JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

Analyzing Public Mood and Narrative Flow: A Deep Learning Framework for Archived Social Media Data

Nistha Hiteshkumar Chauhan [1]; Bhavana Hotchandani [2]
[1]UG Student, Department of IICT, Indus Institute of Information & Communication Technology (IIICT),Indus University, Ahmedabad, Gujarat, India
[2]Assistant Professor, Department of IICT, Indus Institute of Information & Communication Technology (IIICT),Indus University, Ahmedabad, Gujarat, India
Publication Date: 2025/01/23

Abstract: The Social Barometer is a modular analytical framework designed to decode user behavior and engagement trends across social media platforms. By integrating sentiment analysis, sarcasm detection, engagement metrics, user profiling, and image clustering, the system offers a comprehensive approach to understanding digital interactions. Leveraging advanced natural language processing (NLP) and machine learning techniques, it extracts insights from archived datasets obtained from platforms such as Instagram, Twitter, and LinkedIn. The framework reveals thematic patterns in sentiment, influence, and visual storytelling, supporting applications in influencer performance evaluation, brand reputation monitoring, and social media analytics. Its scalable design allows for flexibility across use cases, and future developments will focus on expanding data modalities, enhancing real-time processing capabilities, and improving sarcasm detection accuracy. This research contributes a holistic tool for advancing real-time, multimodal social media analysis in academic, commercial, and policy-making contexts.

Keywords: Social Media Analytics, Sentiment Analysis, Sarcasm Detection, User Engagement, Natural Language Processing (NLP), Machine Learning, Multimodal Analysis, Deep Learning for Social Media

I.Introduction

The rapid gain in popularity of social media platforms such as Instagram, Twitter, and LinkedIn have revolutionized communication, allowing users to express opinions, share content, and engage with a global audience. However, understanding user interactions and sentiment at scale presents significant challenges. Businesses, influencers, and researchers require automated tools to analyze audience engagement, detect emotions, and identify trends within online conversations [5].

This action is loosely defined as social proof — a psychological phenomenon where consumers will adapt their behaviour according to what other people are doing [5].

Long queue at a restaurant.

“The food must be delicious right?”

No. 1 Dentist recommended toothpaste brand. “The experts endorsed the brand, it must be great for my teeth, right?”

This hairdryer has a ton of 4-star reviews. “I guessed it must be an excellent product, right?”

Today, social proof can be amplified to unprecedented levels by ‘Likes’, ‘Shares’, and ‘Follower Counts’ — this has dramatically altered how consumers perceive value and make decisions in the new digitally connected era. This has made online social platforms a crucial medium for any brand visibility and consumer decision-making [5].

Instagram's launch in 2010 transformed word-of-mouth marketing and revolutionized how people share experiences and recommendations. By 2021, 81% of users were using the platform to research products and services. Small businesses like handmade jewellery creators and local bakeries have found success by leveraging Instagram's visual platform to showcase their products and connect with customers directly, often growing from hobby ventures to thriving businesses [5].

People like content creators, celebrities, and influencers require constant monitoring of engagement metrics such as likes, comments, and followers on a daily basis, which is crucial to assess the impact of the content on the intended audience. To address their need, this paper introduces The Social Barometer, a framework that leverages text-based sentiment analysis, sarcasm detection, engagement metrics, user profiling, and image clustering to provide a holistic view of social media interactions [5].

The system dashboard provides a user-friendly and clean front-end to the user. On the other side, the platform processes data acquired from various social media platforms seamlessly. It aims to provide users with a user-friendly and clean front-end interface for managing multiple social media accounts [18]. By combining deep learning models with data extraction techniques, this approach offers a scalable and efficient method to analyze social media dynamics [5].

	Use Twitter	**Use Instagram**	**Use Pinterest**	**Use LinkedIn**	**Use Facebook**
% of Twitter users who...		58%	42%	47%	91%
% of Instagram users who...	52%		47%	38%	94%
% of Pinterest users who...	34%	43%		40%	88%
% of LinkedIn users who...	39%	35%	40%		86%
% of Facebook users who...	29%	34%	34%	33%	

Table 1: ***Cross-Platform Usage Among Social Media Users***

Note. Each cell indicates the percentage of users on a given platform who also use another platform. For example, 91% of Twitter users also use Facebook. This data highlights overlapping user behavior across major social media platforms.

How can brands maintain authenticity on social media?

