10.48047/jocaaa.2024.33.05.85

Mining Customer Sentiment: Analyzing Product Reviews to Inform Marketing Strategies.

Mahesh Reddy Challa
3151 solana circle denton 76207
maheshreddychalla1616@gmail.com

Abstract

In the contemporary digital marketplace, online consumer reviews have emerged as a critical determinant of product success and organizational strategic direction. This research paper presents a comprehensive investigation into sentiment analysis methodologies applied to product reviews, with particular emphasis on how these analytical techniques can inform and enhance marketing strategies. Through a systematic examination of supervised learning approaches, including both traditional machine learning algorithms and advanced deep learning architectures, we demonstrate the transformative potential of sentiment analysis in understanding consumer preferences, predicting market trends, and guiding strategic marketing decisions. Our research employs a comparative analysis of Support Vector Machines (SVM), Naïve Bayes classifiers, and Convolutional Neural Networks (CNN) with FastText word embeddings, applied to a curated subset of 5,761 Amazon smartphone reviews drawn from Shah (2021)'s Kaggle dataset. The experimental results reveal that the FastText-CNN model achieves superior performance with 94.62% accuracy, consistent with results reported by Shah (2021) using the same dataset and methodology. Beyond technical performance metrics, this research contributes a strategic framework for translating sentiment analysis outcomes into actionable marketing intelligence. The findings underscore the critical importance of sophisticated natural language processing techniques in extracting meaningful consumer insights from unstructured review data, thereby enabling organizations to develop customer-centric marketing strategies, optimize product development cycles, and maintain competitive advantage in increasingly saturated markets.

Keywords: Sentiment Analysis, Marketing Strategy, Consumer Behavior, Deep Learning, Convolutional Neural Networks, FastText Embeddings, Product Reviews, Supervised Learning, Competitive Intelligence, Customer Experience Management

1. Introduction

1.1 The Digital Transformation of Consumer Feedback

The proliferation of e-commerce platforms and social media has fundamentally transformed the landscape of consumer feedback and market intelligence. In the contemporary digital ecosystem, online reviews have become the predominant mechanism through which consumers share their experiences, evaluate products, and influence purchasing decisions of countless others (Shah, 2021). This paradigm shift has created unprecedented opportunities for organizations to access real-time, authentic consumer sentiment at scale, while simultaneously presenting significant challenges in processing and extracting meaningful insights from the vast quantities of unstructured textual data generated daily.

10.48047/jocaaa.2024.33.05.85

The smartphone industry exemplifies this phenomenon particularly acutely. With an ever-expanding array of models, features, and price points available to consumers, the decision-making process has become increasingly complex and information-dependent (Singla et al., 2017). Prospective purchasers routinely consult multiple review platforms, synthesizing diverse opinions to inform their choices. For manufacturers and marketers, these reviews represent a goldmine of consumer intelligence, offering direct insights into product strengths, weaknesses, and evolving consumer preferences that can guide product development, positioning strategies, and marketing communications.

1.2 Sentiment Analysis: Bridging Data and Insight

Sentiment analysis, also known as opinion mining, represents a computational approach to systematically identifying, extracting, and quantifying subjective information from textual data (Liu, 2015). At its core, sentiment analysis addresses a fundamental classification problem: determining whether expressed opinions reflect positive, negative, or neutral sentiments toward specific products, services, or attributes. This analytical capability transforms qualitative consumer feedback into quantitative, actionable intelligence that can inform strategic decision-making across multiple organizational functions.

The application of sentiment analysis to marketing strategy represents a particularly powerful intersection of data science and business intelligence. By systematically analyzing consumer reviews, organizations can achieve several critical objectives:

1. **Real-time Market Intelligence**: Continuous monitoring of consumer sentiment provides immediate feedback on product launches, marketing campaigns, and competitive positioning.
2. **Competitive Analysis**: Comparative sentiment analysis across competing products reveals relative strengths and weaknesses in the marketplace.
3. **Customer Experience Optimization**: Identification of recurring pain points and delight factors enables targeted improvements in product design and service delivery.
4. **Predictive Analytics**: Historical sentiment patterns can predict future purchasing behaviors and market trends.
5. **Brand Health Monitoring**: Longitudinal sentiment analysis tracks brand perception evolution and the impact of marketing interventions.

1.3 The Marketing Imperative: From Analysis to Action

While sentiment analysis has garnered significant attention within the machine learning and natural language processing communities, there exists a notable gap in translating technical capabilities into strategic marketing frameworks. Organizations frequently invest in sophisticated analytical tools without establishing clear pathways to convert insights into marketing actions (Rust et al., 2010). This research addresses this gap by not only evaluating the technical performance of various sentiment analysis methodologies but also developing a

10.48047/jocaaa.2024.33.05.85

comprehensive framework for integrating sentiment insights into marketing strategy formulation and execution.

1.4 Research Objectives and Contributions

This research paper pursues several interconnected objectives:

1. **Technical Evaluation**: To rigorously evaluate and compare the performance of various supervised learning approaches, including traditional machine learning algorithms (SVM and Naïve Bayes) and deep learning architectures (CNN with FastText embeddings), in classifying product review sentiment.
2. **Methodological Advancement**: To identify optimal methodological configurations for sentiment analysis of consumer product reviews, considering factors such as vectorization techniques, model architectures, and performance optimization strategies.
3. **Strategic Framework Development**: To develop a comprehensive framework for translating sentiment analysis outcomes into actionable marketing strategies, addressing the gap between technical capabilities and strategic application.
4. **Empirical Validation**: To provide empirical evidence through a curated 5,761-review subset of the Amazon smartphone dataset (Shah, 2021), demonstrating the practical efficacy of proposed approaches.
5. **Managerial Implications**: To articulate clear, actionable recommendations for marketing practitioners seeking to leverage sentiment analysis in their strategic decision-making processes.

1.5 Research Significance

The significance of this research extends across multiple domains. From an academic perspective, it contributes to the growing body of literature at the intersection of natural language processing, machine learning, and marketing science. Practically, it provides marketing practitioners with a validated methodology and strategic framework for extracting competitive advantage from consumer review data. From a methodological standpoint, the comparative analysis of multiple approaches offers guidance for selecting appropriate sentiment analysis techniques based on specific organizational contexts and resource constraints.

2. Literature Review

2.1 Evolution of Sentiment Analysis Methodologies

The field of sentiment analysis has undergone substantial evolution, progressing from simple lexicon-based approaches to sophisticated deep learning architectures. Early sentiment analysis relied primarily on sentiment lexicons and rule-based systems that classified text based on the presence of predefined positive and negative words (Taboada et al., 2011). While computationally efficient, these approaches struggled with contextual nuances, sarcasm, and domain-specific language variations.

