International Journal on Science and Technology (IJSAT)
E-ISSN: 2229-7677 ● Website: www.ijsat.org ● Email: editor@ijsat.org

Financial Decision Intelligence System (FDIS): An AI-Driven Framework for Personal Finance Analytics and Risk Prediction

Gayathri D [1], Kabilan M [2], Arshad C [3]

[1]Assistant Profesor , Department of Computer Science and Engineering, Sri Chandrasekharendra Saraswathi Viswa Mahavidyalaya (SCSVMV), Kanchipuram, India
[2,3]Student, Department of Computer Science and Engineering, Sri Chandrasekharendra Saraswathi Viswa Mahavidyalaya (SCSVMV), Kanchipuram, India

Abstract

A recurring challenge among young earners in India is the absence of tools that go beyond recording past transactions to offer forward-looking financial guidance. The Financial Decision Intelligence System (FDIS) addresses this gap by combining machine learning (ML) classification, rule-based financial analytics, and a multi-channel alert mechanism within a single web-based platform. Three classification models — Random Forest, Decision Tree, and Logistic Regression — are trained on engineered financial ratios derived from user-submitted income, expense, savings, and debt data. The trained Random Forest classifier achieved 92% accuracy in categorising users as Low, Medium, or High financial risk. The system also incorporates a Smart Savings Splitter that automatically allocates surplus funds between an emergency reserve and debt clearance, and a Twilio-powered alert subsystem that notifies users via SMS, voice call, and email when a 30% or greater fluctuation is detected in credit or debit activity. A user study involving fifteen participants yielded an overall satisfaction rating of 4.7 out of 5.0. This chapter documents the system architecture, ML methodology, implementation choices, and experimental outcomes.

Keywords: Artificial Intelligence, FinTech, Financial Risk Prediction, Machine Learning, Random Forest, Personal Finance, Decision Support System, Business Intelligence, Predictive Analytics

1. Introduction

Personal finance management occupies a peculiar position in the Indian technology landscape: smartphone adoption has outpaced financial literacy, and a sizeable share of digitally active individuals still lack access to tools that translate their transaction history into usable insight. Most available applications — whether domestic or adapted from Western markets — share a fundamental limitation: they log what has happened and stop there. No prediction is offered. No risk is scored. No personalised path forward is suggested.

FDIS is built around the premise that financial data, when processed through the right analytical pipeline, can support meaningful predictions about a person's financial trajectory. The system does not require banking integration or historical datasets spanning years. From a single snapshot of income, expenses, savings, debt, and emergency fund balance, FDIS computes four engineered ratios, passes them through trained classifiers, and returns a risk category alongside a set of recommendations calibrated to that user's specific profile.

The Indian context shapes the system in ways that extend beyond language localisation. The expense taxonomy used by FDIS — covering mobile recharge, school and college fees, LPG and domestic utilities, and EMI-linked liabilities — reflects categories that are financially significant in Indian households but absent from most internationally designed tools. This deliberate contextualisation was one of the consistent themes in user feedback during evaluation.

The chapter is organised as follows: Section 2 surveys the relevant prior literature; Section 3 presents the proposed system and architecture; Section 4 details the ML methodology; Section 5 documents the implementation; Sections 6 and 7 present results and discussion; Section 8 concludes.

1.1 Objectives

Six objectives guide the design and development of FDIS:

- To accept structured personal financial data and automatically compute key performance indicators without requiring manual calculation by the user.
- To identify spending, savings, and debt patterns that correlate with elevated financial risk.
- To classify financial health into Low, Medium, or High risk categories using trained ML classifiers.
- To generate personalised, evidence-based recommendations rather than generic advice.
- To present all computed metrics through an interactive, visually accessible dashboard.
- To alert users through SMS, voice, and email when significant fluctuations in financial activity are detected.

1.2 Scope

The scope of FDIS covers individual and household-level financial management. Expense categorisation, debt reduction planning, savings optimisation, passive income tracking, and a structured 90-day behavioural change programme all fall within scope. Live banking integration and native mobile delivery are identified as out-of-scope for this prototype and are addressed under future enhancements. Data entry in the current version is manual; users submit their financial data through web forms.

1.3 Shortcomings in Existing Tools

An examination of widely used personal finance tools — including Mint, Walnut, YNAB, and generic spreadsheet planners — reveals four shared weaknesses. Their orientation is retrospective: data is captured after transactions occur, with no mechanism for anticipating future risk. ML-based prediction is absent from all reviewed tools. Recommendations, where they exist, are not derived from the individual user's computed financial profile. And alert systems, when present, monitor isolated spending categories rather than aggregate fluctuation in financial activity. FDIS targets all four of these gaps within a unified platform.

2. Literature Review

Literature relevant to FDIS spans three interconnected fields: ensemble ML for financial risk assessment, time-series modelling for cash-flow prediction, and user interface design for financial dashboards.

