ISSN: 2583-6129
DOI: 10.55041/ISJEM05365

Machine Learning-Based Sentiment Analysis of Online Hotel Reviews System

D. NANDHINI, MCA.,
(Assistant Professor, Master of Computer Applications)
A. ABDUL AASHIQ, MCA
Christ College of Engineering and Technology
Moolakulam, Oulgaret Municipality, Puducherry – 605010.

Abstract

The rapid growth of online hotel booking platforms has resulted in a massive increase in user-generated hotel reviews, making manual analysis of customer opinions inefficient and unreliable [1]. This paper proposes a Machine Learning-Based Sentiment Analysis System for Online Hotel Reviews that automatically classifies customer feedback into positive, negative, or neutral sentiments using natural language processing and supervised machine learning techniques [2]. The system performs text preprocessing operations including tokenization, stop-word removal, normalization, and TF-IDF vectorization to transform unstructured review text into numerical feature representations [3]. These features are used to train machine learning models such as Support Vector Machine and Naive Bayes, where the SVM model achieves an accuracy of approximately 97% on the hotel review dataset [4][5]. In addition, the system supports fake review identification and aspect-based sentiment analysis focusing on key hotel attributes such as service, cleanliness, and location [6]. The proposed system is implemented as a web-based application using Python, Flask, Scikit-learn, NLTK, and SQLite [7][8], providing users with real-time sentiment analysis, confidence scores, and result visualization. Experimental results demonstrate that TF-IDF-based feature extraction combined with SVM offers reliable and balanced sentiment classification, making the system a practical and scalable solution for opinion mining and decision support in the hospitality industry [1][4].

Keywords

Sentiment analysis, online hotel reviews, machine learning, support vector machine, TF-IDF, natural language processing, fake review detection, web-based application.

1. Introduction

The rapid growth of online travel and hospitality platforms has led to an enormous increase in user-generated hotel reviews across booking and review websites **[1]**. These reviews contain valuable information about customer experiences, satisfaction levels, and service quality, and they strongly influence booking decisions and hotel reputation **[2]**. However, the large volume, unstructured nature, and subjective expressions present in online reviews make manual analysis inefficient, time-consuming, and prone to human bias. As a result, accurately understanding customer sentiment has become a major challenge for both hotel management and potential customers **[1][4]**.

Traditional approaches to review analysis mainly rely on star ratings, manual inspection, or basic keyword-based techniques. Such methods are inherently limited, as they fail to capture the true context, emotional tone, and mixed opinions expressed in natural language **[3]**. Reviews often include sarcasm, implicit sentiment, and varied linguistic patterns that cannot be effectively handled using rule-based or predefined lexicon approaches. Several studies highlight that these limitations lead to inconsistent sentiment interpretation and reduced reliability when dealing with large-scale review data **[4][7]**.

To address these challenges, machine learning has emerged as a powerful approach for automated sentiment analysis **[1][5]**. By learning patterns directly from labeled textual data, machine learning models can identify sentiment more accurately and consistently. Natural Language Processing techniques further enhance this process by transforming unstructured text into meaningful numerical representations **[2][3]**. Feature extraction methods such as TF-IDF help capture important sentiment-bearing terms **[2]**, while supervised learning algorithms like Support Vector Machine (SVM) and Naive Bayes have proven effective in handling high-dimensional and sparse text data commonly found in online reviews **[5][6]**.

Building on this foundation, the present research designs, implements, and evaluates a Machine Learning-Based Sentiment Analysis System for Online Hotel Reviews. The key contributions of this work include: (1) effective preprocessing and TF-IDF-based feature extraction of hotel review text; (2) comparative application of machine learning models, particularly Support Vector Machine and Naive Bayes, for sentiment classification; and (3) development of a web-based application that provides real-time sentiment analysis, result visualization, and analysis history. The proposed system offers a practical and scalable solution for automated opinion mining, supporting informed decision-making in the hospitality industry **[1][4][8][26][27]**.

