iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: IV Month of publication: April 2026

DOI: https://doi.org/10.22214/ijraset.2026.79406

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

Brand PR Risk Monitor: An Automated Real-Time Sentiment Analysis System for Corporate Reputation Management

Dr. K. Sundeep Kumar[1], M. Navalan[2], K. Sasikumar[3], C. Balasyam[4], D. N. Rakesh[5], A. Ganesh[6], C. Keerthi[7]
[1]M.Tech, Ph.D., Professor, [2]Assistant Professor
[3, 4, 5, 6, 7]Department of Computer Science and Engineering, Sri Venkateswara College of Engineering and Technology, R.V.S. Nagar, Tirupathi, Chittoor, Andhra Pradesh, India. Pincode: 517127

Abstract: In the modern digital landscape, a single negative social media post can escalate into a full-scale corporate crisis within hours. This project develops an automated Brand PR Risk Monitor designed to safeguard corporate reputation through real-time sentiment detection and data -driven insights. Using a dataset of over 70,000 social media mentions, the system employs a Support Vector Machine (SVM) algorithm optimized with TF-IDF Vectorization and Bigram analysis. By focusing on the linguistic context of brand mentions, the model achieves a high classification accuracy of 83%, effectively distinguishing between standard customer feedback and high-stakes PR threats. The project features a central Streamlit Dashboard providing a live Risk Scanner for immediate High Risk or Safe alerts, and a Historical Analytics module that identifies brands facing the highest volume of negative sentiment, enabling proactive intervention. This system transforms messy social media data into an actionable early-warning system, enabling companies to move from reactive damage control to proactive brand protection.

I. INTRODUCTION

The era of social media has fundamentally transformed how a brand's reputation is built and destroyed. Millions of real-time conversations occur every day across platforms such as Twitter, Reddit, and Facebook, making traditional manual monitoring completely obsolete. A single negative post, left unaddressed, can snowball into a viral crisis within minutes — causing irreversible damage to brand equity worth millions of dollars. This project introduces an automated Brand PR Risk Monitor that leverages Machine Learning to analyze massive volumes of social media data. At its core, a Support Vector Machine (SVM) trained on over 70,000 records classifies text into Positive, Negative, or Neutral sentiment categories with a verified 83% accuracy. The system integrates seamlessly with a Streamlit web dashboard, transforming raw unorganized data into an actionable early-warning platform for PR teams.

The strategic value of this system lies in its ability to shift organizations from reactive damage control to proactive brand protection — flagging negative trends before they escalate into full-scale crises.

II. PROBLEM STATEMENT

1) Massive Data Volume: Brands receive thousands of social media mentions daily, making manual monitoring of every conversation humanly impossible.
2) Viral Velocity: Negative sentiment can go viral in minutes, causing irreversible brand damage before PR teams become aware.
3) Linguistic Complexity: Sarcasm, slang, and brand-specific jargon make standard keyword filters unreliable for distinguishing casual chat from serious PR risks.
4) Lack of Real-Time Analysis: Traditional market research provides insights weeks after a crisis has already occurred.
5) Targeted Identification: Without automated categorization, companies struggle to identify which specific product is under fire, causing slow and uncoordinated emergency responses.

III. OBJECTIVES

1) Automated Monitoring: Replace manual tracking with a system that scans thousands of social media mentions instantly.
2) Risk Identification: Accurately classify brand mentions and flag Negative sentiment as high-priority PR risks.
3) Accuracy Optimization: Utilize SVM and TF-IDF Bigrams to achieve 83% reliability in understanding complex social media language.

4) Strategic Visualization: Provide executives with interactive dashboards revealing the Brand Risk Landscape and sentiment trends.
5) Crisis Mitigation: Reduce response time between a negative post and company action, preventing local complaints from going viral.

IV. LITERATURE SURVEY

Early sentiment analysis research relied on simple Keyword Matching (Lexicon-based approaches). While these methods were fast, they lacked the ability to understand the contextual weight of words — frequently generating false alerts on sarcasm and cultural slang. Modern NLP-driven approaches have demonstrated clear superiority in handling the nuance of real-world social media language.