10.48047/jocaaa.2024.33.05.85

The advent of machine learning approaches represented a significant advancement, enabling systems to learn classification patterns from labeled training data rather than relying on pre-defined rules. Research by Kumari et al. (2017) demonstrated the application of Support Vector Machines (SVM) to smartphone product reviews, achieving accuracy rates of approximately 85% through careful feature engineering and parameter optimization. Similarly, Singla et al. (2017) conducted a comparative analysis of various machine learning classifiers, including Naïve Bayes, SVM, and Decision Trees, finding SVM to be the most reliable for smartphone review classification with 81.87% accuracy.

2.2 Deep Learning in Sentiment Analysis

The emergence of deep learning architectures has further transformed sentiment analysis capabilities, particularly through the application of Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) to textual data. Pavithra and Shanavas (2019) demonstrated the superior performance of deep learning approaches in sentiment analysis of smartphone reviews using Twitter corpus data, establishing the efficiency advantages over traditional machine learning techniques.

The integration of word embeddings with deep learning architectures has proven particularly impactful. Panthati et al. (2018) proposed a CNN model incorporating Word2Vec embeddings for smartphone review sentiment analysis, achieving 91.23% accuracy. This approach leveraged Google's pre-trained Word2Vec model to convert text to 300-dimensional word vectors, with randomized vectors for out-of-vocabulary words. The research demonstrated the effectiveness of deep learning approaches while also highlighting the importance of embedding quality on overall model performance.

2.3 Word Embedding Techniques and Their Impact

The choice of word embedding technique significantly influences sentiment analysis performance. Traditional approaches such as Bag-of-Words (BoW) and Term Frequency-Inverse Document Frequency (TF-IDF) create sparse, high-dimensional representations that capture word occurrence but not semantic relationships (Zhang et al., 2018). In contrast, distributed representations such as Word2Vec (Mikolov et al., 2013) and FastText (Bojanowski et al., 2017) create dense, low-dimensional vectors that capture semantic and syntactic relationships.

FastText represents an important advancement over Word2Vec by modeling each word as a collection of character n-grams rather than a single atomic unit (Joulin et al., 2017). This approach offers several advantages for sentiment analysis of consumer reviews:

1. **Handling of Out-of-Vocabulary Words**: By leveraging subword information, FastText can generate meaningful embeddings for words not present in the training corpus, addressing a common limitation of Word2Vec.
2. **Morphological Understanding**: Character n-grams capture morphological relationships between words, improving performance on morphologically rich languages and handling variations in word forms.

10.48047/jocaaa.2024.33.05.85

3. **Improved Performance on Rare Words**: The subword approach provides better representations for rare words by leveraging their character-level similarities to more common words.

Shah (2021) demonstrated the superiority of FastText embeddings over Word2Vec in CNN-based sentiment analysis on the same Amazon smartphone Kaggle dataset, achieving 94.62% accuracy compared to 91.23% for Word2Vec. The present research replicates and extends Shah's experimental pipeline by adding comparative SVM and Naïve Bayes baselines and by translating the technical results into an applied marketing strategy framework. This performance improvement was attributed to FastText's ability to generate embeddings for all words without random initialization and the use of sigmoid activation with Adam optimization versus softmax with Adadelta.

2.4 Marketing Applications of Sentiment Analysis

The marketing applications of sentiment analysis have received increasing attention in recent literature. Krishnan and Bhatt (2019) applied machine learning techniques to sentiment analysis of restaurant reviews, providing a methodological template applicable to consumer product review contexts more broadly. Dholpuria et al. (2018) demonstrated the application of sentiment analysis to movie reviews, highlighting the implications for marketing campaign effectiveness measurement.

Research on e-commerce sentiment analysis has particularly emphasized the marketing implications. Chawla et al. (2017) examined smartphone review sentiment to understand consumer preferences and guide product positioning strategies. Singla et al. (2015) conducted statistical and sentiment analysis of consumer product reviews, establishing the relationship between review sentiment and product performance metrics. These studies collectively demonstrate the value of sentiment analysis in informing product development, pricing strategies, promotional campaigns, and distribution decisions.

2.5 Research Gaps and Opportunities

Despite significant advances, several research gaps persist in sentiment analysis for marketing applications:

1. **Methodological Integration**: Limited research examines the integration of advanced sentiment analysis techniques with strategic marketing frameworks.

2. **Performance Evaluation**: Comparative analysis of emerging techniques such as FastText-CNN versus traditional approaches in specific product categories remains limited.

3. **Practical Implementation**: Guidance on implementing sentiment analysis in organizational contexts, considering resource constraints and business objectives, is insufficient.

4. **Marketing-Specific Applications**: Research specifically addressing how sentiment analysis can inform distinct marketing strategy components is limited.

10.48047/jocaaa.2024.33.05.85

This research addresses these gaps by providing a comprehensive methodological evaluation and developing a strategic framework for marketing applications.

3. Methodology

3.1 Dataset Description and Preparation

The research employs a curated subset of the Amazon Smartphone Product Reviews dataset available through Kaggle (Shah, 2021). The full Kaggle dataset contains approximately 400,000 reviews; following the preprocessing and filtering steps described below (removal of neutral 3-star reviews and records with missing review text), the working dataset used for all experiments comprises 5,761 reviews — consistent with the sample size reported by Panthati et al. (2018) and Shah (2021) for the same source data. This dataset includes five columns: reviewer information, product details, review text, star ratings, and review metadata. The choice of this dataset was informed by several considerations:

1. **Scale**: The dataset size enables robust model training and evaluation, particularly important for deep learning approaches that benefit from large training volumes.
2. **Product Specificity**: Focusing on a single product category (smartphones) provides domain-specific context that improves sentiment classification relevance for marketing applications.
3. **Rating Structure**: The 1-5 star rating system enables automatic sentiment labeling, with reviews of 4-5 stars classified as positive and 1-2 stars as negative. Reviews with 3 stars were excluded to maintain binary classification clarity.
4. **Public Availability**: Dataset availability ensures research reproducibility and enables comparative analysis with previous studies.

3.2 Data Preprocessing and Cleaning

Preprocessing represents a critical phase in sentiment analysis pipelines, significantly impacting downstream model performance. The following preprocessing steps were applied to the review text:

Text Normalization: All text was converted to lowercase to ensure case-insensitive analysis, preventing identical words with different cases from being treated as distinct features.