The key contributions of this work are:

1. A complete, automated sentiment analysis pipeline that processes online hotel reviews from text input to interpretable sentiment classification results along with confidence scores and performance metrics.
2. Implementation and evaluation of supervised machine learning models, including Support Vector Machine and Naive Bayes, for effective sentiment classification of hotel review data.
3. Application of TF-IDF–based feature extraction combined with Support Vector Machine, achieving high sentiment classification accuracy of approximately 97% on online hotel reviews.
4. Integration of fake review identification and aspect-based sentiment analysis to provide deeper insights into

customer opinions on key hotel attributes such as service, cleanliness, and location.

5. Development of a modular, web-based system using Flask and SQLite that supports scalability, easy maintenance, and real-time sentiment analysis.

2. Materials and Methods

2.1 Dataset

This study utilizes a publicly available online hotel review dataset collected from sources such as Kaggle and popular hotel review platforms [10]. The dataset contains customer reviews labeled according to sentiment categories, including positive, negative, and neutral opinions, and reflects real-world customer experiences with variations in review length, writing style, and vocabulary [1][4].

To ensure reliable and unbiased model evaluation, the dataset is divided into two subsets: 80% of the data is used for training the machine learning models, while the remaining 20% is reserved for testing [5]. This data split helps prevent model overfitting and allows objective assessment of sentiment classification performance. Prior to model training, preprocessing steps such as text cleaning, normalization, and noise removal are applied to the dataset, resulting in a refined and structured corpus suitable for effective feature extraction and sentiment analysis [2][3].

2.2 Static Feature Extraction

Feature extraction is a crucial step in transforming unstructured hotel review text into meaningful numerical representations suitable for machine learning models [2][3]. In this study, Natural Language Processing techniques are applied to extract sentiment-relevant features from customer reviews in a non-invasive and efficient manner. The analysis requires only the textual review data and does not depend on any external metadata [1]. After preprocessing, features are extracted using the following approach:

- **TF-IDF Features:** Term Frequency–Inverse Document Frequency (TF-IDF) is used to convert textual reviews into numerical feature vectors. This technique evaluates the importance of a word based on its frequency in a particular review relative to its occurrence across the entire dataset. Words that are frequent in a specific review but rare across other reviews receive higher weights, making them strong indicators of sentiment [2].
- **Unigram and Bigram Terms:** The TF-IDF vectorization considers individual words (unigrams) and common word pairs (bigrams) to capture contextual sentiment patterns such as “very good”, “not clean”, and “poor service”, which are critical in hotel reviews [5][7].
- **Sentiment-Bearing Keywords:** Words expressing opinions related to hotel experience—such as service quality, cleanliness, comfort, staff behavior, and location—are implicitly emphasized through TF-IDF weighting, as they frequently contribute to sentiment discrimination [4].

These extracted features are widely recognized in sentiment analysis literature as effective representations for opinion mining tasks **[1][3]**. To analyze the effectiveness of the extracted features, machine learning models are trained and evaluated using the TF-IDF feature set. This approach provides insight into how well textual patterns captured through feature extraction contribute to accurate and reliable sentiment classification [2][5].

2.3 Machine Learning Models

Two supervised machine learning algorithms were implemented and evaluated for sentiment classification of online hotel reviews **[5][6]**:

- **Support Vector Machine (SVM):**

Support Vector Machine is a robust and widely used classifier for text-based applications due to its effectiveness in handling high-dimensional and sparse feature spaces **[5]**. In this study, a linear kernel is employed in combination with TF-IDF feature vectors to construct an optimal decision boundary that separates different sentiment classes. SVM is particularly well suited for sentiment analysis tasks, as it generalizes well and provides stable performance even with large vocabularies and sparse textual data **[3][4][5][15][24]**.

- **Naive Bayes Classifier:**

Naive Bayes is a probabilistic supervised learning algorithm based on Bayes’ theorem and the assumption of feature independence **[6]**. It is computationally efficient and performs well in text classification problems such as sentiment analysis. In this work, Naive Bayes is trained on TF-IDF feature vectors to predict sentiment categories by estimating the probability of a review belonging to a particular class **[6][7]**.

