Grenze International Journal of Engineering and Technology, January Issue

A Geo-Analytical Study of Consumer Sentiment on Apple Products using VADER

Dharanya S[1], Ganga Sree M[2], Joanna Sharon D[3] and Sreemathy J[4]
[1-4]Sri Eshwar College of Engineering/CSE, Coimbatore, India
Email: dharanya.s2023cse@sece.ac.in, gangasree.m2023cse, joannasharon.d2023cse, sreemathy.j@sece.ac.in

***Abstract*— In the digital age, consumer opinions shared on social media platforms offer valuable insights into brand perception and product performance across different regions. This study presents a geographic sentiment analysis of Apple products using Twitter data to understand how public sentiment varies by location. Tweets mentioning Apple products were collected, preprocessed, and analyzed using the VADER sentiment analysis algorithm to classify them into positive, negative, or neutral categories. Sentiment trends were visualized across regions, and average sentiment scores were computed to assess regional differences in consumer attitudes. The findings reveal distinct geographic sentiment patterns, which can inform targeted marketing strategies and product development decisions. Additionally, sentiment distribution and polarity scores offer a snapshot of user satisfaction and public opinion in real time. This research highlights the potential of combining social media mining with sentiment analysis to support data-driven decision-making in brand management and regional market analysis.**

***Index Terms*— Geographic sentiment analysis, Apple products, Twitter data, VADER sentiment analysis, Consumer attitudes.**

I. INTRODUCTION

In today's digital world, social media platforms like Twitter and YouTube have become essential for consumers to share their opinions, emotions, and experiences with products and brands. For global tech companies like Apple Inc., analyzing this user-generated content gives them important insights into public perception, brand sentiment, and market behavior in different regions and demographic groups.

Sentiment analysis is a part of natural language processing (NLP) that helps researchers and marketer measure the emotional tone in social media post. However, most traditional sentiment analysis studies look at textual content in general, ignoring the impact of geographical and demographic contexts. This creates a gap in understanding how consumer sentiment changes across regions like continents or cities and among different demographic traits such as age or device usage.

This study aims to fill that gap by examining the geodemographic factors that shape sentiment toward Apple products. By using datasets from Twitter and YouTube and analyzing user engagement across various locations like North America, Asia, and Bangalore, as well as device types (Apple vs. Android), this research uncovers key sentiment patterns and engagement trends. We specifically look at how sentiments vary based on the user's continent, age group, and mobile operating system, and how these factors affect influencer engagement and regional brand perception. The findings support a more targeted and effective digital marketing strategy by highlighting the need for localized content and tailored outreach.

Grenze ID: 01.GIJET.12.1.507_1

Additionally, the insights gained from over 358 million analyzed tweets and several regional case studies provide a solid basis for future geodemographic sentiment analysis.

II. Literature Survey

The rapid expansion of social media platforms in recent years has led to a significant increase in research interest in sentiment analysis, also referred to as opinion mining. Studies like Kharde and Sonawane (2016) have examined a variety of sentiment analysis approaches using data from Twitter, emphasizing lexicon-based and machine learning approaches. For classifying sentiment on social media, VADER, a lexicon-based tool designed for brief, informal text, has gained widespread use. The use of dashboards, pie charts, and web interfaces to visualize sentiment outputs has been the subject of numerous studies. Nevertheless, deployment issues are frequently encountered when incorporating these models into web applications using frameworks such as Django, particularly in cross-platform settings. This calls for more investigation into the scalable and portable implementation of sentiment analysis systems utilizing contemporary web technologies.[1]

Sentiment analysis has emerged as a major area of study for comprehending consumer behavior as Online Social Networks (OSNs) have grown in popularity. Conventional sentiment analysis methods frequently categorize viewpoints as neutral, negative, or positive, but they fail to take into account the complexity of human emotion and thought. To improve sentiment interpretation, recent research has investigated emotion detection and cognitive computing. The groundwork for classifying emotions has been established by works like Ekman's basic emotions model and Plutchik's Emotion Wheel. Product-based sentiment analysis has made use of crowdsourced data, especially from Twitter. However, there is still a lack of research on integrating cognitive-emotion analytics, which is why the current study focuses on assigning eight human cognition categories to tweets.[2]

