



Prifysgol Cymru
Y Drindod Dewi Sant
University of Wales
Trinity Saint David

Application of Machine Learning

Sentiment Analysis using Traditional ML Algorithms

Submitted To:

Lecturer: Mark Huntly

Module code: ACCA7025

Submitted By:

Name: Shahriar Khan

Student ID: 2507023

Submission Date: 17 April 2026

Contents

1	Abstract	2
2	Introduction	3
3	Related Work in the Literature	4
3.1	Overview of Sentiment Analysis	4
3.2	Naïve Bayes (NB) Classifiers	5
3.3	Support Vector Machines (SVM)	5
3.4	Logistic Regression (LR)	6
3.5	Decision Trees and Random Forests	7
3.6	K-Nearest Neighbors (KNN)	7
3.7	Comparative Analysis of Literature	8
4	Workflow Methodology	9
4.1	Data Collection Understanding	9
4.2	Text Preprocessing	10
4.3	Feature Extraction: TF IDF	11
4.4	Model Selection and Training	11
4.5	Evaluation Metrics	12
4.6	Deployment and Decision-Making	12
4.7	Overall Workflow	13
5	Critical Discussion and Recommendation	14
5.1	Model Comparison	14
5.1.1	Confusion Matrix	15
5.1.2	Model Performance Comparison	15
5.1.3	Accuracy vs Models	15
5.1.4	ROC Comparison	15
5.1.5	SVM Under Computational Cost Constraints	15
5.2	Recommendations	16
5.2.1	Why SVM is the Optimal Choice	17
5.2.2	Why Logistic Regression Performs Well but Not as High as SVM	17

1. Abstract

Sentiment analysis has become a valuable technique for organizations seeking to understand shopper opinions expressed through online platforms and to use these insights to make business decisions. For a marketing firm operating with limited resources, selecting conventional ML approaches that deliver reasonable performance without high computational demands is especially important. This report explores widely used sentiment analysis methods, focusing on traditional algorithms such as Naïve Bayes (NB), Support Vector Machines (SVM), Logistic Regression (LR), K-Nearest Neighbors (KNN), and tree-based decision models. Rather than relying on practical implementation, the discussion is based on a critical evaluation of existing academic literature, in accordance with the assignment requirements. By reviewing prior research and outlining a standard ML workflow, the report seeks to identify the most suitable model for text-based sentiment classification while recommending a straightforward and cost-efficient process that can be easily reproduced. Each technique is compared in terms of accuracy, computational efficiency, scalability, and interpretability. In addition, the report offers practical guidance on structuring the workflow and choosing appropriate evaluation metrics. Through comparative analysis, it ultimately aims to determine the most appropriate non-implemented algorithm for the company described in the assignment, particularly for analyzing shopper review data.

2. Introduction

Although the assignment title refers to image classification, this report focuses on text-based sentiment analysis, as required. In the digital age, organizations increasingly rely on user-generated content—such as shopper reviews, social media posts, and surveys—to understand customer opinions [1]. Sentiment analysis, a key NLP technique, automatically classifies opinions as positive, negative, or neutral, supporting better decision-making and brand monitoring. Manual analysis of large feedback datasets is time-consuming, resource-intensive, and prone to error [2]. As a result, sentiment analysis has emerged as an effective computational approach for extracting meaningful insights from text efficiently [3]. The main aims of this report are:

1. To examine traditional ML algorithms used in sentiment analysis
2. To analyze findings from existing academic research
3. To outline a complete ML workflow suitable for low-resource settings
4. To suggest the most effective model
5. To justify the recommendation based on factors such as accuracy and computational efficiency

The marketing company in question requires a dependable and resource-efficient machine learning approach, avoiding deep learning methods due to their higher computational demands. This report is designed to address that specific requirement

3. Related Work in the Literature

Early research showed that in text categorization tasks, ML techniques perform better than rule-based systems [4] [5].

3.1 Overview of Sentiment Analysis

Sentiment analysis, often referred to as opinion mining, uses computational techniques to detect and classify subjective content within text. This is useful when working with large amounts of unstructured data [6]. With the rapid expansion of e-commerce and social media, the importance of automatic sentiment classification has grown significantly. Conventional ML methods have long been applied in this field because they are efficient and require less computational power. Due to effectiveness and lower resource requirements, traditional ML approaches continue to be frequently employed, despite the rise in popularity of deep learning models. These approaches use methods like Bag-of-Words (BoW) and Term Frequency-Inverse Document Frequency (TF-IDF) to convert textual data into numerical representations [7].