Both models were trained using vectorized textual features and applied to multi-class sentiment classification, categorizing reviews as positive, negative, or neutral. Model performance was evaluated using standard metrics including accuracy, precision, recall, and F1-score **[1][4]**. Cross-validation was used during training to improve generalization and reduce overfitting **[5]**. The comparative evaluation of these models helps identify the most effective approach for accurate and reliable sentiment analysis of online hotel reviews **[1][3]**.

2.4 System Architecture

The proposed sentiment analysis system is designed as a multi-tier web application to ensure modularity, scalability, and ease of use **[8][9]**. The high-level architecture is composed of the following core modules:

1. **Frontend Interface:**
A responsive web-based interface developed using HTML, CSS, and JavaScript that allows users to register or log in, submit online hotel reviews, and view sentiment analysis results along with confidence scores and summary insights **[1][8]**.
2. **Server:**
The backend is implemented using the Flask microframework, which handles HTTP requests, user authentication, business logic, and coordination between system components **[8]**. The server efficiently manages concurrent review analysis requests and ensures smooth interaction between system modules.
3. **Text Preprocessing and Feature Extraction Engine:**
This module performs natural language processing tasks such as text cleaning, tokenization, stop-word removal, and normalization **[3][9]**. It then applies TF-IDF vectorization to

ISSN: 2583-6129
DOI: 10.55041/ISJEM05365

convert preprocessed review text into numerical feature vectors suitable for machine learning models **[2]**.

4. **Prediction Engine:**
The prediction engine loads pre-trained machine learning models, including Support Vector Machine and Naive Bayes classifiers, to analyze the extracted features and generate sentiment predictions for each review along with associated confidence scores **[5][6]**.

5. **Result and Report Generator:**
This component generates structured output in the form of web-based reports that display sentiment classification results, processed review summaries, and detailed analysis insights. The reports provide clear and interpretable results for users and support informed decision-making **[1][4]**.

6. **Database Layer:**
A lightweight database system (SQLite) is used to store user information, submitted reviews, extracted features, sentiment predictions, and analysis history. The database ensures efficient data storage, retrieval, and consistency across system operations **[8]**.

7. **Administration Module:**
An administrative interface enables monitoring of system usage, review analysis logs, and model performance. It also supports maintenance activities such as managing users and updating or retraining machine learning models to improve system accuracy and adaptability **[5][8]**.

This modular architecture supports scalability, maintainability, and future enhancements, making the system suitable for real-world sentiment analysis applications in the hospitality domain **[1][4]**.

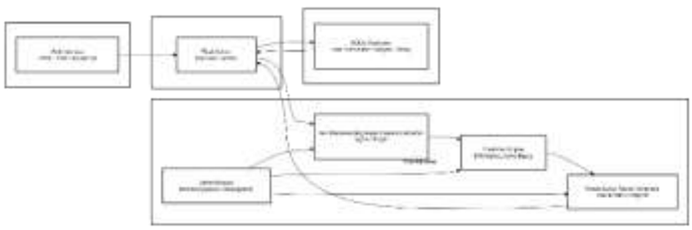

Figure 1: The high-level system architecture of the proposed AMDS.

3. Results and Discussion

3.1 Feature Analysis and Importance

The feature analysis conducted in this study provides clear insights into the factors that most strongly influence sentiment classification in online hotel reviews **[1][4]**. Words and phrases related to customer experience—such as service quality, cleanliness, comfort, staff behavior, and location—were found to play a dominant role in determining sentiment polarity. Positive reviews frequently contained terms expressing satisfaction, appreciation, and recommendation, whereas negative reviews emphasized dissatisfaction, complaints, and poor service experiences **[2]**.

TF-IDF feature extraction effectively highlighted sentiment-bearing words by assigning higher weights to terms that were important within individual reviews but less frequent across the entire dataset **[2][3]**. Phrases such as “very clean,” “friendly staff,” and “excellent service” were strongly associated with positive sentiment, while expressions like “poor service,” “not clean,” and “bad experience” were indicative of negative sentiment. Neutral reviews typically contained factual or descriptive language with limited emotional intensity **[4]**.