With the growth of social media sites like Twitter, where users publish content that is full of opinions in real time, sentiment analysis has become more popular. Previous research has used machine learning techniques to identify sentiments in text, but has frequently encountered difficulties because of slang, informal language, misspellings, and brief messages. Large-scale sentiment datasets have been successfully constructed through research that uses distant supervision, which labels data with emoticons or hashtags. Preprocessing methods like noise reduction, normalization, and stemming have been essential for increasing accuracy. The groundwork for fine-grained opinion mining has been laid by a number of studies that have expanded sentiment classification beyond binary labels to encompass multi-class emotional states.[3]

Due to its potential to enhance user experience and service response, sentiment analysis in customer support communication has emerged as a rapidly expanding field of study. Machine learning techniques like Support Vector Machines (SVM) have been used in earlier research to categorize emotions in textual data. When labeled data is not available, lexicon-based methods such as VADER and domain-specific sentiment lexicons are frequently employed for initial annotation. Prediction accuracy is increased when supervised models are combined with lexicon-based labeling, according to research. Furthermore, research on sentiment forecasting in conversation threads has shown that sentiment frequently exhibits patterns, allowing for the prediction of customer reactions to support responses.[4]

VADER and other lexicon-based sentiment analysis tools have shown promise, particularly for brief, casual texts like posts on social media. However, VADER's direct application in other languages is limited by its English-specific design. Recent studies have looked into language adaptation strategies to address this. The German version of VADER, called GerVADER, combines the language-independent features of VADER with the SentiWS lexicon. Similar studies have demonstrated that lexicon accuracy is improved by user-rated polarity and intensity. Research contrasting GerVADER with datasets such as SB10k and SCARE emphasizes the difficulties in adapting sentiment across languages, but also shows that with the right modifications, multilingual sentiment analysis can be successful.[5]

One of the main tasks in natural language processing (NLP) is sentiment analysis, and there are two main categories of methods for this task: lexicon-based and machine learning-based. VADER is a well-known lexicon-based tool that provides rule-based sentiment scoring without requiring training data and is tailored for social media content. On the other hand, machine learning methods such as Logistic Regression classify sentiment by learning from labeled datasets. The advantages and disadvantages of both strategies have been shown by earlier studies. Although VADER does well in general situations, research shows that supervised models frequently produce better accuracy in domain-specific contexts, like airline reviews, where Logistic Regression has demonstrated better sentiment classification performance.[6]

Sentiment analysis has become a potent tool for gaining insights from online reviews as e-commerce has embraced Big Data. Consumer reviews on websites like Amazon and Flipkart have been found to affect purchasing decisions and assist manufacturers in evaluating the performance of their products. Reviews are frequently categorized as neutral, negative, or positive using traditional sentiment analysis. For greater understanding, recent studies have placed an emphasis on the classification of fine-grained emotions like joy, anger, and trust. Understanding customer preferences and dissatisfaction is enhanced by utilizing such multi-emotion sentiment models. This highlights the unrealized potential of online review mining and allows producers and consumers to make better informed decisions.[7]

Because Twitter can capture public opinion in real time, sentiment analysis has grown in importance as a field of study. Text pre-processing and its effect on classification accuracy receive relatively less attention than feature extraction and model development, which are the subjects of many studies. The difficulties of noisy Twitter data, such as acronyms, negations, URLs, and misspellings, have been highlighted in earlier works. According to research, specific preprocessing can greatly improve sentiment classification performance. Examples of this include negation handling and acronym expansion. Classifier sensitivity also varies; compared to Logistic Regression and SVM models, Naive Bayes and Random Forest models exhibit larger performance changes from preprocessing.[8]

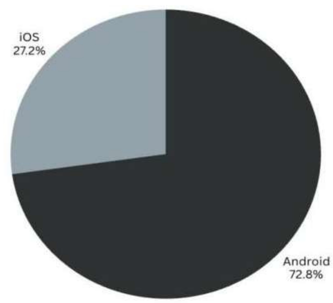

Fig 1: Global Mobile OS Market Share [13]

