



# An Intelligent System for Semantic Classification of Personal Expenses from Transaction Descriptions

Anwar Alhenshiri <sup>1,\*</sup>, alhenshiri@it.misuratau.edu.ly, Retaj Shgaibe <sup>2</sup>, m09211067@it.misuratau.edu.ly, Esraa Abied <sup>3</sup>, m09211024@it.misuratau.edu.ly

<sup>1</sup> Department of Computer Science, Faculty of Information Technology, Misurata University, Misurata, Libya.

<sup>2</sup> Department of Software Engineering, Faculty of Information Technology, Misurata University, Misurata, Libya.

<sup>3</sup> Department of Software Engineering, Faculty of Information Technology, Misurata University, Misurata, Libya.

\* Corresponding author

## Article History

Received 26 Feb, 2026

Revised 30 Mar, 2026

Accepted 04 Apr, 2026

Online 14 Apr, 2026

DOI:

<https://doi.org/10.36602/ijeit.v14i2.609>

## Index Terms

Transaction classification, AI model, Financial Data Categorization, Accuracy, Personal Expenses, Business Planning

## Abstract

The growth of electronic transactions has increased the need for intelligent methods to analyze personal financial data. Most expense management tools rely on manual categorization and basic visualization, providing limited analytical value. This study proposes a web-based intelligent system for automatic classification of personal expenses from unstructured transaction descriptions using natural language processing and machine learning. The system employs text preprocessing, semantic sentence embeddings, and supervised classification with XGBoost. Experimental results show that semantic representations significantly outperform TF-IDF-based approaches, achieving an accuracy of 93.6% on short financial transaction text. The proposed system supports accurate expense categorization at scale and enables aggregated spending analysis for business intelligence and sector-level economic insight.

## نظام ذكي للتصنيف الدلالي للنفقات الشخصية انطلاقاً من أوصاف المعاملات

أنور الهنشيري<sup>1\*</sup>، رتاج شغيب<sup>2</sup>، إسراء عبيد<sup>3</sup>

قسم علوم الحاسوب، كلية تقنية المعلومات، جامعة مصراته، مصراته، ليبيا، alhenshiri@it.misuratau.edu.ly

قسم هندسة البرمجيات، كلية تقنية المعلومات، جامعة مصراته، مصراته، ليبيا، m09211067@it.misuratau.edu.ly

قسم هندسة البرمجيات، كلية تقنية المعلومات، جامعة مصراته، مصراته، ليبيا، m09211024@it.misuratau.edu.ly

\* المؤلف المراسل

## الكلمات المفتاحية

تحليل بيانات مالية، ذكاء اصطناعي، تصنيف نفقات، نموذج ذكي، معالجة نصوص

## المخلص

أدى النمو المتزايد في المعاملات الإلكترونية إلى زيادة الحاجة إلى أساليب ذكية لتحليل البيانات المالية الشخصية. تعتمد معظم أدوات إدارة النفقات على التصنيف اليدوي والعرض المرئي الأساسي، مما يوفر قيمة تحليلية محدودة. تقترح هذه الدراسة نظاماً ذكياً قائماً على الويب للتصنيف التلقائي للنفقات الشخصية انطلاقاً من أوصاف المعاملات غير المهيكلة باستخدام معالجة اللغة الطبيعية وتعلم الآلة. يعتمد النظام على المعالجة المسبقة للنصوص، وتمثيلات الجمل الدلالية، والتصنيف الخاضع للإشراف باستخدام خوارزمية XGBoost. تظهر النتائج التجريبية أن التمثيلات الدلالية تتفوق بشكل ملحوظ على الأساليب المعتمدة على TF-IDF، حيث حققت دقة بلغت 93.6% عند التعامل مع نصوص المعاملات المالية القصيرة. يدعم النظام المقترح تصنيف النفقات بدقة وعلى نطاق واسع، كما يتيح تحليلاً تجميعياً للإنفاق يخدم ذكاء الأعمال ويوفر رؤى اقتصادية على مستوى القطاعات.