Figure 3.1: **Applications of sentiment analysis**

3.2 Naïve Bayes (NB) Classifiers

NB is a popular choice for text classification because of its simplicity and its ability to handle high-dimensional and sparse data effectively. Although it is based on the assumption that features are independent, which is rarely true in real-world language use—it still delivers surprisingly good results. It performs well on large text datasets, especially when efficiency is important [8]. Similarly, it can achieve solid accuracy in classifying movie reviews using a bag-of-words approach [9]. More recent studies also indicate that Multinomial Naïve Bayes works effectively with TF-IDF features for tasks like analyzing sentiment on Twitter [10]. One of the main strengths of Naïve Bayes is its speed and scalability. It requires very little training time and can still perform reasonably well even with smaller datasets. However, its assumption of feature independence can limit its accuracy, particularly when dealing with more complex language patterns. While it may not always match the performance of more advanced models like SVM, its low computational cost and ease of use make it a practical option for smaller organizations. Overall, it is a fast and lightweight method that requires minimal tuning, but its simplifying assumptions can affect its performance in more nuanced text analysis tasks.

3.3 Support Vector Machines (SVM)

SVMs are often considered one of the most reliable algorithms for sentiment analysis, mainly because they perform well in high-dimensional feature spaces. Early research highlighted the effectiveness of SVM in text classification tasks [11]. Later, it was found that SVM produced better results than methods like NB and KNN when classifying sentiment polarity [12]. A broader review of SVM consistently performs strongly across different sentiment analysis datasets [13]. Further studies have shown that SVM often outperforms Logistic Regression, NB, and KNN, particularly when applied to online data [14]. When combined with TF-IDF features, SVM has achieved high accuracy on datasets such as Sentiment140 [15].

It has also been recommended as one of the most suitable traditional ML models for e-commerce sentiment analysis [16]. In addition, SVM is known to be both efficient and robust, especially when working with small to medium-sized datasets [17]. Overall, research published between 2020 and 2026 continues to show that SVM stands out among traditional ML methods. Its strong performance in terms of accuracy, ability to handle sparse and high-dimensional text data, and reasonable computational requirements make it a practical choice for marketing applications with limited resources [15]. In summary, SVM offers high accuracy in text classification, works well with TF-IDF representations, and provides strong generalization performance.

3.4 Logistic Regression (LR)

LR is a linear model that is commonly used as a baseline for sentiment analysis in text classification tasks. When it combined with L2 regularization, it can achieve performance comparable to SVM across several text classification benchmarks [18]. Another advantage of LR is its interpretability, which makes it especially useful for marketing teams that need to understand and explain the results. Unlike Naïve Bayes, Logistic Regression does not assume that features are independent, allowing it to capture more meaningful relationships within the data. It is also computationally efficient and scales well to large datasets, making it suitable for real-world use. It is a sensible option for practical sentiment analysis applications due to its effectiveness, interpretability, and simplicity of implementation. [19] Overall, LR provides a good balance of accuracy, efficiency, and interpretability, although it may fall slightly behind SVM when dealing with more complex datasets.

3.5 Decision Trees and Random Forests

Decision Trees are not commonly used on their own for text classification tasks, mainly because they tend to overfit the data. However, ensemble approaches like Random Forests (RF) can improve performance by combining multiple trees. RF can produce better results when proper feature selection techniques are applied [20]. Even so, these methods generally do not perform as well as SVM or Logistic Regression when dealing with high-dimensional text data. One advantage of Decision Trees is that they are easy to understand and interpret. However, their effectiveness decreases when handling complex and high-dimensional textual features.

3.6 K-Nearest Neighbors (KNN)

K-Nearest Neighbors (KNN) has been used in text classification in the past, but it does not scale well to large datasets because of its lazy learning approach. The method was originally introduced by Cover and Hart (1967), who highlighted its basic strengths [21]. However, later studies in sentiment analysis have shown that KNN usually performs worse than models like SVM and Naïve Bayes. In practice, KNN is a simple, non-parametric algorithm, but it comes with drawbacks such as high memory consumption and slow prediction times, especially when the dataset becomes large.

3.7 Comparative Analysis of Literature

Studies so far suggest there isn't one algorithm that consistently works best in every situation. How well a method performs tends to vary depending on things like how large the dataset is, how the features are structured, and the limits of available computing resources.