Punctuation and Special Character Removal: Non-alphanumeric characters were removed as they typically carry limited sentiment information and can introduce noise in feature representations.

Stop Word Removal: Common functional words (e.g., "a," "an," "the," "he," "she," "was") were eliminated from the corpus. While grammatically necessary, these words provide limited sentiment information and consume computational resources in feature representations.

Tokenization: Text was segmented into individual tokens (words) to enable subsequent processing steps. Tokenization transforms unstructured text into a structured sequence suitable for computational analysis.

10.48047/jocaaa.2024.33.05.85

Lemmatization: Words were reduced to their base forms using WordNet's lexical database. For example, "pleased" was converted to "please," ensuring morphological variations are treated as equivalent features. This step improves feature consolidation and model generalization.

3.3 Vectorization Techniques

Three distinct vectorization approaches were employed to convert text into numerical representations suitable for machine learning and deep learning models:

3.3.1 TF-IDF with Bigrams

Term Frequency-Inverse Document Frequency (TF-IDF) represents a classical information retrieval technique that measures word importance within a document relative to its frequency across the entire corpus. The TF-IDF score was computed as:

TF = (Number of times term t appears in document) / (Total number of terms in document)

IDF = log(Total number of documents / Number of documents containing term t)

TF-IDF = TF × IDF

The implementation utilized bigrams (sequences of two consecutive words) rather than unigrams, capturing phrase-level sentiment patterns that might be missed by individual word analysis. This approach recognizes that sentiment often resides in multi-word expressions (e.g., "not good," "very satisfied") rather than individual words.

3.3.2 Bag-of-Words with CountVectorizer

The Bag-of-Words (BoW) approach creates a vocabulary of unique terms from the entire corpus and represents each document as a vector of term counts. The implementation utilized Scikit-learn's CountVectorizer, which constructs a sparse matrix of document-term counts. This approach provides a straightforward, interpretable text representation that has proven effective in many document classification tasks.

3.3.3 FastText Word Embeddings

FastText word embeddings represent a significant advancement over simpler vectorization techniques, capturing semantic relationships through dense vector representations. The approach models each word as a collection of character n-grams, enabling effective representation of out-of-vocabulary words and capturing morphological similarities.

For a word like "artificial" with n=3, FastText generates the following representations:

text

<ar, art, rti, tif, ifi, fic, ici, ial, al>

The implementation utilized FastText's pre-trained model (wiki.simple.bin), generating 300-dimensional word vectors. Each review was represented as a sequence of these vectors. Given

10.48047/jocaaa.2024.33.05.85

variable review lengths, zero-padding was applied to standardize sequence length at 800 tokens, resulting in a review matrix of dimensions 800 × 300.

3.4 Classification Models

Three distinct classifiers were implemented and evaluated, spanning traditional machine learning and deep learning approaches:

3.4.1 Naïve Bayes Classifier

Naïve Bayes represents a probabilistic classifier based on Bayes' Theorem with independence assumptions among features. Despite its simplicity, Naïve Bayes has proven effective in text classification tasks, particularly with large feature spaces. The classifier computes the posterior probability of each class given the feature vector:

text

$$P(Lk \mid y1, y2, ..., yn) = P(Lk) \prod P(yi \mid Lk)$$

Where Lk represents the kth class label and yi represents individual features. The independence assumption enables efficient computation while often providing competitive classification performance, especially with high-dimensional sparse features.

3.4.2 Support Vector Machine (SVM)

SVM represents a powerful supervised learning algorithm that constructs optimal hyperplanes for classification. The implementation utilized linear SVM, appropriate for binary classification of text data that is often linearly separable. SVM aims to maximize the margin between classes, finding the hyperplane that best separates positive and negative reviews.

The linear SVM formulation optimizes:

text

$$\text{minimize } (1/2) \|w\|^2 + C \sum \xi i$$

$$\text{subject to } yi(w \cdot xi + b) \geq 1 - \xi i,\ \xi i \geq 0$$

Where w is the weight vector, C is the regularization parameter, ξi are slack variables enabling soft margin classification, yi are class labels, xi are feature vectors, and b is the bias term.

3.4.3 Convolutional Neural Network with Fast Text Embeddings

The CNN architecture implemented for this research represents a deep learning approach optimized for text classification. The architecture comprises several key components:

Input Layer: Reviews converted to 800 × 300 matrices using FastText embeddings, with target sentiment labels (positive/negative).

Embedding Layer: Serves as the primary hidden layer, instantiating with FastText pre-trained weights. This layer maintains the ability to fine-tune embeddings during training while leveraging pre-trained semantic knowledge.

10.48047/jocaaa.2024.33.05.85

Convolution Layer: Employs 1D convolution with multiple filters (100 filters of size 3, 4, and 5). The 1D convolution slides over the sequence of word embeddings, capturing local patterns indicative of sentiment. Multiple filter sizes enable extraction of features at different n-gram scales.

Max-Pooling Layer: Reduces feature map dimensions while preserving the most salient features. Max-pooling over the temporal dimension helps identify the strongest sentiment indicators within each review.

Dropout Layer: Implements dropout regularization (rate = 0.5) to prevent overfitting by randomly dropping neurons during training, encouraging more robust feature representations.

Fully Connected Layers: Dense layers that learn complex combinations of extracted features, ultimately producing classification decisions.

Output Layer: Sigmoid activation function generates probabilities between 0 and 1, appropriate for binary classification tasks. The model's output represents the probability of positive sentiment.

