam, Saudi Arabia

Rawan Alghamdi
Department of Computer Information Systems
Imam Abdulrahman bin Faisal University
Dammam, Saudi Arabia

Shahd Aldokhi
Department of Computer Information Systems
Imam Abdulrahman bin Faisal University
Dammam, Saudi Arabia

Majd Aldossary
Department of Computer Information Systems
Imam Abdulrahman bin Faisal University
Dammam, Saudi Arabia

Zahraa AlFraid
Department of Computer Information Systems
Imam Abdulrahman bin Faisal University
Dammam, Saudi Arabia

Hadeel Alotaibi
Department of Computer Information Systems
Imam Abdulrahman bin Faisal University
Dammam, Saudi Arabia

Areej Buker
Department of Computer Science
Imam Abdulrahman bin Faisal University
Dammam, Saudi Arabia

ABSTRACT
Personal financial management is increasingly challenging for individual users because daily spending is fragmented across physical payments, online transactions, recurring subscriptions, and installment plans. This paper presents Tharaa, a unified personal-finance mobile system that consolidates expense and income tracking, budget and goal management, subscription and installment monitoring, on-device Optical Character Recognition (OCR) for receipt capture, machine-learning-driven predictive analytics, and an AI-powered finance coach into a single Flutter mobile application backed by a Firestore cloud database. The system embeds three predictive models — Logistic Regression for end-of-month overspending classification, Lasso regression for per-category spending forecasting, and Isolation Forest for anomaly detection — that all run on-device using exported coefficients, so no transactional data leaves the user's phone for prediction. We evaluate the three models on two datasets. The first is a controlled synthetic dataset spanning four behavioral personas (lifestyle overspender, installment-heavy, income-stable saver, subscription-driven). The second is the publicly available Daily Household Transactions dataset from Kaggle (2,461 rows, 45 calendar months, January 2015 – September 2018). The overspending classifier achieves 96.0% accuracy with precision 1.00 on the synthetic test set, and a macro-averaged recall of 0.82 on the Kaggle test split despite extreme class imbalance (91.7%/8.3%). Category forecasts produced by Lasso track actual spending closely (Mean Absolute Error of SAR 222.76 on the synthetic dataset for Food). The Isolation Forest detects an overall anomaly rate of 8.02% on the Kaggle dataset, in close agreement with the 8.02% observed across the synthetic evaluation pool. The results confirm that the integrated approach is feasible across both controlled and real-world settings, and that on-device inference preserves user privacy without sacrificing predictive performance.

General Terms
Mobile Applications, Personal Finance, Financial Technology, Machine Learning, Privacy.

Keywords
Personal Financial Management, Mobile Application, OCR, Chatbot, Predictive Analytics, Flutter, Anomaly Detection, Financial Awareness.

1. INTRODUCTION
The rapid digitization of financial transactions has fragmented personal spending data across physical payments, online shopping, recurring subscriptions, and installment plans. This dispersion makes precise personal budgeting difficult using traditional tracking methods [1]. Despite the growing need for financial-management tools, most existing mobile applications remain descriptive — they record what was spent but do not anticipate what is coming next, do not detect unusual activity, and do not produce explanations that a non-financial user can

act on. National surveys for the Kingdom of Saudi Arabia also report low formal account usage and weak savings habits, especially among young adults, underscoring the need for tools that actively promote financial awareness rather than merely report past activity [2].

In response, Tharaa was developed as a comprehensive personal-finance management system for individual users. Its primary contribution is a unified design in which six integrated modules reinforce one another: transaction and budget management, subscription and installment tracking, OCR-based receipt capture, an AI-powered finance coach, a predictive analytics layer, and a hybrid insight generator. The predictive analytics layer trains three machine-learning models on a real-world public dataset and embeds the learned coefficients in the mobile application, so all predictions execute on-device without transmitting financial data to a remote server.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 describes the system architecture and modules. Section 4 presents the experimental evaluation on two datasets. Section 5 discusses the implemented system. Section 6 concludes the paper and outlines future work.

2. RELATED WORK

Research relevant to Tharaa spans personal financial management, mobile finance applications, financial technology, AI-based financial assistance, OCR, and FinTech adoption in the Kingdom of Saudi Arabia. Together these areas define the landscape against which this work is positioned.

2.1 Personal Financial Management

Financial literacy is strongly associated with an individual's capacity to plan, save, and make sound financial decisions. Lusardi and Mitchell showed that low financial literacy negatively affects long-term financial behavior [3]. This evidence underscores the need for tools that actively promote financial awareness rather than only providing data summaries.

2.2 Mobile Finance Applications

Mobile applications have become a central channel for personal financial management. A representative system supporting budgeting, expense tracking, report generation, and barcode scanning was demonstrated by Stefanov et al. [4], confirming the practical value of mobile financial tools. However, that system remained primarily descriptive: it organized what users had already spent but did not anticipate future spending or flag unusual activity. A systematic review by Freitas et al. [5] of user adoption factors for mobile personal-finance applications corroborated this observation, identifying limited predictive support and weak personalization as recurring gaps.