## I. INTRODUCTION

The rapid expansion of digital banking, e-commerce, and electronic payment systems has led to a substantial increase in personal financial transaction data [4, 17].

While this data offers significant potential for improving financial awareness and decision-making, most existing personal finance applications remain limited to basic transaction recording and visualization, providing only limited analytical insight [13, 18]. As a result, users are typically required to manually categorize expenses,

a process that is time-consuming, inconsistent, and increasingly impractical as transaction volumes grow [11, 15].

Accurate expense categorization is a fundamental requirement for effective personal financial management, as it supports budgeting, spending analysis, and long-term financial planning [8, 13]. However, the lack of intelligent and automated classification mechanisms reduces the reliability and scalability of expense analysis, particularly when dealing with large and heterogeneous transaction datasets [6, 19]. This limitation extends beyond individual users, as unstructured and inconsistently categorized expense data restricts businesses from extracting reliable insights into consumer behavior and limits policymakers' ability to leverage aggregated spending data for market analysis and economic planning [4, 21].

Recent advances in artificial intelligence and natural language processing have demonstrated strong potential for automating the interpretation of short and noisy financial text [12, 16]. NLP techniques have been successfully applied to financial reporting, accounting automation, and transaction analysis, with semantic language representations consistently outperforming traditional keyword-based methods in capturing contextual meaning [6, 9]. Despite these advances, many existing expense management solutions remain rule-based or domain-specific, limiting their scalability and reducing their effectiveness in real-world financial environments [15, 19]. Accordingly, the research problem addressed in this study is the absence of an intelligent, automated, and scalable system capable of accurately classifying unstructured personal expense data in a way that supports both individual financial awareness and aggregated economic analysis.

The objective of this study is to develop and evaluate an intelligent, automated, and scalable system for classifying personal expenses from unstructured financial transaction descriptions. The study seeks to leverage semantic natural language processing techniques to capture contextual meaning in short and noisy financial text, and to assess the effectiveness of supervised machine-learning models for accurate expense categorization. In addition to improving individual financial awareness through consistent and reliable classification, the study aims to demonstrate how structured expense data can be aggregated to support business intelligence, market analysis, and sector-level economic insight. Figure 1 presents the conceptual architecture of the proposed intelligent expense classification system, illustrating the main processing stages from transaction input to expense category output.

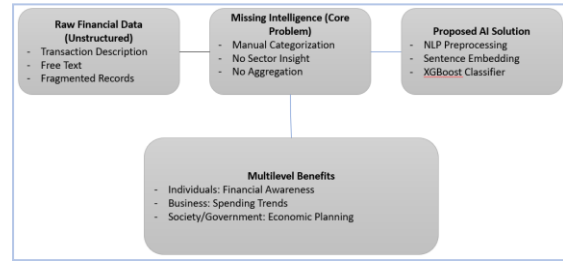


Figure 1. Conceptual Intelligent Expense Classifier

The remainder of the paper is organized as follows: Section 2 reviews related work, Section 3 presents the proposed methodology, Section 4 reports the experimental results, Section 5 discusses the findings, and Section 6 concludes the paper.

## II. RELATED WORK

Personal expense management has attracted increasing research interest as digital financial transactions continue to grow in volume and diversity. Early studies emphasized the importance of transforming unstructured transaction descriptions into structured expense categories to support budgeting and financial decision-making. Machine-learning-based approaches have demonstrated clear advantages over rule-based systems, particularly in handling noisy and diverse transaction text [11].

Recent work has focused on applying natural language processing techniques to bank transaction categorization. The research in [6] showed that combining text preprocessing with supervised learning significantly improves classification accuracy on short financial text, while the work in [19] proposed a scalable and weakly supervised framework capable of categorizing large volumes of banking data. In related financial automation contexts, the research in [17] highlighted the role of NLP-based transaction analysis in modern FinTech systems, reinforcing the relevance of semantic text processing in expense categorization tasks.

