



# INTEGRATED SOCIAL MEDIA SENTIMENT ANALYSIS AND SALES FORECASTING

<sup>1</sup>K. VARSHAA, <sup>2</sup>Dr. P. EDITH LINDA, <sup>3</sup>Dr. R. KIRUTHIKA

<sup>1</sup>Student (MCA), <sup>2</sup>Associate Professor, <sup>3</sup>Assistant Professor

<sup>1</sup>Department of Computer Science,

<sup>1</sup>Dr. G.R.D College of Science, Coimbatore, India

**Abstract :** Nowadays, businesses are using social media to connect with customers and get a feel for market trends, which is pretty amazing, don't you think? Social platforms have become a huge part of shaping how people view products and make buying choices. With so many folks sharing their thoughts online, companies can dive into this treasure trove of data to understand how customers feel and how it ties to their purchasing habits. This study is all about blending sentiment analysis of social media with sales forecasting, giving businesses the tools they need to make smarter, more informed choices based on data. Social media analytics are such a fantastic way to understand what customers are feeling, what they like, and how they behave. The idea here is to combine insights from social media sentiment with sales patterns to help organizations predict market trends and make better decisions. Using Natural Language Processing (NLP) and machine learning, this system will dive into social media data to gauge customer sentiment, then connect those sentiment patterns with past sales data to build a really accurate sales forecasting model.

**Index Terms – Sentiment Analysis, Sales Forecasting, Social Media, Machine Learning, NLP, Data Visualization.**

## I. INTRODUCTION

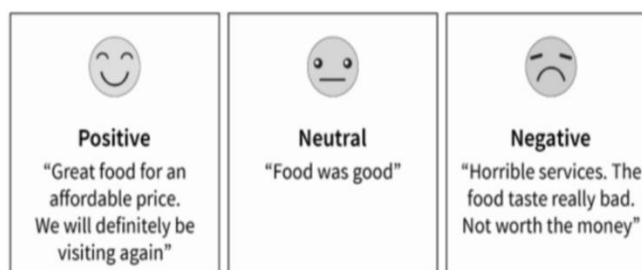
In today's digital world, businesses are leaning into social media to connect with their customers and keep an eye on market trends—it's amazing how much of an impact these platforms have on shaping opinions and influencing what people buy. With so many users sharing their thoughts online, companies can tap into that wealth of information to understand how customers feel and see how it affects sales. This research aims to combine social media sentiment analysis with sales forecasting, giving businesses the tools to make smarter, more informed decisions. Using social media analytics is such a valuable way to get a clearer picture of what customers think, what they prefer, and how they're feeling overall. This initiative merges social media sentiment analysis with sales forecasting to predict market trends and enhance decision-making for organizations. By applying Natural Language Processing (NLP) and machine learning techniques, the system will analyze social media data to gauge customer sentiment. The identified sentiment trends will be linked to historical sales data to create a reliable sales forecasting model. This project aims to deliver a sophisticated analytical tool to businesses that provides real-time insights, facilitating the optimization of marketing strategies, effective inventory management, and improved overall sales performance. The end result will be an interactive dashboard that showcases sentiment trends alongside sales forecasts, enabling the execution of data-driven business strategies.

## 1. DATASET COLLECTION

### 1.1 SENTIMENT DATASET

This dataset is created with sentiment analysis in mind, focusing on user-generated texts and a variety of metadata details. It includes everything from the actual textual content and sentiment classification to user engagement data and time-location attributes. The main highlight, "Text," holds the user-generated content, which is sorted into sentiment categories like Positive, Neutral, and Negative to make the analysis meaningful and insightful.

#### Sentiment Analysis



This dataset is packed with useful details, such as timestamps, user IDs, platforms, and hashtags, giving a rich layer of context to each entry. Plus, metrics like retweets and likes help you see how far and wide the content has reached and its overall impact. There's even geographical data, like the country of origin, combined with time-based elements like year, month, day, and hour, making it easier to uncover trends and explore how sentiments evolve over time. It's a fantastic resource for tasks like analyzing sentiment trends, studying user engagement, or pinpointing the factors that shape public opinion. Whether you're diving into social media insights, tracking brand reputation, or exploring market research, this dataset has plenty to offer.

## 1.2 SALES DATASET

To understand how markets shift, how customers think, and how businesses make money, digging into sales data is super important. This kind of dataset is like a treasure map—it has details like Transaction ID, Date, Product Category, and Product Name that help uncover sales trends across different products. Plus, it includes numbers like Units Sold, Unit Price, and Total Revenue, which are key for figuring out how much money's coming in and how profitable things are. On top of that, with info about Region and Payment Method, you can peek into where sales are happening and how people prefer to pay. This dataset is perfect for a variety of needs, like forecasting sales, analyzing demand, fine-tuning pricing strategies, and understanding customer buying habits. It's especially helpful for businesses looking to boost their sales by making smarter, data-backed decisions.