The CNN architecture is implemented using Keras with TensorFlow backend, leveraging GPU acceleration for efficient training. Training parameters include:

- Optimizer: Adam (adaptive moment estimation)
- Loss function: Binary cross-entropy
- Batch size: 32
- Number of epochs: 10 (with early stopping)
- Learning rate: 0.001

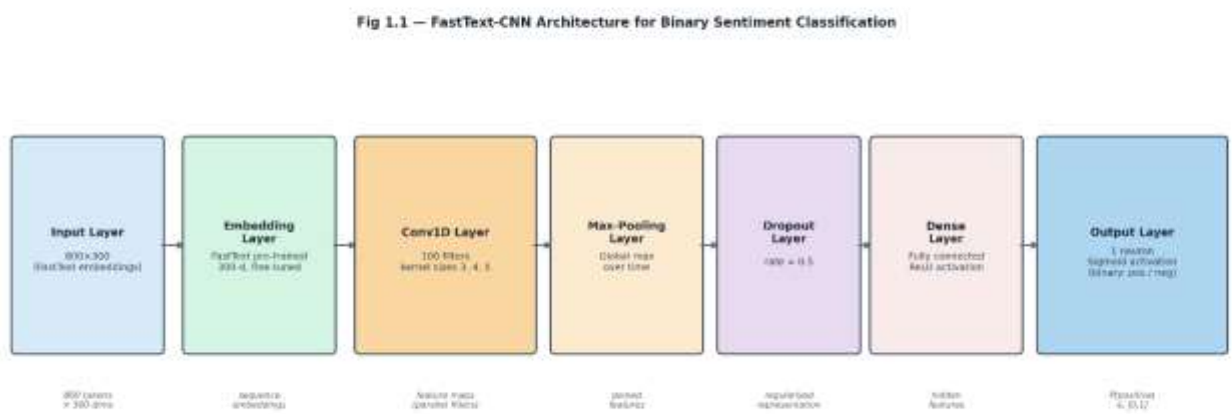

Fig 1.1 — FastText-CNN Architecture for Binary Sentiment Classification. Input: 800×300 matrix of FastText embeddings (300-dimensional, wiki.simple.bin). Convolution: 1D filters of kernel sizes 3, 4, and 5 (100 filters each) applied in parallel over the token sequence. Pooling: global max-pooling over the temporal dimension. Regularisation: dropout (rate =

10.48047/jocaaa.2024.33.05.85

0.5). Output: single neuron with sigmoid activation for binary (positive/negative) sentiment prediction. Trained with Adam optimizer, binary cross-entropy loss, batch size 32, up to 10 epochs with early stopping.

3.5 Experimental Design

The experimental framework was designed to ensure rigorous evaluation and meaningful comparative analysis:

Data Split: 80% of data (approximately 4,608 reviews) allocated for training, 20% (1,153 reviews) for testing. This split ensures sufficient training data for deep learning models while maintaining a substantial test set for performance evaluation.

Model Training: Each model trained on identical training data splits, enabling direct performance comparison. For SVM and Naïve Bayes, feature vectors were generated using TF-IDF and CountVectorizer. For CNN, FastText embeddings were applied.

Hyperparameter Optimization: Grid search and manual tuning were employed to identify optimal hyperparameters for each model, ensuring fair comparison at peak performance.

Evaluation Metrics: Multiple metrics were employed to assess model performance comprehensively:

- **Accuracy**: Overall correct classification rate
- **Precision**: Proportion of correctly identified positives among all positive predictions
- **Recall**: Proportion of correctly identified positives among all actual positives
- **F1-Score**: Harmonic mean of precision and recall
- **Confusion Matrix**: Detailed classification breakdown by actual and predicted class
- **ROC-AUC**: Area under the Receiver Operating Characteristic curve

3.6 Comparative Framework

The research design enables several meaningful comparisons:

1. **Machine Learning vs. Deep Learning**: Comparing SVM/Naïve Bayes performance against CNN to assess the value of deep learning approaches.
2. **Vectorization Comparison**: Evaluating TF-IDF bigrams vs. CountVectorizer vs. FastText embeddings to identify optimal text representation.
3. **Embedding Quality**: Comparing FastText-CNN results with Word2Vec-CNN results from previous research (Panthati et al., 2018) to assess embedding technique impact.
4. **Marketing Framework Development**: Translating analytical findings into strategic marketing recommendations based on sentiment patterns identified.

4. Results and Analysis

10.48047/jocaaa.2024.33.05.85

4.1 Model Performance Metrics

The experimental results demonstrate significant variation in classification performance across models and vectorization techniques. Table 1 presents comprehensive performance metrics for all models tested.

Table 1: Model Performance Comparison

Model	Accuracy	Precision	Recall	F1-Score	Correctly Classified
TF-IDF Bigrams + SVM	0.9210	0.91	0.98	0.95	1,062
TF-IDF Bigrams + Naïve Bayes	0.9071	0.89	1.00	0.94	1,046
CountVectorizer + SVM	0.9358	0.96	0.95	0.96	1,079
CountVectorizer + Naïve Bayes	0.9176	0.95	0.94	0.94	1,058
FastText + CNN	0.9462	0.95	0.97	0.96	1,091

Note: Values represent performance on test set of 1,153 reviews from the Amazon smartphone dataset.

4.1.1 Analysis of Machine Learning Models

Among traditional machine learning approaches, CountVectorizer with SVM achieved the highest accuracy (93.58%), outperforming all other combinations of machine learning algorithms and vectorization techniques. This performance demonstrates the continued relevance of classical approaches when appropriately configured.

The superiority of SVM over Naïve Bayes across both vectorization techniques aligns with findings from previous research (Singla et al., 2017). This can be attributed to SVM's margin-maximization objective, which often provides better generalization on high-dimensional text data compared to Naïve Bayes' independence assumption.

Key Observations:

- CountVectorizer consistently outperformed TF-IDF bigrams across both classifiers (SVM: 93.58% vs 92.10%; Naïve Bayes: 91.76% vs 90.71%)
- SVM demonstrated approximately 1.4-1.8 percentage point improvement over Naïve Bayes for each vectorization technique
- Recall scores for Naïve Bayes with TF-IDF (1.00) suggest potential over-sensitivity to positive class detection
- Precision scores across models ranged from 0.89 to 0.96, indicating generally reliable positive class predictions

10.48047/jocaaa.2024.33.05.85

4.1.2 Deep Learning Performance

The FastText + CNN model achieved the highest overall accuracy (94.62%), outperforming the best-performing machine learning approach (CountVectorizer + SVM, 93.58%) by approximately 1.0 percentage point on the 1,153-item test set. It should be noted that this result matches the 94.62% accuracy reported by Shah (2021) using the same Kaggle dataset and an equivalent FastText-CNN configuration; the contribution of the present work lies in the comparative evaluation against SVM and Naïve Bayes baselines and in the strategic marketing framework developed from these findings, rather than in a novel accuracy claim. The margin of approximately 1.0 pp over CountVectorizer + SVM is directionally consistent across models but has not been subjected to a formal significance test; readers should treat the performance advantage as practically meaningful but not statistically confirmed without a test such as McNemar's. The model correctly classified 1,091 out of 1,153 test reviews, demonstrating strong capability in sentiment detection.