Because deep learning can directly learn hierarchical representations from unprocessed text, it has become increasingly popular in sentiment analysis. Twitter sentiment classification has made extensive use of Convolutional Neural Networks (CNNs), which are frequently improved by pre-trained word embeddings. Previous studies show that CNN weights can be initialized with embeddings trained on large unlabeled corpora (through unsupervised neural language models) to increase accuracy and decrease the need for handcrafted features. For fine-tuning embeddings, distant supervision—such as labeling tweets with emoticons or hashtags—provides a large amount of training data. Such pre-trained and optimized CNN models achieve state-of-the-art sentiment classification performance, as demonstrated by studies conducted on shared tasks such as SemEval.[9]

Because Twitter is widely used for opinion sharing and presents special linguistic challenges, sentiment analysis has become a hot topic of research. Because tweets are brief and frequently contain informal grammar, slang, and abbreviations, text preprocessing and feature extraction are essential. Lexicon-based methods, machine learning classifiers, and, more recently, deep learning models have all been investigated in previous research on sentiment classification. Accuracy has increased with hybrid approaches that combine sentiment lexicons, feature engineering, and preprocessing. Effective sentiment extraction and analysis from massive Twitter datasets are made possible by Python-based implementations that make use of libraries like NLTK, scikit-learn, and TextBlob.[10]

Research on business innovation frequently looks at how top tech firms create and maintain competitive advantages through innovations in their ecosystems, processes, and products. Previous research on Apple Inc. emphasizes the company's integrated strategy, which combines services, software, and hardware to provide a flawless customer experience. Apple's success has been attributed by academics to its strategic market positioning, user-centric development, and emphasis on design. Case studies also show how Apple's ecosystem encourages repeat business and brand loyalty. The literature, however, highlights the difficulties in sustaining innovation momentum in the face of market saturation and competition, highlighting the necessity of ongoing adaptation and long-term strategic vision.[11]

From straightforward lexicon-based methods to sophisticated machine learning and optimization strategies, sentiment analysis has changed over time. Prior research has shown that classifiers like Naive Bayes, Support Vector Machines, and Logistic Regression are effective at classifying customer opinions. In order to improve model performance through pipeline optimization and hyperparameter tuning, Genetic Algorithms (GA) and automated optimization tools such as TPOT have been investigated more and more. For increased accuracy, research also shows how useful it is to combine polarity-based sentiment scoring with other rating-based techniques. Combining GA optimization with conventional classifiers has demonstrated encouraging outcomes in terms of improved predictive performance, providing frameworks that can be applied to a variety of product domains.[12]

III. METHODOLOGIES

The process for performing a geographic sentiment analysis of Apple product reviews is explained in detail in this section. Data collection, preprocessing, sentiment classification using the VADER algorithm, geographic-based aggregation, and result visualization are all steps in the process. Python libraries like Pandas, NLTK, Matplotlib, Seaborn, and the vaderSentiment module were used to implement the methodology.

A. Dataset Description

The apple_reviews_by_region.csv dataset was used in this investigation. It includes user-generated text reviews of Apple products in addition to a column that displays the user's location. The reviews are based on informal, brief comments that are frequently posted on social media sites like Twitter.
The dataset was designed to mimic Twitter content's tone, brevity, and variability. The following essential characteristics are present in every record in the dataset:
review_text: The raw textual content submitted by the user
region: A categorical field indicating the user's geographic location (e.g., North America, Europe, Asia-Pacific, etc.)
Both linguistic analysis and regional sentiment comparisons are made easier by this structure. To enable statistically significant sentiment trend analysis across various regions, the dataset contains hundreds to thousands of such entries.

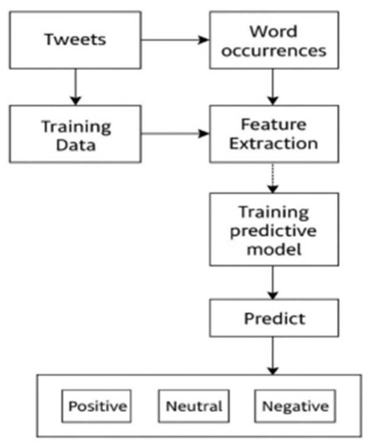

Fig 2: Flow Diagram

B. Data Preprocessing

The dataset was subjected to a thorough preprocessing pipeline to clean and normalize the review text prior to sentiment analysis. This step is crucial for enhancing analysis quality and making sure noise doesn't impede sentiment detection.