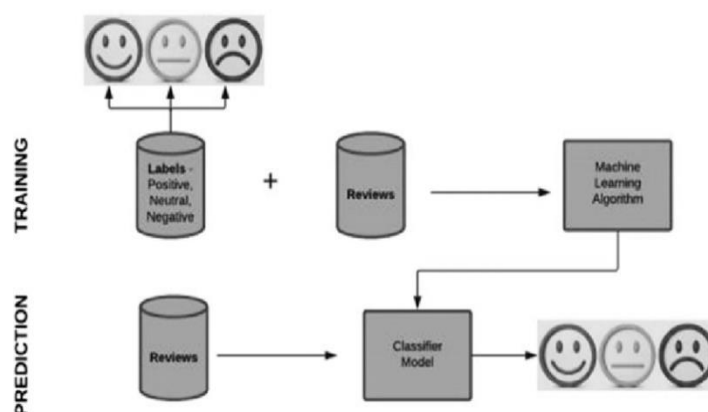
## 2. METHODOLOGY

### 2.1 DATA PREPROCESSING

Data preprocessing is a crucial step in any data-driven project, ensuring that raw data is transformed into a clean and structured format suitable for analysis. This process involves multiple steps, including data cleaning, integration, transformation, reduction, and feature selection, all of which improve data quality and enhance the accuracy of analytical models. In the context of this project, data preprocessing is applied to both sales data and sentiment data, preparing them for forecasting and sentiment analysis. The first step, data cleaning, involves handling missing values, removing duplicate records, correcting inconsistent formatting, and addressing outliers. This ensures that the dataset is complete and reliable. For sales data, missing entries in product names, sales amounts, or dates are either imputed using statistical methods or removed if necessary. In sentiment data, raw text is pre-processed using Natural Language Processing (NLP) techniques, including tokenization, stop word removal, stemming, and lemmatization, which help standardize text for further analysis. Next, data transformation involves normalizing numerical values, encoding categorical variables, and structuring text data for machine learning models. In sales forecasting, normalization techniques such as Min-Max Scaling or Z-score normalization ensure that numerical features, such as price and sales volume, are on a comparable scale. Similarly, text-based sentiment data is converted into numerical representations using techniques like TF-IDF (Term Frequency-Inverse Document Frequency) or Word Embeddings, enabling machine learning models to process textual content effectively. Another essential step is data reduction, which involves dimensionality reduction techniques like Principal Component Analysis (PCA) to eliminate redundant features while retaining critical information. This helps improve computational efficiency and reduces the risk of overfitting in predictive models. Additionally, feature selection techniques such as correlation analysis, chi-square tests, and mutual information scores are applied to retain only the most relevant variables for sales prediction and sentiment classification. Finally, data integration ensures that multiple datasets, such as sales records and customer reviews, are merged correctly for joint analysis. This step is particularly important in this project, as sentiment data needs to be aligned with sales trends over time. By integrating structured sales data with unstructured textual data, the model can analyze the relationship between customer opinions and product demand.

### 2.2 SENTIMENT ANALYSIS

Sentiment analysis is like a helpful tool for getting a clearer picture of what customers think, feel, and prefer, as well as spotting market trends—it's super valuable for this project! By diving into things like reviews, social media posts, and feedback, businesses can understand how the public feels about their products and services. In this case, sentiment analysis focuses on product-related text data to explore how customer emotions influence sales. Looking at sentiment over time helps uncover patterns in how people perceive things—for example, a spike in negative sentiment about a product might signal quality issues affecting sales, while more positive vibes could hint at growing demand. By blending sentiment analysis with sales forecasting models, businesses can get a better sense of market shifts, enhance their products, and create more focused and effective marketing plans.



The diagram represents the sentiment analysis process using a machine learning-based approach, divided into two main phases: training and prediction. In the training phase, a dataset containing user reviews is labeled with sentiment categories such as Positive, Neutral, and Negative. These labeled reviews serve as input to a machine learning algorithm, which learns patterns in the text to distinguish between different sentiment classes. The algorithm processes the textual data, extracting relevant features and associations that enable it to classify new reviews accurately. In the prediction phase, the trained model is used to analyze unseen reviews. New reviews are fed into the classifier model, which applies the learned patterns to determine the sentiment of the text. The final output categorizes the review as either positive, neutral, or negative, helping businesses understand customer opinions. In the context of this project, sentiment analysis plays a crucial role in linking customer feedback to sales performance. By analyzing sentiments from reviews, businesses can identify trends in consumer satisfaction, detect potential issues, and anticipate fluctuations in demand. This integration of sentiment analysis with sales forecasting enables data-driven decision-making, helping businesses refine their strategies and improve customer engagement.

