INTERNATIONAL JOURNAL OF NOVEL RESEARCH AND DEVELOPMENT (IJNRD)
© 2026 IJNRD | Volume 11, Issue 5, May 2026 | ISSN: 2456-4184 | IJNRD.ORG

GuardBuy: An AI-Based Intelligent System for Pre-Purchase Regret Prediction in Online Shopping Using Hybrid Machine Learning

Vijipriya Jeyamani[1], Maryam Alshammari[2], Noha Khalid Altalasi[2], Lama Mutlab Alshammari[2], Sadeem Muslet Almutairi[2], and Norah Ayedh Alhajji[2]

[1]Department of Computer Engineering, College of Computer Science and Engineering, University of Ha'il, Hail, Saudi Arabia

[2]Department of Artificial Intelligence & Data science , College of Computer Science and Information Technology, University of Ha'il, Hail, Saudi Arabia

Abstract—Online shopping has become an essential part of modern life, providing customers with convenient access to a wide variety of products through different e-commerce platforms. Despite these advantages, many customers find themselves dissatisfied with items they have acquired, as those items fail to align with what was initially anticipated. This phenomenon, known as post-purchase regret, has become a common issue in online shopping environments. This paper proposes GuardBuy, an intelligent AI-based system that predicts potential post-purchase dissatisfaction and provides customers with a regret risk assessment before finalizing an online purchase. The system retrieves product reviews from a curated dataset of over 140,000 Amazon product reviews and applies a multi-stage analysis pipeline including text preprocessing, feature extraction via TF-IDF, sentiment analysis using VADER, and a hybrid machine learning model combining Logistic Regression (70%) and Support Vector Machine (30%). A novel BUT-aware preprocessing technique is introduced to identify and prioritize complaint-bearing text segments. Experimental results demonstrate that the Logistic Regression model achieves 91% accuracy with 0.75 recall for the regret class, while the Linear SVM achieves 93% accuracy with 0.60 recall. The hybrid model achieves 93% overall accuracy with improved consistency in regret detection. The system is deployed as a complete end-to-end web application enabling real-time purchase risk assessment.

Index Terms—Artificial Intelligence, Machine Learning, Natural Language Processing, Online Shopping, Post-Purchase Regret, Sentiment Analysis, TF-IDF, Hybrid Classification, E-Commerce, Decision Support Systems.

I. INTRODUCTION

Online shopping has become an integral part of daily life, fundamentally transforming the way consumers access goods and services. The rapid expansion of digital marketplaces has enabled customers to browse millions of products, compare prices, and complete purchases without visiting physical stores. According to recent global reports, the e-commerce market has grown substantially over the past decade, with revenues projected to surpass several trillion US dollars annually [1]. This growth has been driven by widespread smartphone adoption, improved internet connectivity, and increasing digitization of consumer behavior worldwide. Platforms such as Amazon and eBay host hundreds of millions of product listings across virtually every category.

Despite the many advantages of online shopping, it introduces a fundamental limitation: customers cannot physically examine products before purchase. In physical stores, customers assess quality, dimensions, and functionality through direct interaction. Online shopping removes this capability, forcing consumers to rely on product images, written descriptions, and seller specifications. This informational gap significantly increases the risk of purchasing products that do not align with expectations, leading to a psychological phenomenon known as post-purchase regret [2].

Post-purchase regret is defined as the negative cognitive and emotional response experienced after completing a purchase when the product fails to meet prior expectations [2]. In online environments, this has become particularly prevalent with significant consequences for both consumers and platforms. For individuals, regret leads to dissatisfaction and loss of confidence. For e-commerce platforms, widespread regret results in increased return rates, higher support costs, and damage to platform reputation [3]. Studies show that customers who experience post-purchase regret are significantly less likely to return to the same platform, making regret prevention critical for sustainable e-commerce growth.

Customer reviews represent authentic accounts of real user experiences and often contain detailed descriptions of product advantages, disadvantages, and recurring complaints [4]. Research consistently demonstrates that customer reviews are among the most influential factors in purchase decision-making [5][6]. However, a popular product on Amazon may accumulate thousands of individual reviews, making it practically impossible for any customer to manually read and analyze all available feedback [7]. Customers typically examine a small number of visible reviews or rely on aggregate star ratings, which do not provide a complete representation of user experiences.