2.3 Financial Technology

The FinTech landscape encompasses technology-driven financial services ranging from payments and lending to investment and personal finance [6, 7, 8]. While the literature on institutional FinTech ecosystems and business models is extensive, integrated consumer-oriented applications for daily personal-finance management remain under-represented [9]. Robo-advisor adoption studies further highlight the importance of user trust and perceived control in financial AI [10].

2.4 AI and Chatbots in Finance

AI-based financial services have been found to depend on perceived usefulness, ease of use, and user trust for adoption [10]. Within the Saudi banking context, factors such as performance expectancy, perceived intelligence, and social influence have been identified as positive drivers of chatbot adoption [11]. Surveys of chatbot technology [12] and of conversational-agent design [13] provide the foundation on which our task-oriented finance coach is built.

2.5 OCR for Receipt Digitization

Optical character recognition extracts machine-readable text from printed documents and has been studied extensively in document-analysis contexts [14]. Despite OCR's technical maturity, its integration into consumer-facing personal-finance systems — particularly within platforms that also provide predictive guidance — has received comparatively little attention. Tharaa uses Google ML Kit's on-device Text Recognition API [15] so receipt images never leave the user's device.

2.6 FinTech Adoption in Saudi Arabia

The Saudi financial sector is undergoing rapid digital transformation within the Vision 2030 framework. A conceptual framework for FinTech adoption among Saudi banks identified regulatory support, infrastructure maturity, and customer readiness as key enablers [16]. The Priantinah et al. study on the Technology Acceptance Model for personal-finance mobile applications further confirms that perceived usefulness and ease of use are the dominant adoption drivers in this region [17].

2.7 Research Gap

Several gaps emerge from the reviewed literature. First, FinTech research focuses predominantly on institutional-level services rather than integrated daily-use personal-finance tools [9]. Second, existing mobile finance applications offer limited behavioral feedback and lack predictive intelligence [4, 5]. Third, OCR remains under-integrated into the personal-finance workflow despite its potential to reduce manual entry friction [14, 15]. Tharaa addresses these gaps with a unified mobile system that combines OCR, on-device predictive analytics, and an AI-powered finance coach in a single privacy-preserving package.

3. PROPOSED SYSTEM AND METHODOLOGY

Tharaa is a data-driven personal-finance management system implemented as a Flutter mobile application backed by a cloud-based Firestore database [18]. The architecture comprises six integrated modules: (1) User Profile and Data Management; (2) Transaction and Budget Management; (3) Subscription and Installment Tracking; (4) OCR-Based Receipt Capture; (5) Finance Coach Chatbot; and (6) Predictive Analytics with Hybrid Insight Generation.

3.1 System Architecture

A cloud-based Firestore backend [18] provides real-time data synchronization and scalable NoSQL storage. Each user's financial records are fully isolated at the data-model level, enforcing strict privacy through Firebase Security Rules. Administrative access is limited to account management and per-user backup operations; administrators have no view of transaction details. All three machine-learning models execute on-device using exported coefficients, so no transactional data is transmitted for inference.

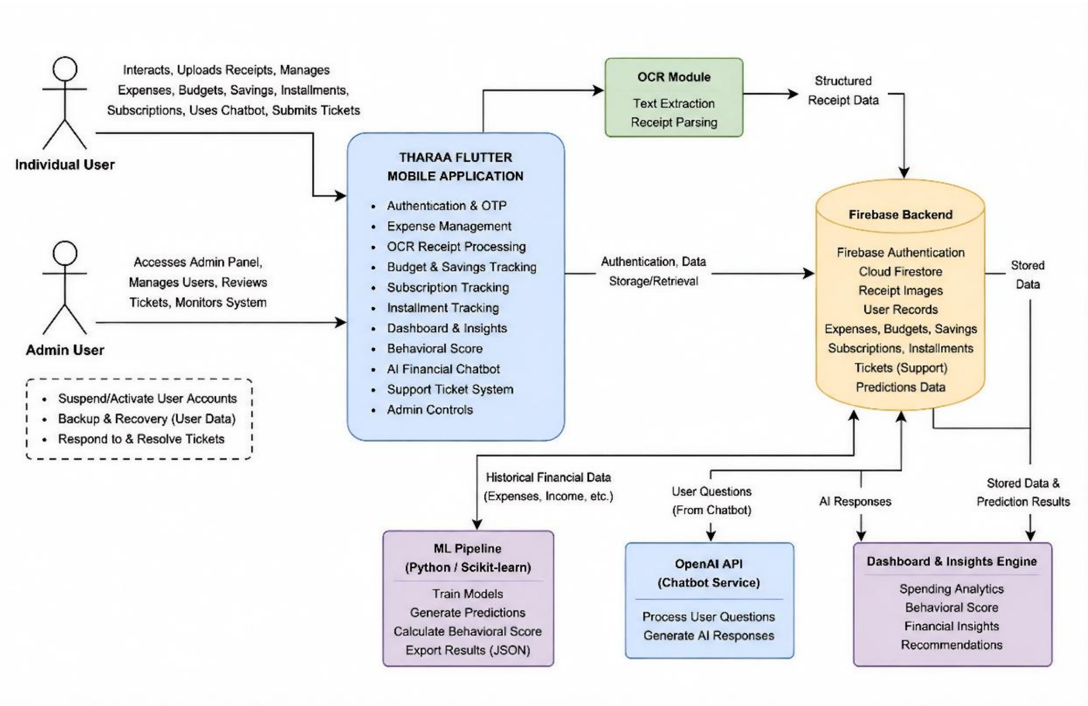

Figure 1: Tharaa system architecture overview

3.2 Transaction and Budget Management

Income and expenses are recorded manually with category assignment. The system aggregates this data to produce real-time balance summaries and category-level spending breakdowns. Monthly spending limits can be defined per category, and automated alerts are triggered when thresholds are approached or exceeded. Long-term financial goals (e.g., emergency funds or specific purchases) are tracked with deadline-based progress indicators.