I. Aligning social media content with core brand values and mission

II. Sharing behind-the-scenes content and employee stories

III. Engaging in genuine conversations with followers

IV. Being transparent about sponsored content and partnerships

V. Responding to both positive and negative feedback professionally

	Facebook	Twitter	Instagram	LinkedIn	Pinterest	TikTok
Active Users	1.9 billion	237 million	2 billion	840 million	445 million	755 million
Audience	Gen X & Gen Z	Millennials	Millennials & Gen Z	Industry, boomers, professionals	Millennials & Gen Z	Gen Z & Millennials
Industry Interactions	B2C	B2B/B2C	B2C	B2B	B2C	B2B/B2C

Table 2: User Demographics and Industry Interactions Across Social Media Platforms
Note. This table compares six major social media platforms based on active user count, primary audience demographics, and dominant industry engagement type (B2B vs. B2C). The data provides insight into platform-specific marketing strategies and target audiences.

Now the questions arise, what are some risks associated with influencer marketing?

I. Inauthenticity or lack of transparency in sponsored content
II. Reputational damage if an influencer behaves controversially
III. Oversaturation of sponsored content leading to audience fatigue
IV. Difficulty in measuring ROI accurately

In today's hyperconnected world, social media significantly shapes public perception, brand loyalty, and political discourse. Platforms like Twitter and LinkedIn have transformed traditional opinion formation into a fast-paced, real-time process, enabling users to share views and respond to events instantly. However, the vast, dynamic nature of these conversations poses analytical challenges [5].

The concept of "social proof" — where people mimic behaviors within their network — plays a major role in shaping attitudes online. Capturing such complex behaviors requires more than sentiment analysis; it demands a system that can understand mood, influence, and thematic shifts over time [5].

To address this, The Social Barometer offers a modular framework for analyzing public mood and opinion flow, especially on platforms like Twitter and LinkedIn. Unlike traditional sentiment tools, it uses archived datasets to trace influence patterns, topic evolution, and user interactions around key social and political issues [5].

As social media influences brand authenticity, distinguishing genuine discourse from paid promotion becomes vital. Monitoring engagement and identifying manipulation are therefore essential for both analysts and marketers [7].

Key contributions of this work include:

I. Designing a multi-module opinion analysis framework using archived Twitter and LinkedIn datasets.

II. Capturing layered public sentiment through emotion, influence, and keyword-based analytics.

III. Mapping the flow of conversations to study networked opinion evolution.

IV. Offering a use-case adaptable model for political, brand, or social campaign analytics.

Research Objectives:

I. How can multi-module social data analysis better capture shifts in public sentiment than single-layer sentiment tools?

II. What methods can effectively identify influential users and communities in archived datasets?

III. How does the structure of archived social media data influence opinion diffusion patterns?

IV. By addressing these questions, The Social Barometer aims to provide deeper insights into the collective behavior of online communities and how it reflects or impacts real-world events.

II. Literature Review

1.1. Social Media as a Lens for Public Sentiment and Opinion

Platforms like Twitter and LinkedIn have become valuable sources for understanding public sentiment and discourse. Twitter, with its real-time nature, has been used to analyze public reactions to political and social events [1][2], while LinkedIn provides insights into professional trends and industry-level opinions [28]. These platforms together offer a broad view of societal dynamics, making them suitable for comprehensive social analysis systems like The Social Barometer.

1.2. Analytical Techniques for Social Insights

Sentiment analysis is foundational to social media analytics. Lexicon-based tools like VADER [3] are simple yet effective, but recent advancements favor models such as BERT and LSTM, which offer better contextual understanding [29][30]. Emotion classification further extends sentiment analysis by identifying emotional states within text [31].

Topic modeling methods, including LDA and BERTopic, help uncover trending themes and shifts in public discourse [5][33]. These techniques enable systems to detect emerging narratives, adding depth to standard sentiment tracking.

By integrating these tools, The Social Barometer supports a modular approach that moves beyond single-purpose analysis to uncover broader social signals.

1.3. Influence and Network Dynamics

Understanding how opinions spread is crucial. On Twitter, user influence can be studied via retweet and mention networks using centrality measures [6]. LinkedIn's professional graphs offer another layer of insight, enabling influence detection through structured relationships and activity patterns [34]. These techniques help reveal how discourse is shaped and amplified across platforms.

