International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue:2 Volume No.10 March - April – 2026
DOI:10.46647/ijetms.2026.v10i02.072 ISSN: 2581-4621

SpendWise AI: An Intelligent Personal Finance Assistant

Dr. Babu Rao[1], N. Anvith[2], G. Aravind[3], J. Nandhini[4], M. Tharuni[5]
Department of Artificial Intelligence and Data Science, Vidya Jyothi Institute of Technology, Hyderabad, India

Abstract—*The growing complexity of digital financial transactions has made effective personal finance management an increasingly difficult exercise for the average consumer. Conventional approaches such as paper ledgers and basic spreadsheets are largely passive: they record what has already been spent but rarely produce actionable insight, predictive analysis, or personalized guidance. This paper introduces SpendWise AI, an intelligent personal finance assistant that brings together Machine Learning, Natural Language Processing, and interactive data visualization within a single, accessible platform. The proposed architecture is organized as four interacting tiers — a user interface layer for visual financial dashboards, an application layer responsible for transaction processing and budget calculation, an intelligence layer that hosts the ML and NLP modules, and a data layer that securely stores user financial records. The ML component analyzes historical transactions to identify behavioral patterns, forecast upcoming expenses, and flag anomalies, while the NLP component allows users to query their financial state through plain conversational language rather than rigid menu navigation. The visualization tier presents category-wise breakdowns, monthly trends, and budget progress indicators that update dynamically as new transactions are added. A comparative evaluation against widely deployed tools such as YNAB and Walnut highlights the additional capabilities delivered by SpendWise AI, particularly its conversational interface, predictive intelligence, and explainable outputs. The prototype establishes a strong foundation for future extensions including secure banking API integration, IoT-based real-time updates through MQTT, and cloud-hosted multi-device synchronization.*

Index Terms**—***Personal Finance Management, Artificial Intelligence, Machine Learning, Natural Language Processing, Expense Tracking, Predictive Analytics, Data Visualization, Anomaly Detection, MQTT, Smart Assistant.*

I. INTRODUCTION

Managing personal finances has always been a basic responsibility of adult life, but in the modern digital era this task has become unexpectedly intricate. Although individuals today have access to a wide assortment of banking applications, online wallets, and payment services, many still find it difficult to maintain an accurate ongoing picture of their financial situation. Behavioral economics research has repeatedly shown that limited financial awareness, when combined with weak budgeting habits, tends to result in impulsive purchases, accumulating debt, and a failure to meet longer-term goals such as savings, education funds, or emergency reserves.

Traditional methods of money management — handwritten ledgers, periodic statement reviews, and basic spreadsheet tools — have only been transformed superficially by the introduction of digital alternatives. Most existing budgeting applications behave as passive recorders rather than active assistants. They capture transactions after the fact but rarely provide interpretation, prediction, or personalized recommendations. The user is left to identify trends manually, draw their own conclusions, and decide on appropriate actions without any form of intelligent guidance. For users who lack formal training in finance, this analytical burden is often overwhelming and tends to produce poor financial outcomes.

The maturation of Artificial Intelligence (AI) and its associated subfields — particularly Machine Learning (ML), Natural Language Processing (NLP), and predictive analytics — has substantially expanded what software systems are capable of offering. Capabilities once confined to specialized research environments are now accessible enough to be embedded into consumer-grade applications. Personal finance is a domain particularly well-suited to such integration: AI-driven systems can

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue:2 Volume No.10 March - April – 2026
DOI:10.46647/ijetms.2026.v10i02.072 ISSN: 2581-4621

learn the spending behavior of individual users, generate forward-looking forecasts, and explain financial information in accessible, conversational language.

This work presents SpendWise AI, a smart personal finance assistant designed to bridge the gap between conventional budgeting tools and the intelligent capabilities now achievable through modern AI techniques. The system pursues five core objectives: (1) to provide a unified platform for income tracking, budget allocation, and expenditure monitoring; (2) to deploy ML algorithms for spending pattern analysis, anomaly detection, and forecasting; (3) to integrate NLP to support plain-language financial queries; (4) to present financial data through interactive visual dashboards; and (5) to maintain a modular and extensible architecture that supports the future integration of banking APIs, cloud-based synchronization, and IoT-enabled financial data ingestion through MQTT.