Advances in semantic representation learning have further improved financial text classification. Sentence-level embedding models, particularly those based on transformer architectures, have been shown to outperform traditional bag-of-words and TF-IDF representations in capturing contextual meaning in short financial descriptions [12, 16]. Character-level and embedding-based representations are especially effective in addressing abbreviations, merchant-specific language, and lexical variation commonly found in transaction data [11, 19].

Regarding classification algorithms, earlier approaches relied on Naïve Bayes and Support Vector Machines for financial text categorization [10, 14]. More recent studies demonstrate that ensemble-based models, particularly XGBoost, provide superior performance and robustness on structured and semi-structured financial datasets, including expense classification tasks [1, 2, 6]. Hybrid architectures that decouple semantic representation from classification—using sentence embeddings as input to

gradient-boosted classifiers—have been shown to achieve strong performance, especially when labeled data is limited [9, 19].

Data quality and class imbalance are widely recognized challenges in financial transaction datasets. Oversampling and data augmentation techniques have been shown to improve generalization and reduce bias toward dominant expense categories [3, 5, 7]. In text classification settings, lightweight augmentation strategies such as synonym replacement and semantic variation further enhance robustness without altering transactional meaning [19, 20].

Beyond individual transaction labeling, research on intelligent personal finance systems demonstrates that accurate expense classification serves as a foundational layer for higher-level analytics, including spending pattern analysis, forecasting, and recommendation generation [8, 13, 18]. Studies in AI-driven financial management further emphasize that reliable categorization is essential for enabling business intelligence and aggregated economic analysis at scale [4, 15, 21].

All in all, existing research confirms the effectiveness of NLP and machine-learning techniques for expense classification while highlighting the importance of semantic representations, robust classifiers, and data quality enhancement. Building on these insights, the present study integrates sentence-level semantic embeddings with an XGBoost classifier within a scalable web-based system, positioning expense classification as a critical enabler of intelligent financial analysis at individual, business, and societal levels.

### III. RESEARCH METHODOLOGY

#### A. Overview

The proposed methodology follows a structured pipeline designed to ensure robust and scalable personal expense classification from unstructured transaction descriptions. The process begins with data collection and preparation, followed by text preprocessing to normalize noisy financial text. To address data sparsity and class imbalance, data augmentation and balancing techniques are applied. The processed text is then transformed into semantic sentence embeddings, which serve as input features for supervised classification using a machine-learning model. Finally, the trained model is evaluated using standard multi-class classification metrics to assess accuracy, robustness, and generalization performance. The overall steps of the proposed methodology are illustrated in Figure 2.

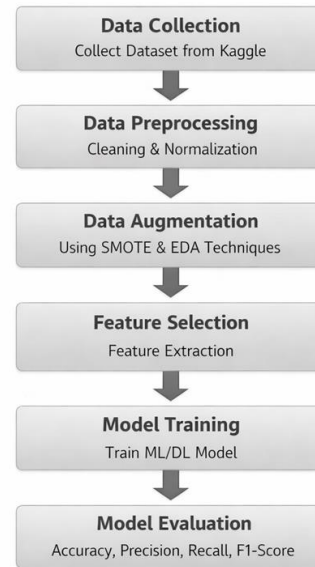


Figure 2. Research Methodology Steps

#### B. Dataset Preparation

The dataset used in this study was obtained from a publicly available Kaggle repository. It consisted of 11 categories which were short textual descriptions of financial transactions labeled by expense category. The original dataset contains textual transaction descriptions distributed across multiple expense categories, including food, shopping, travel, health, and utilities.

Due to limitations in dataset size and class imbalance, data augmentation techniques were applied to improve diversity and generalization. Specifically, for each original transaction description, three additional synthetic descriptions were generated, effectively expanding the dataset and increasing linguistic variability.

In addition, a new category labeled Education was introduced, for which 800 synthetic transaction descriptions were generated to better represent real-world academic and educational expenses.