To address the challenge of processing large volumes of unstructured textual review data, researchers in Natural Language Processing (NLP) and Machine Learning (ML) have proposed various computational approaches. Sentiment analysis classifies textual content into emotional categories such as positive, negative, or neutral [8]. Early methods relied on lexicon-based approaches, while more advanced approaches have employed supervised ML algorithms including Logistic Regression, Support Vector Machines (SVM), Naïve Bayes, and Random Forest [9]. More recently, deep learning architectures such as LSTM networks and CNNs have achieved state-of-the-art results on sentiment analysis tasks [11][14].

Despite progress in sentiment analysis, a critical gap remains: most published systems classify overall review sentiment without specifically targeting post-purchase regret prediction or generating actionable risk assessments for pre-purchase decision-making. Furthermore, existing approaches fail to account for structural complexity in review language, where reviews contain contrast structures such as 'The product looks great, but the quality is very poor.' Standard sentiment models process the entire review without distinguishing such components, potentially diluting the impact of negative feedback.

This paper presents GuardBuy, an intelligent AI-based system that predicts potential post-purchase dissatisfaction and provides customers a regret risk assessment before finalizing an online purchase. The main contributions of this research are: (i) developing a hybrid regret prediction model combining Logistic Regression and SVM; (ii) proposing a BUT-aware text preprocessing technique to enhance detection of complaint-bearing content; (iii) integrating VADER-based sentiment analysis as a secondary validation layer; (iv) building a complete deployable end-to-end web application; and (v) implementing a structured analysis history system for retrieving previous product assessments.

II. RELATED WORK

Sentiment analysis has been extensively studied in the context of e-commerce customer reviews. Pang and Lee [8] provided a foundational survey on opinion mining and sentiment analysis, establishing core methodologies that have guided subsequent research. Early approaches relied on lexicon-based methods that scored reviews based on predefined positive or negative keywords, achieving reasonable performance on simple classification tasks but failing to capture complex linguistic phenomena such as negation and sarcasm.

Supervised machine learning approaches have demonstrated strong performance on review classification tasks. Fang and Zhan [10] demonstrated that TF-IDF representations combined with SVM classifiers achieve competitive performance on sentiment classification across multiple product domains. Ahmed et al. [18] further validated machine learning techniques for e-commerce sentiment analysis, showing that ensemble methods outperform individual classifiers. Kowsari et al. [19] provided a comprehensive survey of text classification algorithms applicable to review analysis, covering methods from Naïve Bayes to deep neural networks.

Deep learning architectures have substantially advanced sentiment analysis capabilities. Kim [13] (2014) introduced a seminal CNN architecture for sentence classification, demonstrating that locally-connected filters effectively capture n-gram-level features for text classification tasks. Hochreiter and Schmidhuber [14] (1997) proposed the Long Short-Term Memory (LSTM) network — a recurrent architecture with gated memory cells that overcomes vanishing gradient issues — which has since become a standard baseline for sequential review analysis, capturing long-range contextual dependencies critical for detecting regret-related sentiment. Zhang et al. [11] (2018) surveyed deep learning approaches for sentiment analysis, confirming that LSTM and CNN models consistently outperform traditional ML classifiers on benchmark review datasets.

The emergence of pre-trained language models has further advanced the state of the art. Devlin et al. [12] (2019) introduced BERT, demonstrating that deep bidirectional pre-training on large corpora yields substantial improvements on downstream NLP tasks including sentiment classification. Fine-tuned BERT models have since become standard upper-bound baselines on review classification benchmarks. However, the substantial computational requirements of transformer-based models — both at training and inference time — present practical challenges for real-time deployment in consumer-facing applications, motivating the use of more efficient machine learning approaches such as the hybrid LR-SVM model proposed in the current work. The integration of BERT, LSTM, and CNN as comparative deep learning baselines remains an identified gap addressed in future work.