The remainder of this paper is organized as follows. Section II reviews related work in personal finance applications and AI-based financial intelligence. Section III defines the underlying problem and analyzes the system requirements. Section IV describes the proposed system architecture. Section V presents the methodology and implementation details. Section VI discusses experimental results together with a comparative evaluation. Section VII concludes the paper and outlines directions for future research.

II. RELATED WORK

Research in computational personal finance management has expanded considerably over the past decade, with several distinct lines of inquiry contributing to the broader foundation of intelligent financial assistants. Early systems concentrated primarily on the digitization of expense recording, replacing handwritten ledgers with mobile applications and lightweight spreadsheet templates. While such tools improved the convenience of record keeping, they offered limited analytical depth and provided no mechanism to actively support user decision making.

Subsequent studies examined the application of supervised learning to financial transaction data. Approaches based on Decision Trees, Support Vector Machines, and Random Forest classifiers have shown that automated categorization of expenses using merchant descriptions and metadata is both feasible and reasonably accurate. Other works have explored regression models — including Linear Regression and Random Forest Regression — to forecast monthly expenditures from historical patterns. Anomaly detection methods, ranging from simple statistical thresholding to clustering-based techniques such as K-Means and Isolation Forests, have been employed to flag irregular spending events that may signal fraud or unusual user behavior.

A separate strand of research has focused on the role of Natural Language Processing in financial assistant systems. Conversational interfaces built around intent classification, named entity recognition, and dialogue management have been investigated as a way of reducing the cognitive friction associated with traditional menu-driven applications. Tools such as Cleo and Plum have demonstrated that NLP-based interfaces can make financial information more approachable, particularly for users who are unfamiliar with technical financial terminology.

Studies on data visualization have consistently emphasized the importance of clear, dynamic representations of financial information. Comparative work on visualization techniques has shown that interactive charts — particularly category-wise pie charts and time-series line plots — provide significantly stronger support for pattern recognition than static tabular reports. Human-computer interaction research has further argued that AI-generated insights such as predictions, classifications, and behavioral summaries must be accompanied by transparent explanations in order to earn user trust.

Despite the breadth of this prior work, several gaps remain. Many existing studies are conducted on financial datasets drawn from Western economies, where transactional patterns and currency conventions differ noticeably from those of emerging economies such as India. Equally important, most prior systems treat ML, NLP, and visualization as independent modules rather than as components of a unified, end-to-end intelligent platform. SpendWise AI directly addresses these

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue:2 Volume No.10 March - April – 2026
DOI:10.46647/ijetms.2026.v10i02.072 ISSN: 2581-4621

gaps by combining all three capabilities within a single architecture and by tailoring the system to a more diverse and financially heterogeneous user base.

III. PROBLEM DEFINITION AND REQUIREMENT ANALYSIS

The central problem addressed in this work is the absence of an integrated, intelligent, and accessible personal finance management system. Existing applications generally fall into one of two categories. The first comprises passive expense trackers that record transactions but offer minimal analytical insight. The second includes more sophisticated platforms that, while feature-rich, present complex interfaces requiring a steep learning curve and a baseline level of financial literacy. As a consequence, users are typically forced to choose between simplicity and intelligence, with neither option fully satisfying their needs.

A. Functional Requirements

The functional requirements of SpendWise AI follow directly from this problem statement. The system must support secure user authentication and profile management; permit manual entry, modification, and deletion of expense records; allow users to define and monitor budgets across multiple expenditure categories such as food, transport, utilities, and entertainment; analyze historical transactions to identify recurring patterns and anomalies; forecast future expenditures together with associated confidence values; classify transactions automatically into appropriate categories; render an interactive financial dashboard with dynamically updating visualizations; and emit alerts when spending approaches or exceeds defined thresholds.

B. Non-Functional Requirements

The non-functional requirements emphasize accuracy, scalability, security, usability, and reliability. The ML and NLP components must maintain consistent performance across users with diverse spending profiles. Sensitive financial information has to be stored under strong encryption and protected by authenticated access. The system must remain responsive as the user's transaction history grows over time, and the interface must remain comprehensible to users with no formal background in either finance or computing. Color-coded visual indicators and conversational queries are used to lower the cognitive cost of interacting with the system.