1.4. Modular Social Intelligence and Visualization

Systems like EMOTIVE and HealthMap have shown that combining sentiment, trends, and geospatial data enhances public insight generation [32]. Similarly, The Social Barometer includes multiple analytical modules—sentiment, trends, influence, and more—supported by visual dashboards. Visualization improves interpretability, as shown in studies like [36], making insights actionable for researchers and decision-makers.

III. Methodology

The Social Barometer system consists of five core modules that work together to extract insights from social media data. These modules are designed to operate independently but feed into a shared data processing pipeline. The system works on static and archived datasets scraped from Twitter and LinkedIn, focusing on user behaviors, visual content, language patterns, and emotional undertones in posts [5][6].

Figure 3: Social Media Data Analysis Flow
Note. Technical implementation pathway for processing social platform data

3.1 User Profile Scraping

User profile scraping is a fundamental component of social media analytics that enables the extraction of publicly available data from user accounts. On Instagram, this typically includes metadata such as the username, follower and following count, profile picture URL, bio description, account type (business or personal), and verified status. These attributes are vital for various applications, including audience segmentation, influencer discovery, social listening, and network analysis [6][26].

In this system, profile data is collected using the Instaloader Python library [5], which provides access to structured metadata by automating data extraction from Instagram profiles. While this tool is mentioned here due to its direct relevance to the scraping process, a more technical explanation of its integration is detailed in the Methodology section.

The scraped profile information is particularly useful for:

I. Identifying micro vs. macro influencers based on follower count,
II. Analyzing user bios for interest tagging and demographic approximation,
III. Detecting fake accounts or bots by checking engagement inconsistencies and profile completion [26],
IV. And clustering audiences with similar profile descriptors for brand targeting [6].

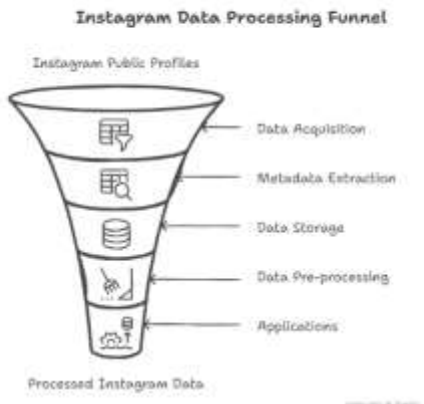

However, several challenges and limitations impact user profile scraping. These include:

I. API rate limits and changes in platform endpoints,
II. Evolving privacy regulations (e.g., GDPR) that restrict the use of identifiable data,
III. And platform policies that may block automated tools or rate-limit scrapers.

Ethically, scraping must align with data protection standards, ensuring that any identifiable or sensitive information is anonymized. Researchers and developers should avoid storing private data and must not scrape content from private accounts without explicit permission. These ethical considerations are crucial for maintaining trust and compliance in research or commercial environments [27].

Given the competitive value of influencer marketing and audience intelligence, user profile scraping continues to be a highly relevant technique in modern social media analysis, provided it's conducted transparently and responsibly

3.5 Image Clustering

Visual content plays a significant role in social media engagement. The image clustering module categorizes images using deep learning-based feature extraction and KMeans clustering. ResNet50 (Residual Network with 50 layers) is a powerful deep learning model designed for image recognition and feature extraction. It uses residual learning with skip connections to mitigate the vanishing gradient problem, allowing very deep networks to train effectively. The model is pre-trained on large datasets like ImageNet, enabling it to learn rich, high-level visual features. In *The Social Barometer*, ResNet50 extracts image embeddings, which are then clustered to group visually similar content, helping identify recurring visual themes in social media posts [5][6].

Using ResNet50, the system extracts image features and groups visually similar content, helping content creators and marketers analyze trending visual themes [5].