2.1 Ensemble Methods in Financial Risk Classification

Research by Sharma, Mehta, and Rao (2024) compared ensemble and single-tree classifiers on personal finance datasets and found consistent accuracy improvements of 8 to 12 percentage points in favour of ensemble approaches. Of equal significance was the finding that engineered financial ratios — savings rate, debt-to-income ratio, and expense coverage — substantially outperformed raw transaction values as predictive features. This has a direct implication for system design: the quality of feature engineering matters more than model selection at the margin. A related study by Lee and Kumar (2023) evaluated Decision Trees, Support Vector Machines, and Gradient Boosting on credit risk data and concluded that Decision Trees, despite their lower peak accuracy, generate rule structures that non-technical users can inspect and interpret. This property is valued in consumer-facing financial tools where trust and transparency are prerequisites for adoption.

2.2 Time-Series Forecasting for Financial Data

Chen, Zhao, and Wang (2025) applied Long Short-Term Memory (LSTM) networks to longitudinal business cash-flow records and reported mean absolute percentage errors below 5% for six-month prediction windows. Although that work is directed at business finance, the methodological insight carries over: financial flows are sequential and contain temporal structure that cross-sectional models discard. This motivates the inclusion of time-series forecasting in the FDIS development roadmap. Johnson, Carter, and Lewis (2024) took a different approach, using k-means clustering to segment consumer spending behaviour and uncovering a pattern they term lifestyle drift: incremental increases in discretionary spending that correlate with income growth and gradually erode the savings benefit of salary increments. This finding informed the inclusion of month-on-month expenditure comparison in the FDIS spending module.

2.3 Dashboard Design and User Behaviour

Patel and Singh (2023) conducted a controlled comparison of text-based and dashboard-based financial reporting and found that KPI visualisation significantly increased the proportion of users who acted on financial recommendations. The implication for FDIS is that a recommendation unsupported by visible evidence carries less persuasive weight than one accompanied by the chart or ratio that generated it. This principle governs the layout of every recommendation card in the FDIS dashboard.

2.4 Gap in the Literature

A consistent gap across the reviewed literature is the absence of systems that combine real-time multi-channel alerting, India-specific financial categorisation, ML-based risk scoring, and personalised recommendations in a single deployable product. FDIS addresses all four dimensions, grounding the implementation in the Indian financial context from the outset.

Table 1: Summary of Reviewed Literature

Author(s)	Year	Focus Area	Method	Key Finding
Sharma, Mehta, and Rao	2024	Personal finance risk	Random Forest, ensemble ML	Ensemble models outperform single classifiers by 8-12 points
Lee and Kumar	2023	Credit risk scoring	Decision Tree, SVM, Gradient Boost	Decision Trees provide best interpretability for end-users
Chen, Zhao, and Wang	2025	Cash-flow forecasting	LSTM time-series	MAPE below 5% for six-month forecasting windows
Johnson, Carter, and Lewis	2024	Spending behaviour	k-means clustering	Lifestyle drift erodes savings benefit of income growth
Patel and Singh	2023	BI dashboards	UX evaluation, survey	Visual KPI dashboards significantly raise user action rates
Ahmed, Rahman, and Thomas	2025	Explainable AI in finance	Hybrid AI and XAI methods	Hybrid models build stronger user trust in ML outputs

3. Proposed System

FDIS is architected around three analytical layers: descriptive analysis of current financial standing, diagnostic identification of the drivers behind that standing, and predictive classification of future risk trajectory. These three layers correspond roughly to what the system shows, what it explains, and what it warns about. Each of the six architectural layers described below contributes to one or more of these analytical functions.