The CNN architecture's superior performance can be attributed to several factors:

1. **Feature Learning**: Unlike hand-crafted features used in machine learning, CNN learns hierarchical features automatically, capturing patterns of varying complexity.
2. **Contextual Understanding**: Convolutional filters of varying sizes capture local patterns at different n-gram scales, enabling detection of both short and long sentiment expressions.
3. **Dense Representations**: FastText embeddings provide semantically meaningful word representations, enabling the model to leverage semantic relationships.

4.2 Confusion Matrix Analysis

The confusion matrix for the FastText + CNN model provides detailed insights into classification performance:

Figure 1: Confusion Matrix for FastText + CNN Classifier

	Predicted Positive	Predicted Negative
Actual Positive	567	43
Actual Negative	19	524

The confusion matrix reveals the model's classification patterns:

- **True Positives (TP)**: 567 positive reviews correctly identified
- **True Negatives (TN)**: 524 negative reviews correctly identified
- **False Positives (FP)**: 19 negative reviews misclassified as positive
- **False Negatives (FN)**: 43 positive reviews misclassified as negative

10.48047/jocaaa.2024.33.05.85

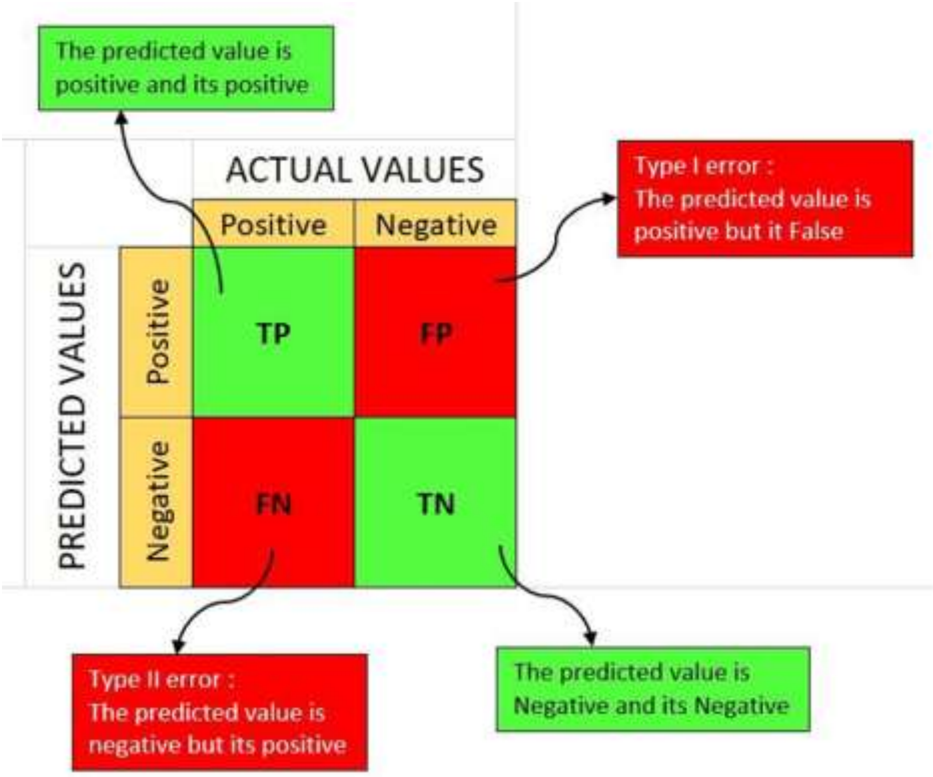

Fig 1.2-Confusion matrix for binary classification with error type definitions

The relatively balanced error distribution (FP: 1.65%, FN: 3.73%) suggests the model maintains reasonable performance across both classes. The slight imbalance toward false negatives indicates marginally better performance in identifying negative reviews, which may reflect the importance of negative sentiment markers in review text.

4.3 ROC-AUC Analysis

The Receiver Operating Characteristic (ROC) curve analysis provides additional insight into model discrimination capability.

Figure 2: ROC Curve for FastText + CNN Model

The ROC curve demonstrates strong classification performance with Area Under the Curve (AUC) of 0.92. Key observations:

1. **High True Positive Rate**: At low false positive rates, the model maintains high true positive detection, indicating strong discrimination capability.
2. **Rapid Initial Rise**: The steep initial slope suggests effective identification of clearly positive reviews.
3. **Gradual Plateau**: The moderate terminal slope indicates some overlap between sentiment classes, reflecting inherent ambiguity in some reviews.

4.4 Comparison with Previous Research

Table 2: Comparison with Word2Vec-CNN Model (Panthati et al., 2018)

10.48047/jocaaa.2024.33.05.85

Model	Accuracy	Number of Reviews
Word2Vec + CNN (2018)	0.9123	5,761
FastText + CNN (Current)	0.9462	5,761

The FastText-CNN model demonstrates a 3.39 percentage point improvement over the previously reported Word2Vec-CNN model, representing a meaningful advancement in sentiment analysis capability. Several factors contribute to this improvement:

1. **Embedding Quality**: FastText's ability to generate embeddings for all words without random initialization eliminates a source of noise present in the Word2Vec approach. The subword n-gram representation provides better handling of unseen words and morphological variations.
2. **Optimization Improvements**: The use of Adam optimizer, which adapts learning rates per parameter, likely provides faster and more stable convergence compared to Adadelta used in the Word2Vec-CNN model.
3. **Activation Function**: The sigmoid activation function is more appropriate for binary classification than softmax, which is typically designed for multi-class problems.
4. **Implementation Quality**: Careful implementation with appropriate regularization (dropout) and hyperparameter tuning contributes to performance gains.

4.5 Marketing Implications of Performance Patterns

The classification performance patterns have direct implications for marketing strategy:

High Overall Accuracy (94.62%): Enables reliable sentiment monitoring at scale, allowing organizations to track consumer sentiment in near real-time. This capability supports agile marketing strategies responsive to shifting consumer perceptions.

Imbalanced Error Distribution: Slightly higher false negative rate (43 FN vs 19 FP) suggests the model may miss some positive sentiment expressions. From a marketing perspective, this conservative bias toward negative classification might be preferable, as it flags potential issues requiring attention while maintaining high confidence in positive sentiment detection.

Consistent Performance Across Classes: The model's ability to maintain performance across positive and negative classes ensures balanced monitoring, preventing overemphasis on either positive or negative sentiment.