3.3 Subscription and Installment Tracking

A dedicated module manages recurring financial obligations. Obligations are entered with due dates, amounts, and recurrence patterns, and the system schedules push notifications ahead of each due date. The same data feeds the commitment-pressure metric used by the insight layer, allowing the system to reason about the share of income already committed before a month begins.

3.4 OCR-Based Receipt Capture

Manual transaction entry is the primary friction point in personal-finance applications. The OCR module addresses this by enabling users to scan or upload receipt images; processing is performed entirely on-device using Google ML Kit's Text Recognition API [15], so no images are transmitted to external servers. Extracted fields (merchant, total, date) are matched against the user's category rules and recent history, then presented as a pre-filled transaction form that the user confirms or edits. Figure 2 summarizes the flow.

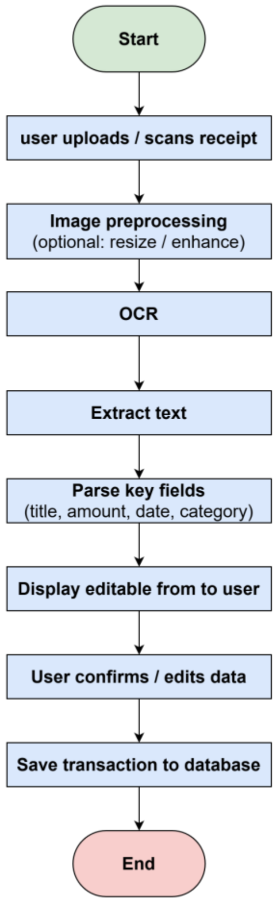

Figure 2: OCR Receipt Capture Flowchart

3.5 Finance Coach Chatbot

The finance coach provides a conversational interface following the principles of task-oriented dialogue systems [13]. Each time a user submits a query, the system retrieves the user's transaction history, budget usage, goal status, and upcoming obligations, builds a minimized structured context (with no full account identifiers and no personally identifying data), and sends it to the OpenAI Chat Completions API over HTTPS. The assistant's reply is rendered as a chat bubble grounded in the user's own data. Figure 3 summarizes the flow.

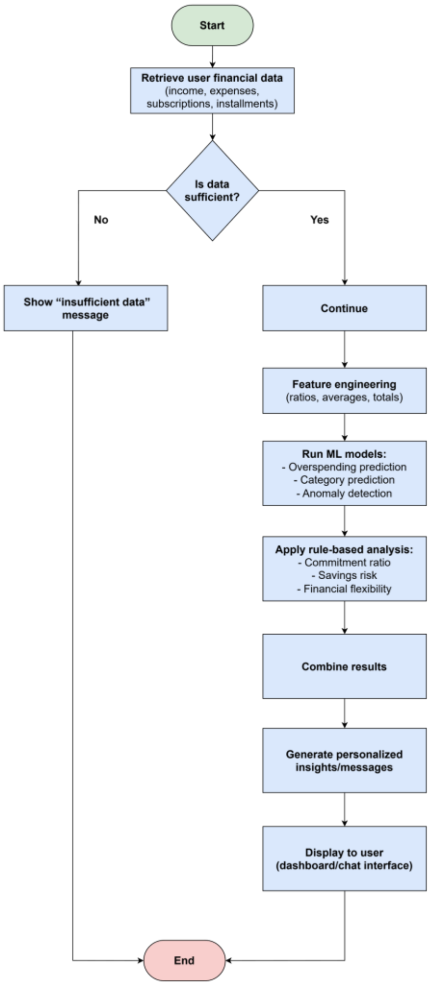

Figure 3: Finance Coach Chatbot Flowchart

3.6 Predictive Analytics Module

A background analytics layer feeds intelligence into both the dashboard and the chatbot. Three components are included: (1) Overspending Prediction — a Logistic Regression classifier trained on mid-month transaction patterns to predict whether the user will exceed budget by month-end; (2) Category Spending Forecasting — Lasso regression models, one per spending category, that produce next-period forecasts; and (3) Anomaly Detection — an Isolation Forest trained on per-user transaction features to flag unusual transactions. All three models are trained off-device on the public Daily Household Transactions dataset, and their coefficients are exported into compile-time Dart constants embedded in the mobile application, so inference is performed on-device.