## VADER

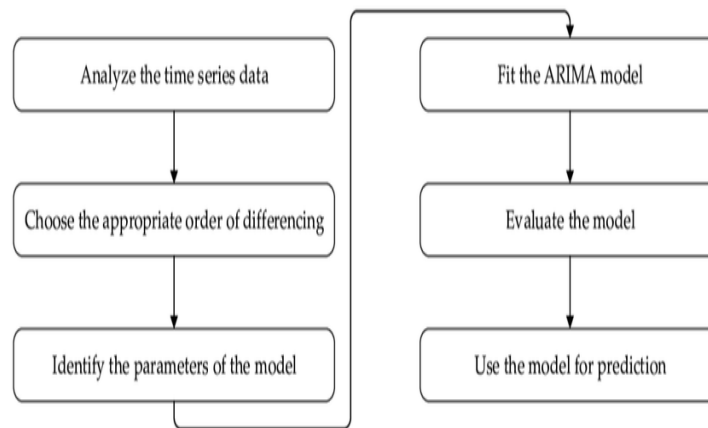
VADER (Valence Aware Dictionary and sEntiment Reasoner) is a lexicon-based sentiment analysis tool specifically designed to analyze text from social media, customer reviews, and other informal language sources. It is particularly effective in handling slang, emoticons, capitalization, and punctuation, making it well-suited for analyzing sentiment in real-world customer feedback. In this project, VADER is used to extract sentiment scores from customer reviews and opinions, which are then correlated with historical sales data to enhance sales forecasting accuracy. VADER assigns a compound sentiment score to each piece of text based on its polarity (positive, neutral, or negative). It relies on a predefined sentiment lexicon that contains words labeled with their respective sentiment intensities. The algorithm takes into account things like word modifiers, such as intensifiers and negations, along with punctuation—think multiple exclamation marks—and even capitalization, like "GREAT" versus "great," to measure how strong a sentiment is. For instance, saying "This product is amazing!!!" would score higher on positivity than simply "This product is amazing," thanks to those added exclamation marks. In this project, VADER is used to process and make sense of customer reviews by assigning each one a sentiment score, which is then grouped into positive, neutral, or negative categories based on specific benchmarks. By grouping customer sentiments over time and connecting them to sales data, we can get a clearer picture of how people feel about products and how that impacts demand. For example, a boost in positive sentiment might hint at growing interest and lead to higher sales, while an uptick in negative sentiment could point to decreasing demand. This project takes things a step further by combining VADER-based sentiment analysis with sales forecasting, going beyond just looking at past trends. It gives businesses the ability to use real-time customer feedback to make smarter, more flexible decisions about predicting demand. Being able to understand sentiment on a large scale gives businesses a great way to learn about what customers are thinking, helping them fine-tune their marketing, make their products even better, and stay ahead of changing trends.

## 2.3 SALES FORECASTING

Sales forecasting is all about anticipating future sales by looking at past trends, market movements, and other key factors. Think of it as a helpful tool for making smarter business decisions—like fine-tuning inventory, planning budgets, aligning production, or crafting spot-on marketing strategies. Getting it right can prevent hiccups like stocking too much or too little while boosting cash flow and guiding smart investments. By getting a clearer picture of future demand, businesses can make confident choices and stay ahead in the game. To tackle forecasting, you can use either qualitative insights or crunch numbers with quantitative methods—it's all about finding what works best for you. Qualitative methods focus on tapping into expert opinions, market studies, and survey insights to get a feel for future sales. Approaches like the Delphi method are all about collecting thoughtful feedback from a group of experts, refining ideas step by step. Meanwhile, quantitative methods dive into numbers, using statistical tools and machine learning to dig into past data and uncover trends. When it comes to predicting demand, there are some tried-and-true methods like time series analysis, moving averages, exponential smoothing, and regression analysis—all great tools for spotting patterns in past sales. On top of that, advanced machine learning models can take things up a notch by factoring in multiple variables and uncovering hidden connections. Of course, several things influence these forecasts, like seasonal shifts during holidays or weather changes, economic ups and downs like inflation and interest rates, and even what competitors are up to, whether it's tweaking prices or running big promotions. Additionally, marketing efforts, including advertisements and social media engagement, contribute to shifts in consumer behavior and purchasing decisions. An emerging trend in sales forecasting is the integration of sentiment analysis, which analyses customer opinions from reviews, surveys, and social media posts. By assessing sentiment trends, businesses can gauge consumer perception and anticipate demand shifts. Positive sentiments often indicate rising interest in a product, potentially leading to increased sales, while negative sentiments may signal declining demand or customer dissatisfaction. This integration enhances traditional forecasting models, allowing businesses to adjust their strategies proactively. Sales forecasting is widely applied across industries such as retail, manufacturing, finance, and healthcare, where accurate predictions drive operational efficiency and revenue growth. In retail and e-commerce, businesses use forecasting to optimize stock levels and prepare for seasonal demand fluctuations. Manufacturing industries rely on forecasts to align production schedules with market demand, reducing waste and improving efficiency. In finance, revenue predictions aid in investment planning and credit risk assessment. Similarly, the healthcare sector uses forecasting to anticipate demand for medical supplies, ensuring adequate inventory levels.