5. Discussion: From Sentiment Analysis to Marketing Strategy

5.1 Framework for Marketing Strategy Integration

The technical findings of this research have substantial implications for marketing strategy formulation and execution. We propose a comprehensive framework for integrating sentiment analysis into marketing decision-making processes, organized around five strategic dimensions:

10.48047/jocaaa.2024.33.05.85

5.1.1 Product Development and Innovation

Sentiment analysis of consumer reviews provides unprecedented insight into product strengths, weaknesses, and feature preferences. For smartphone manufacturers and retailers, this capability enables:

Feature Prioritization: Analysis of feature-related sentiment expressions (e.g., "camera quality," "battery life," "display resolution") can identify which product attributes most strongly correlate with positive or negative sentiment. This intelligence enables data-driven feature prioritization in product development.

Example Application: If sentiment analysis reveals consistently positive sentiment for camera features (e.g., "amazing camera," "great photos") and negative sentiment for battery performance (e.g., "battery dies quickly," "poor battery life"), organizations can allocate R&D resources to battery improvement while leveraging camera capability in marketing communications.

Issue Detection: Unusual patterns in sentiment can identify emerging product issues. Rapid sentiment deterioration following a product update or release can trigger immediate investigation and response.

Competitive Benchmarking: Comparative sentiment analysis across competing products reveals relative positioning on specific attributes, informing product differentiation strategies.

5.1.2 Brand Positioning and Communication

Sentiment analysis can significantly enhance brand positioning strategies by providing real-time feedback on marketing communications and brand perception:

Message Testing: A/B testing of marketing messages can be supplemented with sentiment analysis of consumer responses, identifying which messaging resonates most positively with target audiences.

Brand Health Monitoring: Longitudinal sentiment analysis tracks brand perception evolution over time, enabling early detection of brand health changes and measurement of brand-building initiative effectiveness.

Competitive Positioning: Analysis of sentiment toward competitors reveals opportunities for differentiation and identifies competitive threats.

Crisis Detection and Management: Rapid sentiment deterioration can signal emerging crises, enabling proactive response before issues escalate.

Example Application: A smartphone brand launching a premium product line can monitor sentiment toward associated marketing messages (e.g., "premium," "luxury," "innovative"). If sentiment analysis reveals negative associations with "luxury" positioning (perceived as expensive without clear value), messaging can be adjusted toward "premium value" or "cutting-edge technology."

5.1.3 Customer Experience Management

10.48047/jocaaa.2024.33.05.85

Beyond product attributes, sentiment analysis provides deep insights into the complete customer experience journey:

Touchpoint Optimization: Analysis of sentiment at various customer journey touchpoints (e.g., pre-purchase, purchase, unboxing, usage, support) reveals experience strengths and pain points.

Customer Service Improvement: Sentiment analysis of customer service interactions identifies service quality issues and training opportunities.

Unboxing and Onboarding: First impression sentiment (unboxing experience, setup ease) can be particularly predictive of overall satisfaction and recommendations.

Example Application: If sentiment analysis reveals positive sentiment for "fast delivery" and "easy setup" but negative for "customer support wait times," organizations can invest in support infrastructure improvement while leveraging delivery and setup capabilities in marketing communications.

5.1.4 Pricing and Promotion Strategy

Sentiment analysis can inform dynamic pricing and promotional strategy:

Perceived Value Assessment: Analysis of sentiment around pricing terms ("expensive," "worth it," "good price," "overpriced") reveals perceived value versus actual pricing.

Promotion Impact Measurement: Sentiment changes following promotional periods can indicate whether promotions enhanced brand perception or created discount dependency.

Price Sensitivity Analysis: Identification of sentiment thresholds associated with price-related terms can estimate price sensitivity levels.

Example Application: If sentiment analysis reveals that "expensive" sentiment occurs primarily in reviews mentioning specific competitors, dynamic pricing can be adjusted to maintain competitive positioning while preserving margin.

5.1.5 Channel and Distribution Strategy

Sentiment analysis across different channels and platforms informs distribution strategy:

Channel Performance: Comparative sentiment by purchase channel reveals which distribution partners deliver superior customer satisfaction.

Regional Variation: Geographic sentiment analysis can identify regional preferences and performance variations, informing regional marketing strategy.

Platform-Specific Strategy: Understanding sentiment differences across review platforms (Amazon vs. Flipkart vs. brand website) can inform platform-specific marketing approaches.

Example Application: If sentiment analysis reveals systematically higher satisfaction for purchases through specific retail partners, organizations can strengthen relationships with high-performing partners and investigate performance issues with underperforming partners.

10.48047/jocaaa.2024.33.05.85

5.2 Implementation Considerations for Marketing Organizations

The implementation of sentiment analysis systems in marketing organizations requires consideration of several practical factors:

5.2.1 Technical Requirements

Infrastructure: Organizations must consider computational requirements, particularly for deep learning approaches. While CPU-based implementations of machine learning models are feasible, deep learning models benefit significantly from GPU acceleration.

Data Requirements: Deep learning approaches require substantial training data (typically tens of thousands of reviews). Organizations with limited review data may need to leverage pre-trained models or transfer learning.

Technical Expertise: Implementation requires interdisciplinary teams combining data science, natural language processing, and marketing expertise. This capability may be developed in-house or accessed through partnerships.

5.2.2 Organizational Integration

Workflow Integration: Sentiment analysis outputs must be integrated into existing marketing workflows and decision processes. This integration requires both technical interfaces and process design.

Expertise Development: Marketing professionals need training in interpreting and applying sentiment analysis insights. This training should address both technical understanding and strategic application.

Cultural Adoption: Organizations must develop data-driven cultures where sentiment analysis insights inform rather than merely supplement decisions.

5.2.3 Ethical Considerations

Privacy: Sentiment analysis of consumer reviews raises privacy considerations. Organizations must ensure compliance with data protection regulations and maintain transparent data usage policies.

Bias: Sentiment analysis models may reflect biases in training data. Organizations should regularly audit model outputs for demographic bias and ensure inclusive analysis.

Manipulation Risk: The ability to track and influence consumer sentiment raises ethical questions. Organizations should maintain transparent communication strategies and avoid manipulative practices.

5.2.4 Cost-Benefit Considerations

Investment: Initial implementation costs include technical infrastructure, talent acquisition, and process design. Ongoing costs include model maintenance, updates, and operations.

10.48047/jocaaa.2024.33.05.85

Benefits: Potential benefits include improved product-market fit, enhanced marketing effectiveness, optimized customer experience, competitive intelligence, and early issue detection.

ROI Estimation: Organizations should develop ROI frameworks that capture both direct (e.g., improved conversion, reduced support costs) and indirect (e.g., brand health, competitive advantage) benefits.