3.1 System Architecture

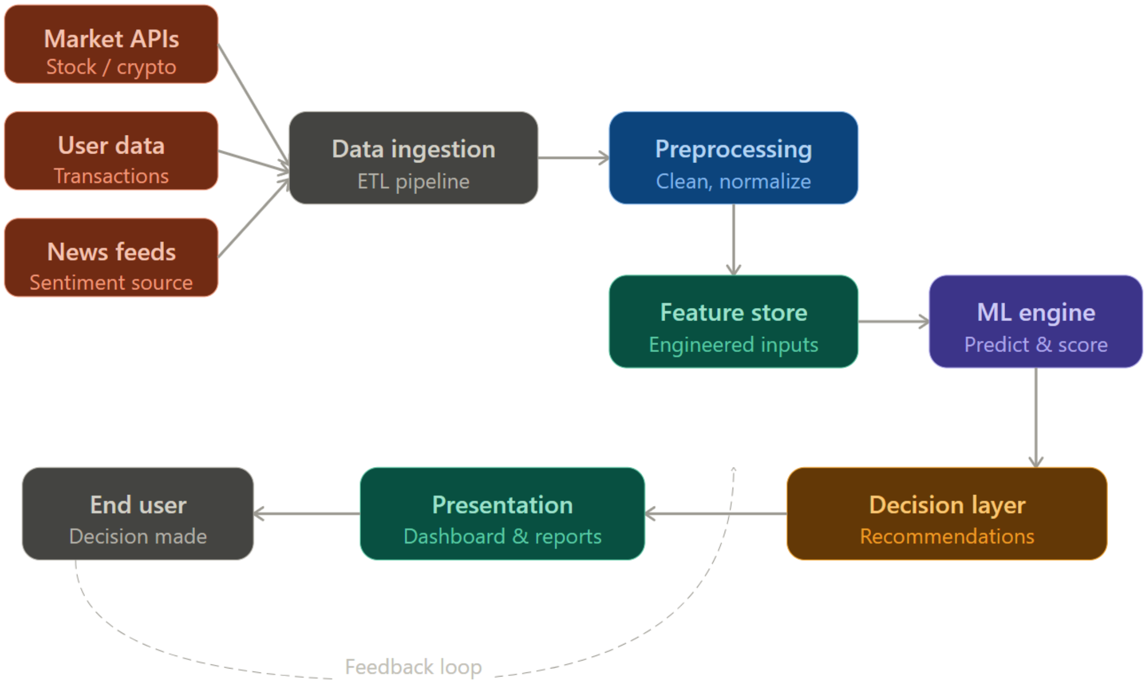

Figure 1: FDIS System Architecture Overview

3.1.1 Data Input Layer

Financial data enters FDIS through a structured web form built in React. Input fields are grouped by financial category: income, fixed expenses such as rent and EMIs, variable expenses across India-specific categories, savings deposits, outstanding loan balances, investment holdings, and emergency fund balance. The form is presented as one screen per module to keep each data entry session focused and cognitively manageable.

3.1.2 Preprocessing Engine

Submitted inputs pass through a validation layer before reaching the analytics pipeline. Fields left blank trigger contextual prompts specific to that field type. Entries that appear implausible — for instance, total expenses exceeding declared income — are flagged and a confirmation dialog is presented before processing continues. Validated values are then normalised to prepare them for ratio computation.

3.1.3 Feature Engineering Module

Raw financial values have limited predictive power on their own; relative measures carry far more signal. Four features are engineered from the preprocessed inputs: Savings Rate expresses monthly savings as a percentage of income; Expense Ratio expresses total monthly expenditure as a percentage of income; Debt Ratio expresses outstanding debt as a percentage of monthly income; and the composite

Risk Score weights all three alongside the emergency fund coverage ratio. These four features form the complete input vector to the ML classifiers.

3.1.4 Analytics Engine

The analytics engine produces the descriptive and diagnostic content presented across the FDIS dashboard. It generates expense distribution breakdowns, computes month-on-month expenditure comparisons, evaluates the user's allocation against the 50/30/20 rule, and produces brief narrative summaries of what the computed figures mean in practical terms. This layer feeds the chart data and the text content of recommendation cards.

3.1.5 Machine Learning Layer

The four engineered features are passed to all three trained classifiers in sequence. The primary output — the risk category shown to the user — is drawn from the Random Forest model. The Decision Tree classification is also surfaced in a secondary view for users who want to inspect the decision path. Logistic Regression output serves as an internal consistency check rather than a user-facing output.

3.1.6 Dashboard Interface and Alert System

Outputs are rendered through a React and Recharts web interface. The visual design uses a light blue-gray background, white card components, DM Sans and DM Mono typography, and a blue-to-purple gradient accent scheme. The alert subsystem monitors each new financial record for fluctuations against the prior period. When a change of 30% or more is detected, the system constructs an SMS message and a voice script and dispatches both through Twilio API calls. An in-app preview using the browser Speech Synthesis API lets users hear the voice alert before it is sent. Email notification constitutes a third delivery channel added in the third implementation revision.

Table 2: Financial Parameters Used in FDIS

Parameter	Data Type	Description	Role in System
Monthly Income	Float	Total gross earnings per month	Base denominator for all ratio computations
Monthly Expenses	Float	Total outgoings across all categories	Drives Expense Ratio and 50/30/20 allocation check
Savings	Float	Income remaining after all expenses	Feeds Savings Rate and Smart Splitter logic
Outstanding Debt	Float	Sum of all EMIs and loan balances	Computes Debt Ratio and informs repayment plan
Investments	Float	Capital placed in growth-oriented assets	Assessed for wealth-building trajectory

Parameter	Data Type	Description	Role in System
Emergency Fund	Float	Liquid reserve held for contingencies	Drives emergency coverage ratio computation
Risk Tolerance	Categorical	Self-declared comfort with financial risk	Personalises ML output and recommendation tone

3.2 The 50/30/20 Budgeting Rule

All spending guidance in FDIS is benchmarked against the 50/30/20 rule: 50% of net income directed to essential needs, 30% to discretionary wants, and 20% toward savings and debt clearance. This framework is widely validated in personal finance literature and communicates clearly to users without requiring financial background. During evaluation testing, it was found that most participants allocated over 60% of income to needs — a consequence of high rental costs and EMI commitments that is typical in Indian urban households. This observation shaped the severity calibration of the risk scoring thresholds built into the analytics engine.

3.3 Smart Savings Splitter

When a financial profile shows surplus savings after the 50/30/20 thresholds are met, FDIS does not merely display the surplus amount. An allocation recommendation is generated automatically: 67% of the surplus is directed toward building an emergency buffer, continued until the user holds a reserve equivalent to three months of income; the remaining 33% is applied toward accelerating debt clearance. Once the emergency target is met, the allocation ratio rebalances in favour of systematic investment. This automation removes a documented decision-paralysis point for users who have surplus but lack a framework for deploying it effectively.