IV. SYSTEM ARCHITECTURE

The proposed architecture follows a layered design composed of four interacting tiers, supplemented by an external interaction subsystem. Fig. 1 presents an overview of the architecture and the data flow that connects its components.

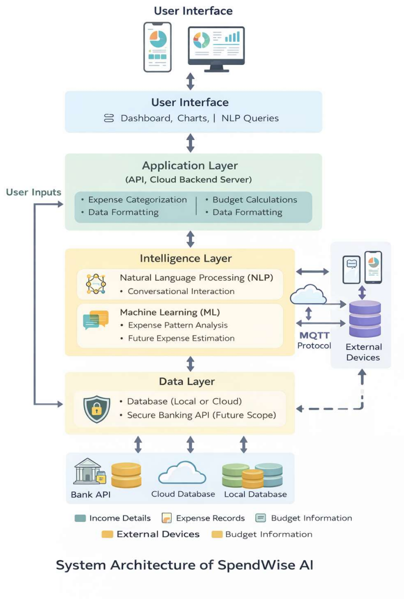

Fig. 1. Layered system architecture of SpendWise AI.

Website: ijetms.in Issue:2 Volume No.10 March - April – 2026
DOI:10.46647/ijetms.2026.v10i02.072 ISSN: 2581-4621

A. User Interface Layer

The User Interface Layer serves as the primary point of contact between the user and the system. It is implemented as a web-based dashboard that presents charts, summary metrics, and forms for both manual and natural-language input. Through this layer, users access category breakdowns, monthly trend visualizations, and budget status indicators, and they submit expense entries either through structured forms or by issuing conversational queries to the AI assistant.

B. Application Layer

The Application Layer mediates between the user interface and the intelligent modules. It exposes a set of REST APIs and is responsible for transaction validation, budget calculation, data formatting, and routing requests to the appropriate downstream component. Conceptually, this layer encapsulates the application logic and ensures that user inputs are translated into the structured representations required by the intelligence and data tiers.

C. Intelligence Layer

The Intelligence Layer constitutes the analytical core of the system. It contains two primary components: an NLP module and an ML module. The NLP module interprets free-form user queries — for example, "How much did I spend on food last month?" — and converts them into structured operations executed against stored data. The ML module is responsible for spending pattern analysis, expenditure forecasting, anomaly detection, and predictive budgeting. Together these components transform the system from a passive recorder into an active assistant.

D. Data Layer

The Data Layer manages the persistent storage of user information. It maintains income records, expense entries, and budget definitions in either a local database (used in the current prototype) or a cloud-based store (planned for future deployment). The layer is also designed to integrate with secure banking APIs in subsequent versions, which will allow transactions to be ingested automatically without manual user intervention.

E. External Device Interaction

External devices — including mobile clients and IoT-enabled appliances — interact with the platform through the MQTT protocol, an event-driven, lightweight publish-subscribe messaging mechanism well-suited to real-time financial updates across multiple devices. This design choice anticipates a future in which financial activity is generated and consumed by a broader ecosystem of connected devices.

V. METHODOLOGY AND IMPLEMENTATION

Development of SpendWise AI followed an Agile Software Development Life Cycle, in which the project was decomposed into a sequence of short, focused sprints. Each sprint produced a working increment of the system, beginning with the foundational user interface and budgeting module and progressively incorporating the ML, NLP, visualization, and integration components. Agile was selected because it accommodates iterative refinement, supports modular development, and aligns naturally with academic project timelines.

A. Data Collection

Two complementary data acquisition strategies were employed. In the prototype phase, users manually enter transaction details — amount, category, date, merchant, and an optional note — through the expense entry form. To supplement user-generated data, publicly available open-source personal finance datasets were used to train and validate the ML models. In future versions, data ingestion will be augmented through secure banking APIs, removing the need for manual entry and enabling true real-time monitoring.

B. Data Preprocessing

Raw financial data is rarely suitable for direct use in modeling. The preprocessing pipeline begins with a cleaning stage that removes duplicate transactions, fills or discards missing fields, and normalizes inconsistent formats such as varied date representations, divergent category labels, and mixed currency symbols. Numerical attributes are then scaled to prevent dominant features from

DOI:10.46647/ijetms.2026.v10i02.072 ISSN: 2581-4621

biasing the model, and categorical fields are encoded into numerical representations. For the NLP component, user queries are tokenized, lowercased, lemmatized, and stripped of stop words before being passed into the intent classification model.