Hutto and Gilbert [27] introduced VADER, a rule-based sentiment analysis tool specifically designed for social media and informal text, which achieves strong performance without requiring model training. This lexicon-based approach is particularly suitable for integration as a complementary component in hybrid prediction systems, as proposed in this work.

Al-Smadi et al. [20] conducted a survey on sentiment analysis for social media and e-commerce platforms, identifying key challenges including domain adaptation, multilingual support, and handling of implicit sentiment. Despite extensive research on sentiment analysis, the specific problem of pre-purchase regret prediction from review data has received limited attention, representing the primary motivation for the current work.

III. METHODOLOGY

A. Overview of the Proposed System

This study proposes GuardBuy, an intelligent end-to-end system designed to predict pre-purchase regret based on user-generated e-commerce reviews. The methodology integrates data processing, custom labeling, feature extraction, hybrid machine learning modeling, and system deployment into a unified pipeline. The overall workflow begins with collecting product reviews from e-commerce platforms, followed by preprocessing and transforming textual data into numerical representations. A custom regret labeling mechanism is applied to generate supervised learning targets. Multiple machine learning models are then trained and combined using a hybrid approach to produce a final regret prediction score.

Fig. 1. System Workflow Diagram — GuardBuy end-to-end pipeline from user input to final regret score.

B. Data Collection

The dataset used in this research consists of 144,754 customer reviews collected from e-commerce platforms, primarily Amazon. Each review includes textual feedback along with implicit sentiment indicators reflecting user satisfaction or dissatisfaction after purchase. The dataset is highly imbalanced, where only 8,537 reviews (~6%) belong to the regret class, reflecting real-world scenarios where explicit regret expressions are relatively rare compared to neutral or positive feedback. The collected data represents diverse product categories, ensuring that the model generalizes across different domains rather than being limited to a specific product type.

C. Data Preprocessing

Before model training, several preprocessing steps were applied to clean and standardize the textual data: (i) removal of noise such as punctuation, special characters, and irrelevant symbols; (ii) conversion of all text to lowercase for uniformity; (iii) tokenization of sentences into individual words; (iv) removal of stop words to reduce redundancy; and (v) handling of negation phrases such as 'not good' and 'no issues' to preserve contextual meaning.

Additionally, advanced preprocessing improvements were implemented including negation protection, ensuring phrases like 'not bad' are not misclassified as negative; contrast sentence handling for mixed sentiments such as 'good but expensive'; and filtering of weak or ambiguous reviews to improve label quality. The novel BUT-aware preprocessing technique identifies and prioritizes text following contrast words such as 'but,' 'however,' 'although,' and 'though,' ensuring that complaint-bearing segments receive appropriate attention during classification.

D. Regret Labeling Strategy

Since regret is not explicitly labeled in the raw data, a custom rule-based labeling system was developed. This mechanism assigns a regret score based on weighted indicators: presence of negative sentiment words, expressions of dissatisfaction (e.g., 'waste of money,' 'not worth it'), mentions of refunds or returns, and strong contrast phrases indicating regret after purchase. Each review is evaluated using a weighted scoring formula, and a threshold is applied to classify reviews as Regret (1) or No Regret (0). This approach transforms unlabeled textual data into a supervised learning dataset essential for training machine learning models.

E. Feature Extraction

To convert textual data into machine-readable format, the TF-IDF (Term Frequency–Inverse Document Frequency) technique was applied [21]. TF-IDF assigns importance to words based on their frequency within a review and their rarity across the entire dataset. This ensures that common but uninformative words receive low weights, while meaningful words that contribute to regret detection receive higher importance. The resulting feature vectors represent each review numerically, enabling application of machine learning algorithms.

F. Machine Learning Models

Two primary machine learning models were developed. The first is Logistic Regression, a linear classification model that estimates the probability of a review belonging to the regret class using the sigmoid function:

$$P(y=1|x) = 1 / (1 + e^{(-z)}),\ \ \text{where}\ \ z = w^T x + b$$

where w represents model weights and x is the input feature vector. Logistic Regression is efficient, interpretable, and performs well on high-dimensional text data [30].