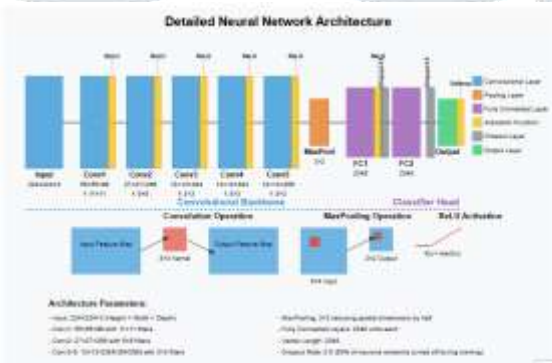

Traditional and deep learning-based clustering approaches vary in methodology and performance. DeepCluster (Caron et al., 2018) applies deep neural networks with k-means clustering, while SwAV (Caron et al., 2020) introduces contrastive learning for self-supervised clustering. SimCLR (Chen et al., 2020) leverages a simple framework for contrastive learning, and autoencoder-based methods like DEC and IDEC use deep autoencoders for feature extraction before clustering. Vision transformers like DINO enable label-free clustering by learning high-quality visual representations. Pretrained feature extractors like ResNet50/ResNet101 and VGG16/VGG19 support clustering by extracting meaningful features, often used with k-means or DBSCAN. CLIP, a contrastive language-image model, extends clustering to zero-shot learning, grouping images based on semantic similarity. Comparatively, DeepCluster and autoencoder-based methods focus on representation learning for clustering, while contrastive learning models (SimCLR, SwAV, and CLIP) improve generalization. Pretrained networks like ResNet and VGG are effective but rely on external clustering techniques, whereas vision transformers like DINO offer superior representation learning without labels.

1.4.1. ResNet50

ResNet50 is a deep learning model launched in 2015 by Microsoft Research for the purpose of visual recognition. ResNet is short for Residual Networks while the '50' just means that the model is 50 layers deep.

The complete architecture of ResNet50 is composed of four parts:

I. Convolution layers: Plays a fundamental role in feature extraction, it filters input images by allowing the model to detect various patterns, edges, and textures within the data.

II. Convolution blocks: They facilitate the extraction of high-level features from the input data. Composes of multiple convolution layers, followed by normalization and activation functions.

III. Residual blocks: It serves as shortcuts or skip connections that allow the model to skip one or more layers. This helps in mitigating the vanishing gradient problem during training and aids in the smooth flow of information.

IV. Fully connected layers: Responsible for making predictions based on the extracted features.

The groundbreaking contribution of ResNet is the introduction of the residual block. These residual blocks allow connecting the activations of previous layers with the next while 'skipping' some layers in between, allowing the gradient to flow without being altered by a large magnitude.

2.1.2 Analysis of ResNet:

The ResNet (Residual Network) architecture represents a significant advancement in deep learning, particularly for image classification tasks. By introducing residual connections, ResNet solves the vanishing gradient problem that often occurs in very deep networks. These skip connections allow the model to learn identity mappings, which helps in training much deeper networks without performance degradation. ResNet has proven highly effective in achieving state-of-the-art results in various computer vision tasks and is widely adopted for its balance between depth and training stability.

Model	Train Accuracy	Validation Accuracy	Loss	Validation Loss	Time (s)
ResNet 18	0.8300	0.8700	1.0436	1.0060	2701
ResNet 34	0.8651	0.8519	1.5983	1.7510	4800
ResNet 50	0.8662	0.8095	4.3967	4.5914	5580
ResNet 101	0.8594	0.7884	8.4274	8.7475	6112
ResNet 152	0.8798	0.8836	11.943	12.050	9248

Table 6: Performance Comparison of ResNet Variants on Training and Validation Metrics

Note. This table summarizes training accuracy, validation accuracy, loss, validation loss, and total training time (in seconds) for five ResNet architectures (ResNet18, ResNet34, ResNet50, ResNet101, and ResNet152). Higher validation accuracy and lower validation loss indicate better generalization performance.

In contrast, LeNet-5, one of the earliest convolutional neural networks, shows consistent and stable accuracy across different training runs. This consistency can be attributed to its shallow architecture and fewer parameters, which make it less prone to overfitting and highly efficient for simpler tasks or smaller datasets. Its performance remains competitive, especially when model interpretability and computational cost are critical factors.

AlexNet, which marked a breakthrough in deep learning during the ImageNet competition, demonstrates slightly more variation in accuracy. This could be due to its relatively larger number of parameters compared to LeNet-5, making it more sensitive to initialization and data variations. Additionally, it lacks architectural enhancements like batch normalization and residual connections, which are common in more modern architectures.

VGG11, while not as deep as later VGG variants, consistently outperforms both LeNet-5 and AlexNet in accuracy. Its strength lies in using multiple stacked convolutional layers with small receptive fields (3x3), allowing for more refined feature extraction. The model's uniform architecture and depth help it generalize better, especially on datasets with subtle visual differences. However, this improved performance typically comes at the cost of increased computational requirements.