Model	Accuracy	Computational Cost	Notes
Naïve Bayes	Medium	Very Low	Best for very small datasets
Logistic Regression	High	Low	Good baseline; interpretable
SVM	Very High	Medium	Strongest performance on text
Random Forest	Medium	Medium-High	Less effective on sparse text
KNN	Low-Medium	High	Impractical for large datasets

Table 3.1: Summary of Literature Findings

4. Workflow Methodology

This section describes the approach taken to build a sentiment analysis system based on classic ML methods. The goal is to create a model that is both easy to understand and efficient to run, while still being able to correctly sort customer reviews into their respective sentiment groups.



Figure 4.1: Sentiment Analysis workflow

4.1 Data Collection Understanding

The process begins by gathering text data, usually shopper reviews, sourced either from open datasets or internal company records. These collections are generally labelled to reflect different sentiments, such as positive, nega-

tive, or neutral. It's important to examine how the data is organized, how evenly the classes are represented, and whether there is any noise, since issues like imbalance or poor-quality entries can have a strong impact on how well the model performs [22].

4.2 Text Preprocessing

Text data in its raw form is unstructured, so it needs to be cleaned and prepared before it can be used in machine learning models.




Figure 4.2: **Process of Sentiment Analysis**

This stage usually involves a few important steps,

- **Lowercasing:** Converting all text to lowercase to keep it consistent
- **Removal of punctuation:** Eliminates unnecessary symbols
- **Stop-word removal:** Removes common words (e.g., “the”, “is”) that do not contribute to sentiment
- **Tokenisation:** Breaking sentences into individual words or tokens
- **Stemming/Lemmatisation:** Simplifying words to their base forms

Applying these techniques helps cut down unnecessary complexity and noise in the data, which in turn allows the model to perform better on new

inputs. Because of this, careful preprocessing is often seen as a key step in building a reliable sentiment analysis system [23].

4.3 Feature Extraction: TF IDF

TF-IDF, short for Term Frequency–Inverse Document Frequency, is a commonly used method for turning text into numerical form in traditional ML tasks like text classification and sentiment analysis. After the text has been cleaned and prepared, this technique is applied to measure how important each word is within a document compared to the entire dataset. Words that appear frequently in a specific document but not across all documents are given higher importance. This makes it easier for standard ML models to work with text data without relying on more complex approaches [24].




Figure 4.3: **Frequency–Inverse Document Frequency (TF-IDF)**

TF-IDF is preferred over bag-of-words because it:

- Reduces weight of common terms
- Helps the classifier pick up more meaningful differences in the data
- Works extremely well with linear models

4.4 Model Selection and Training

Several approaches, including NB, SVM, and Logistic Regression, are commonly used. These models are trained on labeled datasets so they can identify patterns linked to different sentiment categories. They have shown

good effectiveness in sentiment analysis while still being computationally efficient [23], [11]. After feature extraction, each model is trained on the dataset, with hyperparameters adjusted when needed to improve overall performance.

4.5 Evaluation Metrics

Models are assessed using performance measures such as

- **Accuracy:** overall correctness
- **Precision/Recall:** essential for marketing contexts
- **F1 Score:** balances precision and recall
- **Confusion Matrix:** helps diagnose misclassification trends

These metrics are derived from the confusion matrix and provide a comprehensive evaluation of classification performance, particularly in cases of class imbalance [6].

4.6 Deployment and Decision-Making

The selected model can be integrated into business systems to support decision-making, such as tracking trends in shopper satisfaction. In a real-world setting, it can process incoming shopper reviews and predict sentiment in real time.

4.7 Overall Workflow




Figure 4.4: **Proposed Sentiment Analysis Workflow**

The overall workflow follows a structured pipeline that includes data collection, preprocessing, feature extraction, model training, evaluation, and deployment. This systematic process helps ensure that the sentiment analysis model is both efficient and effective, particularly in resource-constrained environments.




Figure 4.5: **Overall Workflow**

Overall, the figure offers a clear visual summary of how sentiment analysis can be implemented in a marketing context using efficient and well-established ML methods.

5. Critical Discussion and Recommendation

Several widely used classification methods differ in terms of accuracy and computational demand. For instance, NB is efficient and lightweight, but its simplifying assumptions can limit performance in complex sentiment tasks. LR tends to produce more stable predictions, although it may face challenges when sentiment categories are not clearly separable. The KNN approach places most of the computational burden at prediction time, leading to higher memory usage and slower inference, which reduces its suitability for large-scale or real-time applications. In contrast, SVM, particularly with a linear kernel, strike a strong balance between efficiency and predictive performance. They handle high-dimensional text features effectively, require relatively low resources during deployment, and consistently deliver reliable classification results.