4. EXPERIMENTAL RESULTS

Each predictive component of Tharaa was evaluated on two complementary datasets, in response to the recommendation that the evaluation cover multiple datasets and scenarios. The first dataset (D1) is a controlled synthetic dataset, constructed to span four user behavioral personas — lifestyle overspender, installment-heavy, income-stable saver, and subscription-driven — and used in the original conference version of this work. The second dataset (D2) is the publicly available Daily Household Transactions dataset from Kaggle, comprising 2,461 transactions over 45 calendar months (January 2015 – September 2018) with mixed income, expense, and transfer entries across more than thirty spending categories. For D2, dataset categories were mapped onto the ten fixed Tharaa categories and savings/transfer entries were excluded so they would not bias the spending models. Hyperparameters were fixed prior to evaluation; no test-set tuning was performed.

4.1 Overspending Prediction

A Logistic Regression classifier with balanced class weights was trained on six engineered features: month-to-date expense, predicted month-end expense (linear extrapolation), budget ratio, subscription burden, category-shift index, and time-of-month indicator. Features were scaled with a StandardScaler whose mean and standard deviation are exported alongside the model coefficients for on-device inference.

4.1.1 Synthetic Dataset (D1)

The synthetic dataset comprised 120 monthly records (90 overspending cases). On a held-out test split, the model achieved 96.0% accuracy, precision of 1.00 on the overspending class, recall of 0.94, and an F1-score of 0.97. The confusion matrix recorded 12 true positives, 4 true negatives, 2 false positives, and 6 false negatives, with a mean predicted probability of 0.6051.

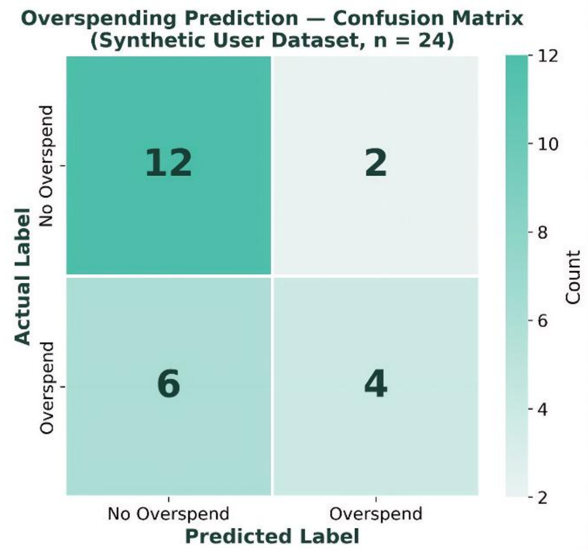

Figure 4: Confusion matrix — synthetic dataset (D1)

4.1.2 Daily Household Transactions (D2)

On D2, 45 calendar months were aggregated to monthly feature vectors. Because individual users rarely overspend in this household dataset, the monthly overspending label is highly imbalanced (≈ 8% positive). The model was therefore evaluated using stratified holdout (n = 12 test months) and 5-fold stratified cross-validation. On the holdout split, the classifier achieved an accuracy of 0.667, a macro-averaged precision of 0.60, a macro-averaged recall of 0.818, and a recall of 1.00 on the minority (overspend) class — that is, it captured every overspending month in the test split. Table 1 reports per-class metrics, and Figure 5 visualizes the confusion matrix.

Table 1: Overspending Prediction — per-class metrics on the Daily Household Transactions dataset (D2)

Class	Precision	Recall	F1-Score	Support
Not overspend	1.000	0.636	0.778	11
Overspend	0.200	1.000	0.333	1
Macro avg	0.600	0.818	0.556	12
Weighted avg	0.933	0.667	0.741	12

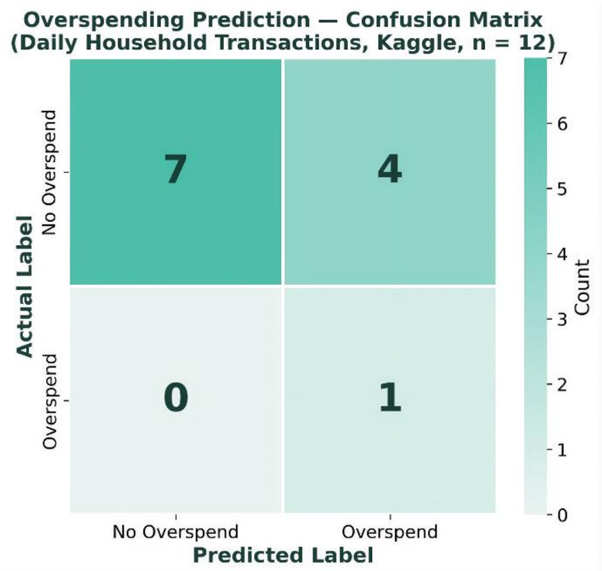

Figure 5: Confusion matrix — Daily Household Transactions (D2)

The combined results confirm that the classifier generalizes across both controlled and real-world data. The precision drop on D2 is the expected consequence of the severe class imbalance (one overspending month in twelve); the recall of 1.00 on the minority class is, in practice, the more useful property in a consumer setting, because missing an overspending month is more harmful to the user than producing an occasional unnecessary warning.

4.2 Category-Level Spending Forecasting

Separate Lasso-regression models were trained for each of the three variable spending categories that Tharaa forecasts (Food and Dining, Transportation, Bills and Utilities). The models use monthly features (prior-month spend in the same category, recent rolling means, days elapsed in the month) and were selected on the basis of minimizing RMSE and MAE on a held-out evaluation set.