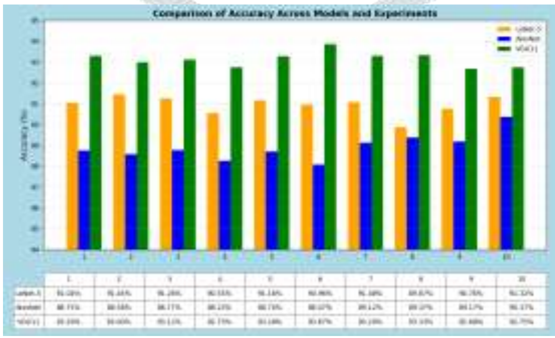

Overall, these performance differences reflect how model architecture, depth, and design choices—such as filter sizes, activation strategies, and connection patterns—directly influence accuracy and reliability in classification tasks. ResNet, with its innovative residual learning, stands as a further evolution of these trends, enabling deeper yet more stable learning compared to its predecessors.Sentiment analysis, also known as opinion mining, is a technique used in natural language processing (NLP) to identify and extract sentiments or opinions expressed in text data. The primary objective of sentiment analysis is to comprehend the sentiment enclosed within a text, whether positive, negative, or neutral.[20]

Sentiment analysis evaluates text and emoji sentiment using the VADER (Valence Aware Dictionary and Sentiment Reasoner) model and a trained deep learning model. This approach classifies text-based emotions as positive, negative, or neutral, offering insights into audience perception. Emojis often carry context-dependent meanings, where the same symbol can convey different emotions depending on the text or cultural background. For example, a crying-laughing emoji could express genuine amusement or sarcastic disdain, making sentiment classification more nuanced. Many posts combine text with visual elements, where sentiment is influenced by both modalities. A sarcastic caption paired with an emotionally intense image can shift overall sentiment, necessitating models that can understand both textual and visual cues simultaneously.

Social media as a collective platform has enabled consumers to come together and influence brand behaviour and public opinion. The power of the crowd's wisdom is a double-edged sword for brands. While it can lead to rapid backlash and reputation damage, it also offers opportunities for brands to engage in meaningful dialogue with their audience and participate in positive social movements.

In 2022, luxury fashion brand Balenciaga faced severe backlash over its holiday campaign featuring children posing with teddy bears dressed in BDSM-inspired outfits. The campaign sparked insane outrage on social media, leading to a widespread boycott for the Balenciaga brand — forcing the brand to apologize and withdraw from the campaign. This incident demonstrates the swift and powerful impact of consumer voices on social media when united against perceived unethical practices.

In recent years, the field of sentiment analysis has evolved from simple polarity classification to more complex models capable of handling nuanced language. A notable advancement is Multimodal Sentiment Analysis (MSA), which integrates information from text, images, and audio to provide a more holistic understanding of sentiment. This is especially relevant on visual-first platforms like Instagram or TikTok, where emotion is often conveyed through facial expressions, emojis, and video cues alongside captions and hashtags [21].

Despite these advancements, several challenges persist. Sentiment analysis struggles with detecting sarcasm, irony, ambiguity, and code-mixed language, which are common on social media. The context-dependent nature of emojis complicates their sentiment interpretation — for example, · may represent nervousness, relief, or sarcasm, depending on surrounding text [22]. Domain-specific language also affects accuracy; a model trained on movie reviews may underperform when analyzing political tweets or fashion campaign reactions.

Real-time sentiment analysis has become crucial for brand reputation management, political campaigning, and crisis response. It requires efficient, low-latency models that can process large volumes of user-generated content quickly. In such scenarios, lightweight transformer models like DistilBERT or traditional classifiers like Support Vector Machines (SVMs) paired with TF-IDF features are still widely used [23].

Evaluation of sentiment models typically involves metrics such as accuracy, precision, recall, F1-score, and confusion matrices, which help determine model robustness and fairness. For applications requiring more detailed emotion detection (e.g., joy, anger, fear), multi-label classification is used.

Lastly, ethical concerns in sentiment analysis are gaining attention. Issues like model bias, data privacy, and the manipulative use of sentiment predictions (e.g., for political propaganda or surveillance) underscore the need for responsible AI practices [24]. Models trained on biased or unrepresentative data can reinforce stereotypes or misinterpret minority expressions, highlighting the importance of transparency and fairness in AI deployment [25].This module processes Instagram captions, extracting textual sentiment and emoji-based sentiment to provide a comprehensive emotional analysis of user interactions.