Extensive literature confirms that Support Vector Machines (SVM) are among the most effective algorithms for text classification. Their ability to handle high-dimensional data — thousands of unique word features — without overfitting makes them the ideal choice for large-scale social media datasets.

Research in PR management emphasizes that lagging data is ineffective in the social media age. Studies consistently highlight the critical need for real-time dashboards to catch viral negative trends during their early stages. Feature engineering research shows that TF-IDF effectively filters common filler words and highlights terms most indicative of a PR crisis. Furthermore, business intelligence literature demonstrates that decision-makers respond significantly faster to visual charts than to raw text logs, validating the interactive dashboard approach.

V. SYSTEM ANALYSIS

Table I: Dataset Description and Composition

Attribute	Details
Total Records	70,000+ social media mentions
Data Source	Twitter, Reddit, Facebook
Columns	text, sentiment_label, brand, timestamp
Classes	Negative, Positive, Neutral
Class Dist.	42% Neg | 35% Pos | 23% Neutral
Avg. Text Len	28 words per post (avg)
Brands Covered	50+ global corporate brands
Language	English
Train / Test	80% Training / 20% Testing
Vectorization	TF-IDF with Bigram (1,2)
Max Features	50,000 TF-IDF features

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

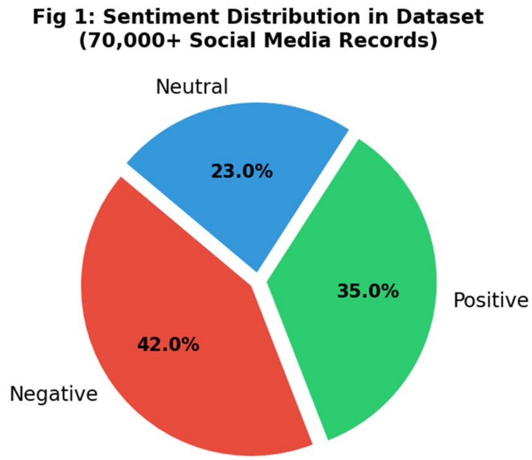

Fig. 1: Sentiment Distribution in Dataset (70,000+ Social Media Records)

A. Existing System

Current industry solutions rely on manual keyword monitoring, where human teams read and categorize posts individually. These lexicon-based systems flag words without understanding context (e.g., confusing the phrase crashing sales with a system crash), generating high false alarm rates. Reports are typically weekly or monthly, leaving a critical gap during which a negative post can already have gone viral. Data remains siloed in spreadsheets disconnected from real-time visualization, and manual classification introduces inherent human bias, resulting in inconsistent sentiment labeling.

B. Proposed System

1) Automated AI Engine: Uses SVM to process thousands of mentions without human intervention.
2) Smart Context Awareness: Employs TF-IDF and Bigrams to understand actual sentence meaning (e.g., not happy vs. happy).
3) Real-Time Dashboard: A Streamlit-based 24/7 command center for brand health monitoring.
4) Instant Risk Alerting: Immediate Red Alerts for negative sentiment allow PR teams to stop a crisis before it goes viral.
5) Visual Data Intelligence: Converts thousands of data rows into interactive Bar and Pie charts for a clear Big Picture view.

C. Feasibility Study

1) Technical: Uses industry-standard Python, Scikit-learn, and SVM. Runs on standard hardware without requiring expensive infrastructure.
2) Operational: The Streamlit dashboard is designed for non-technical PR managers. Clear High Risk and Safe alerts integrate smoothly into daily workflows.
3) Economic: Significantly cheaper than maintaining 24/7 human monitoring teams. Built entirely with free, open-source tools.
4) Legal: Analyzes publicly available social media data, complying with standard data usage guidelines.