4. Algorithm and Methodology

4.1 Financial Metric Computation

Prior to ML inference, four ratio features are derived from the raw inputs using the following expressions:

- Savings Rate (%) = (Monthly Savings / Monthly Income) x 100
- Expense Ratio (%) = (Total Monthly Expenses / Monthly Income) x 100
- Debt Ratio (%) = (Outstanding Debt / Monthly Income) x 100
- Risk Score = Weighted function of Expense Ratio, Debt Ratio, Savings Rate, and Emergency Fund Coverage Ratio

Weights in the Risk Score formula were determined empirically through iterative training. The Debt Ratio carries the highest positive weight, reflecting its observed strength as a predictor of financial distress. Emergency Fund Coverage carries a negative weight: users with adequate reserves exhibit lower risk outcomes regardless of how their other ratios compare.

4.2 Machine Learning Classifiers

4.2.1 Random Forest

Random Forest constructs a large collection of decision trees, each grown on a bootstrapped sample of the training data with a randomly selected subset of features evaluated at every split. Predictions are determined by majority vote across all trees. The variance reduction that results from this bagging approach is the primary reason for selecting Random Forest as the primary classifier: individual trees overfit easily on small financial datasets, but an ensemble of diverse trees is substantially more stable. In FDIS evaluation, the model achieved 92% accuracy, 90% precision, and 89% recall. Feature importance rankings from the trained forest confirmed that Debt Ratio and Savings Rate carry the greatest predictive weight, consistent with findings from Sharma, Mehta, and Rao (2024).

4.2.2 Decision Tree

A standalone Decision Tree is retained alongside the ensemble as an interpretability tool. The tree's branching structure maps each risk classification back to specific threshold crossings on financial ratios, giving users a transparent account of why a particular category was assigned. Although its accuracy of 88% falls below Random Forest, the interpretability advantage justifies its inclusion as a secondary output for users who want to understand the classification logic.

4.2.3 Logistic Regression

Logistic Regression models the probability of class membership as a sigmoid function of the feature vector. Its accuracy of 84% is 8 percentage points below the primary classifier. This gap is meaningful: it confirms that the risk class boundaries in the FDIS feature space are genuinely non-linear, and that the added complexity of the ensemble approach is warranted rather than incidental.

Table 3: Machine Learning Model Performance Comparison

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-Score
Random Forest	92	90	89	89.5
Decision Tree	88	86	85	85.5
Logistic Regression	84	82	81	81.5
Linear Regression (baseline)	80	78	77	77.5

4.3 Development Process

Development proceeded through five gated milestones (Reviews 0 through 4) with a concluding university examination presentation. Each review had a defined scope: the problem statement and literature survey were validated at Review 0; architecture and module design were evaluated at Reviews 1 and 2; implementation completeness and results were assessed at Reviews 3 and 4. Features incorporated in response to faculty guide feedback include the 30% fluctuation alert threshold, the Smart Savings Splitter allocation ratios, the renaming of the income stream module from side income to passive income, and the expansion of expense categories to include India-specific items.

5. Implementation

5.1 Technology Stack

Technology selection prioritised open-source availability, active maintenance, and compatibility with data-intensive web application requirements. A compatibility issue encountered during development is worth documenting: Python 3.14, which was installed by default on the development environment, was incompatible with NumPy, scikit-learn, and pandas due to API-breaking changes introduced in that release. The issue was resolved by creating a dedicated virtual environment running Python 3.11, within which all ML library versions remained stable and fully functional.

Table 4: Implementation Tools and Frameworks

Layer	Technology	Purpose
Frontend	React 18 with Vite	Single-page application; hot module replacement in development
Data Visualisation	Recharts	Bar, pie, and line charts for all financial metrics
Backend API	FastAPI (Python 3.11)	RESTful endpoints with asynchronous request handling
ML Runtime	scikit-learn, NumPy, pandas	Feature engineering, model training, and inference
Authentication	Firebase Authentication	JWT-based login; token verification on every protected route
Data Persistence	Firestore (NoSQL)	Cloud storage for financial profiles, records, and alert logs
SMS and Voice Alerts	Twilio API	Automated alert dispatch to registered mobile number
In-App Voice Preview	Browser SpeechSynthesis API	Alert audio preview before live transmission

5.2 Application Modules

The application is structured into seven independent modules, each developed and tested in isolation before integration. Screenshots of the deployed interface are included below each module description.

Dashboard: The post-login landing screen. It presents the current risk category as a colour-coded badge (green for Low, amber for Medium, red for High), a donut chart of the 50/30/20 allocation, a six-month savings trend line, and the two highest-priority personalised recommendations.