Text Cleaning Steps
The following preprocessing tasks were applied sequentially:
URL Removal: To prevent skewing sentiment scores, any hyperlinks that were in the text were eliminated using regular expressions.
Lowercasing: To ensure consistency, all text was changed to lowercase, particularly for lexicon-based sentiment analysis where case affects word matching.
Punctuation Removal: Since punctuation usually lacks emotion in brief reviews, it was removed from the text.

Stop Word Removal: The English stopword corpus from the NLTK library was used to eliminate common words like "the," "is," and "and." These words might make extra noise and don't add much to the sentiment.
Tokenization: For more detailed processing, the text was divided into discrete words, or tokens. This step serves as a prelude to optional activities like stemming and lemmatization.
Text Reconstruction: For the sentiment classifier to analyze each review, cleaned tokens were reassembled into a single string.

Cleaned Text Storage
A new column, clean_text, was added to the dataset to store the preprocessed version of each review. This preserves the original data while preparing the cleaned version for analysis.

C. Sentiment Analysis Algorithm

For sentiment classification, the VADER (Valence Aware Dictionary and Sentiment Reasoner) algorithm was chosen. VADER is a lexicon- and rule-based sentiment analysis tool created especially for social media text by Hutto and Gilbert (2014).

Justification of VADER Selection
VADER is ideal for this project for several reasons:
Handles informal language: It can identify capitalization, emojis, slang, and intensifiers (like "very good").
Lightweight and fast: It is quick and lightweight, and because it can analyze sentiment in real time, it can be integrated into web applications.

Output Metrics
For each input sentence, VADER produces four sentiment scores:
Positive (pos): Proportion of text conveying positive sentiment.
Negative (neg): Proportion of text conveying negative sentiment.
Neutral (neu): Proportion of text conveying neutral sentiment.
Compound (compound): A normalized score that goes from +1 (most positive) to -1 (most negative). When it comes to sentiment classification, this is the most important value.

Compound Score Normalization Formula
The total sentiment intensity of a sentence is combined into a normalized metric called the compound score in VADER. The sigmoid normalization function that follows is used to calculate it.

$$\text{Compound} = \frac{x}{\sqrt{x^2 + \alpha}}$$

Where:
x is the sum of all valence scores (i.e., individual positive and negative word scores in a sentence)
α is a normalization constant (default value is 15)
With the help of this function, the final compound score is guaranteed to fall between -1 (the most negative) and +1 (the most positive). The main metric for classification in this study is normalization, which facilitates the comparison of sentiment intensity across texts with different lengths and sentiment strengths.

Sentiment Classification Thresholds
The compound score was used to assign sentiment labels to each review using the following thresholds:
Positive: compound $\geq$ 0.05
Negative: compound $\leq$ -0.05
Neutral: -0.05 < compound < 0.05
Each record in the clean_text column had this logic applied to it using a Python function, creating a new column called predicted_sentiment.

D. Region-wise Sentiment Aggregation and Visualization

We categorized the sentiment results by region and produced a number of visualizations to examine the geographic variations in sentiments.

Aggregation Technique
For each region, we counted the number of positive, negative, and neutral reviews using the Pandas groupby() function. This made it easier to see how sentiment distributions vary around the world.region_
sentiment=df.groupby(['region','predicted_sentiment']).size().unstack(fill_value=0)

Visualization Tools

Grouped Bar Chart: Showed sentiment type comparison per region using Matplotlib. Each region had three bars representing the sentiment counts.

Stacked Bar Chart: Provided a proportion-based view to assess dominant sentiment type in each region.

Average Sentiment Score Plot: Using Seaborn, we plotted the average compound sentiment score per region, giving insight into overall sentiment polarity in each geographic area.

IV. RESULT ANALYSIS

The sentiment scores of Apple product reviews from six regions around the world—Africa, Asia, Australia, Europe, North America, and South America—are seen to have consistent prevalence of positive sentiment with significant regional variation in neutral and negative reviews. By applying the VADER sentiment analysis tool, each review was labeled as Positive, Neutral, or Negative, and the counts were summarized by region.