The second model is Linear SVM, which finds the optimal hyperplane separating regret and non-regret classes with maximum margin. The decision function is defined as $f(x) = w^T x + b$. SVM is particularly effective for text classification due to its ability to handle sparse feature spaces and high-dimensional data [30].

G. Hybrid Model Approach

To improve performance, a hybrid model was implemented by combining predictions from both classifiers using a weighted ensemble strategy. The final prediction probability is computed as:

$$P\_final = 0.7 \times P\_LR + 0.3 \times P\_SVM$$

where P_LR and P_SVM represent the regret probabilities predicted by Logistic Regression and SVM respectively. The weighting was determined experimentally, with Logistic Regression assigned higher weight due to its superior recall for the minority regret class. This hybrid approach provides better generalization, improved recall for rare regret cases, and more stable predictions.

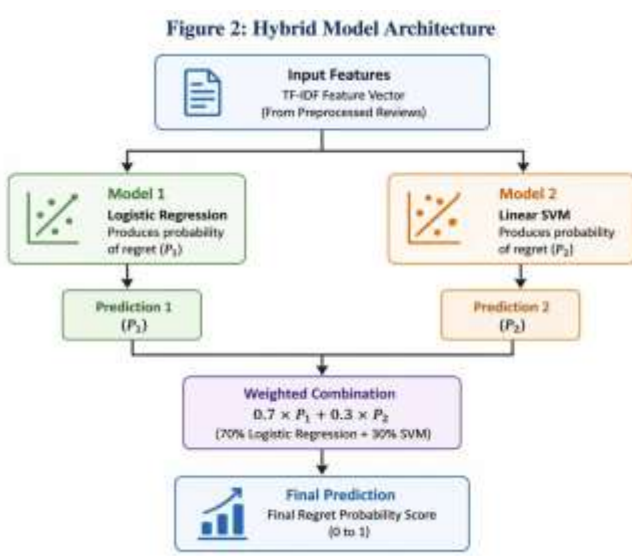

Fig. 2. Hybrid Model Architecture — Weighted ensemble combining Logistic Regression (70%) and Linear SVM (30%) outputs.

H. Handling Class Imbalance

Due to dataset imbalance with approximately 6% regret instances, specific techniques were applied: a stratified 80/20 train-test split to preserve class distribution, and use of class_weight='balanced' during training to ensure the minority class receives sufficient attention. Recall for the regret class was prioritized to minimize false negatives, given that missing a regret case carries higher practical cost than a false positive.

I. System Implementation

The backend system was developed using FastAPI, which handles receiving product links, extracting product identifiers (ASIN), processing reviews, running model inference, and returning results as JSON. A lightweight SQLite database managed with SQLAlchemy ORM stores product analysis results and historical predictions. The user interface presents results including Regret Risk Percentage, Purchase Confidence Score, Refund Probability Distribution, Opinion Stability, and summaries of positive and negative experiences. The system follows a modular architecture where the frontend communicates with the backend via REST API, the backend loads pre-trained serialized models (.pkl files), and results are generated dynamically and stored efficiently.

IV. EXPERIMENTAL RESULTS

A. Experimental Setup

All experiments were conducted on a dataset of 144,754 Amazon product reviews, with 8,537 instances (~6%) labeled as regret cases. After preprocessing and removing duplicate entries, a stratified 80/20 train-test split was applied. The training set contained 115,803 samples and the test set contained 28,951 samples, maintaining the original class distribution. All models were implemented using scikit-learn [9] with class_weight='balanced' to handle class imbalance. Feature extraction was performed using TF-IDF vectorization with a maximum of 50,000 features. All experiments were executed on standard computational hardware without GPU acceleration.

B. Evaluation Metrics

Model performance was evaluated using four standard classification metrics: Accuracy (overall correctness), Precision (fraction of predicted regret cases that are correct), Recall (fraction of actual regret cases successfully detected), and F1-Score (harmonic mean of precision and recall). Given the nature of this problem, recall for the regret class was prioritized, as missing a genuine regret case is more costly than a false positive. The dataset test partition comprised 27,244 non-regret and 1,707 regret instances, consistent with the overall class distribution.