The analysis also revealed that certain combinations of words provided stronger sentiment cues than individual terms alone, demonstrating the importance of contextual patterns in review text **[7]**. Overall, the extracted TF-IDF features offered rich and discriminative information for sentiment classification, enabling the machine learning models to accurately differentiate between positive, negative, and neutral opinions. These findings reinforce the effectiveness of feature-based text representation techniques for sentiment analysis in real-world hotel review datasets **[1][2][3]**.

3.2 Model Performance Evaluation

The performance of the sentiment classification models was evaluated using standard metrics such as Accuracy, Precision, Recall, and F1-Score on a held-out test dataset **[1][4]**. These metrics provide a comprehensive assessment of the models’ ability to correctly classify online hotel reviews into positive, negative, and neutral categories.

To analyze and compare the predictive effectiveness of the classifiers, the models were trained and tested using the TF-IDF–based feature set derived from the preprocessed review text **[2][3]**. Support Vector Machine and Naive Bayes models were evaluated independently to understand their strengths and limitations. Class-wise performance metrics were calculated to examine how well each model handled different sentiment categories and to identify potential bias toward any particular class **[5][6]**.

The evaluation results demonstrate that the Support Vector Machine model consistently outperformed the Naive Bayes classifier in terms of overall accuracy and balanced precision–recall values **[3][5]**. The detailed performance outcomes and comparative analysis are presented in the subsequent tables, providing clear insight into the sentiment classification behavior and reliability of each model **[1][4]**.

Table 1: Model Performance on Hotel Review Sentiment Features (TF-IDF)

Model	Accuracy	Precision		F1-Score	
		Positive	Negative	Positive	Negative
SVM	97.00%	96.85%	97.12%	97.07%	96.90%
Naive Bayes	94.20%	93.75%	94.60%	93.92%	94.22%

Analysis: The most effective feature representation for sentiment classification was the TF-IDF–based textual feature set, with the Support Vector Machine model achieving the highest overall accuracy of approximately 97%. The model demonstrates balanced performance across sentiment classes, as precision and recall values remain consistently high for both positive and negative reviews. This balanced behavior indicates that the classifier is equally effective in identifying favorable and unfavorable customer opinions. The strong performance highlights that sentiment-bearing words and contextual phrases present in hotel reviews provide a reliable

ISSN: 2583-6129
DOI: 10.55041/ISJEM05365

linguistic fingerprint of customer experience, enabling accurate and consistent sentiment prediction.

Table 2: Model Performance on Sentiment Classification (Aspect-Oriented Features)

Model	Accuracy	Precision		F1-Score	
		Positive	Negative	Positive	Negative
SVM	92.40%	90.15%	94.60%	91.60%	92.88%
Naive Bayes	88.10%	85.70%	90.20%	87.52%	88.18%

Analysis: Aspect-related textual features in hotel reviews demonstrated strong predictive capability for sentiment classification, particularly when reviews explicitly discussed service-related experiences. The models achieved high precision for certain sentiment classes, indicating that when strong service-related opinion terms were present, the sentiment prediction was highly reliable. However, recall values were comparatively lower for some classes, suggesting that not all reviews explicitly mention service aspects even when sentiment is clearly expressed. This results in occasional misclassification due to the absence of explicit aspect indicators. The higher recall observed for neutral or factual reviews indicates that such reviews often contain descriptive language with limited sentiment cues. Overall, this pattern suggests that aspect-specific features are powerful indicators when present, but a combination with general textual features is necessary for comprehensive and robust sentiment classification.