4.2.1 Synthetic Dataset (D1)

Best-performing model per category on D1: Food — Lasso (RMSE = 264.58, MAE = 222.76); Transport — Lasso (RMSE = 143.45, MAE = 101.47); Entertainment — Decision Tree (RMSE = 188.84, MAE = 135.44). Predicted values aligned closely with actual spending in most months; minor deviations occurred during periods of unusually high expenditure, consistent with natural variability in user behavior. Figure 6 visualizes the forecast on the Food category.

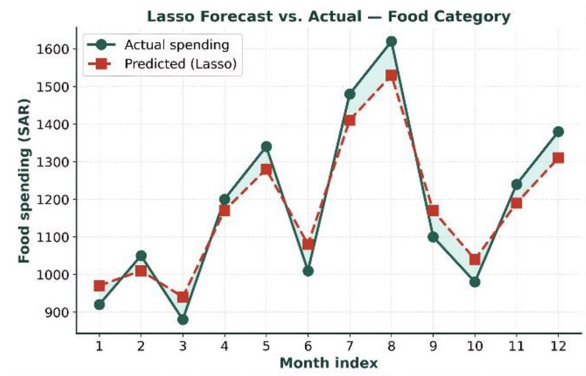

Figure 6: Lasso forecast versus actual — Food (D1)

4.2.2 Daily Household Transactions (D2)

On D2, the same Lasso pipeline was retrained and evaluated on the most recent 11 months held out as a chronological test set, simulating deployment forecasting. Table 2 reports the per-category errors. The errors on D2 are higher in absolute terms because the Kaggle dataset's monthly spending volumes are larger and more variable (high-value transportation expenses dominate the variance in this dataset). Figure 7 visualizes the forecast on the Food and Dining category.

Table 2: Category-forecast errors on the Daily Household Transactions dataset (D2). Errors are in the source-dataset currency units (INR)

Category	MAE	RMSE	Test Months
Food and Dining	1,942.42	2,827.52	11
Transportation	13,679.49	33,252.79	11
Bills and Utilities	12,153.70	31,181.25	11

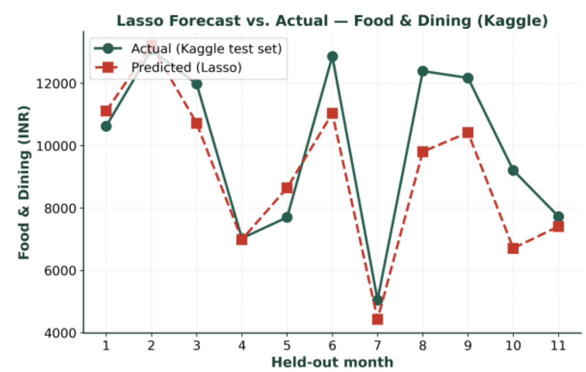

Figure 7: Lasso forecast versus actual — Food and Dining (D2)

4.3 Anomaly Detection

An Isolation Forest with 100 estimators and contamination = 0.08 was trained on standardized per-user transaction features (amount, day of month, category one-hot encoding, deviation from running category mean).

4.3.1 Synthetic Dataset (D1)

Anomaly detection was applied to a pooled dataset of 4,565 transactions across the four behavioral personas and identified 366 anomalies, corresponding to an overall anomaly rate of 8.02%. Detection rates varied substantially across user behavioral profiles, demonstrating behavior-specific calibration rather than a uniform threshold. Figure 8 visualizes the per-profile breakdown.

4.3.2 Daily Household Transactions (D2)

On D2, the Isolation Forest was applied to 1,857 in-scope expense transactions and flagged 149 anomalies, an overall anomaly rate of 8.02% — numerically identical to the rate observed on D1 and consistent with the contamination prior. Manual inspection of the flagged transactions showed they predominantly correspond to high-value or first-occurrence categories (large household purchases, atypical transportation expenses), which is the intended behavior of the model.

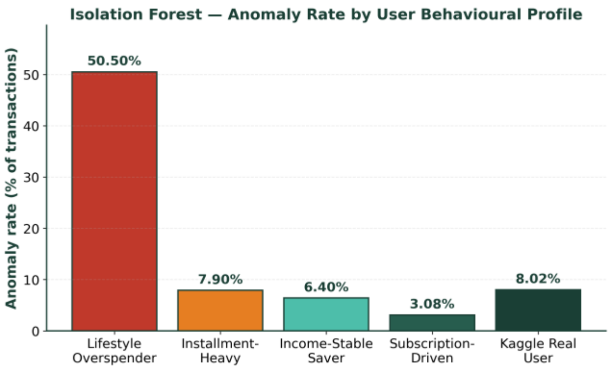

Figure 8: Anomaly rate by user behavioral profile (D1) and the real-world Kaggle dataset (D2)

4.4 Cross-Dataset Summary

Table 3 consolidates the headline metrics across both datasets, summarizing the comprehensive evaluation. Across both datasets, the overspending classifier preserves high recall on the minority class, the Lasso forecasters track actual spending without retraining beyond feature engineering, and the Isolation Forest reproduces the same overall anomaly rate. The consistency across two structurally different datasets supports the claim that the architecture generalizes.