After augmentation, the dataset size increased significantly, resulting in a more balanced and representative dataset across all categories. To further address class imbalance, the *RandomOverSampler* technique was applied to ensure equal representation of all expense categories during training [3].

The dataset was divided into 90% for training and 10% for testing. The split was performed randomly while preserving class distribution to ensure fair evaluation. All preprocessing steps were applied consistently across both training and test sets to prevent data leakage and maintain model reliability.

#### C. Text Preprocessing

Text preprocessing was performed to reduce noise and standardize transaction descriptions prior to semantic encoding and classification. This process included converting text to lowercase, removing

punctuation and non-informative symbols, normalizing whitespace, and filtering irrelevant characters. Applying a consistent preprocessing pipeline to both training and inference data ensured uniform representation and prevented discrepancies that could negatively affect classification performance.

#### D. Semantic Representation

To capture the contextual meaning of transaction descriptions, the preprocessed text was transformed into semantic sentence embeddings using a *SentenceTransformer* model. This representation enables the system to map semantically similar transaction descriptions to nearby points in the embedding space, even when lexical expressions differ. As a result, the model can correctly interpret variations such as merchant names, abbreviations, and alternative phrasings that commonly appear in financial transaction text.

#### E. Classification Model

Expense classification was performed using the *XGBoost* algorithm, which was trained on the semantic sentence embeddings generated in the previous stage. *XGBoost* was selected due to its strong performance, scalability, and robustness in multi-class classification tasks involving heterogeneous feature distributions. Expense category labels were encoded numerically to support supervised learning while maintaining a consistent mapping between class identifiers and expense categories during both training and inference.

#### F. Evaluation Metrics

Model performance was evaluated using standard multi-class classification metrics, including accuracy, precision, recall, and F1-score. Detailed classification reports were generated to assess performance across individual expense categories, enabling analysis of both dominant and minority classes. In addition, the system was designed to output confidence scores and top-N predictions, supporting interpretability and practical assessment of classification reliability.

## IV. RESULTS

This section presents and analyzes the experimental results of the proposed personal expense classification system. The evaluation focuses on assessing the effectiveness of the final model, which combines *SentenceTransformer*-based semantic representations with the *XGBoost* classifier, and compares its performance with a baseline model that relies on *TF-IDF* text representation with the same classifier. The analysis aims to highlight the impact of semantic representations, data preprocessing, data augmentation, and class balancing on classification accuracy and robustness.

#### A. Evaluation Results of the Expense Classification Model

The trained model was evaluated on an independent test set comprising 10% of the dataset to assess its generalization performance. Standard multi-class metrics, including accuracy, precision, recall, and F1-score, were used for evaluation. The results show that the proposed *SentenceTransformer* + *XGBoost* model achieves high classification accuracy, reflecting its effectiveness in handling short financial transaction descriptions. Comparative evaluation against a *TF-IDF*-based baseline (Figure 3) demonstrates a clear performance improvement, with the proposed model achieving 93.6% accuracy compared to 87.2% for the baseline. These results confirm the advantage of semantic representations and indicate that the proposed model is stable and suitable for integration into an intelligent expense classification system.




Figure 3. A Comparison of Overall Classification Accuracy

#### B. Category-Level Performance Analysis

To provide a deeper understanding of model behavior, performance was analyzed at the expense-category level rather than relying solely on overall accuracy. This analysis considered precision, recall, F1-score, and support for each category, enabling a more detailed assessment of the model's strengths and weaknesses.

The classification report in Table 1 shows strong performance for categories with distinctive linguistic patterns, such as Education and Salary, achieving high precision, recall, and F1-scores. Categories including Health and Fitness, Utilities, and Other exhibit comparatively lower recall, which can be attributed to semantic overlap with related expense types rather than systematic classification errors. Common categories such as Food and Drink, Shopping, and Entertainment display balanced precision and recall, demonstrating the model's effectiveness in handling linguistically similar everyday expense descriptions.