## ARMIA

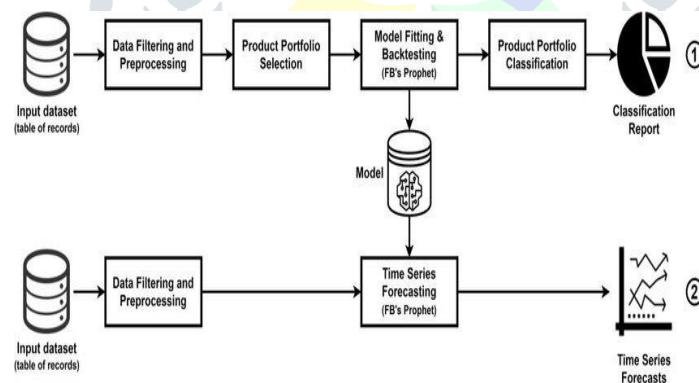
ARIMA, which stands for AutoRegressive Integrated Moving Average, is such a fantastic statistical tool for forecasting time series data! In this project, it's especially handy for predicting future sales trends. By analyzing past sales, ARIMA uncovers patterns like trends, seasonal shifts, and other changes, helping businesses gain insights that make it easier to manage inventory, plan for demand, and adjust marketing strategies effectively.



The ARIMA model is defined by three key components: the AutoRegressive (AR) component (p), which captures the relationship between past and present sales data; the Integrated (I) component (d), which ensures the data is stationary by differencing; and the Moving Average (MA) component (q), which accounts for past forecast errors to smooth out fluctuations. The ARIMA model is expressed as ARIMA (p, d, q), where the parameters are selected based on historical data characteristics. In this project, ARIMA is applied to forecast sales by first preprocessing and cleaning the sales data, ensuring that missing values are handled and the time series is formatted correctly. The Augmented Dickey-Fuller (ADF) test is then used to check for stationarity, and if necessary, differencing is applied to stabilize the data. The values of p, d, and q are determined using Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots, ensuring that the optimal model configuration is selected based on criteria such as the Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC). Once trained on historical sales data, the ARIMA model generates future sales predictions, which are evaluated using error metrics such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE) to ensure accuracy.

## PROPHET

Prophet is an advanced time series forecasting model developed by Facebook (Meta), designed to handle complex seasonal patterns and trends in data. It is particularly effective for business applications such as sales forecasting, making it a valuable tool in this project. Prophet follows an additive model approach, where trends, seasonality, and holiday effects are modeled separately and combined to generate accurate predictions. Unlike traditional models such as ARIMA, Prophet is highly flexible and capable of handling missing data, outliers, and irregular time series patterns, making it well-suited for real-world business scenarios. This project utilizes Prophet to analyze historical sales data and generate forecasts of future sales trends.



The model decomposes the time series into three main components: trend, seasonality, and holidays/events. The trend component captures long-term growth patterns in sales, while the seasonality component identifies recurring patterns, such as weekly, monthly, or yearly fluctuations. The holiday effect accounts for special occasions or promotional events that influence sales. Prophet's ability to automatically detect and adjust for these components enhances the accuracy of its predictions. The implementation of Prophet begins with thorough data preprocessing, which includes addressing missing values, identifying and rectifying anomalies within the time series, and ensuring the dataset is structured to meet specified requirements. The framework necessitates the dataset to include two primary columns: 'ds' for date/time and 'y' representing the target variable, such as sales volume. Once the data is ready, the model undergoes training to examine historical sales patterns and generate future forecasts. Moreover, Prophet includes built-in capabilities for integrating external regressors, such as customer sentiment metrics, to improve the accuracy of sales predictions. A key advantage of Prophet is its ability to automatically detect changepoints, allowing it to adjust for sudden shifts in trends, such as market fluctuations or demand spikes. Additionally, it offers hyperparameter tuning options to control growth rate, seasonal components, and uncertainty intervals, ensuring robust and reliable forecasts. The performance of the Prophet model evaluation using error metrics such as Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE) to assess its accuracy.

## LSTM

Long Short-Term Memory (LSTM) is an advanced variance of Recurrent Neural Network (RNN) designed specifically for time series prediction and sequential data analysis. In contrast to conventional RNNs, which often suffer from the vanishing gradient issue, LSTMs effectively capture long-term dependencies thanks to their distinctive memory cell architecture. This feature renders LSTM a robust model for predicting future sales based on historical data, which is essential for this project. In this initiative, LSTM

is employed for sales forecasting by examining historical sales figures and recognizing patterns over time. The model excels at identifying seasonal variations, abrupt increases in sales, and long-term trends. Unlike statistical models such as ARIMA and Prophet, which depend on established mathematical frameworks, LSTMs derive patterns directly from the data, making them highly responsive to fluctuating market dynamics. The key strength of LSTM lies in its capacity to retain relevant past data while discarding extraneous information through its gating mechanisms, which include the input gate, forget gate, and output gate. These gates regulate the append and elimination of historical data, allowing LSTM models to focus on the most critical patterns efficiently. The application of LSTM in this project commences with data preprocessing. Historical sales data undergoes cleaning, normalization, and transformation into a format suitable for deep learning models. Given that LSTMs require input sequences, the data is reshaped into time steps to facilitate effective learning of temporal dependencies. The LSTM model architecture consists of multiple LSTM layers, dense (fully connected) layers, and various activation functions to improve learning efficiency. The model is trained through backpropagation through time (BPTT) and optimized using advanced algorithms, including Adam and RMSprop, to minimize prediction errors effectively. Standard error metrics, including Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE), were employed to evaluate the performance of the LSTM model.