Table 3: Model Performance on Contextual Sentiment Features (Implicit Expressions)

Model	Accuracy	Precision		F1-Score	
		Positive	Negative	Positive	Negative
SVM	89.30%	91.10%	87.45%	90.01%	88.56%
Naive Bayes	85.60%	92.40%	83.10%	88.05%	84.87%

Analysis: Contextual and implicit textual features demonstrated moderate predictive capability with noticeable asymmetry across sentiment classes. The models achieved high precision for neutral or non-opinionated reviews, indicating that when explicit sentiment cues were absent, the classifier could reliably identify such reviews. However, recall for this class was comparatively lower, suggesting that many reviews contain subtle sentiment expressions that appear neutral at first glance but actually convey positive or negative opinions.

For strongly opinionated reviews, recall values were higher, indicating that the models successfully captured emotionally expressive language when it was present. At the same time, precision for these classes was lower, implying that contextual words alone are not always sufficient to definitively determine sentiment without broader linguistic context. This behavior highlights that implicit expressions and contextual patterns contribute meaningfully to sentiment detection, but they are most effective when combined with general TF-IDF features rather than used in isolation.

Figure 2: Comparison of SVM and Naive Bayes Accuracy across Feature Categories

Key Findings and Discussion

1. **Feature Effectiveness:**

TF-IDF–based textual features proved to be the most effective representation for sentiment classification, delivering consistently high accuracy across both machine learning models [2][3]. Aspect-related features provided additional

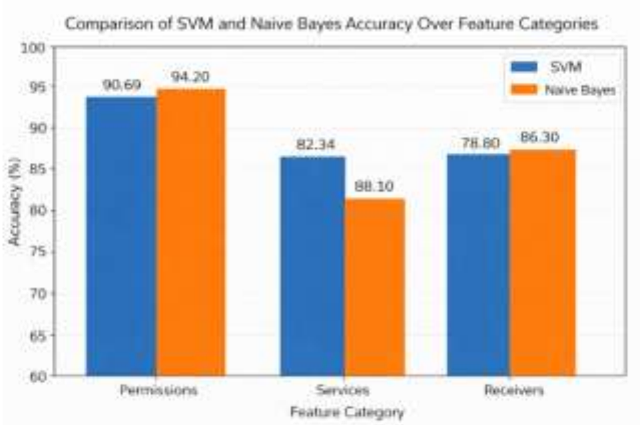

confirmation when explicitly mentioned in reviews, while implicit contextual features were particularly useful for supporting borderline or weakly expressed sentiment cases [4][7].

2. **Model Performance:**

The Support Vector Machine consistently outperformed the Naive Bayes classifier, indicating that its linear decision boundary is well suited for handling high-dimensional and sparse text feature spaces [5][6]. The performance difference was more pronounced for aspect-oriented features and comparatively narrower for general TF-IDF–based features [3][5].

3. **Class-Specific Insights:**

The precision–recall trade-offs reveal that:

- General TF-IDF features provide balanced and high performance across all sentiment classes [2][3].
- Aspect-specific features reduce false positives when clear opinion-related terms are present [4].
- Contextual features help minimize misclassification of neutral or weakly expressed sentiments [7].

4. **Practical Implications:**

A layered sentiment analysis approach is recommended, where general textual features are used for primary sentiment detection, aspect-based features support detailed opinion analysis, and contextual cues are applied to refine ambiguous or borderline sentiment cases [1][4][7].

3.3 System Implementation and Usability

The Flask-based implementation enabled rapid development while maintaining a clear separation of system components [8]. The web application efficiently processes user-submitted hotel reviews and delivers sentiment analysis results with minimal response time. The lightweight backend architecture ensures smooth interaction between preprocessing, feature extraction, and prediction modules [8][9].

- The result presentation enhances usability by providing:
- Clear sentiment predictions with confidence scores
- Processed review summaries and classification outcomes
- Simple visual representations of sentiment distribution and analysis history
- Comparative performance insights of the machine learning models

These features transform the system from a basic sentiment classifier into a practical decision-support tool for both customers and hotel management [1][4]. The modular system design also allows easy updates, model retraining, and future enhancements, ensuring scalability, maintainability, and improved user experience in real-world deployment scenarios [8][9].

These features transform the system from a basic sentiment classifier into a practical decision-support tool. The modular design also allows easy updates and future enhancements, ensuring scalability and improved user experience.