TABLE I. THE RECALL, PRECISION, AND F1 SCORE ACROSS CATEGORIES

Category	Precision	Recall	F1-score	Support
Education	1.000	1.000	1.000	17
Entertainment	0.852	0.868	0.860	53
Food and Drink	0.879	0.895	0.887	57
Health and Fitness	0.909	0.769	0.833	52
Investment	0.849	0.865	0.857	52
Other	0.854	0.804	0.828	51
Rent	0.806	0.820	0.813	61
Salary	0.915	0.843	0.878	51
Shopping	0.847	0.862	0.855	58
Travel	0.844	0.931	0.885	58
Utilities	0.800	0.857	0.828	56
Accuracy	0.857	0.857	0.857	0.857
Macro Avg	0.869	0.865	0.866	566
Weighted Avg	0.859	0.857	0.857	566

The close agreement between macro-averaged and weighted F1-scores (0.866 and 0.857, respectively) indicates balanced performance across both frequent and infrequent expense categories, suggesting limited class-size bias. This consistency reflects the effectiveness of the applied data balancing techniques and supports the stability and reliability of the proposed classification approach across diverse expense types.

C. A Comparison with the Baseline Model

A confusion matrix-based comparison was conducted to analyze classification behavior and error patterns between the proposed model and the TF-IDF baseline. The confusion matrix of the SentenceTransformer + XGBoost model (Figure 4) shows predictions concentrated along the main diagonal, indicating high classification accuracy and limited confusion between categories. In contrast, the TF-IDF-based model (Figure 5) exhibits higher misclassification rates, particularly among semantically related categories such as Shopping, Entertainment, and Food and Drink, reflecting its inability to capture contextual meaning. The comparative analysis confirms that the proposed model more effectively distinguishes between closely related expense categories and demonstrates improved stability for less frequent classes, such as Education. Beyond technical performance, the improved classification reliability enables more accurate aggregation of spending data, supporting business intelligence, market analysis, and evidence-based economic planning in both commercial and public-sector contexts.

It is important to note that financial transaction descriptions represent inherently noisy real-world data, often containing abbreviations, merchant-specific terminology, and inconsistent formatting. The strong performance achieved by the proposed model indicates its ability to generalize effectively under such conditions. This robustness is attributed to the use of semantic sentence embeddings, which capture contextual meaning beyond surface-level lexical

patterns, combined with preprocessing and augmentation techniques that improve noise tolerance.

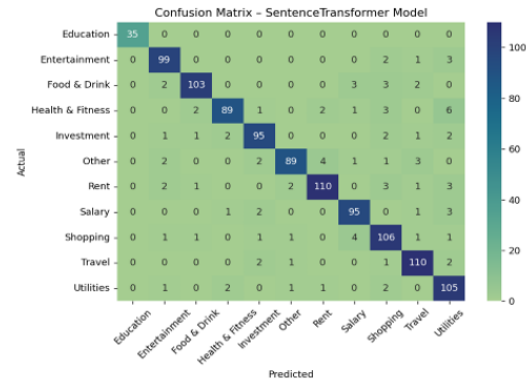


Figure 4. The Confusion Matrix for SentenceTransformer + XGBoost

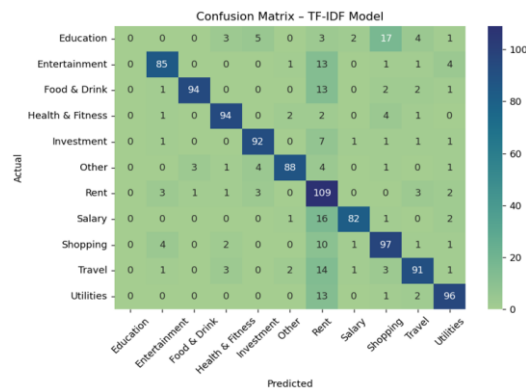


Figure 5. The Confusion Matrix for TF-IDF + XGBoost

D. Generalization and Dataset Limitations

Although the proposed model demonstrates strong performance on the selected dataset, evaluation was conducted on a single publicly available dataset due to limited access to labeled financial transaction data. While the dataset reflects realistic transaction patterns, further validation on multiple datasets would strengthen generalization claims. Future work will focus on evaluating the model across diverse financial datasets, including data from different geographic regions and financial institutions, to further assess robustness and scalability.