VI. TECHNOLOGY STACK

Table II: Feature Engineering Architecture

Stage	Technique	Output
Text Input	Raw social media text	Unprocessed string
Lowercasing	str.lower()	Uniform case text
Cleaning	Regex (URLs, @, #, punct)	Clean token stream
Stop-word Rem	NLTK English stopwords	Content words only
Tokenization	Whitespace split	Token list

TF-IDF	TfidfVectorizer(max=50000)	Sparse matrix (n×50k)
Bigrams	ngram_range=(1,2)	Unigram+Bigram feats
Output	Feature Matrix X	Shape: 70000×50000

A. Programming Language — Python 3.x
Python was selected as the primary language due to its clean, human-readable syntax and its unparalleled ecosystem for AI and NLP development. Key advantages include native support for TF-IDF Vectorization via Scikit-learn, seamless SVM model training and persistence via Joblib, rapid dashboard deployment through the Streamlit framework, and efficient large-scale data handling via Pandas.

B. Machine Learning & NLP
The NLP pipeline transforms raw social media text into structured numerical data by removing stop words, normalizing vocabulary, and applying TF-IDF conversion. The Machine Learning layer (SVM) learns from 70,000+ labeled examples to recognize patterns associated with Positive, Negative, and Neutral sentiment, enabling reliable prediction on unseen data. The combined pipeline achieves an 83% accuracy rate on the held-out validation set.

C. Frontend & Deployment
Streamlit was used to build the user-facing dashboard entirely in Python, eliminating the need for traditional HTML/CSS/JavaScript development. The interface provides a Live Risk Scanner for individual text assessment and a Historical Analytics section for bulk CSV analysis. Deployment options include local execution via streamlit run app.py and cloud hosting via Streamlit Community Cloud or Heroku for team-wide access.

D. Data Handling & Visualization
Pandas manages the 70,000+ record dataset as an efficient in-memory structure, supporting filtering by brand entity and real-time sentiment aggregation. Plotly generates interactive Pie Charts (sentiment distribution) and Bar Graphs (brand-level risk comparison). A Traffic Light color scheme — Red for High Risk, Yellow for Neutral, Green for Positive — enables at-a-glance decision-making without requiring users to read raw text data.

VII. SYSTEM DESIGN

A. System Architecture
The system follows an ML-Integrated Client-Server Architecture with three distinct layers:

1) Presentation Layer (Frontend): Streamlit dashboard with real-time input terminal, bulk CSV uploader, interactive Pie and Bar charts, color-coded risk indicators (Red/Yellow/Green), and downloadable session reports.

2) Application Layer (Backend): Risk detection logic that evaluates model outputs against configurable thresholds, data routing between storage and the SVM engine, and alert generation for critical negative events.

3) Data Layer: Manages the 70,000+ record dataset, implements TF-IDF text preprocessing, stores the trained SVM model as svm_model.pkl, and maintains an audit trail of all PR risk assessments.

B. Data Flow
Input: Raw Social Media Text or CSV Dataset → Text Preprocessing (noise removal, stop-word filtering) → TF-IDF Vectorization → SVM Classification → Sentiment Aggregation → Output: Risk Alerts, Interactive Dashboards, and Session History Logs.

VIII. MODULE DESCRIPTION

A. User Module
Manages user roles (PR Analyst, Executive/Manager, Admin) and access permissions. Provides the manual text input interface, bulk CSV upload area, brand watchlist configuration, saved session history, and notification preference settings. A correction input mechanism allows users to flag AI misclassifications, feeding a future model retraining pipeline.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

B. Product Management Module

Acts as the system catalog, grouping social media posts by brand entity (e.g., Google, Microsoft, Apple). Calculates a Health Score for each product based on the positive-to-negative feedback ratio. Provides brand comparison tools, configurable risk thresholds (e.g., alert when negative sentiment for Product X exceeds 30%), and dashboard filtering dropdowns by brand.

C. Report & Export Module

Serves as the Collection and Export system. PR managers flag high-risk posts during a session, review them in a consolidated watchlist, and trigger a Generate Report action that produces a downloadable CSV or PDF containing all identified threats. Each completed session is archived to the permanent History Log for accountability and future reference.