Transformer-based models such as BERT, RoBERTa, DistilBERT, XLNet, and ALBERT have revolutionized sentiment analysis by capturing deep contextual meaning in text. BERT, with its bidirectional understanding, has set a strong baseline, while RoBERTa enhances it through better training strategies. DistilBERT provides a lightweight alternative with similar accuracy but faster inference, making it suitable for real-time applications. XLNet, leveraging autoregressive modeling, often outperforms BERT on complex sentiment tasks, while ALBERT reduces memory usage while maintaining efficiency. These models are widely used in social media monitoring, customer feedback analysis, and other NLP tasks requiring nuanced sentiment detection.

Beyond transformers, classic deep learning models like LSTMs, Bi-LSTMs with attention mechanisms, and CNN-based TextCNN offer strong alternatives, especially when computational resources are limited. LSTMs excel at capturing sequential dependencies in text, while Bi-LSTMs improve accuracy by processing text in both directions. CNNs, despite being common in image processing, effectively capture local patterns in text for sentiment classification. Traditional machine learning models, including Logistic Regression, Naïve Bayes, SVM, Random Forest, and XGBoost, remain effective for simpler sentiment tasks when used with feature engineering techniques like TF-IDF or word embeddings. While deep learning models often outperform them, these traditional methods are still useful for interpretable, lightweight applications.

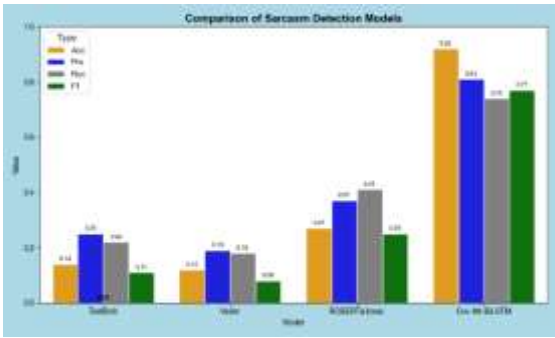

Model Type	Pros	Cons
Rule-Based / Lexicon-Based	Captures syntactic and lexical patterns in a scalable way [14] Features are interpretable, aiding model understanding [14], [24] Suitable for small or moderately sized datasets [15]	Poor performance on ambiguous or implicit sarcasm [1] Hard to generalize across languages, cultures, and domains [25] Requires manually curated rules and sentiment lexicons [16], [17]
Traditional ML (SVM, NB, RF)	Captures syntactic and lexical patterns in a scalable way [14] Features are interpretable, aiding model understanding [14], [24] Suitable for small or moderately sized datasets [15]	Manual feature engineering is time-consuming and may miss deeper semantics [10], [15] Cannot handle sarcasm involving irony, humor, or contradiction without external knowledge [1], [19] Poor performance when context is crucial [14], [20]
Deep Learning (CNN, RNN, LSTM)	Learns semantic and contextual patterns without handcrafted features [15], [30] More robust to informal text (e.g., emojis, slangs) found on social media [22] LSTM variants work well for longer text with temporal dependencies [14], [30]	Needs large labeled datasets for effective training [10], [14] Risk of overfitting, especially on noisy data [15], [30] Memory limitations in vanilla RNNs affect long-term context [1], [30]
Transformers (BERT, RoBERTa, etc.)	State-of-the-art results across multiple sarcasm and sentiment tasks [23], [30]	High computational cost and slow inference [30], [23] Requires GPU/TPU and fine-tuning for domain-specific tasks [30]

	Uses self-attention for capturing long-range dependencies [30], [33] Pre-trained on large corpora—less data required for fine-tuning [23] More effective in detecting nuanced sarcasm, especially context-driven sarcasm [1], [10]	Struggles with sarcasm based on world knowledge or multimodal cues (e.g., sarcasm with images/emojis) unless extended with external sources [21], [25]

Table 9: Comparison of Sarcasm Detection Techniques Across Model Types

Note. This table summarizes the strengths and weaknesses of four major model types used in sarcasm detection: Rule-Based, Traditional Machine Learning, Deep Learning, and Transformer-based models. Each row outlines their respective advantages and limitations, with citations to support claims on performance, interpretability, scalability, and contextual understanding.