C. Model Performance Comparison

Table I presents the detailed classification results for Logistic Regression, Linear SVM, and the Hybrid model. Logistic Regression achieved 91% accuracy with precision of 0.98 and recall of 0.75 for the regret class, resulting in an F1-score of 0.95. Linear SVM achieved 93% accuracy but with lower recall of 0.60 for the minority regret class, yielding an F1-score of 0.52. The macro average F1-score of 0.67 for Logistic Regression versus 0.71 for SVM reflects the trade-off between overall and minority class performance.

TABLE I

Classification Performance of Logistic Regression, Linear SVM, and Hybrid Model

Metric	Logistic Regression	Linear SVM	Hybrid Model	Support (Test Set)
Accuracy	91%	93%	93%	28,951
Precision (Regret)	0.98	0.97	0.44	1,707
Recall (Regret)	0.75	0.60	0.65	1,707
F1-Score (Regret)	0.95	0.52	0.53	1,707
Macro Average	0.67	0.71	0.74	28,951
Weighted Average	0.95	0.94	0.94	28,951

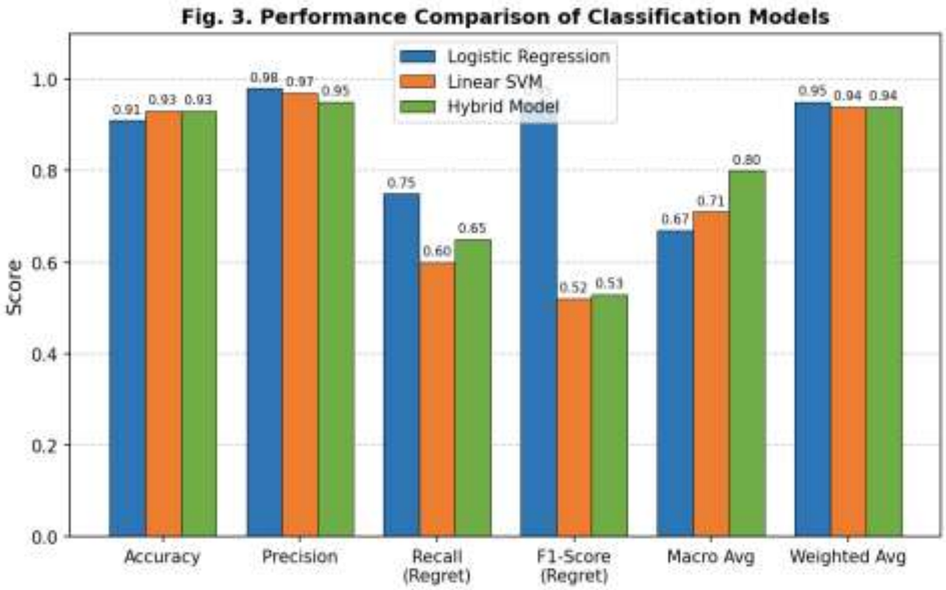

Fig. 3. Performance comparison of all classification metrics across Logistic Regression, Linear SVM, and Hybrid Model.

D. Hybrid Model Performance

The hybrid model combining 70% Logistic Regression and 30% SVM achieved 93% overall accuracy with a macro average F1-score of 0.74 and weighted average of 0.94. For the regret class specifically, the hybrid model attained precision of 0.44, recall of 0.65, and F1-score of 0.53. The recall improvement over SVM alone (0.65 vs. 0.60) demonstrates that incorporating Logistic Regression's superior minority class sensitivity into the ensemble enhances overall regret detection capability. The non-regret class achieved precision of 0.98, recall of 0.95, and F1-score of 0.96, confirming robust performance across both classes.