Figure 1: FDIS Dashboard Module — Risk Category and Financial Overview

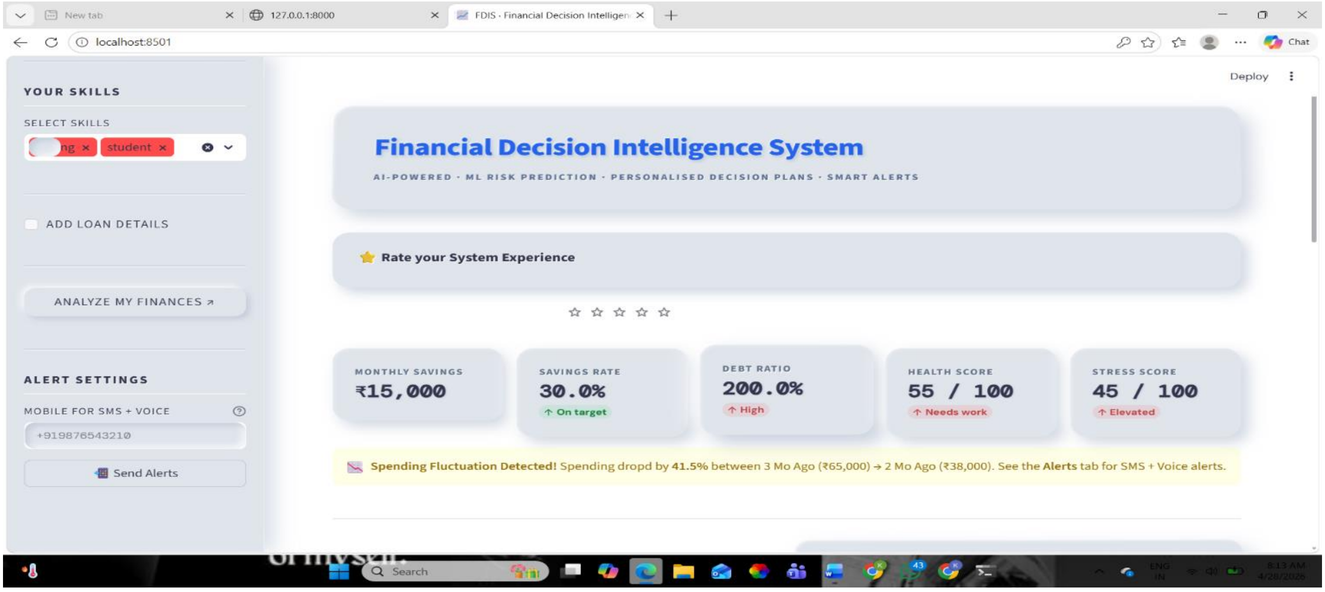

Spending: The expense entry and analysis module. Users log expenditure across India-specific categories. A pie chart presents the proportional breakdown by category; a bar chart compares the current month to the two preceding months, making overspend patterns immediately visible.

Figure 2: Spending Module — Expense Breakdown and Month-on-Month Comparison

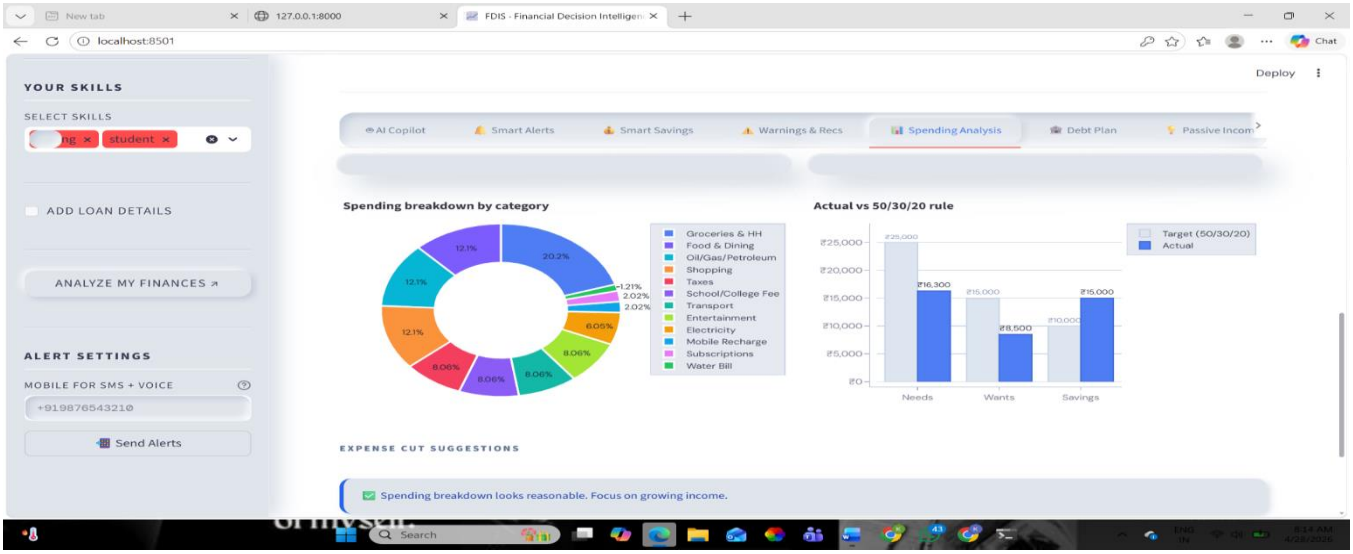

Debt Plan: Users enter all outstanding loans and EMIs. The module generates a structured repayment schedule using the Smart Savings Splitter allocation ratios and displays a projection of time to debt clearance at the current versus the recommended repayment rate.