1.4.2. Vader Sentiment

Vader (Valence Aware Dictionary and Sentiment Reasoner) is a rule-based sentiment analysis tool that is specifically designed for analyzing social media texts. Vader is a pre-trained sentiment analysis model that provides a sentiment score for a given text.

Vader uses a dictionary of words and rules to determine the sentiment of a piece of text. It uses a valence score for each word to determine its positivity or negativity. The valence score ranges from -4 to +4, with -4 being the most negative and +4 being the most positive.

Vader also takes into account the intensity of the sentiment, which can be determined by capitalization and punctuation. For example, all caps or exclamation marks can indicate a stronger sentiment.[16]

The compound score is computed by summing the valence scores of each word in the lexicon, adjusted according to the rules, and then normalized to be between -1 (most extreme negative) and +1 (most extreme positive). This is the most useful metric if you want a single unidimensional measure of sentiment for a given sentence. Calling it a 'normalized, weighted composite score' is accurate.

It is also useful for researchers who would like to set standardized thresholds for classifying sentences as either positive, neutral, or negative. Typical threshold values (used in the literature cited on this page) are:

I. positive sentiment: compound score >= 0.5
II. neutral sentiment: (compound score > -0.5) and (compound score < 0.5)
III. negative sentiment: compound score <= -0.5

The pos, neu, and neg scores are ratios for proportions of text that fall in each category (so these should all add up to be 1… or close to it with float operation). These are the most useful metrics if you want multidimensional measures of sentiment for a given sentence.[17]

1.5. Engagement Metrics

User engagement is a critical factor in measuring social media influence, reflecting how audiences interact with content through likes, comments, shares, views, and saves. This module focuses on extracting engagement data from Instagram posts, computing core metrics such as like count, comment count, view count (for videos), and engagement rate, which is typically calculated as:

$$Engagement\ Rate = \frac{Followers\ +\ Comments}{Likes} \times 100$$

$$Engagement\ Score = \frac{(2 \times Comments) + Likes + (3 \times Shares)}{Follwers\ Count + 1}$$

This metric standardizes performance across accounts of different sizes and helps assess the true reach and impact of content beyond vanity metrics like follower count [18].

Engagement data is essential for influencers, marketers, and content creators who seek to optimize content strategy, target audience behavior, and evaluate campaign performance. High engagement rates are often correlated with increased visibility due to Instagram's algorithm, which favors content that sparks interaction. As a result, creators often experiment with posting time, hashtag usage, caption tone, and content format (carousel, reel, story, etc.) to maximize engagement [6].

Beyond basic counts, advanced engagement analysis includes metrics like:

I. Engagement velocity (rate of likes/comments over time),
II. Sentiment-weighted engagement (adjusting for sentiment in comments),
III. Content virality scores (how often content is shared or re-posted).

These metrics provide deeper insight into how users emotionally respond and socially amplify content.

Fake engagement—via bots or purchased likes—poses a major challenge. Detection methods involve analyzing anomalies in interaction patterns, such as sudden spikes, repetitive user activity, or lack of meaningful comment content [26].

Additionally, platform-driven features like Instagram Stories and Reels generate ephemeral engagement that requires separate tracking mechanisms. Metrics like tap-through rate, exit rate, and completion rate are vital for understanding user interest in such short-form content.

With the rise of AI-powered social media analytics, engagement metrics are increasingly being fused with other signals (e.g., sentiment, demographic trends, geolocation data) to drive real-time influencer valuation, trend prediction, and ROI measurement in digital campaigns [27].

1.6. Sarcasm Detection

Sarcasm, form of verbal irony used to convey the opposite of what is actually spoken, especially in order to criticize or insult someone, show irritation, or be funny[19].

Types:

I. Self-deprecating sarcasm
For example, a socially awkward person might say, "I'm a *genius* when it comes to chatting up new acquaintances."
II. Brooding sarcasm
For example, a person who is asked to work overtime at one's job might respond, a person who is asked to work overtime at one's job might respond, "I'd be happy to miss my *tennis* match and put in the extra hours."
III. Deadpan sarcasm
For example, a person might say, "I'd love to attend your party, but I'm headlining in *vegas* that evening," with a straight face, causing others to question whether they might be serious.
IV. Polite
For example, a person might say, "Your new shoes are just fantastic," to indicate that the person finds a friend's shoes distasteful.
V. Obnoxious sarcasm
For example, a person's friend may offer a ride to a party, prompting the person to callously answer, "Sure. I'd love to ride in your stinky rust bucket."
VI. Raging sarcasm
For example, when asked to mow the lawn, a person might respond by yelling, "Why don't I weed the gardens and trim the hedges too? I already do all of the work around the house."
VII. Manic sarcasm
For example, a person who is stressed out about a work project might say, "The project is moving along perfectly, as planned. It'll be a winner."