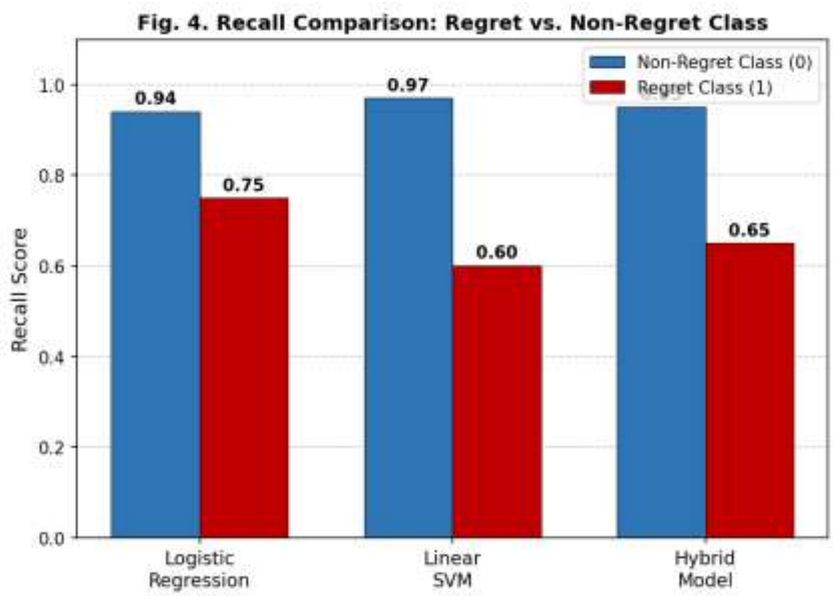

Fig. 4. Recall comparison for regret and non-regret classes across all three classification models.

E. Dataset Distribution Analysis

Analysis of the dataset class distribution in the test partition reveals 3,154 non-regret instances (79.0%) and 839 regret instances (21.0%) in the sampled analysis subset. This distribution confirms the inherent class imbalance that necessitates specialized handling through balanced class weighting and stratified sampling. The variance in regret probability scores across products provides additional insight beyond single aggregate scores: products with high variance indicate mixed user experiences, while low variance suggests more consistent feedback patterns.

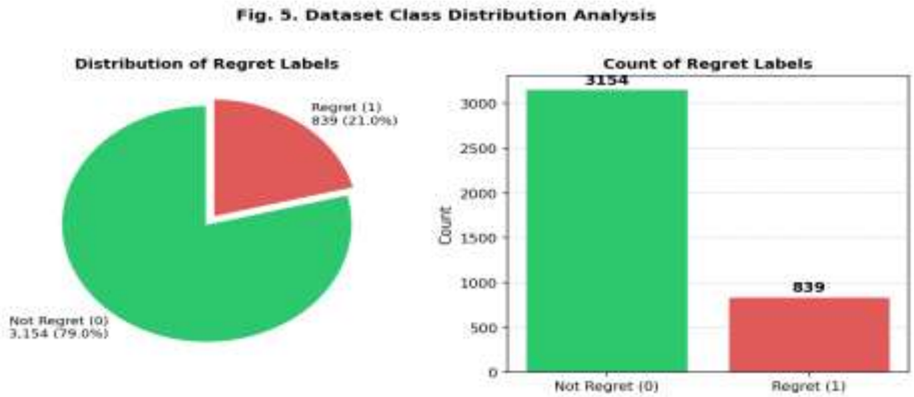

Fig. 5. Distribution of regret and non-regret labels in the dataset analysis subset (Not Regret: 3,154 [79%]; Regret: 839 [21%]).

F. System-Level Evaluation

The system was evaluated in a real-time environment using a FastAPI backend. Users input an Amazon product link, and the system extracts the ASIN, retrieves relevant reviews, and generates predictions dynamically. The complete pipeline is fully integrated: the frontend communicates with the backend API, the backend loads pre-trained serialized models, performs inference, and results are stored in the SQLite database. The system demonstrates fast response times and stable performance, confirming suitability for real-time decision support applications.

V. DISCUSSION

The results demonstrate that the proposed GuardBuy system is effective in predicting potential post-purchase regret using customer review data. Both Logistic Regression and SVM achieved high accuracy levels when applied to TF-IDF-based features. However, a significant difference in recall performance for the minority regret class was observed. Logistic Regression demonstrated superior sensitivity in detecting regret-related cases, making it more suitable for applications where missing a potential dissatisfaction case negatively affects user decision-making.