A key insight from the literature is that increasing model complexity does not always lead to proportionate improvements in performance. In many real-world applications, simpler models provide sufficient accuracy while offering greater efficiency and reliability.

5.1 Model Comparison

The performance of each classifier is compared using standard evaluation metrics. Model selection is not based solely on accuracy but also considers the balance between precision and recall, depending on the application needs. In marketing analytics, for example, correctly detecting negative sentiment can be more important, as it helps organizations respond effectively to shopper dissatisfaction. The final model is selected based on its ability to deliver consistent and reliable performance while maintaining

computational efficiency.

5.1.1 Confusion Matrix

The confusion matrix is an important evaluation tool for classification models. It compares actual and predicted labels using true/false positives and negatives and is used to compute metrics such as accuracy, precision, recall, and F1-score for sentiment analysis [25], [26]. Unlike a single metric, it provides deeper insight into prediction errors and is especially useful for imbalanced datasets [27].

5.1.2 Model Performance Comparison

Model comparison evaluate different classifiers using metrics like accuracy and F1-score. In sentiment analysis, NB, Logistic Regression, and SVM are commonly compared to select the best-performing model [27]. Research also shows that multiple metrics provide a more reliable evaluation than accuracy alone, particularly for imbalanced datasets [26], [28].

5.1.3 Accuracy vs Models

SVM often outperform other conventional classifiers in text-based sentiment analysis tasks [9].

5.1.4 ROC Comparison

SVM exhibits better class separability than Logistic Regression and Naïve Bayes, consistent with the results in [29].

5.1.5 SVM Under Computational Cost Constraints

Linear SVM offers a strong balance between efficiency and performance, satisfying low resource requirements. Using TF-IDF features, sentiment data is often approximately linearly separable, enabling SVM to learn an effective boundary without expensive kernel computations, which keeps training and prediction efficient for large-scale and real-time use [29], [30].

Studies also show that linear SVM typically achieves better accuracy and F1-score than Naïve Bayes, k-NN, and Random Forest while maintaining low inference cost [31] – [30]. Although Naïve Bayes is computationally faster, its lower accuracy can lead to more misclassifications, which is not ideal for marketing applications. Overall, linear SVM provides the best balance between performance and computational efficiency for sentiment analysis.

Algorithm	Advantages	Disadvantage	Suitability
SVM (Linear)	High accuracy, robust to high-dimensional data	Needs tuning, not ideal for very large datasets	Best choice overall
Logistic Regression	Interpretable, fast, strong baseline	Slightly less accurate than SVM	Good baseline model
Naïve Bayes	Very fast, handles small data well	Lower accuracy, independence assumption	Good for simple tasks
Random Forest	Handles non-linear data	Poor with sparse text vectors	Not ideal
KNN	Simple	Slow prediction, memory heavy	Not recommended

Table 5.1: Summary of Literature Findings

5.2 Recommendations

LR can be an appropriate option for sentiment analysis due to its balance of performance and efficiency, while Naïve Bayes is often used as a fast baseline approach. From a business standpoint, model selection should account not only for predictive performance but also for practical factors such as processing speed, scalability, and ease of deployment. A linear SVM combined with TF-IDF features is the most appropriate choice for this use case, as it offers scalability, efficiency, and strong accuracy while meeting the constraints of limited resources. LR is recommended as a secondary option where interpretability is required.

5.2.1 Why SVM is the Optimal Choice

1. High Accuracy Across Benchmarks:

SVM repeatedly outperforms NB, KNN, RF, and LR for sentiment analysis in major studies [9], [12], [13].

2. Works Extremely Well with TF-IDF:

Linear SVM and TF-IDF are de facto standards in classical NLP pipelines.

3. Efficiency:

SVM trains quickly on text even with modest hardware. This suits the marketing company's computational constraints.

4. Strong Generalization Ability:

SVM manages sparse text vector spaces better than tree-based or distance-based models.

5. Minimal Hyperparameter Tuning

Only "C" (regularization) typically needs tuning.

SVMs have been widely reported as one of the most effective models for customer sentiment analysis, often outperforming other traditional ML techniques in terms of accuracy and F-measure [32].