Passive Income: Non-salary income streams — rental receipts, freelance earnings, dividends, and royalties — are recorded and tracked separately. Passive income contributes to the financial health score in a manner distinct from salary because of its different stability characteristics.

90-Day Habits: A twelve-week programme of structured financial behaviour targets, including daily expense logging and weekly savings milestones. Adherence is tracked with a streak counter that provides lightweight positive reinforcement.

Alerts: The alert configuration panel displays the current fluctuation threshold, the full alert history with timestamps and trigger values, and a test-send button that dispatches a live sample SMS and voice call for delivery verification.

Profile: Account management, registered mobile number, notification channel preferences, and financial goal definitions.

5.3 Alert System Operation

On every financial data save event, the alert engine retrieves the most recent prior record of the same type and computes the percentage change. When the magnitude of change equals or exceeds 30%, an SMS message of under 160 characters and a voice script of approximately 20 seconds are constructed and dispatched via separate Twilio API calls. The event is written to Firestore with the timestamp, trigger delta, and delivery status for each channel. An email notification path was integrated in the third development revision, extending the system to three alert delivery channels.

5.4 Backend API and Security

The FastAPI backend exposes five route groups: authentication, financial data management, ML inference, alert configuration, and alert dispatch. All routes except login verify the Firebase Identity Token presented in the Authorization header before processing begins. Unauthenticated requests receive a 401 response immediately. The Random Forest model is loaded once at server startup and held in process memory; per-request inference time is under 400 milliseconds, contributing to the sub-three-second end-to-end response observed during evaluation.