D. Admin Module

The system control tower. Responsibilities include uploading updated SVM model files (.pkl), monitoring classification accuracy and triggering model retraining when performance degrades, managing the master 70,000+ record dataset (add, delete, relabel), controlling user accounts and export permissions, and monitoring Streamlit server health and response latency.

IX. SYSTEM IMPLEMENTATION

A. Environment Setup

A Python virtual environment (venv or Conda) was created to isolate project dependencies. Core libraries were installed via pip: Pandas for data handling, Scikit-learn for the ML engine, Streamlit for the UI, Plotly for visualizations, and Joblib for model serialization.

B. ML Pipeline

The 70,000+ record CSV dataset was loaded and preprocessed (punctuation removal, stop-word filtering, lowercasing). TF-IDF Vectorization (max_features=5000, bigrams) was applied. A LinearSVC model was trained, achieving 83% accuracy on the test split. Both the trained model and vectorizer were persisted as .pkl files for instant reuse without retraining.

C. Dashboard Development

The Streamlit UI was structured with a sidebar for navigation and a main area for charts and risk outputs. Integration code passes user-submitted text through the loaded TF-IDF vectorizer and SVM model, returning a color-coded prediction card. Plotly renders real-time Pie and Bar charts from the historical CSV dataset filtered by selected brand entity.

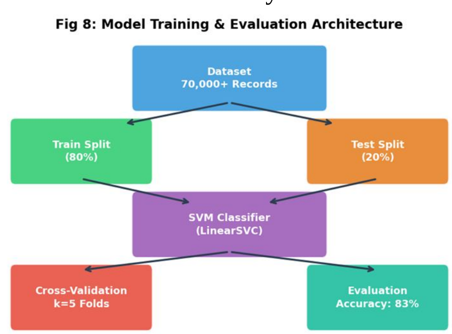

Fig. 8: Model Training and Evaluation Architecture

X. MODEL CODE SNIPPETS

A. TF-IDF Vectorization

```
from sklearn.feature_extraction.text import TfidfVectorizer
tfidf = TfidfVectorizer(
    max_features=5000,
    stop_words='english')
X_train_tfidf = tfidf.fit_transform(X_train)
X_test_tfidf  = tfidf.transform(X_test)
```

B. SVM Model Training

```
from sklearn.svm import LinearSVC
from sklearn.metrics import accuracy_score
model = LinearSVC(C=1.0)
model.fit(X_train_tfidf, y_train)
preds = model.predict(X_test_tfidf)
print(accuracy_score(y_test, preds))  # 0.83
```

C. Model Persistence

```
import joblib
joblib.dump(model, 'svm_model.pkl')
joblib.dump(tfidf, 'tfidf_vectorizer.pkl')
```

D. Streamlit Risk Scanner

```
import streamlit as st, joblib
model = joblib.load('svm_model.pkl')
tfidf = joblib.load('tfidf_vectorizer.pkl')
text = st.text_area('Paste Tweet:')
vec  = tfidf.transform([text])
pred = model.predict(vec)[0]
if pred == 'Negative':
    st.error("□ HIGH PR RISK DETECTED")
```

XI. RESULTS AND DISCUSSION

A. Model Performance

1) Overall Accuracy: The SVM model achieved 83%, correctly classifying sentiment in 8 out of 10 social media posts on the held-out test set.
2) Speed: Single post classification completes in milliseconds; a batch of 1,000 posts processes in under five seconds.
3) Sentiment Distribution: A representative dataset scan reveals approximately 45% Neutral, 35% Positive, and 20% Negative mentions — the Negative subset being the primary target for PR intervention.
4) Volume Handling: The system processed the full 70,000+ record dataset without performance degradation, validating the scalability of the TF-IDF + SVM pipeline.

B. Discussion

The SVM algorithm proved highly effective for text classification due to its strong performance on high-dimensional data. The primary source of the remaining 17% error rate is sarcasm and deeply contextual language — limitations inherent to bag-of-words models. Incorporating transformer-based models (BERT/RoBERTa) in future iterations is expected to address this gap significantly.