### 3. SYSTEM IMPLEMENTATION

The implementation of this project follows a structured approach that integrates sentiment analysis and sales forecasting to provide valuable business insights. The process begins with data collection, which includes historical sales records sourced from databases, Excel spreadsheets, or APIs. Customer sentiment information is gathered through online reviews, feedback submissions, and social media channels. After data collection, it is subjected to preprocessing procedures designed to ensure its quality and maintain consistency. Sales data is cleaned by addressing missing values, removing outliers, and converting it into a time-series format suitable for forecasting models. Sentiment data is processed through text cleaning, tokenization, and lemmatization. Sentiment scores are then extracted using the VADER sentiment analysis tool, which categorizes them as positive, negative, or neutral. These scores are merged with the sales data to establish a correlation between customer sentiment and sales trends. After preprocessing, exploratory data analysis (EDA) is conducted to visualize sales patterns, seasonal trends, and sentiment distributions. Stationarity checks, such as the Augmented Dickey-Fuller (ADF) test, help determine whether time-series transformations are necessary. With a clear understanding of the data patterns, different forecasting models are trained. The ARIMA model is applied for statistical time-series forecasting, with optimal parameters selected using Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) plots. The Prophet model, developed by Facebook (Meta), is employed to capture long-term trends and seasonality while adjusting for external factors such as holidays and promotions. A Long Short-Term Memory (LSTM) model, which is a deep learning-based approach, is also trained on historical sales data to learn complex non-linear relationships and long-term dependencies. The sentiment scores are integrated into these forecasting models as an additional feature to analyze their impact on future sales. To evaluate the performance of the models, various error metrics, including Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), and Mean Absolute Percentage Error (MAPE), are used to measure prediction accuracy. The model with the best performance is then selected for deployment. The final dashboard provides businesses with actionable insights by displaying interactive sales forecasts, sentiment impact analysis, and trend comparisons. Users can adjust parameters, filter data, and explore various visualization options in real-time. This implementation empowers businesses with data-driven decision-making, allowing them to optimize inventory management, marketing strategies, and pricing models based on forecasted trends. The integration of Plotly Dash enhances the usability of the system, making it an intuitive and powerful tool for business intelligence and predictive analytics.

#### 3.1 Results and Discussions

This project underscores the significance of sentiment analysis in enhancing sales forecasting accuracy. A systematic methodology was employed to rigorously assess various machine learning and deep learning models, aiming to determine their effectiveness in forecasting future sales trends through sentiment score analysis.

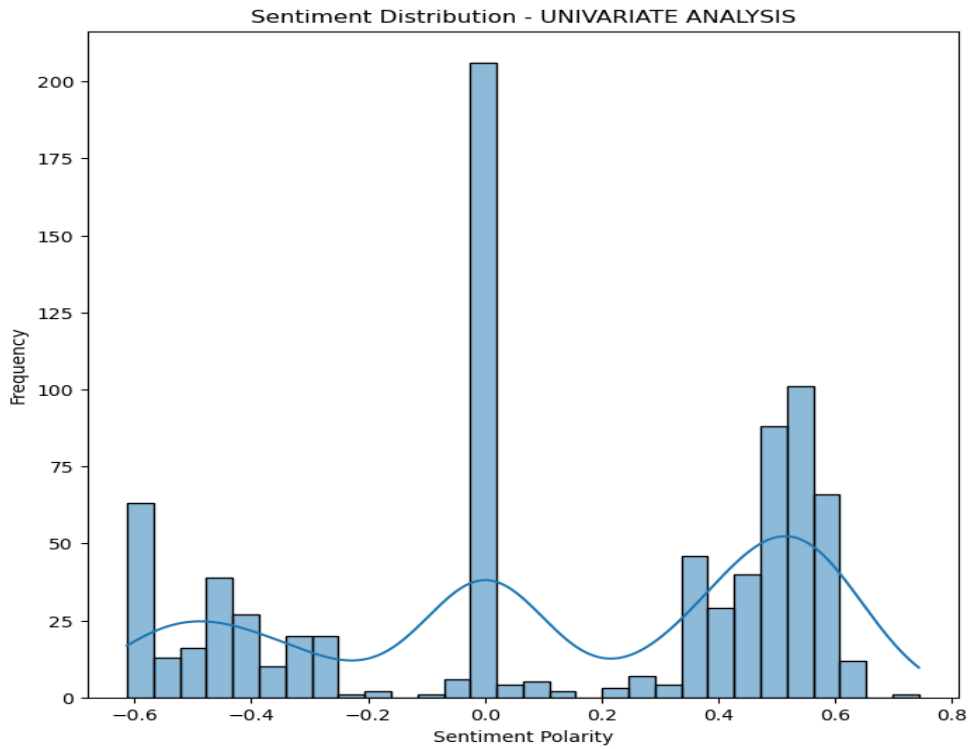


Figure 1: Sentiment Distribution

The sentiment analysis findings indicate a polarization in customer opinions, characterized by a noticeable decline in neutral responses and an increase in extreme sentiment expressions. The high frequency of neutral sentiments might be due to ambiguous or factual statements that do not carry strong emotions. The insights gained from this analysis can help businesses better understand customer sentiment trends, identify areas of improvement, and make informed decisions in sales and marketing strategies. To gain deeper insights, further analysis can be performed by segmenting the data based on product categories, customer demographics, or time-based trends.