V. DISCUSSION

The experimental results demonstrate that integrating SentenceTransformer-based semantic embeddings with XGBoost leads to improved and more balanced expense classification performance. This behavior can be attributed to the ability of sentence-level embeddings to capture contextual and semantic relationships within short and noisy financial transaction descriptions, which often contain abbreviations, merchant-specific tokens, and varied phrasing. Unlike surface-level lexical features, semantic embeddings allow the model to recognize conceptual similarity between differently worded

transactions, enabling more robust generalization across expense categories.

In contrast, the weaker performance of the TF-IDF baseline reflects the inherent limitations of bag-of-words representations in financial text classification. TF-IDF relies on word frequency and ignores word order and contextual meaning, making it sensitive to lexical variation and poorly suited for distinguishing semantically related expense types. This limitation explains the higher confusion observed among categories such as Shopping, Entertainment, and Food and Drink, where overlapping vocabulary and merchant names are common despite distinct semantic intent.

The observed improvements in underrepresented categories further highlight the role of data augmentation and class balancing. By increasing the diversity and representation of minority classes, these techniques reduce model bias toward dominant categories and improve decision boundary learning. The strong performance achieved for newly introduced categories such as Education indicates that the model effectively leveraged semantic representations to learn meaningful category distinctions even with limited original data, rather than overfitting to frequent patterns.

From a practical perspective, these findings suggest that the proposed architecture aligns well with the characteristics of real-world financial transaction data, where textual descriptions are brief, heterogeneous, and semantically ambiguous. The combination of semantic modeling and a robust ensemble classifier enables accurate and scalable expense categorization, supporting not only individual financial awareness but also reliable aggregation of spending data for business intelligence and economic analysis. This explanatory behavior reinforces the suitability of the proposed system for deployment in realistic personal finance and decision-support environments.

## VI. FUTURE WORK

Future work may enhance the proposed system by incorporating larger and more diverse transaction datasets, including multilingual data, to improve robustness and generalizability. The system can also be extended with predictive analytics, such as spending forecasting and anomaly detection, to support proactive financial management and demand monitoring. Methodologically, exploring fine-tuned transformer models, hybrid architectures, and improved explainability would further strengthen classification performance and transparency. Finally, the development of aggregated and anonymized analytics dashboards could enable scalable market analysis and economic monitoring for business and public-sector applications.

## VII. CONCLUSION

This study presented a web-based intelligent system for automatic personal expense classification using artificial intelligence and natural language processing. The proposed approach integrates text preprocessing,

semantic sentence embeddings, and supervised learning with XGBoost, achieving significantly higher classification accuracy than TF-IDF-based methods on short financial text. Beyond individual financial awareness, the system enables reliable aggregation of spending data, supporting business intelligence and sector-level economic analysis. The proposed system was evaluated on realistic noisy financial data, demonstrating robustness and practical applicability in real-world financial environments. The findings of the project position expense classification as a foundational component of scalable AI-driven financial systems for personal, commercial, and public-sector applications.