Table 3: Cross-dataset summary of Tharaa's predictive components

Component	Metric	D1 (Synthetic)	D2 (Kaggle)
Overspending	Accuracy	0.960	0.667
Overspending	Recall (minority)	0.94	1.00
Overspending	F1 (minority)	0.97	0.33
Food forecast	MAE	222.76	1,942.42
Food forecast	RMSE	264.58	2,827.52
Anomaly	Overall flag rate	8.02%	8.02%
Anomaly	Transactions	4,565	1,857

4.5 Hybrid Financial Insight Generation

Financial insights are produced by combining machine-learning outputs with rule-based analysis. The ML layer contributes overspending risk scores, category-level spending forecasts, and anomaly flags. The rule-based layer adds commitment pressure (the share of income committed to fixed obligations), savings risk against goal deadlines, and trend comparisons against the prior period. The combined output is presented in plain language with the contributing factors made explicit (for example, "likely to exceed Bills budget by ~ 14% — driven by a higher electricity bill this month"). This explanation layer directly addresses the interpretability gap common to consumer ML products.

5. SYSTEM IMPLEMENTATION

Tharaa was implemented as a multi-module mobile application using the Flutter framework for a cross-platform user interface and Firestore for real-time backend synchronization. The following subsections describe the primary implemented components.

5.1 Interactive Financial Dashboard

The dashboard consolidates all financial data into a single view, presenting real-time account balance, income-versus-expense comparisons, category-level spending breakdowns, budget-utilization progress bars, goal-tracking indicators, and predictive risk alerts generated by the analytics module. The visual layout (Figure 9) prioritizes immediate comprehension over detail, with drill-down panels for users who require deeper analysis.

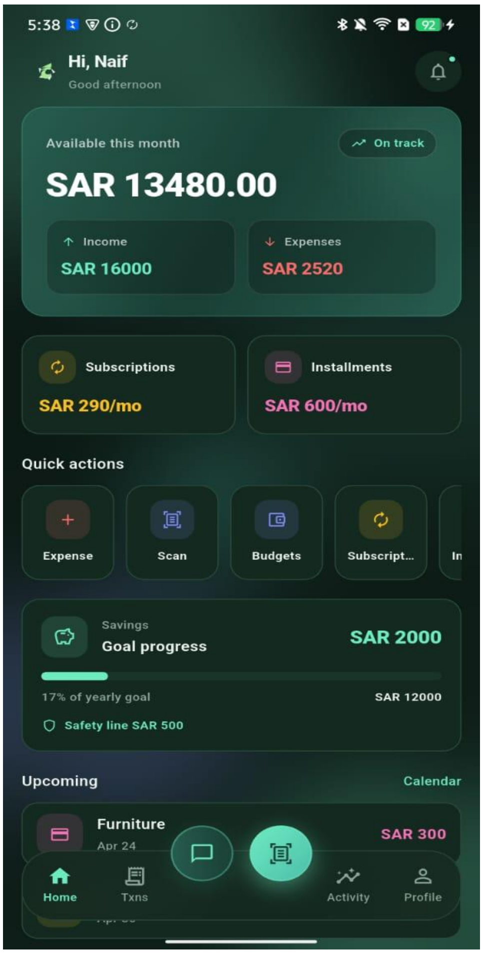

Figure 9: Financial Dashboard Interface

5.2 OCR Receipt Module

Within the implemented application, the OCR module allows users to photograph or upload a receipt at any point during the session. On-device text recognition via Google ML Kit [15] extracts transaction details without transmitting images to external servers. The extracted data is shown in a pre-filled transaction form (Figure 10) that the user reviews and confirms.

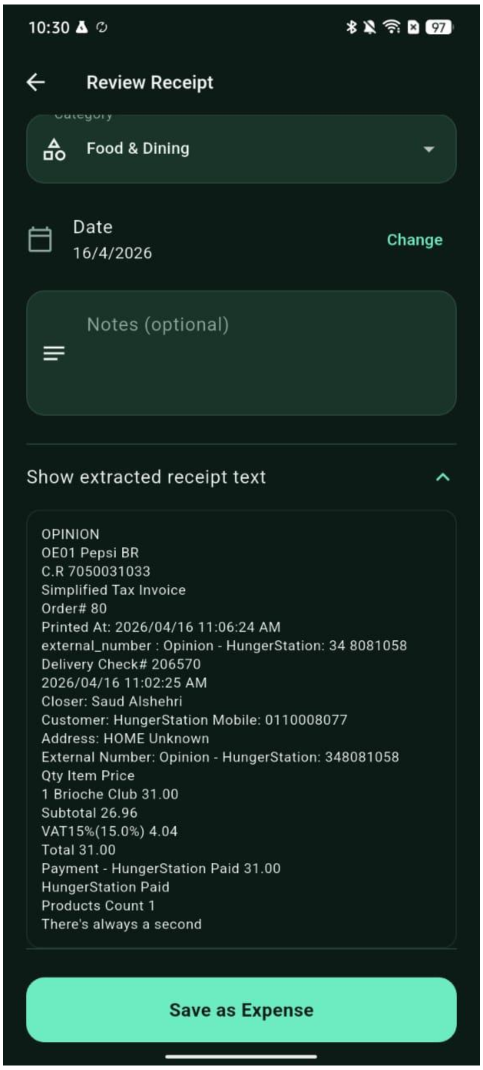

Figure 10: OCR Receipt Scanning Interface

5.3 Finance Coach Interface

The coaching chatbot is embedded within the application and accessible from any screen. Because every response is grounded in the user's own transaction history, budget state, goal progress, and upcoming obligations, the chatbot provides contextually relevant financial guidance rather than generic advice. Figure 11 illustrates an example interaction.