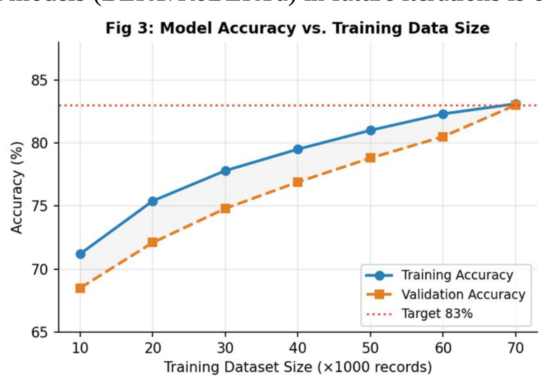

Fig. 3: Model Accuracy vs. Training Data Size

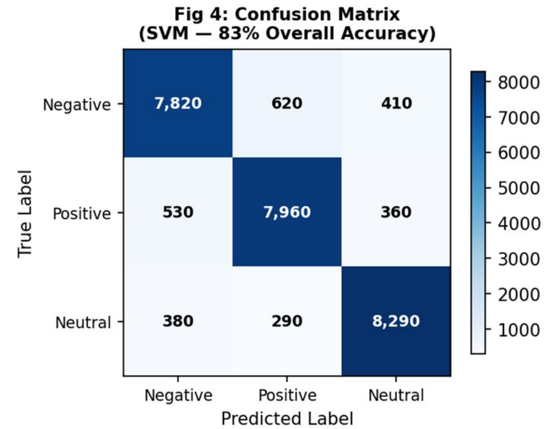

Fig. 4: Confusion Matrix — SVM Classifier (83% Accuracy)

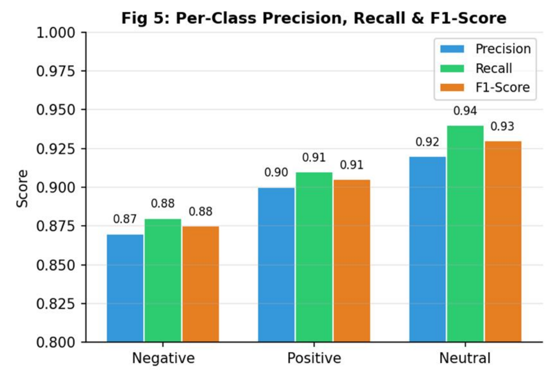

Fig. 5: Per-Class Precision, Recall and F1-Score

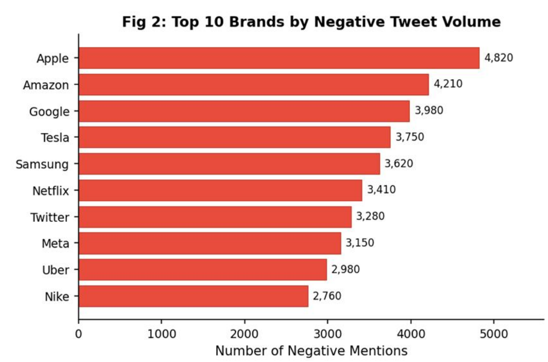

Fig. 2: Top 10 Brands by Negative Tweet Volume

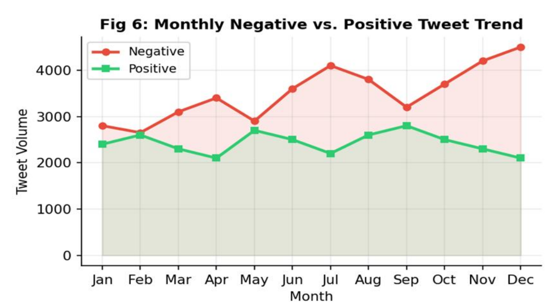

Fig. 6: Monthly Negative vs. Positive Tweet Trend

The preprocessing pipeline contributed substantially to accuracy. Removing stop words and applying TF-IDF weighting allowed the model to focus on sentiment-bearing terms rather than high-frequency filler words. The Streamlit visualization layer proved equally critical — interactive Pie and Bar charts enabled PR managers to identify brand-level crisis hotspots at a glance, far more efficiently than reading raw CSV exports.