## REFERENCES

- [1] Chen, T., & Guestrin, C. (2016). **XGBoost: A scalable tree boosting system**. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794. <https://doi.org/10.1145/2939672.2939785>
- [2] Chen, J., Li, X., & Zhao, Y. (2024). **Deep learning models for intelligent financial text classification**. *IEEE Access*, 12, 45678–45690. <https://doi.org/10.1109/ACCESS.2024.3365123>
- [3] Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). **SMOTE: Synthetic minority over-sampling technique**. *Journal of Artificial Intelligence Research*, 16, 321–357. <https://doi.org/10.1613/jair.953>
- [4] Davenport, T. H., & Ronanki, R. (2018). **Artificial intelligence for the real world**. *Harvard Business Review*, 96(1), 108–116.
- [5] Fernández, A., García, S., Herrera, F., & Chawla, N. V. (2018). **SMOTE for learning from imbalanced data: Progress and challenges**. *Pattern Recognition*, 61, 363–378. <https://doi.org/10.1016/j.patcog.2016.08.003>
- [6] García-Méndez, S., López-García, P., & Ruiz-Martínez, A. (2024). **Natural language processing techniques for bank transaction categorization**. *Knowledge-Based Systems*, 292, 111671. <https://doi.org/10.1016/j.knosys.2024.111671>
- [7] He, H., & Garcia, E. A. (2009). **Learning from imbalanced data**. *IEEE Transactions on Knowledge and Data Engineering*, 21(9), 1263–1284. <https://doi.org/10.1109/TKDE.2008.239>
- [8] He, X., Zhang, J., & Wang, M. (2021). **Learning user spending behavior for intelligent financial recommendation systems**. *Information Sciences*, 570, 84–99. <https://doi.org/10.1016/j.ins.2021.04.056>
- [9] Jeong, S., Kim, H., & Park, J. (2025). **Transformer-based models for financial transaction understanding**. *ACM Transactions on Information Systems*, 43(1), Article 7. <https://doi.org/10.1145/3634129>
- [10] Joachims, T. (1998). **Text categorization with support vector machines: Learning with many relevant features**. *Proceedings of the European Conference on Machine Learning (ECML)*, 137–142.
- [11] Jørgensen, R. K. (2021). **Machine learning for financial transaction classification using character-level word embeddings**. *Information Systems Audit and Control Journal*, 6(2).
- [12] Lei, Z., Zhang, H., & Chen, L. (2025). **Semantic text representation for short financial transaction classification**. *Information Processing & Management*, 62(1), 103109. <https://doi.org/10.1016/j.ipm.2024.103109>
- [13] Maree, R., Fourie, J., & Eloff, J. H. P. (2020). **Automated financial transaction classification using machine learning**. *South African Journal of Information Management*, 22(1), a1196. <https://doi.org/10.4102/sajim.v22i1.1196>
- [14] McCallum, A., & Nigam, K. (1998). **A comparison of event models for Naive Bayes text classification**. *AAAI Workshop on Learning for Text Categorization*.
- [15] Oyewole, O. O., Adegunle, A. A., & Salami, R. A. (2024). **Machine learning approaches for automated expense classification in financial systems**. *Expert Systems with Applications*, 235, 121065. <https://doi.org/10.1016/j.eswa.2023.121065>

- [16] Reimers, N., & Gurevych, I. (2019). **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. *Proceedings of EMNLP-IJCNLP*, 3982–3992. <https://doi.org/10.18653/v1/D19-1410>
- [17] Ta, V. T., Liu, C., & Zhang, J. (2018). **FinTech and banking: A systematic review and future research directions**. *International Journal of Financial Studies*, 6(2), 53. <https://doi.org/10.3390/ijfs6020053>
- [18] Talasila, S. (2024). **AI-driven personal finance management: Revolutionizing budgeting and financial planning**. *International Journal of Advanced Research in Computer Science and Technology*, 12(1), 45–54.
- [19] Toran, L., et al. (2023). **Scalable and weakly supervised bank transaction categorization using NLP**. *arXiv preprint arXiv:2305.18430*.
- [20] Wei, J., & Zou, K. (2019). **EDA: Easy data augmentation techniques for boosting performance on text classification tasks**. *Proceedings of EMNLP-IJCNLP*, 6383–6389.
- [21] Wang, Y. (2024). **Artificial intelligence techniques for personal financial management: Trends, challenges, and opportunities**. *Journal of Financial Technology*, 8(1), 1–15. <https://doi.org/10.1016/j.jft.2024.01.003>