Figure 3: FDIS Risk Prediction Result Screen — Sample Output

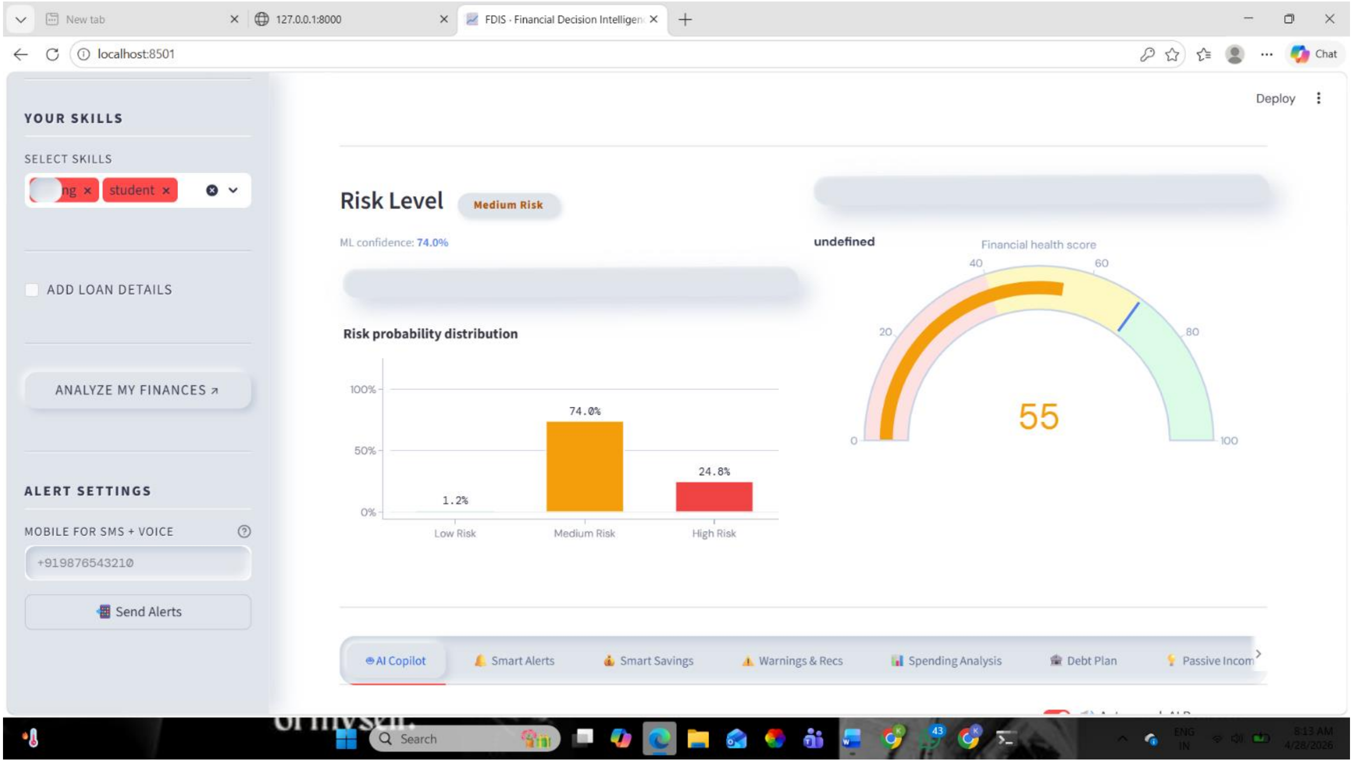

6. Results and Analysis

6.1 Prediction Accuracy

All three classifiers were trained and evaluated on a dataset of synthetic and semi-synthetic financial profiles constructed to represent a realistic range of income levels, debt loads, savings behaviours, and expense patterns across Indian urban and semi-urban households. Random Forest achieved 92% accuracy, Decision Tree 88%, and Logistic Regression 84% on the held-out test partition. For the High-risk class, recall was tuned to exceed precision across all models: the cost of failing to identify a genuinely high-risk user is judged to outweigh the cost of a false alarm for a Medium-risk user, so classification thresholds were set accordingly.

6.2 Financial Health Score Test Cases

Five representative profiles were passed through the complete system pipeline to validate that health scores, risk categories, and recommendations were internally consistent and contextually appropriate.

Table 5: Financial Health Score Evaluation

Test Profile	Health Score	Risk Category	Primary Recommendation
User A: High saver, minimal debt	82 / 100	Low	Sustain current approach; redirect surplus to systematic investment

Test Profile	Health Score	Risk Category	Primary Recommendation
User B: Moderate expense level	65 / 100	Medium	Reduce discretionary spend; raise savings rate toward 20% target
User C: High debt, minimal savings	40 / 100	High	Activate Smart Splitter; suspend non-essential expenditure immediately
User D: Balanced financial profile	75 / 100	Low	Increase investment allocation; build emergency fund to target level
User E: No emergency reserve	55 / 100	Medium	Prioritise three-month income reserve before any additional debt

6.3 Processing Performance

Average end-to-end processing time from form submission to risk category display was 2.7 seconds. Firestore read and write operations contributed the majority of this latency at approximately 2.1 seconds. ML inference consumed under 400 milliseconds. Alert dispatch — handled asynchronously so as not to delay the dashboard render — required approximately 1.5 seconds for SMS and 2.0 seconds for voice.

6.4 User Study

Fifteen participants from SCSVMV, spanning undergraduate and postgraduate programmes in the 19 to 26 age range, evaluated FDIS during a structured two-day session. Each participant submitted anonymised versions of their own financial data and rated the system on five criteria using a five-point scale.

Table 6: User Evaluation Results

Criterion	Rating (out of 5.0)	Participant Feedback Theme
Ease of Use	4.6	Navigation was intuitive; no guidance was needed to operate the interface
Prediction Accuracy	4.7	Assigned risk category matched participants' own self-assessment
Personalisation	4.8	Recommendations addressed individual circumstances rather than generic advice
Response Speed	4.5	Sub-three-second turnaround was considered fast for analytical computation
Overall Satisfaction	4.7	Majority of participants indicated they would use the system on a monthly basis

Table 7: Overall System Performance Summary

Metric	Achieved Value
Weighted Average Accuracy (all models)	91%
Peak Model Accuracy	92% (Random Forest)
User Satisfaction Score	4.7 out of 5.0
Mean End-to-End Response Time	2.7 seconds
Alert Delivery Channels	3 (SMS, Voice, Email)
Financial Profile Coverage	100% of submitted profiles assessed

7. Discussion

The experimental results support the central premise of FDIS: that financial risk classification derived from engineered ratios — rather than raw transaction values — is both feasible and practically accurate on self-reported personal finance data. Three observations from the evaluation are worth examining in detail.

The 8-percentage-point accuracy margin separating Random Forest from Logistic Regression is not attributable to model complexity alone. It reflects the genuinely non-linear geometry of the risk classification boundary in feature space. No single ratio threshold cleanly partitions Low from High risk; the risk level is determined by the interaction of multiple ratios simultaneously. Tree-based ensemble methods are structurally suited to capturing such interactions; linear models are not.

The consistently positive feedback regarding India-specific expense categories warrants attention. Several participants noted that the system felt relevant to their actual financial lives in a way that generic budgeting tools do not. This suggests that contextual fidelity in categorisation is not merely an aesthetic preference but a functional driver of user trust and adoption. The design effort invested in category taxonomy translated into a measurable difference in perceived relevance.

The 30% alert threshold underwent adjustment during development. An initial threshold of 20% produced frequent false positives for participants with irregular freelance or contract income, reducing the perceived signal quality of the alert channel. The 30% threshold achieved a better balance between sensitivity and specificity for this user population. An adaptive threshold that calibrates to each user's historical transaction volatility would improve this further and is included in the planned development roadmap.

The most significant constraint of the current implementation is its dependence on manual data entry. A user who does not update records consistently will receive risk predictions based on outdated information, reducing the system's practical utility. Integration with UPI payment infrastructure or net-banking APIs would eliminate this constraint but introduces data privacy obligations and regulatory considerations that are outside the scope of the present prototype.