5.2.2 Why Logistic Regression Performs Well but Not as High as SVM

LR is frequently the second-best model:

- Performs strongly on linearly separable sparse text data.
- Has straightforward interpretability.
- But lacks the robust margin optimization of SVM, causing slight drops in accuracy and F1.

Based on the reviewed literature and system requirements, SVM with a linear kernel emerges as the most suitable candidate model for the client. A linear SVM combined with TF-IDF features is recommended. This approach offers high accuracy [29], efficient computation, strong generalization, and practical scalability.

Conclusion

Based on extensive evidence from the reviewed literature, SVMs emerge as the most effective traditional ML model for sentiment analysis, particularly in resource-constrained environments. SVMs demonstrate consistently high accuracy when applied to high-dimensional and sparse text data produced by feature extraction techniques such as TF-IDF. Compared to Naïve Bayes, SVM offers superior classification boundaries and improved generalization while avoiding the unrealistic independence assumption between features. Although LR provides strong interpretability and competitive performance, it typically underperforms SVM due to the lack of margin maximization. Tree-based methods and instance-based learners such as Random Forest and KNN are less suitable due to overfitting, scalability issues, and inefficiency with sparse text vectors. Therefore, when balancing accuracy, computational efficiency, scalability, and implementation simplicity, a linear SVM combined with TF-IDF feature representation represents the most appropriate and academically justified choice for sentiment analysis using traditional ML techniques.

Future work may focus on exploring hybrid approaches, such as combining multiple traditional models or incorporating incremental learning techniques to adapt to evolving data patterns. Such enhancements could further improve model performance and ensure long-term effectiveness in dynamic business environments.

Bibliography

- [1] A. Maroof, S. Wasi, S. I. Jami, and M. S. Siddiqui, “Aspect-based sentiment analysis for service industry,” *IEEE Access*, vol. 12, pp. 109702–109713, 2024.
- [2] B. Rahman and Maryani, “Optimizing customer satisfaction through sentiment analysis: A bert-based machine learning approach to extract insights,” *IEEE Access*, vol. 12, pp. 151476–151489, 2024.
- [3] A. Qazi, N. Hasan, R. Mao, M. Elhag Mohamed Abo, S. Kumar Dey, and G. Hardaker, “Machine learning-based opinion spam detection: A systematic literature review,” *IEEE Access*, vol. 12, pp. 143485–143499, 2024.
- [4] K. Aurangzeb, N. Ayub, and M. Alhussein, “Aspect based multi-labeling using svm based ensembler,” *IEEE Access*, vol. 9, pp. 26026–26040, 2021.
- [5] Y. Dong, X. Yang, Z. Deng, T. Pang, Z. Xiao, H. Su, and J. Zhu, “Black-box detection of backdoor attacks with limited information and data,” in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 16462–16471, 2021.
- [6] B. Pang and L. Lee, *Opinion mining and sentiment analysis*. Now Publishers Inc, 2008.
- [7] H. Liu, X. Chen, and X. Liu, “A study of the application of weight distributing method combining sentiment dictionary and tf-idf for text sentiment analysis,” *IEEE Access*, vol. 10, pp. 32280–32289, 2022.

- [8] A. McCallum, “A comparison of event models for naive bayes text classification,” 1998.
- [9] B. Pang, L. Lee, and S. Vaithyanathan, “Thumbs up? sentiment classification using machine learning techniques,” in *Proceedings of the 2002 conference on empirical methods in natural language processing (EMNLP 2002)*, pp. 79–86, 2002.
- [10] M. Khan and A. Srivastava, “Sentiment analysis of twitter data using machine learning techniques,” *International Journal of Engineering and Management Research*, vol. 14, no. 1, pp. 196–203, 2024.
- [11] T. Joachims, “Text categorization with support vector machines: Learning with many relevant features,” in *European conference on machine learning*, pp. 137–142, Springer, 1998.
- [12] T. Mullen and N. Collier, “Sentiment analysis using support vector machines with diverse information sources,” in *Proceedings of the 2004 conference on empirical methods in natural language processing*, pp. 412–418, 2004.
- [13] K. Ravi and V. Ravi, “A survey on opinion mining and sentiment analysis: tasks, approaches and applications,” *Knowledge-based systems*, vol. 89, pp. 14–46, 2015.
- [14] X. He, “Sentiment classification of social media user comments using svm models,” in *2024 5th International Seminar on Artificial Intelligence, Networking and Information Technology (AINIT)*, pp. 1755–1759, IEEE, 2024.
- [15] K. Matondkar, “Sentiment analysis on social media using support vector machines (svm),” *International Journal of Trend in Scientific Research and Development*, vol. 9, pp. 704–708, Apr. 2025.