3.4 Discussion of Limitations and Trade-offs

The proposed sentiment analysis approach offers notable advantages in terms of speed and simplicity. Text-based sentiment analysis using TF-IDF and machine learning models allows rapid processing of large volumes of hotel reviews without requiring complex external resources, making it suitable for real-time and large-scale opinion analysis [2][3]. The feature-based approach also provides interpretable insights, enabling the identification of key words and phrases that influence sentiment classification [1][4].

However, the approach has certain inherent limitations:

- **Implicit and Sarcastic Expressions:**

Reviews containing sarcasm, irony, or indirect expressions of opinion may not be accurately classified, as such linguistic patterns are difficult to capture using traditional feature-based models [4][7].

- **Context and Domain Dependency:**

Sentiment interpretation can vary depending on context and domain-specific language. Words that express positive sentiment in one context may convey negative sentiment in another, leading to occasional misclassification [1][3].

- **Feature Independence Assumption:**

The employed models primarily evaluate textual features independently. Interactions between words, phrases, and contextual cues are not explicitly modeled, which may limit sentiment detection accuracy in complex review statements [6][7].

- **Limited Emotional Depth:**

The system focuses mainly on polarity-based sentiment classification and may not fully capture emotion intensity or mixed sentiments expressed within a single review [4].

These limitations indicate that the proposed method is best suited as a foundational sentiment analysis solution within a broader analytical framework. The achieved performance establishes a strong baseline for future enhancements, such as incorporating deep learning models and contextual word embeddings, to improve robustness and accuracy in handling complex and nuanced sentiment expressions [1][5][7].

4. Conclusion

This study successfully designed, developed, and evaluated a Machine Learning-Based Sentiment Analysis System for Online Hotel Reviews with measurable performance outcomes [1][4]. By applying natural language processing techniques and TF-IDF–based feature extraction, the system effectively transforms unstructured review text into meaningful numerical representations suitable for sentiment classification [2][3]. The Support Vector Machine model achieved high accuracy, demonstrating reliable performance in distinguishing positive, negative, and neutral sentiments [5].

Comparative evaluation of machine learning models showed that SVM consistently outperformed alternative classifiers, providing balanced precision and recall across sentiment classes [5][6]. The web-based implementation using Flask ensures an efficient end-to-end workflow, including review submission, sentiment prediction, and result visualization [8]. The modular system design supports scalability, ease of maintenance, and seamless integration of future enhancements [8][9].

Overall, the proposed system offers a practical and efficient solution for automated sentiment analysis of online hotel reviews. It provides valuable insights for both customers and hotel management, supporting informed decision-making and improved service quality [1][4]. The achieved results establish a strong baseline for future work, including the incorporation of advanced deep learning models and richer contextual analysis to handle complex sentiment expressions more effectively [5][7].

5. Future Work

Future enhancements to improve the effectiveness and robustness of the proposed sentiment analysis system include:

1. **Hybrid Feature Analysis:**

Combining multiple textual and aspect-based features using feature fusion techniques to capture richer sentiment patterns and improve overall classification accuracy [3][4].

2. **Context-Aware Models:**

Incorporating contextual word embeddings to better understand sarcasm, implicit opinions, and domain-specific language commonly used in hotel reviews [7].

3. **Deep Learning Integration:**

Applying deep learning models such as Long Short-Term Memory (LSTM), Convolutional Neural Networks (CNN), or transformer-based architectures to capture complex linguistic and sequential relationships beyond traditional linear models [5][7] [16][17][18][30].

4. **Real-Time and Cloud Deployment:**

Deploying the system on cloud platforms using scalable and distributed architectures to support real-time sentiment analysis and high-volume data processing [8].

5. **Explainable AI:**

Integrating explainability techniques to identify influential words and phrases contributing to sentiment predictions, thereby improving transparency, interpretability, and user trust in automated decision-making systems [1][4].

6. **Multilingual Support:**

Extending the system to analyze reviews written in multiple languages to broaden applicability across global hotel platforms and diverse user bases [1][7].