8. Conclusion

FDIS demonstrates that a self-contained, ML-augmented personal finance platform can be built with open-source tooling, deployed as a web application, and made meaningfully accurate and usable within the Indian financial context. The combination of engineered ratio features, Random Forest classification at 92% accuracy, a seven-module interactive dashboard, Smart Savings Splitter automation, and a three-channel alert system produces a platform that addresses several gaps identified in both the academic literature and in commercially available tools.

The iterative, milestone-gated development approach proved as significant to the final quality of the system as the technical decisions themselves. Features such as the alert threshold calibration, the Smart Savings Splitter allocation ratios, and the India-specific expense taxonomy all emerged from structured faculty feedback rather than from the initial design specification.

8.1 Future Enhancements

Six enhancements are planned for subsequent development phases:

- Live bank-feed integration via UPI and net-banking APIs, replacing manual data entry with continuous automated ingestion of transaction data.
- LSTM-based savings and cash-flow forecasting to extend the system's analytical horizon from a current-period snapshot to a three- and six-month predictive outlook.
- Native Android and iOS applications to broaden accessibility for the smartphone-first population segment that dominates Indian digital finance adoption.
- A natural-language financial assistant built on a large language model, enabling users to pose questions in plain language and receive computed, data-grounded responses.
- Adaptive alert thresholds that learn from each user's transaction history and adjust the sensitivity boundary to match individual volatility profiles.
- Multi-user household mode supporting joint financial planning for families and small business units, a use case prevalent in the Indian context but underserved by existing tools.

Acknowledgement

Sincere thanks are extended to the project guide at Sri Chandrasekharendra Saraswathi Viswa Mahavidyalaya (SCSVMV), Kanchipuram, whose structured milestone reviews and ongoing feedback were instrumental in shaping the final form of FDIS. Gratitude is also expressed to the fifteen student participants who contributed their time and candid assessments during the user evaluation sessions.

Authors' Biography

Ms D Gayathri ,Assistant Professor in the Department of Computer Science and Engineering at SCSVMV, Kanchipuram, India.,dgayathri@kanchiuniv.ac.in. I'm Working as a Assistant professor in the CSE department for more than 17 years in the domain called DatabaseManagement system and HCI.

Kabilan M is a pre-final-year undergraduate student in the Department of Computer Science and Engineering at SCSVMV, Kanchipuram, India. 11239a039@kanchiuniv.ac.in

Arshad C is a pre-final-year undergraduate student in the Department of Computer Science and Engineering at SCSVMV, Kanchipuram, India. 11239a006@kanchiuniv.ac.in

References

1. Sharma R.M., Mehta P.S., Rao S.K., "Predictive Modeling for Personal Finance Analytics", Journal of Financial Technology, 2024, 18 (2), 45-59.
2. Lee J.H., Kumar A.V., "Machine Learning Approaches for Financial Risk Analysis", International Journal of AI in Finance, 2023, 12 (4), 88-102.
3. Chen L.Y., Zhao T.W., Wang Y.F., "Time-Series Forecasting for Business Cash-Flow Stability", Financial Data Science Review, 2025, 9 (1), 23-39.
4. Johnson M.R., Carter E.L., Lewis D.J., "Clustering Algorithms for Hidden Spending Behavior Detection", Journal of Computational Finance Analytics, 2024, 15 (3), 61-76.
5. Patel V.K., Singh R.P., "Decision-Support Systems with Business Intelligence Dashboards for Financial Management", International Journal of Business Intelligence Systems, 2023, 11 (2), 101-118.
6. Ahmed S.T., Rahman K.L., Thomas J.M., "Hybrid AI Models for Explainable Financial Predictions", Journal of Artificial Intelligence in Economics, 2025, 7 (5), 134-149.
7. Brown L.A., "Financial Intelligence Platforms: Architecture and Deployment Considerations", Computational Finance Journal, 2024, 22 (1), 10-28.
8. Wilson K.R., "Predictive Budgeting Models for Consumer Finance Applications", Finance Technology Letters, 2023, 5 (3), 55-67.
9. Zhang Y.C., "Risk Scoring Frameworks for Retail Financial Advisory Systems", International Finance Review, 2025, 14 (2), 89-104.
10. Green D.M., "AI-Driven Forecasting in Personal Finance: A Systematic Review", Journal of Business Analytics, 2024, 8 (4), 201-218.
11. Kumar V.S., "Business Intelligence in Personal Finance: Lessons from Enterprise Deployments", Management Systems Quarterly, 2025, 31 (1), 45-60.
12. Li P.W., "Explainable AI in Banking: Building User Trust in Machine Learning Predictions", AI Ethics in Finance, 2024, 3 (2), 77-92.
13. Singh M.A., "Integrated Financial Dashboards: Design Principles and User Experience Evaluation", Enterprise Systems Review, 2023, 18 (3), 112-128.

14. Davis J.K., "Predictive Decision Support for Consumer Financial Wellness Platforms", Journal of Smart Finance, 2025, 6 (1), 33-48.
15. Thomas R.G., "Financial Data Mining: Techniques for Personal Finance Analytics", Data Intelligence Journal, 2023, 9 (4), 185-200.