Table 1: Comparison

	MODEL	RMSE	MAE	R <sup>2</sup> SCORE
0	ARIMA	0.462505	0.365412	0.732612
1	PROPHET	0.822168	0.762045	0.155049
2	LSTM	0.307053	0.304933	0.858577

The table outlines the evaluation metrics, including Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and R<sup>2</sup> Score, for three forecasting models ARIMA, Prophet, and LSTM. These methodologies are employed to rigorously evaluate the ability of each model to accurately predict future sales trends and the efficient operation of business processes undergo thorough and critical evaluation.

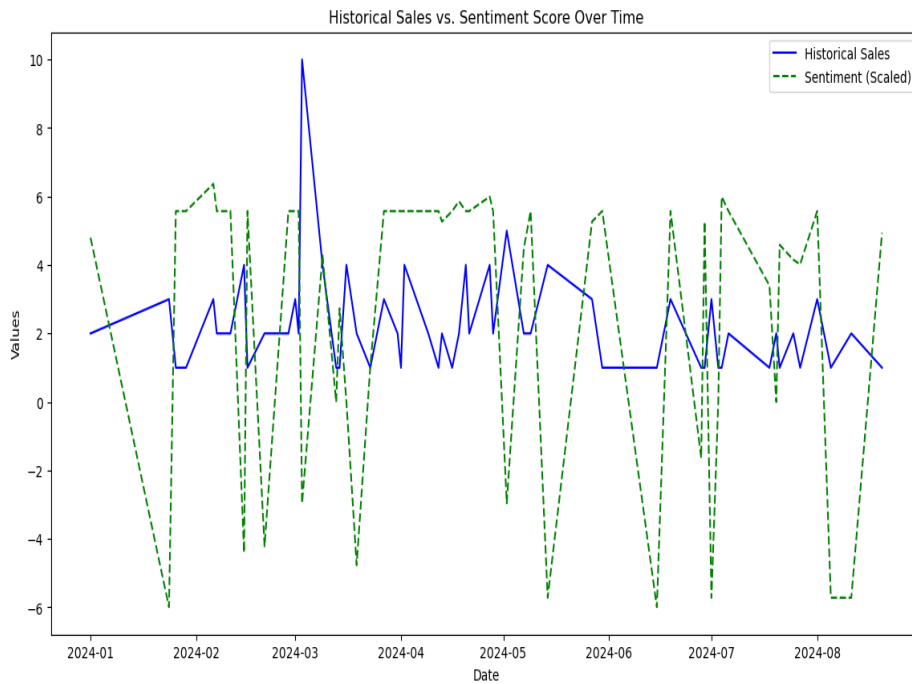


Figure 2: Historical Sales vs Sentiment Score

The "Historical Sales vs. Sentiment Score Over Time" graph provides a detailed representation of the correlation between sales performance and sentiment scores across a specified timeframe. In the visualization, the blue solid line illustrates the historical sales data and the green dashed line depicts sentiment scores data that has been adjusted to enable a more accurate and effective comparison.

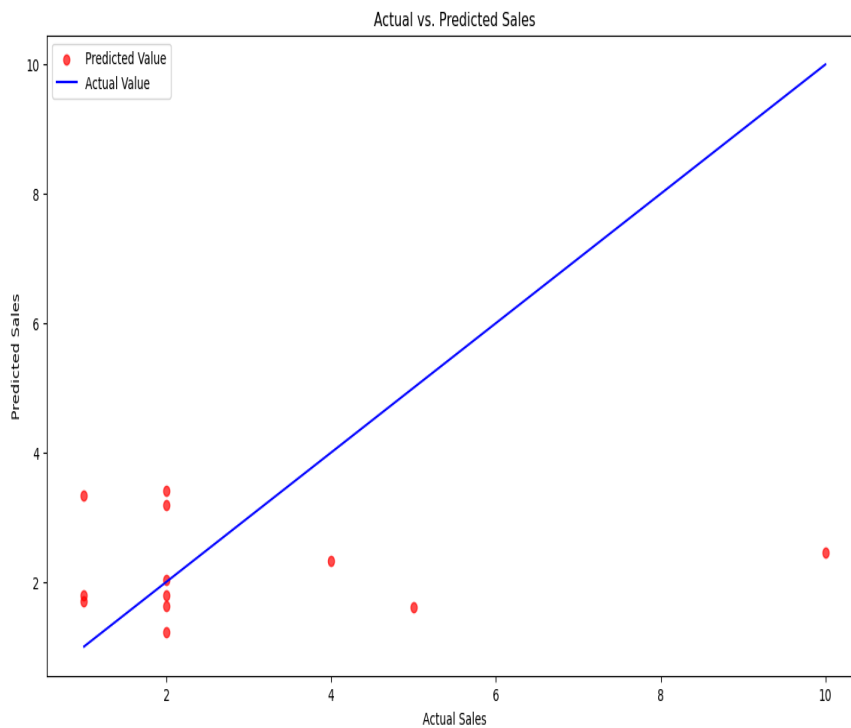


Figure 3: Actual vs Predicted Sales

The Actual versus Predicted Sales graph offers a professional analysis, contrasting forecasted sales projections with actual performance metrics. Representing predicted sales figures, the red dots are plotted against the blue line, which depicts actual sales results. Accurate forecasting is indicated when the predicted values align closely with the blue diagonal line, reflecting a strong correspondence between projected and realized sales outcomes.