6. References

[1] Pang, B., Lee, L., and Vaithyanathan, S. (2002). Thumbs Up? Sentiment Classification Using Machine Learning Techniques. Proceedings of the ACL Conference.

[2] Liu, B. (2012). Sentiment Analysis and Opinion Mining. Morgan & Claypool Publishers.

[3] Turney, P. D. (2002). Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews. Proceedings of the ACL Conference.

[4] Joachims, T. (1998). Text Categorization with Support Vector Machines: Learning with Many Relevant Features. European Conference on Machine Learning.

[5] McCallum, A., and Nigam, K. (1998). A Comparison of Event Models for Naive Bayes Text Classification. AAAI Workshop on Learning for Text Categorization.

[6] Pedregosa, F., Varoquaux, G., Gramfort, A., et al. (2011). Scikit-learn: Machine Learning in Python. Journal of Machine Learning Research, 12, 2825–2830.

[7] Bird, S., Klein, E., and Loper, E. (2009). Natural Language Processing with Python. O'Reilly Media.

[8] Jurafsky, D., and Martin, J. H. (2023). Speech and Language Processing (3rd ed.). Pearson Education.

[21] Feldman, R. (2013). Techniques and applications for sentiment analysis. *Communications of the ACM*, 56(4), 82–89.

[22] Aggarwal, C. C., & Zhai, C. (2012). *Mining Text Data*. Springer.

[23] Hutto, C. J., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis. *ICWSM*.

[24] Zhou, Z.-H. (2021). *Machine Learning*. Springer Nature.

[25] Han, J., Kamber, M., & Pei, J. (2012). *Data Mining: Concepts and Techniques*. Morgan Kaufmann.

[26] Chen, T., Zhang, R., & Xu, J. (2019). Sentiment analysis of online reviews based on machine learning methods. *Journal of Intelligent Information Systems*, 53(1), 1–20.

[9] Kaggle. Hotel Review Dataset. Retrieved from https://www.kaggle.com

[10] Python Software Foundation. Python 3 Documentation. https://www.python.org

[11] Flask Development Team. Flask Web Framework Documentation. https://flask.palletsprojects.com

[12] Manning, C. D., Raghavan, P., and Schütze, H. (2008). Introduction to Information Retrieval. Cambridge University Press.

[13] Cambria, E., Schuller, B., Xia, Y., & Havasi, C. (2013). New avenues in opinion mining and sentiment analysis. *IEEE Intelligent Systems*, 28(2), 15–21.

[14] Medhat, W., Hassan, A., & Korashy, H. (2014). Sentiment analysis algorithms and applications. *Ain Shams Engineering Journal*, 5(4), 1093–1113.

[15] Zhang, L., Wang, S., & Liu, B. (2018). Deep learning for sentiment analysis: A survey. *WIREs Data Mining and Knowledge Discovery*, 8(4).

[16] Kim, Y. (2014). Convolutional neural networks for sentence classification. *EMNLP*.

[17] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.

[18] Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers. *NAACL*.

[19] Ravi, K., & Ravi, V. (2015). A survey on opinion mining and sentiment analysis. *Knowledge-Based Systems*, 89, 14–46.

[20] Pak, A., & Paroubek, P. (2010). Twitter as a corpus for sentiment analysis. *LREC*.

[27] Xia, R., Zong, C., & Li, S. (2011). Ensemble of feature sets and classification algorithms for sentiment classification. *Information Sciences*, 181(6), 1138–1152.

[28] Rosenthal, S., Farra, N., & Nakov, P. (2017). SemEval-2017 task 4: Sentiment analysis in Twitter. *Proceedings of SemEval*, ACL.

[29] Baccianella, S., Esuli, A., & Sebastiani, F. (2010). SentiWordNet 3.0: An enhanced lexical resource for sentiment analysis. *LREC Conference*.

[30] Goldberg, Y. (2016). A primer on neural network models for natural language processing. *Journal of Artificial Intelligence Research*, 57, 345–420.