C. Exploratory Data Analysis

Prior to model construction, descriptive statistics — means, medians, variances, and quantile distributions — were computed for transaction amounts, total category spend, and monthly outflows. Correlation analysis was used to examine the relationships between income and expenditure, while temporal analysis searched for seasonal patterns such as recurring month-end spending peaks. The insights derived from this stage informed both feature selection and the broader design of the modeling pipeline.

D. Model Development

Several algorithms were investigated for the analytical tasks of the system. For expense forecasting, Linear Regression and Random Forest Regression were trained on historical spending sequences. For automated transaction classification, Decision Trees, Support Vector Machines, and Gradient Boosting classifiers were evaluated. K-Means clustering was applied to group users into behavioral segments based on their spending patterns. The NLP module relied on intent recognition trained over a curated corpus of finance-specific user utterances, accompanied by lightweight entity extraction for fields such as category, time range, and amount.

E. Model Evaluation

Regression performance was measured using Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and the coefficient of determination (R^2). Classification models were assessed using precision, recall, F1 score, and confusion matrices, with particular attention to minority categories where training data is scarce. The NLP module's intent classification accuracy and entity extraction precision were evaluated against a held-out validation set drawn from a diverse mix of phrasings and complexity levels.

F. Software Stack

The system was implemented in Python (3.8 or later) with Scikit-learn for machine learning, NLTK and spaCy for natural language processing, and Matplotlib together with Chart.js for data visualization. The web interface was built using HTML5 and CSS3 with a Flask backend. MySQL was selected for structured persistent storage. Visual Studio Code served as the primary development environment, while Git and GitHub provided version control and collaboration support throughout the project.

VI. RESULTS AND DISCUSSION

The implemented prototype demonstrates the feasibility of integrating ML, NLP, and interactive visualization within a single personal finance assistant. The dashboard provides users with immediate visibility into their financial state by displaying total spend, the top expense category, and the number of recorded entries (Fig. 2). The spending distribution is rendered as a category-wise donut chart, while monthly trends are visualized as a bar chart, allowing users to detect temporal variation in their behavior at a glance.

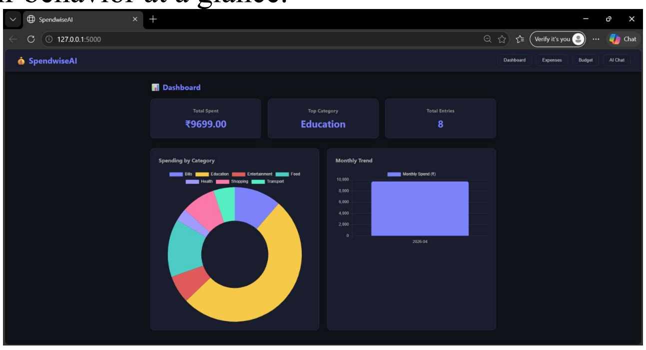

Fig. 2. Main dashboard showing total spend, top category, and the spending-by-category and monthly-trend visualizations.

International Journal of Engineering Technology and Management Sciences
Website: ijetms.in Issue:2 Volume No.10 March - April – 2026
DOI:10.46647/ijetms.2026.v10i02.072 ISSN: 2581-4621

The Budget vs. Spending module offers a continuous visual representation of the user's adherence to their predefined budget limits (Fig. 3). Each category is displayed with a color-coded progress bar — green indicating on-track, yellow signaling proximity to the limit, and red flagging an overage. Textual indicators such as "On Track", "Near Limit", and "Over Budget" complement the color cues to ensure that the interface remains accessible to color-impaired users.

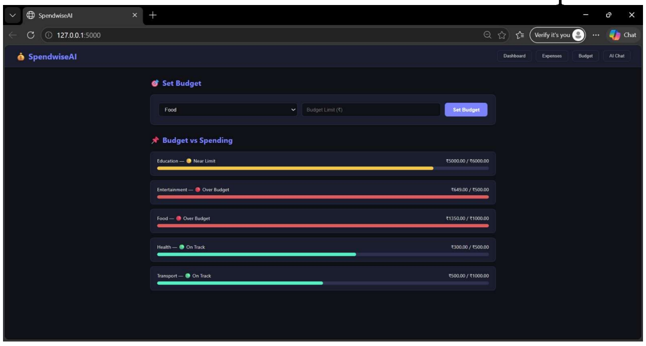

Fig. 3. Budget versus spending view with category-level color-coded progress indicators.