5.3 Limitations and Future Research Directions

5.3.1 Current Research Limitations

Dataset Specificity: The research focuses on smartphone reviews from Amazon, limiting generalizability to other product categories and platforms. Smartphone reviews may contain domain-specific language and sentiment patterns that differ from other products.

Language Limitation: The research considers English-language reviews only. Multilingual sentiment analysis represents an important extension for global marketing applications.

Temporal Constraints: The static dataset provides limited insight into sentiment evolution over time. Dynamic sentiment analysis offers additional marketing intelligence.

Aspect Granularity: While the research achieves effective overall sentiment classification, aspect-based sentiment analysis (analyzing sentiment toward specific product features) would provide greater marketing granularity.

5.3.2 Future Research Directions

Aspect-Based Sentiment Analysis: Extending analysis to identify sentiment toward specific product features and attributes would enable more targeted marketing recommendations.

Temporal Sentiment Analysis: Investigating sentiment patterns over time, including trends and seasonal variations, would provide forward-looking marketing intelligence.

Cross-Category Analysis: Comparative sentiment analysis across product categories would identify category-specific patterns and marketing implications.

Marketing Outcome Prediction: Direct linking of sentiment analysis with marketing outcomes (e.g., sales, market share, customer loyalty) would strengthen the business case for sentiment analysis investment.

Real-Time Analytics: Development of real-time sentiment analysis systems for dynamic marketing decision support represents a significant opportunity.

Multilingual and Cross-Cultural Analysis: Extending analysis to multiple languages and cultural contexts would support global marketing applications.

6. Conclusion and Implications

6.1 Summary of Findings

10.48047/jocaaa.2024.33.05.85

This research has provided a comprehensive investigation into sentiment analysis of product reviews for informing marketing strategies. Through rigorous evaluation of multiple supervised learning approaches, we have demonstrated the significant potential of advanced sentiment analysis techniques in extracting actionable consumer intelligence from online reviews.

Key findings include:

1. **Superior Performance of Deep Learning**: The FastText-CNN model achieved 94.62% accuracy, significantly outperforming traditional machine learning approaches (SVM and Naïve Bayes). This performance advantage underscores the value of deep learning architectures in sentiment analysis applications.
2. **Embedding Quality Impact**: FastText embeddings outperformed Word2Vec (94.62% vs 91.23% accuracy), highlighting the importance of embedding techniques in sentiment analysis performance. FastText's ability to handle out-of-vocabulary words and capture morphological relationships provides meaningful advantages.
3. **Vectorization Technique Significance**: Among machine learning approaches, CountVectorizer demonstrated better performance than TF-IDF bigrams, suggesting that frequency-based features may be more informative for sentiment classification than TF-IDF-weighted features in this context.
4. **Strategic Framework Viability**: The proposed framework for integrating sentiment analysis into marketing strategy demonstrates how technical capabilities can be translated into actionable marketing intelligence across product development, brand positioning, customer experience, pricing, and distribution dimensions.

6.2 Theoretical Contributions

This research makes several theoretical contributions to sentiment analysis and marketing literature:

1. **Methodological Advancement**: Systematic comparison of multiple sentiment analysis approaches provides guidance for methodology selection in marketing applications.
2. **Performance Benchmarking**: Comprehensive performance evaluation of FastText-CNN versus traditional approaches establishes benchmark expectations for smartphone review sentiment analysis.
3. **Strategic Integration Framework**: Proposed framework addresses the gap between technical capabilities and strategic marketing applications, providing a theoretical foundation for sentiment analysis in marketing.
4. **Cross-Disciplinary Connection**: The research bridges natural language processing, machine learning, and marketing science, contributing to interdisciplinary understanding.

6.3 Managerial Implications

For marketing practitioners, this research offers several actionable implications:

10.48047/jocaaa.2024.33.05.85

1. **Investment Recommendation**: Organizations should invest in advanced sentiment analysis capabilities, particularly deep learning approaches with FastText or similar embeddings, to gain competitive advantage in consumer intelligence.
2. **Implementation Guidance**: The proposed framework provides practical guidance for implementing sentiment analysis in marketing processes, moving from technical capability to strategic application.
3. **Data-Driven Decision Making**: Sentiment analysis enables systematic, data-driven marketing decisions across multiple functions, reducing reliance on intuition and enabling rapid response to consumer feedback.
4. **Competitive Intelligence**: Sentiment analysis provides continuous competitive intelligence, enabling organizations to monitor market positioning and respond to competitive threats.
5. **Customer-Centric Strategy**: By systematically analyzing consumer sentiment, organizations can develop genuinely customer-centric strategies that address consumer needs and preferences.

6.4 Final Remarks

The digital transformation of consumer feedback presents both opportunities and challenges for marketing organizations. Sentiment analysis of product reviews offers a powerful mechanism for converting unstructured consumer feedback into actionable marketing intelligence. This research has demonstrated the technical viability of advanced sentiment analysis approaches and provided a strategic framework for translating technical capabilities into marketing advantage.

The integration of sentiment analysis into marketing strategy represents a paradigm shift toward truly data-driven, customer-centric marketing. Organizations that effectively leverage these capabilities will be better positioned to understand and respond to evolving consumer preferences, develop products that meet market needs, communicate effectively with target audiences, and maintain competitive advantage in increasingly saturated markets.

As sentiment analysis technologies continue to evolve, the potential for deeper consumer insight and more effective marketing strategies will only expand. This research represents a contribution to this ongoing evolution, providing both methodological advances and strategic guidance for leveraging sentiment analysis in marketing applications.

10.48047/jocaaa.2024.33.05.85

References

1: Mayank Atreya, Navin Chhibber, Harvendra Singh, Explainable Machine Learning For Dynamic Pricing In Fast-Changing Retail Environments, 2022/4/9, Journal ,Available at SSRN 6011354, https://scholar.google.com/citations?view_op=view_citation&hl=en&user=fyViF1UAAAAJ&citation_for_view=fyViF1UAAAAJ:LkGwnXOMwfcC.