Asia had the largest percentage of positive reviews (≈273 positives) that reflect high customer satisfaction with the product quality, innovation, and post-purchase service of Apple. South America ranked second (≈260 positives), showing a very positive brand image. North America and Africa also exhibited high positive scores (≈259 and ≈256 positives, respectively) but with relatively higher neutral numbers, which may signify more moderate evaluation tendencies. Australia had 253 positive reviews along with a greater proportion of neutral sentiment (≈36) revealing mixed views potentially attributed to pricing and availability. Europe, with ≈249 positive reviews, witnessed the lowest ratio of positive sentiment, accompanied by moderate neutral (≈47) and negative (≈204) reviews, pointing towards regional take-up issues.

The Sentiment Distribution by Region bar chart (Fig. 4) identifies the dominance of positive responses across all regions, and the relative sizes of neutral and negative bars expose slight but significant variations in outlook. These findings are relevant to Apple's product and marketing teams to modify strategies for regions where neutral and negative proportions are greater.

TABLE I: SENTIMENT DISTRIBUTION BY REGION

Region	Positive	Neutral	Negative
Africa	256	48	196
Asia	273	49	178
Australia	253	36	211
Europe	249	47	204
North America	259	41	200
South America	260	46	194

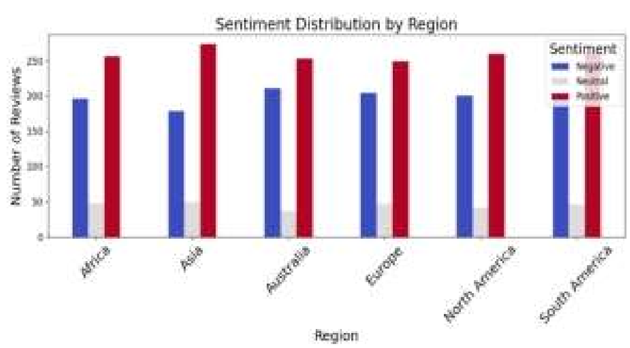

Fig. 3: Sentiment Distribution by Region

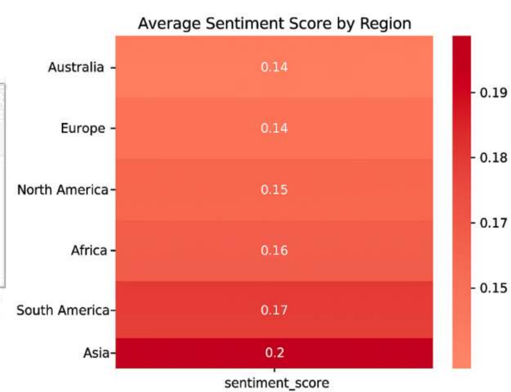

Fig. 4: Average Sentiment Score by Region

The Average Sentiment Score by Region chart offers an insight into the intensity of customer sentiment across various regions on a continuous scale, thus complementing the discrete classification view within Fig. 4. The sentiment intensity expressed in each review is measured preserving the polarity and power of that sentiment by the VADER Compound Score whose maximum and minimum values can be taken as -1 for extremely negative and +1 for extremely positive.

Asia received the highest average sentiment score across the regions, implying that it not only enjoys the greatest share of positive reviews but also a higher emotional intensity within customer response. South America stands next; its reflection suggests equally favorable attitudes toward the brand and strong brand loyalty. North America ranks third and shows a generally positive tone.

Africa takes the middle stage, with the leave-one-out score indicating a balance of enthusiasm and constructive criticism. Europe records a poor score in sentiment intensity, with less positive reviews reflecting a restrained or moderate tone on the part of customers. South America shows moderate positivity but does not record the lowest average sentiment, as its neutral proportion remains close to other regions. The data suggest that positive sentiment is dominant, but the strength of positivity varies; Asia demonstrates the most enthusiastic customer tone, followed by North America, while Australia exhibits a more polarized mix of positive and negative reviews.

V. Future Work

The current study used a lexicon-based approach to classify sentiment. To better capture contextual meaning, this work can be expanded in the future by integrating deep learning models like transformer-based architectures, BERT, or LSTMs. Furthermore, by associating opinions with particular emotions like joy, anger, or fear, incorporating emotion detection frameworks can enhance analysis. Using Twitter APIs to stream tweets in real time can yield dynamic insights into changing consumer sentiment. Lastly, organizations will find it easier to track regional brand perception if the system is implemented as a scalable web application with interactive dashboards.