The conversational AI assistant (Fig. 4) allows users to interact with their financial data through plain-language queries such as "Where am I overspending?". The NLP module parses the query, extracts the relevant intent and entities, and produces a structured response generated from the underlying data store. This conversational pathway is intended to serve as a low-friction alternative to traditional menu-based navigation, particularly for users who prefer to ask questions rather than sift through dashboards.

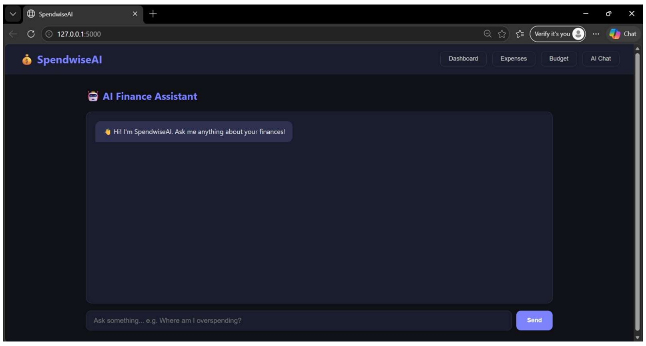

Fig. 4. AI Finance Assistant conversational interface.

The Expense Tracker (Fig. 5) presents a tabular view of recorded transactions and supports adding, viewing, and deleting entries. Each transaction includes a title, amount, category, date, and an optional note, ensuring that the dataset feeding the ML models remains rich and well structured. This view also serves as the primary mechanism through which users curate the historical record on which the predictive models depend.

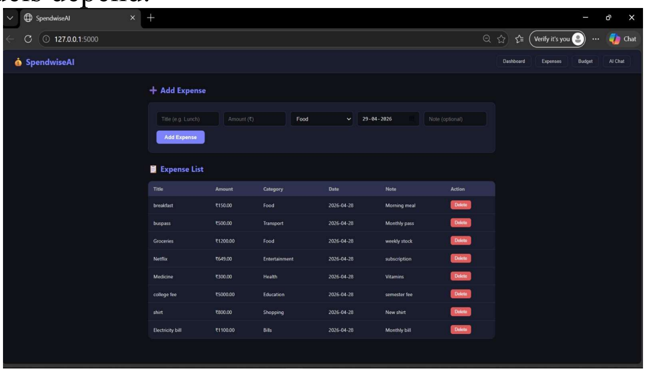

Fig. 5. Expense Tracker view with transaction entry and tabular history.

A. Comparative Analysis

A comparative evaluation of SpendWise AI against two widely used personal finance applications — YNAB (You Need A Budget) and Walnut — was conducted to highlight the distinguishing features of the proposed system. The results are summarized in Table I and indicate that SpendWise

DOI:10.46647/ijetms.2026.v10i02.072 ISSN: 2581-4621

AI provides a substantially broader feature set, particularly in the areas of ML-driven analytics, conversational NLP support, and interactive visualization.

TABLE I

COMPARISON OF SPENDWISE AI WITH EXISTING PERSONAL FINANCE TOOLS

Parameter	YNAB	Walnut	SpendWise AI
Approach	Manual zero-based budgeting	SMS-based transaction tracking	Integrated AI-driven platform
Machine Learning	Not supported	Limited to SMS parsing	Pattern analysis & anomaly detection
NLP Interface	Form-based only	Not available	Conversational queries
Anomaly Detection	Not supported	Not supported	Supported
Visualization	Basic, low interactivity	Simple pie charts	Interactive dashboard with dynamic charts
Accessibility	Requires financial literacy	Easy for Indian users	Accessible via NLP interface

The comparative evidence indicates that none of the surveyed tools simultaneously delivers ML-based behavioral analysis, predictive forecasting, conversational querying, and interactive visualization within a single, cohesive product. SpendWise AI uniquely consolidates these capabilities and, through its modular architecture and India-specific contextual awareness, positions itself as a particularly relevant solution for a diverse and heterogeneous user base.