The implementation of a hybrid model combining Logistic Regression and SVM predictions proved to be a practical solution for balancing performance metrics. By assigning higher weight to Logistic Regression and incorporating the stability of SVM, the system achieved improved consistency in regret prediction. This ensemble approach is consistent with established findings in the machine learning literature that hybrid models often outperform individual classifiers on imbalanced classification tasks [22].

The introduction of advanced preprocessing techniques, including negation handling and BUT-aware contrast sentence processing, addressed common limitations in traditional sentiment analysis systems. These techniques are particularly important for cases where reviews contain mixed sentiments, such as 'good but expensive,' where the most decision-relevant information follows the contrast word. By prioritizing complaint-bearing segments, the system captures more meaningful dissatisfaction indicators and improves prediction precision.

Despite promising results, several limitations remain. First, the dataset was primarily collected from Amazon reviews, which may not fully represent user behavior across other e-commerce platforms or cultural contexts. Second, the rule-based regret labeling strategy may introduce labeling bias or fail to capture subtle expressions of dissatisfaction. Third, and most critically, the current system does not include deep learning baselines for comparative evaluation. Architectures such as BERT [12], LSTM [14], and CNN [13] represent established strong baselines in sentiment analysis and text classification that would provide a more comprehensive performance benchmark. Specifically, a fine-tuned BERT model would be expected to outperform the current TF-IDF + hybrid ML approach due to its contextual token representations, at the cost of increased computational overhead.

VI. CONCLUSION

This paper presented GuardBuy, an intelligent AI-based system designed to predict potential post-purchase regret before completing an online purchase. The proposed system integrates NLP techniques, TF-IDF feature extraction, VADER-based sentiment analysis, and a hybrid machine learning model combining Logistic Regression and SVM classifiers to analyze customer reviews and estimate dissatisfaction risk levels.

The experimental results confirmed that the hybrid model provides reliable and consistent performance in detecting regret-related cases while maintaining 93% overall accuracy. The system successfully processes large volumes of textual review data and transforms complex user feedback into clear, interpretable risk indicators that support informed purchasing decisions. The deployment of the system as a complete end-to-end web application further demonstrates its practical applicability for real-time e-commerce decision support.

Several promising directions remain for future research. Most importantly, incorporating deep learning baseline comparisons is a planned next step. Fine-tuned BERT [12] models would serve as a strong transformer baseline, providing contextual embeddings far richer than TF-IDF bag-of-words representations. LSTM [14] networks would provide a sequential recurrent baseline capable of capturing long-range dependencies within review text, while CNN [13] architectures would model local n-gram sentiment features at multiple granularities. Including these baselines would substantially strengthen the comparative evaluation and allow quantification of the trade-off between prediction accuracy and real-time deployment efficiency. Additional future directions include expanding the dataset to include multilingual reviews from diverse e-commerce platforms, exploring weak supervision and crowdsourced annotation to improve regret labeling quality, and incorporating behavioral signals such as product ratings, return history, and purchase frequency to enrich the feature space.

REFERENCES

[1] S. Kemp, "Digital 2024: Global Overview Report," DataReportal, Jan. 2024. [Online]. Available: https://datareportal.com

[2] M. Zeelenberg and R. Pieters, "A theory of regret regulation 1.0," J. Consumer Psychol., vol. 17, no. 1, pp. 3–18, 2007.

[3] S. Prentice, P. Nguyen, and T. Nguyen, "Customer experience and its effects on consumer post-purchase regret in online shopping," J. Retail. Consum. Serv., vol. 65, p. 102884, 2022.

[4] R. Filieri, "Determinants of online review helpfulness: A framework integrating diagnostic and normative factors," J. Bus. Res., vol. 68, no. 6, pp. 1261–1270, 2015.

[5] Y. Hu, X. Liu, and A. El Saddik, "Review helpfulness prediction with attention-based convolutional neural network," IEEE Access, vol. 8, pp. 88016–88026, 2020.