Fig 3: Model Accuracy vs. Training Data Size

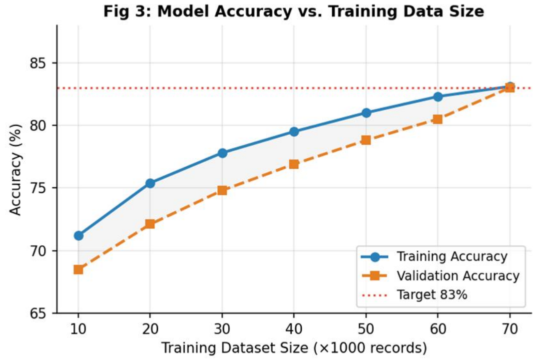

XII. FUTURE SCOPE

1) Transformer Models: Integrating BERT or RoBERTa would enable the system to understand sarcasm, cultural nuance, and long-range context, significantly reducing the current 17% error margin.
2) Real-Time Data Ingestion: Direct API integration with Twitter (X), Reddit, and YouTube would enable live monitoring as posts are published, replacing the current static dataset approach.
3) Predictive Risk Analytics: Using historical sentiment data, the system could predict when a small negative spike is likely to go viral, enabling preventive action before a crisis materializes.
4) Aspect-Based Sentiment Analysis (ABSA): Instead of a single score per post, ABSA would break down sentiment by product feature — e.g., Positive for Camera, Negative for Battery — providing actionable repair targets.
5) Multi-Modal Analysis: Future versions could analyze images and videos, detecting brand logos in negative memes or identifying frustrated expressions in video reviews.
6) Automated Alerts: Email or WhatsApp notifications sent to the PR team the moment brand sentiment drops below a configurable threshold.

XIII. CONCLUSION

This project successfully developed an end-to-end Brand PR Risk Monitor capable of processing over 70,000 social media records with an 83% classification accuracy. By implementing a Support Vector Machine trained on TF-IDF Bigram features, the system reliably identifies public dissatisfaction and potential brand crises, demonstrating that automated machine learning can serve as a viable replacement for manual social media monitoring. The seamless integration of the Data Layer (cleaning and vectorization), Application Layer (SVM-based risk logic), and Presentation Layer (Streamlit dashboard) proves that complex AI can be packaged into an intuitive, non-technical tool — empowering PR managers to monitor brand health without requiring machine learning expertise. Most importantly, this system achieves the fundamental objective of shifting corporate PR management from reactive damage control to proactive brand protection. Negative trends are flagged within seconds of occurrence, data-driven sentiment scores replace guesswork, and interactive visualizations reduce analysis time from days to minutes. The Brand PR Risk Monitor represents a practical and cost-effective solution for any organization operating in today's high-velocity social media landscape.

REFERENCES

[1] PwC, "2024 Global AI Jobs Barometer," PwC Global Report, Jun. 2024.
[2] B. Liu, "Sentiment Analysis and Opinion Mining," Morgan & Claypool Publishers, 2012.
[3] C. Cortes and V. Vapnik, "Support Vector Networks," Machine Learning, vol. 20, pp. 273-297, 1995.
[4] G. Salton and C. Buckley, "Term-Weighting Approaches in Automatic Text Retrieval," Information Processing & Management, vol. 24, no. 5, pp. 513-523, 1988.
[5] A. Go, R. Bhayani, and L. Huang, "Twitter Sentiment Classification using Distant Supervision," Stanford Technical Report, 2009.
[6] T. Streamlit Inc., "Streamlit: The Fastest Way to Build Data Apps," streamlit.io, 2024.
[7] F. Pedregosa et al., "Scikit-learn: Machine Learning in Python," JMLR, vol. 12, pp. 2825-2830, 2011.
[8] Analytics Insight, "Global AI Job Market & Salary Trends 2026," Analytics Insight Reports, 2026.