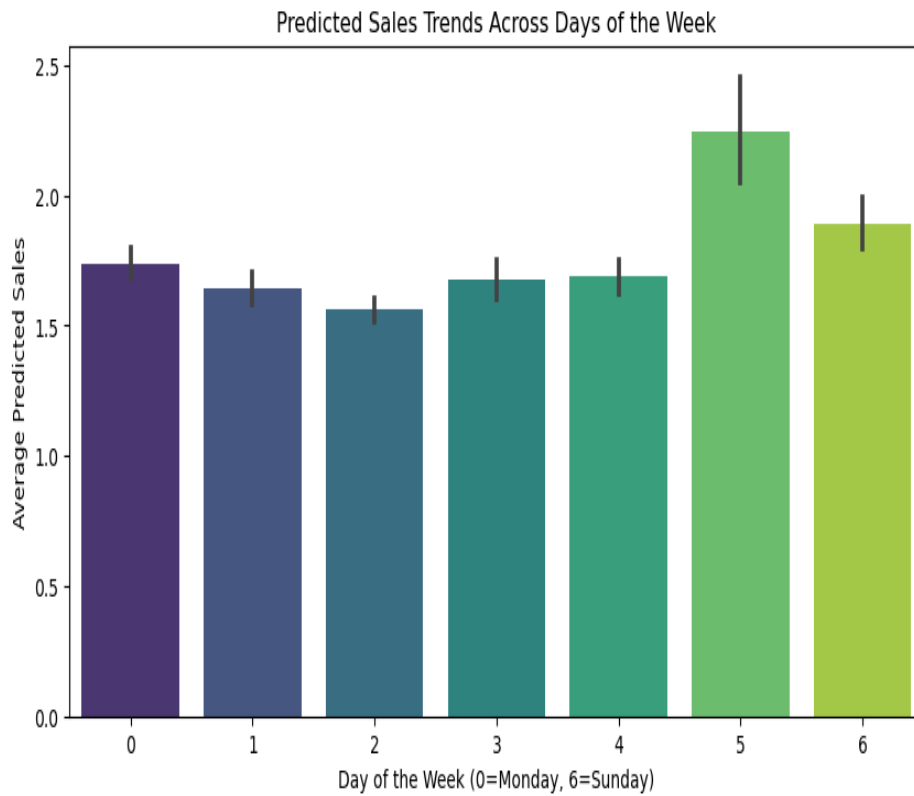


Figure 4: Weekly Distribution of Predicted Sales

The "Predicted Sales Trends Across Days of the Week" graph presents the average forecasted sales for each day, using 0 to represent Monday and 6 to represent Sunday. The height of each bar reflects the predicted sales values, while the accompanying error bars denote the range of variability in these predictions.