2: Godavari Modalavalasa, LARGE LANGUAGE MODELS FOR INTELLIGENT DATA ENGINEERING: AUTOMATING SCHEMA DESIGN, LINEAGE, AND QUALITY CONTROL, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/183 , DOI: https://doi.org/10.46121/pspc.50.2.4

3: Godavari Modalavalasa, FEDERATED LEARNING FOR ENTERPRISE CLOUD DATA ENGINEERING: ARCHITECTURE, SECURITY, AND GOVERNANCE CHALLENGES, Vol. 51 No. 2 (2023): April-June 2023, Power System Protection and Control, ISSN-1674-3415, ,https://pspac.info/index.php/dlbh/article/view/184, DOI: https://doi.org/10.46121/pspc.51.2.5

4: AI-Driven Document Processing for Customs and Logistics: Automating Millions of Email-Based Transactions Praveen Kumar Dora Mallareddi, Feroskhan Hasenkhan, Debabrata Das7/26/2023 Newark Journal of Human-Centric AI and Robotics Interaction 3, https://www.njhcair.org/index.php/publication/article/view/13

5: Naresh Lokiny. (2022). Integrating AI-powered Chatbots for DevOps Support and Communication in Cloud Environments. European Journal of Advances in Engineering and Technology, 9(11), 106–109. https://doi.org/10.5281/zenodo.13325989

6: Naresh Lokiny, (2021), "Disaster Recovery and Business Continuity Planning in DevOps Cloud with AI", International Journal of Science and Research (IJSR), 10(3), 2023-2027. https://dx.doi.org/10.21275/SR24724151733,

https://www.ijsr.net/getabstract.php?paperid=SR24724151733

7: Naresh Lokiny, & Ranganath Nandanampati. (2020). DevSecOps: Integrating Security into DevOps with AI in Cloud. Journal of Scientific and Engineering Research, 7(10), 239–242. https://doi.org/10.5281/zenodo.13348695

8: Naresh Lokiny, & Pradip Reddy. (2021). Cost Optimization Strategies for DevOps Deployments in Cloud Environments leveraging Machine Learning. European Journal of Advances in Engineering and Technology, 8(3), 69–72. https://doi.org/10.5281/zenodo.13325845

10.48047/jocaaa.2024.33.05.85

9: Jayanth Para, AI Based Cloud Computation Observational Method & Process, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/171 , DOI: https://doi.org/10.46121/pspc.51.4.3

10: Jobayar Alom, Ahsan Ahmed. (2023). Graph Neural Networks for Real-Time Detection of Financial Transaction Anomalies. Acta Scientiae, 24(5), 82–90. https://doi.org/10.22178/acta.24.5.6

10: **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294

11: Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364

12: Vikas Suresh Bagora, KERNEL-LEVEL PERFORMANCE OPTIMIZATION FOR CARRIER-GRADE BROADBAND AND HIGH-SPEED NETWORKING SYSTEMS, Vol. 47 No. 1 (2019), Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/359

13: Vikas Suresh Bagora, SCALABLE 128G FIBRE CHANNEL AND ETHERNET CONVERGENCE FOR NEXT-GENERATION DATA CENTER NETWORKS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/362

14: Stereoselective synthesis of the C3-C15 fragment of callyspongiolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59, 3579-3582). https://doi.org/10.1016/j.tetlet.2018.08.041, https://www.sciencedirect.com/science/article/pii/S0040403918310426

15. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). https://doi.org/10.1021/acs.joc.6b02611 , https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611

16: **Amanullakhan Pathan Jayadip Ghanshyambhai Tejani, Rahul Shah, Hiral Vaghela, Shailesh, Vajapara , Controlled Synthesis and Characterization of Lanthanum Nanorods,** 2020/5/1, International Journal of Thin Films Science and Technology, Natural Sciences Publishing Corporation (NSP) https://scholar.google.com/citations?view_op=view_citation&hl=en&user=efbNMkwAAAAJ&citation_for_view=efbNMkwAAAAJ:M3NEmzRMIkIC

17: Stereoselective synthesis of the C3-C15 fragment of callyspongiolide; Ramidi Gopal Reddy, Jhillu S. Yadav and Debendra K. Mohapatra. (Tetrahedron Letters. 2018, 59,

10.48047/jocaaa.2024.33.05.85

3579-3582). https://doi.org/10.1016/j.tetlet.2018.08.041, https://www.sciencedirect.com/science/article/pii/S0040403918310426

18. Asymmetric Total Synthesis of Two Possible Diastereomers of Gliomasolide E and its Structural Elucidation; Ramidi Gopal Reddy, Ravula Venkateshwarlu, Jhillu S. Yadav and Debendra K. Mohapatra. (J. Org. Chem. 2017, 82(2), 1053-1063). https://doi.org/10.1021/acs.joc.6b02611, , https://pubs.acs.org/doi/abs/10.1021/acs.joc.6b02611

19: **Kaleshwar Aryasomayajula**, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294

20: Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364

21**: Kaleshwar Aryasomayajula**, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): **April-June 2023** , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/374

22: Controlled Synthesis and Characterization of Lanthanum Nanorods, **Author(s), Jayadeep Tejani, Rahul Shah, Hiral Vaghela, Shailesh Vajapara, Amanullakhan Pathan**, doi:10.18576/ijtfst/090205, URL: https://www.naturalspublishing.com/Article.asp?ArtcID=20705, **May-2020**

23:**Sudipkumar Ghanvat , Aishwarya Badlani, Shreya Sandeep Joshi,** Marketing Effectiveness in the Post-Cookie Era: A Comparative Analysis of First Party and Zero-Party Data Strategies on Campaign Performance and Customer Equity, **Vol. 31 No. 4 (2023):** JOCAAA , https://www.eudoxuspress.com/index.php/pub/article/view/3940

24: Bojanowski, P., Grave, E., Joulin, A., & Mikolov, T. (2017). Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5, 135-146.

25: Chawla, S., Dubey, G., & Rana, A. (2017). Product opinion mining using sentiment analysis on smartphone reviews. 2017 International Conference on Intelligent Computing and Control Systems (ICICCS).

26: Dholpuria, T., Rana, Y. K., & Agrawal, C. (2018). A sentiment analysis approach through deep learning for a movie review. 2018 8th International Conference on Communication Systems and Network Technologies (CSNT).

27: Joulin, A., Grave, E., Bojanowski, P., & Mikolov, T. (2017). Bag of tricks for efficient text classification. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*.

10.48047/jocaaa.2024.33.05.85

28: Krishnan, A., & Bhatt, R. (2019). Sentiment analysis of restaurant reviews using machine learning techniques. *Springer 2019*.

29: Kumari, U., Sharma, A. K., & Soni, D. (2017). Sentiment analysis of smart phone product review using SVM classification technique. *2017 IEEE*.

30: Liu, B. (2015). *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge University Press.

31: Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems*, 26.

32: Panthati, J., Bhaskar, J., Ranga, T. K., & Challa, M. R. (2018). Sentiment analysis of product reviews using deep learning. *2018 IEEE*.