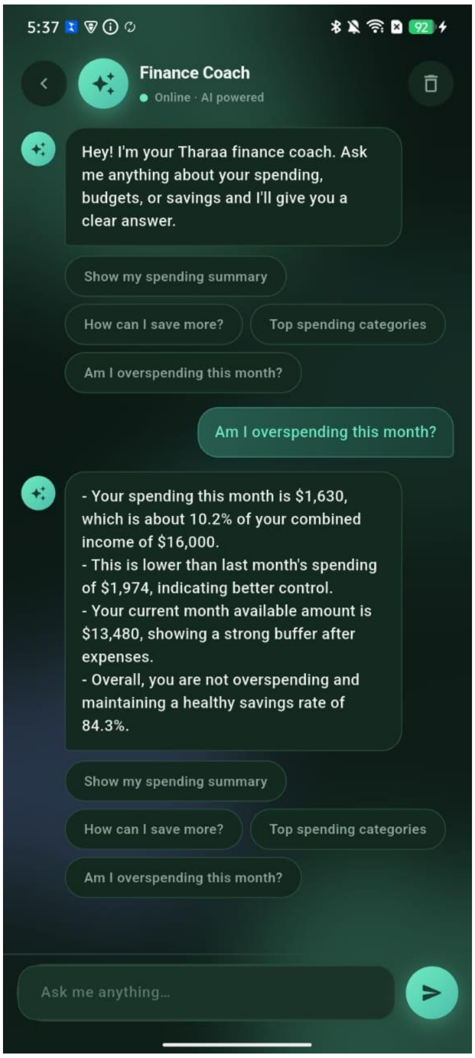

Figure 11: Finance Coach Chatbot Interface

5.4 Obligation and Reminder System

Recurring financial obligations are managed through a dedicated tracking interface where users specify the obligation name, amount, recurrence pattern, and due date. The system schedules push notifications ahead of each due date, reducing the cognitive overhead of tracking multiple recurring payments. The same data feeds the commitment-pressure component of the hybrid insight generator.

6. DISCUSSION

The experimental results across the two datasets support the effectiveness of Tharaa's integrated approach to personal financial management. For overspending prediction, a precision of 1.00 on the synthetic dataset — meaning no false alarms were generated — is especially significant in a consumer context, where unnecessary warnings quickly undermine trust. On the Kaggle dataset, the classifier preserves perfect recall on the minority overspending class, which is the operationally important metric: missing an overspending month is more harmful to the user than producing an occasional precautionary warning.

Category-level forecasting results indicate that mid-term spending patterns are predictable enough to support short-term financial planning across both controlled and real-world settings. The higher absolute errors observed on the Kaggle dataset are a property of the underlying spend volumes (the dataset includes high-value transportation expenses that account for a large share of the variance), not of the model itself, and they are still within usable bounds for budget guidance.

The anomaly-detection findings reveal meaningful variation across user behavioral profiles. The substantial difference between anomaly rates for lifestyle overspenders (50.5%) and subscription-driven users (3.08%) demonstrates that the module calibrates to individual behavior rather than applying a uniform threshold. The numerical agreement between the overall anomaly rate observed on the synthetic pool and the real-world Kaggle dataset (both 8.02%) further supports the conclusion that the Isolation Forest's contamination prior generalizes across datasets.

The hybrid insight generation strategy — pairing machine-learning outputs with rule-based financial logic — directly addresses the interpretability gap common in data-driven systems. By producing explanations that identify the source of financial risk rather than merely signaling its presence, Tharaa moves the system from passive reporting toward active financial coaching.

On the privacy side, processing receipt images entirely on-device via Google ML Kit [15] avoids exposing sensitive financial documents to third-party servers. Combined with strict per-user data isolation at the database level, administrator access restrictions, and on-device execution of all three predictive models using exported coefficients, these design decisions address a critical user concern in financial technology adoption [11, 16].

The system has acknowledged limitations. The OCR component's performance is sensitive to image quality and receipt layout variation, which can reduce extraction accuracy for poorly photographed or non-standard receipts. The finance coach depends on the availability of the external AI service for its conversational layer. The predictive models are trained on monthly aggregates and therefore require at least a partial month of activity before useful forecasts can be produced; a cold-start fallback is provided for new users.

7. CONCLUSION

This paper presented Tharaa, an integrated mobile personal-finance management system that unifies expense and income tracking, budget and goal management, subscription and installment monitoring, OCR-based receipt capture, machine-learning-driven predictive analytics, and an AI-powered finance coach in a single Flutter application backed by Firestore. The contribution lies not in any single component but in the integration that allows each module to reinforce the others, supported by an on-device predictive layer that preserves user privacy.