VII. CONCLUSION AND FUTURE WORK

This paper has presented SpendWise AI, an intelligent personal finance assistant that integrates Machine Learning, Natural Language Processing, and interactive data visualization within a unified, accessible platform. The system addresses the longstanding gap between simple expense trackers and feature-heavy financial management tools by combining analytical intelligence with a conversational, user-friendly interface. The prototype demonstrates that a layered architecture — comprising a user interface tier, an application tier, an intelligence tier, and a data tier — can support both the immediate needs of consumer-grade budgeting and the longer-term goals of scalable, AI-enabled financial services.

Several promising directions remain for future development. The integration of secure banking APIs would automate transaction ingestion and remove the need for manual entry, while the adoption of more advanced ML models — including deep learning architectures for time-series forecasting and transformer-based models for NLP — would substantially enhance predictive accuracy and conversational fluency. Voice-enabled interaction, mobile application packaging, cloud-based deployment for multi-device synchronization, and MQTT-driven IoT connectivity are all natural next steps. Additional functional extensions include goal-based savings tracking, investment recommendations, and personalized financial planning tools, all of which would broaden the scope of the platform from reactive expense tracking toward proactive financial coaching.

In summary, SpendWise AI illustrates that intelligent, AI-augmented financial tools need not be intimidating or technically demanding. With careful design and the considered application of modern AI techniques, it is possible to build personal finance systems that meet users where they are, learn from their behavior over time, and guide them toward more informed and confident financial decisions.

DOI:10.46647/ijetms.2026.v10i02.072 ISSN: 2581-4621

ACKNOWLEDGMENT
The authors gratefully acknowledge the guidance of Dr. Babu Rao, Assistant Professor, Department of Artificial Intelligence and Data Science, Vidya Jyothi Institute of Technology, Hyderabad, for his continuous support throughout this work. The authors also thank Dr. K. S. R. K. Sarma, Head of the Department, Dr. A. Srujana, Principal, and Dr. A. Padmaja, Dean of Accreditation and Rankings, for providing the institutional support and infrastructure that made this project possible.

REFERENCES
[1] D. Kahneman, Thinking, Fast and Slow. New York, NY, USA: Farrar, Straus and Giroux, 2011.
[2] T. M. Mitchell, Machine Learning. New York, NY, USA: McGraw-Hill, 1997.
[3] D. Jurafsky and J. H. Martin, Speech and Language Processing, 3rd ed. Stanford, CA, USA: Stanford University, 2023.
[4] L. Breiman, "Random forests," Machine Learning, vol. 45, no. 1, pp. 5–32, Oct. 2001.
[5] C. Cortes and V. Vapnik, "Support-vector networks," Machine Learning, vol. 20, no. 3, pp. 273–297, 1995.
[6] J. MacQueen, "Some methods for classification and analysis of multivariate observations," in Proc. 5th Berkeley Symp. Math. Statist. Probab., Berkeley, CA, USA, 1967, pp. 281–297.
[7] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.
[8] S. Bird, E. Klein, and E. Loper, Natural Language Processing with Python. Sebastopol, CA, USA: O'Reilly Media, 2009.
[9] OASIS, "MQTT Version 5.0," OASIS Standard, 2019. [Online]. Available: https://docs.oasis-open.org/mqtt/mqtt/v5.0/mqtt-v5.0.html
[10] M. Grinberg, Flask Web Development, 2nd ed. Sebastopol, CA, USA: O'Reilly Media, 2018.
[11] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining, San Francisco, CA, USA, 2016, pp. 785–794.
[12] H. Wickham, "A layered grammar of graphics," J. Comput. Graph. Stat., vol. 19, no. 1, pp. 3–28, 2010.
[13] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," ACM Comput. Surveys, vol. 41, no. 3, pp. 1–58, Jul. 2009.
[14] M. Honnibal and I. Montani, "spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing," 2017. [Online]. Available: https://spacy.io
[15] You Need A Budget (YNAB), "Personal budgeting software," 2024. [Online]. Available: https://www.ynab.com
[16] Walnut, "Walnut – Money manager," Capital Float Holdings, 2024.
[17] J. Brownlee, Deep Learning for Time Series Forecasting. Machine Learning Mastery, 2018.
[18] A. Vaswani et al., "Attention is all you need," in Adv. Neural Inf. Process. Syst., vol. 30, 2017, pp. 5998–6008.