VI. Conclusion

This study used the VADER sentiment analysis algorithm and a dataset modeled after Twitter to present a geographic sentiment analysis of Apple product reviews. The suggested system used a structured text preprocessing pipeline to classify opinions into three categories: positive, negative, and neutral. This classification was based on the compound sentiment score. Significant regional differences in sentiment were shown by visualizations such as average sentiment score plots and stacked and grouped bar charts. The results point to regional variations in consumer perceptions that may be impacted by user expectations, product availability, and pricing. The method shows how lexicon-based sentiment analysis can be applied for quick, comprehensible insights. Real-time data collection, emotion detection integration, and sophisticated natural language processing techniques for better contextual comprehension are examples of future improvements.

References

[1] A. Sarlan, C. Nadam and S. Basri, "Twitter sentiment analysis," Proceedings of the 6th International Conference on Information Technology and Multimedia, Putrajaya, Malaysia, 2014, pp. 212-216, doi: 10.1109/ICIMU.2014.7066632.

[2] A. Abdul Rasheed, "Cognitive Analytics of Apple Products from Crowdsourced Tweets," 2021 5th International Conference on Computing Methodologies and Communication (ICCMC), Erode, India, 2021, pp. 1543-1549, doi: 10.1109/ICCMC51019.2021.9418302.

[3] Dupinder Kaur | IJCSET(www.ijcset.net) | 2017 | Volume 7 Issue 9, 76-78

[4] Anton Borg, Martin Boldt,Using VADER sentiment and SVM for predicting customer response sentiment,Expert Systems with Applications,Volume 162,2020,113746,ISSN 0957-4174,https://doi.org/10.1016/j.eswa.2020.113746.

[5] K. M. Tymann, M. Lutz, P. Palsbröker, and C. Gips, "GerVADER - A German adaptation of the VADER sentiment analysis tool for social media texts," presented at the Lernen, Wissen, Daten, Analysen - LWDA 2019, Humboldt-Universität zu Berlin, 2019.

[6] P. Dhanalakshmi, G. A. Kumar, B. S. Satwik, K. Sreeranga, A. T. Sai and G. Jashwanth, "Sentiment Analysis Using VADER and Logistic Regression Techniques," 2023 International Conference on Intelligent Systems for Communication, IoT and Security (ICISCoIS), Coimbatore, India, 2023, pp. 139-144, doi: 10.1109/ICISCoIS56541.2023.10100565.

[7] Z. Singla, S. Randhawa and S. Jain, "Statistical and sentiment analysis of consumer product reviews," 2017 8th International Conference on Computing, Communication and Networking Technologies (ICCCNT), Delhi, India, 2017, pp. 1-6, doi: 10.1109/ICCCNT.2017.8203960.

[8] Z. Jianqiang and G. Xiaolin, "Comparison Research on Text Pre-processing Methods on Twitter Sentiment Analysis," in IEEE Access, vol. 5, pp. 2870-2879, 2017, doi: 10.1109/ACCESS.2017.2672677.

[9] Aliaksei Severyn and Alessandro Moschitti. 2015. Twitter Sentiment Analysis with Deep Convolutional Neural Networks. In Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '15). Association for Computing Machinery, New York, NY, USA, 959–962. https://doi.org/10.1145/2766462.2767830

[10] Gupta, Bhumika & Negi, Monika & Vishwakarma, Kanika & Rawat, Goldi & Badhani, Priyanka. (2017). Study of Twitter Sentiment Analysis using Machine Learning Algorithms on Python. International Journal of Computer Applications. 165. 29-34. 10.5120/ijca2017914022.
[11] M. Siahaan, “Analysis and Evaluation of the Business Innovation Strategy: A Case Study of Apple Inc.”, economy, vol. 1, no. 2, pp. 42-48, Dec. 2023.
[12] Sugunan, Sruthy. "Enhancing Apple Product Review Sentiment Analysis using Machine Learning and Genetic Algorithm Optimization." PhD diss., Dublin Business School, 2023.
[13] https://www.digitalsilk.com/digital-trends/iphone-vs-android-user-stats/