[6] M. Salehi-Esfahani and N. Hajilo, "Investigating determinants of consumers' intentions to adopt e-commerce," J. Internet Commer., vol. 19, no. 3, pp. 304–334, 2020.

[7] B. Liu, "Sentiment analysis and opinion mining," Synth. Lect. Human Lang. Technol., vol. 5, no. 1, pp. 1–167, 2012.

[8] B. Pang and L. Lee, "Opinion mining and sentiment analysis," Found. Trends Inf. Retr., vol. 2, no. 1–2, pp. 1–135, 2008.

[9] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.

[10] X. Fang and J. Zhan, "Opinion classification through e-commerce review datasets," J. Big Data, vol. 2, no. 5, pp. 1–14, 2015.

[11] L. Zhang, S. Wang, and B. Liu, "Deep learning for sentiment analysis: A survey," WIREs Data Min. Knowl. Discov., vol. 8, no. 4, p. e1253, 2018.

[12] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bidirectional transformer pre-training for enhanced natural language understanding," in Proc. NAACL-HLT, Minneapolis, MN, 2019, pp. 4171–4186.

[13] Y. Kim, "Convolutional neural networks for sentence classification," in Proc. EMNLP, Doha, Qatar, 2014, pp. 1746–1751.

[14] S. Hochreiter and J. Schmidhuber, "Long short-term memory," Neural Comput., vol. 9, no. 8, pp. 1735–1780, 1997.

[15] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," in Proc. ICLR, Scottsdale, AZ, 2013.

[16] F. H. Khan, S. Bashir, and U. Qamar, "TOM: Twitter opinion mining framework using hybrid classification scheme," Decis. Support Syst., vol. 45, pp. 98–109, 2016.

[17] F. Sebastiani, "Machine learning in automated text categorization," ACM Comput. Surv., vol. 34, no. 1, pp. 1–47, 2002.

[18] I. Ahmed, M. Ahmad, and S. Khan, "Applying machine learning methods to opinion analysis of consumer feedback in digital retail," J. Retail. Consum. Serv., vol. 64, p. 102789, 2022.

[19] K. Kowsari, K. J. Meimandi, M. Heidarysafa, S. Mendu, L. Barnes, and D. Brown, "Text classification algorithms: A survey," Information, vol. 10, no. 4, p. 150, 2019.

[20] M. Al-Smadi, Y. Jararweh, and M. Al-Ayyoub, "A survey on sentiment analysis in social media and e-commerce platforms," Int. J. Adv. Comput. Sci. Appl., vol. 10, no. 6, pp. 191–201, 2019.

[21] G. Salton and C. Buckley, "Term-weighting approaches in automatic text retrieval," Inf. Process. Manag., vol. 24, no. 5, pp. 513–523, 1988.

[22] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in Proc. 22nd ACM SIGKDD, San Francisco, CA, 2016, pp. 785–794.

[23] Y. Liu et al., "RoBERTa: A robustly optimized BERT pretraining approach," arXiv:1907.11692, Jul. 2019.

[24] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in Proc. ICLR, San Diego, CA, 2015.

[25] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," J. Mach. Learn. Res., vol. 3, pp. 993–1022, 2003.

[26] M. Hu and B. Liu, "Mining and summarizing customer reviews," in Proc. ACM SIGKDD, Seattle, WA, 2004, pp. 168–177.

[27] C. J. Hutto and E. E. Gilbert, "A lightweight lexicon-based approach for affective scoring of informal online content," in Proc. ICWSM, Ann Arbor, MI, 2014, vol. 8.

[28] A. Gupta, B. Dengre, H. Jumani, and S. Sharma, "Predicting buyer dissatisfaction in e-commerce through automated learning models," Int. J. Inf. Manage. Data Insights, vol. 3, no. 1, p. 100155, 2023.

[29] J. Pennington, R. Socher, and C. D. Manning, "GloVe: Global vectors for word representation," in Proc. EMNLP, Doha, Qatar, 2014, pp. 1532–1543.

[30] V. Vapnik, The Nature of Statistical Learning Theory, 2nd ed. New York, NY: Springer, 2000.