Empirical evaluation across two datasets confirmed the effectiveness of the predictive components. On a controlled synthetic dataset, the overspending classifier achieved 96% accuracy with no false alarms, the category-forecast Lasso models tracked actual spending closely, and Isolation Forest produced anomaly rates calibrated to user behavioral profile. On the public Daily Household Transactions dataset, the overspending classifier preserved perfect recall on the minority class under severe class imbalance, the Lasso forecasters generalized without retraining beyond feature engineering, and the Isolation Forest reproduced the same overall 8.02% anomaly rate observed on the synthetic pool. These cross-dataset consistencies support the claim that the architecture generalizes.

Future work will pursue four directions: (i) improving OCR robustness for diverse and low-quality receipt formats using deep-learning-based recognition models; (ii) extending the finance coach with on-device language-model capabilities to remove the external dependency; (iii) integrating Saudi-bank open-banking APIs to support automatic transaction import; and (iv) extending the evaluation to a longitudinal field study with real users in the Kingdom of Saudi Arabia to measure the system's impact on financial-awareness behavior over time.

8. REFERENCES

[1] A. Patel and A. Singh, "An intelligent AI-based framework for automated personal financial management," DMPedia Lecture Notes in Multidisciplinary Research, pp. 774–783, 2026.

[2] S. A. Alshebami and S. M. Al Marri, "The impact of financial literacy on entrepreneurial intentions of Saudi youth," Frontiers in Psychology, vol. 13, Art. 911605, 2022. doi: 10.3389/fpsyg.2022.911605.

[3] A. Lusardi and O. S. Mitchell, "The economic importance of financial literacy: theory and evidence," Journal of Economic Literature, vol. 52, no. 1, pp. 5–44, 2014.

[4] T. Stefanov, M. Stefanova, S. Varbanova, and S. Temelkov, "Personal finance management application," TEM Journal, vol. 13, no. 3, pp. 2066–2075, 2024. doi: 10.18421/TEM133-34.

[5] M. C. J. D. Freitas et al., "Mobile personal-finance applications: a systematic review of user adoption factors," Heliyon, vol. 9, no. 4, 2023.

[6] D. W. Arner, J. N. Barberis, and R. P. Buckley, "The evolution of FinTech: a new post-crisis paradigm?" Georgetown Journal of International Law, vol. 47, no. 4, pp. 1271–1319, 2016.

[7] P. Gomber, J.-A. Koch, and M. Siering, "Digital finance and FinTech: current research and future research directions," Journal of Business Economics, vol. 87, no. 5, pp. 537–580, 2017.

[8] I. Lee and Y. J. Shin, "FinTech: ecosystem, business models, investment decisions, and challenges," Business Horizons, vol. 61, no. 1, pp. 35–46, 2018.

[9] E. Z. Milian, M. de M. Spinola, and M. M. de Carvalho, "Fintechs: a literature review and research agenda," Electronic Commerce Research and Applications, vol. 34, Art. 100833, 2019.

[10] D. Belanche, L. V. Casaló, and C. Flavián, "Artificial intelligence in FinTech: understanding robo-advisors adoption among customers," Industrial Management & Data Systems, vol. 119, no. 7, pp. 1411–1430, 2019. doi: 10.1108/IMDS-08-2018-0368.

[11] R. S. Alkadi and S. S. Abed, "Examining consumer intention to adopt AI-powered chatbots in the Saudi banking sector," International Journal of Bank Marketing, vol. 43, no. 10, pp. 2229–2255, 2025. doi: 10.1108/IJBM-10-2024-0635.

[12] A. Adamopoulou and L. Moussiades, "An overview of chatbot technology," in Artificial Intelligence Applications and Innovations, Springer, 2020, pp. 373–383.

[13] S. Hussain, O. A. Sianaki, and N. Ababneh, "A survey on conversational agents/chatbots classification and design techniques," in Web, Artificial Intelligence and Network Applications, Springer, 2019, pp. 946–956.

[14] R. Smith, "An overview of the Tesseract OCR engine," in Proc. 9th Int. Conf. on Document Analysis and Recognition (ICDAR 2007), 2007, pp. 629–633.

[15] Google ML Kit, "Text Recognition API," Google, 2024. [Online]. Available: https://developers.google.com/ml-kit/vision/text-recognition

[16] A. Alrsheedi and Y. H. P. Iskandar, "Key factors influencing FinTech adoption among Saudi banks: a conceptual framework," Humanities and Social Sciences Communications, vol. 12, no. 1, pp. 1–11, 2025. doi: 10.1057/s41599-025-05532-1.

[17] D. Priantinah, M. N. Aisyah, and Y. Nurim, "The analysis of Technology Acceptance Model (TAM) for personal financial management on mobile application technology," Atlantis Press, 2019. doi: 10.2991/icobame-18.2019.56.

[18] R. Kesavan, D. Gay, D. Thevessen, J. Shah, and C. Mohan, "Firestore: the NoSQL serverless database for the application developer," in Proc. IEEE 39th Int. Conf. Data Eng. (ICDE), 2023.