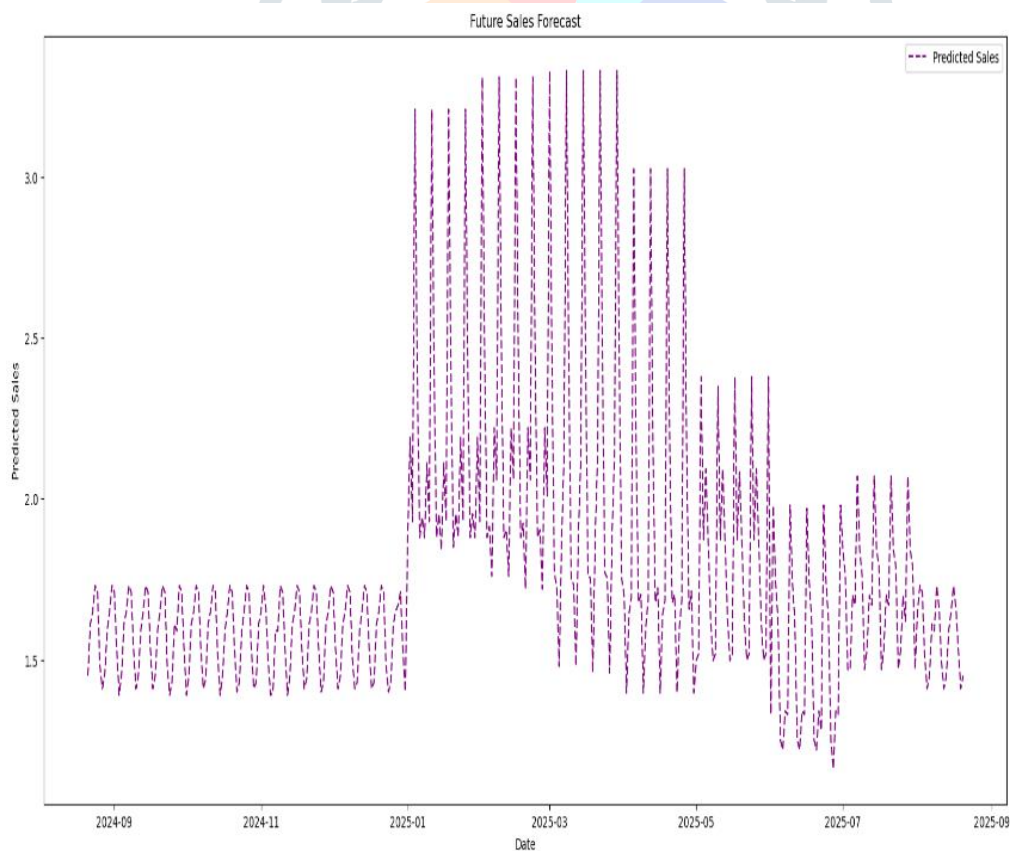


Figure 5: Future Sales Forecast

The Future Sales Forecast graph displays the predicted sales trends over an extended period, showing fluctuations in sales patterns. The forecasted data presents noticeable periodic variations, suggesting seasonal trends in sales activity.



Figure 6: Realtime Sentiment and Sales Dashboard

The Real-Time Sentiment and Sales Dashboard combines sentiment analysis with sales data to deliver actionable insights. By visualizing sentiment trends alongside sales performance, it facilitates informed, data-driven decisions and provides a comprehensive understanding of how customer sentiment influences sales dynamics.

**4. FUTURE SCOPE**

This project has significant potential for further enhancements and real-world applications. A key focus for improvement lies in the integration of real-time data streaming to enable continuous updates to forecasts, thereby enhancing the dynamism and

responsiveness of predictions about shifting market conditions. Advancing this further, the adoption of sophisticated deep learning models, such as Transformer-based architectures, holds potential for achieving greater accuracy. Expanding datasets through the inclusion of external variables—such as economic indicators, competitor pricing, and seasonal trends—can yield more comprehensive insights into sales patterns. Additionally, the development of an interactive dashboard using Plotly Dash offers a practical solution for real-time monitoring and predictive analytics, empowering businesses to make informed and efficient data-driven decisions. Future research avenues may include the creation of automated recommendation systems designed to propose optimal sales strategies based on anticipated trends, thereby fostering strategic alignment and improved performance. Overall, the scalability and adaptability of this project make it a valuable tool for businesses aiming to optimize sales performance and strategy.

## 5. CONCLUSION

This initiative adeptly combines sentiment analysis with sales forecasting to deliver a thorough evaluation of the correlation between social media sentiment and sales performance. Utilizing sophisticated machine learning models, including ARIMA, Prophet, and LSTM, it accurately projects future sales trends by integrating historical data and sentiment insights. The findings confirm that deep learning models, notably LSTM, surpass conventional statistical approaches, providing superior forecasting precision. Additionally, the integration of Plotly Dash augments the project by presenting an interactive and intuitive dashboard for effective data visualization and analysis. The results underscore the significant influence of consumer sentiment on sales, stressing the critical role of real-time data analysis in informed business decision-making. By incorporating advancements like real-time streaming data and additional external variables, this system has the potential to become an invaluable decision-support tool for businesses aiming to enhance their sales strategies.

## REFERENCES

1. Dash, P., & Dash, R. (2020). Real-Time Data Visualization Using Plotly Dash. *Springer Nature*.
2. Chakraborty, I., & Ghosh, S. (2021). Impact of Social Media on Sales Prediction Using Machine Learning. *Journal of Retailing and Consumer Services*, 61, 102570.
3. Zhang, T., & Liu, X. (2022). Sentiment Analysis-Based Sales Forecasting Using Machine Learning Techniques. *International Journal of Data Science and Analytics*, 14(2), 215–230.
4. Brown, C., & Gupta, R. (2023). Enhancing ARIMA-Based Time Series Forecasting with Social Media Sentiment Analysis. *IEEE Transactions on Computational Social Systems*, 10(1), 45–58.
5. Kim, D., & Park, H. (2024). Leveraging LSTM for Predicting Sales Trends from Social Media Sentiments. *Neural Computing and Applications*, 36(4), 1256–1270.
6. Wang, Y., & Lee, J. (2022). A Hybrid Prophet-LSTM Model for Sales Forecasting Using Online Reviews. *Expert Systems with Applications*, 196, 116540.
7. Kumar, S., & Singh, P. (2023). Sentiment-Driven Demand Forecasting Using VADER and Deep Learning. *Applied Artificial Intelligence*, 37(2), 89–108.
8. Taylor, S. J., & Letham, B. (2020). Advancements in Forecasting with Prophet: Applications and Performance Analysis. *Journal of Business Analytics*, 5(3), 127–145.
9. Chen, X., & Zhao, L. (2021). Predicting Future Sales Trends with Social Media Data and Machine Learning Algorithms. *Data Mining and Knowledge Discovery*, 35(1), 221–245.