- [16] E. Park, R. Dave, and M. Bhavsar, “Comparative research in sentiment analysis using machine learning technique,” *J. Data Anal. Inf. Process*, vol. 13, no. 3, pp. 269–280, 2025.
- [17] S. Cao, Z. Liu, Z. Wang, and H. Tang, “Text sentiment analysis based on word2vec: A comparative study of svm and deep learning models,” in *International Conference on Computational Linguistics and Natural Language Processing*, pp. 266–272, Springer, 2025.
- [18] A. Genkin, D. D. Lewis, and D. Madigan, “Large-scale bayesian logistic regression for text categorization,” *technometrics*, vol. 49, no. 3, pp. 291–304, 2007.
- [19] F. Shadabi and D. Sharma, “Comparison of artificial neural networks with logistic regression in prediction of kidney transplant outcomes,” in *2009 International Conference on Future Computer and Communication*, pp. 543–547, 2009.
- [20] H. Kim, Y. Jo, and S. Lee, “Sentiment classification using random forests,” *Information Sciences*, vol. 502, pp. 324–337, 2019.
- [21] V. Vapnik, *The nature of statistical learning theory*. Springer science & business media, 2013.
- [22] B. Liu, “Sentiment analysis and opinion mining,” 2012.
- [23] F. Sebastiani, “Machine learning in automated text categorization,” *ACM computing surveys (CSUR)*, vol. 34, no. 1, pp. 1–47, 2002.
- [24] C. D. Manning, *Introduction to information retrieval*. Syngress Publishing,, 2008.
- [25] H. M. Omran, K. Ibrahim, G. T. Abdel-Jaber, and A.-N. Sharkawy, “Brain tumor classification from mri images using hybrid deep learning approaches: Vgg19 with softmax and svm classifiers,” *International Journal of Robotics and Control Systems*, vol. 6, no. 1, pp. 16–35, 2026.

- [26] O. Rainio, J. Teuho, and R. Klén, “Evaluation metrics and statistical tests for machine learning,” *Scientific reports*, vol. 14, no. 1, p. 6086, 2024.
- [27] S. K. Singh *et al.*, “Sentiment analysis: A survey on design framework and applications,” 2022.
- [28] J. C. Obi, “A comparative study of classification metrics and their performances,” *World Journal of Advanced Engineering Technology and Sciences*, 2023.
- [29] M. Danyal, S. Khan, M. Khan, M. Ghaffar, B. Khan, and M. Arshad, “Sentiment analysis based on performance of linear support vector machine and multinomial naïve bayes using movie reviews with baseline techniques,” *Journal on Big Data*, vol. 5, p. 1, 2023.
- [30] K. Suresh Kumar, A. Radha Mani, T. Ananth Kumar, A. Jalili, M. Gheisari, Y. Malik, H.-C. Chen, and A. Jahangir Moshayedi, “Sentiment analysis of short texts using svms and vsms-based multiclass semantic classification,” *Applied Artificial Intelligence*, vol. 38, no. 1, p. 2321555, 2024.
- [31] T. A. Khan, R. Sadiq, Z. Shahid, M. M. Alam, and M. B. M. Su’ud, “Sentiment analysis using support vector machine and random forest,” *Journal of Informatics and Web Engineering*, vol. 3, no. 1, pp. 67–75, 2024.
- [32] L. Ashbaugh and Y. Zhang, “A comparative study of sentiment analysis on customer reviews using machine learning and deep learning,” *Computers*, vol. 13, no. 12, p. 340, 2024.

Appendix

Appendix A: Conceptual Workflow Summary

This appendix presents a brief overview of the sentiment analysis workflow proposed in the main report. The workflow reflects a traditional ML approach designed for deployment in resource-constrained organizational environments. The process begins with the collection of shopper reviews data, followed by essential text preprocessing steps such as tokenization, lowercasing, and stopword removal to reduce noise. The cleaned text is then transformed into numerical features using the Term Frequency–Inverse Document Frequency (TF-IDF) method. A linear SVM is trained using these features, and model performance is evaluated using precision, recall, and F1 score. This workflow prioritizes efficiency, scalability, and ease of maintenance while maintaining effective sentiment classification performance.