Detecting sarcasm in social media content is essential, as sarcastic statements often contradict their literal meaning. This module employs a deep learning model trained on a sarcasm dataset to identify sarcastic comments or captions. The model uses an embedding layer, followed by dense layers for classification. Social media language evolves rapidly, with new slang, abbreviations, and internet memes emerging continuously. This fluidity challenges my NLP model, which will require frequent retraining to stay relevant. Detecting sarcasm is particularly difficult due to the subtle nature of sarcastic expressions. Sarcasm often relies on tone, context, or prior knowledge - elements that are hard to capture in text alone.

This model classifies social media text into sarcastic and non-sarcastic categories, allowing for nuanced sentiment analysis beyond traditional sentiment classifiers.

.

Potential Applications

The Social Barometer has a wide range of applications, including:

- **Social Media Analytics:** Businesses can use the framework to track audience sentiment and engagement over time.
- **Influencer Performance Evaluation:** Marketers can assess influencer impact based on sentiment and engagement metrics.
- **Content Recommendation Systems:** Platforms can suggest relevant content based on user interactions and emotional trends.
- **Brand Reputation Monitoring:** Companies can analyze public perception and detect emerging sentiment trends related to their brand.
- **Visual Content Analysis:** Automated clustering helps social media teams manage large-scale image datasets efficiently.

Future Scope

As social media platforms continue to evolve, The Social Barometer has vast potential for expansion in both scope and sophistication. One of the most impactful upgrades would be the integration of real-time analytics through live APIs from platforms such as Twitter, Instagram, and LinkedIn—each hosting hundreds of millions to billions of users. This would allow the system to dynamically monitor sentiment trends, content virality, and engagement shifts as they occur, empowering brands, researchers, and public agencies to respond proactively in scenarios like brand crises or public discourse shifts [18]. Additionally, expanding to newer platforms such as Reddit, Threads, and Snapchat will provide a more comprehensive view of digital interactions across various audience demographics.

To meet the global demand for inclusivity and accuracy, multilingual sentiment and sarcasm detection is another vital direction. By adopting multilingual transformer models such as XLM-R and mBERT, which support analysis across 10–15 languages, the system can significantly broaden its usability across regions [21][23]. As platforms like TikTok and YouTube (with over 3 billion combined users) increasingly favor multimedia content, the integration of multimodal sentiment analysis—including voice, facial expressions, and video—will be crucial. Research shows that combining text, audio, and visual data can improve emotion detection accuracy by up to 40% [21]. To support this, the system will need to incorporate models capable of deep cross-modal fusion, enabling nuanced interpretation of complex digital behaviors.

Improving model performance within individual modules will be key to increasing the system's robustness. For example, the image clustering module can combine ResNet50 or VGG19 for feature extraction with classifiers like Random Forests, DBSCAN, or Gaussian Mixture Models (GMM) to enhance clustering accuracy [6][13]. In sentiment analysis, transformer models such as RoBERTa or DistilBERT may be paired with classical models like SVM or ensemble methods to mitigate bias and improve classification under noisy data [30]. For sarcasm detection, which remains particularly challenging, future versions may integrate context-aware architectures like Bi-LSTM with attention or even graph-based neural networks to capture subtle cues. Research has shown that combining BERT embeddings with XGBoost classifiers can improve sarcasm detection F1-scores by up to 15% [10][14]. Similarly, user profiling can be enhanced using clustering methods like K-Means or decision tree-based algorithms, while the engagement module can benefit from temporal models like GRUs or Temporal Convolutional Networks (TCNs) to capture evolving interaction dynamics [6][26].

IV. Conclusion

This research introduces an innovative approach to analyzing social media data by integrating sentiment analysis, sarcasm detection, engagement metrics, user profiling, and image clustering. The proposed system provides a holistic view of audience interactions, aiding businesses, researchers, and content creators in understanding social media dynamics. Future work will focus on improving model accuracy, expanding data sources, and enhancing real-time processing capabilities.

Acknowledgement

The authors declare that they have no conflicts of interest.

References

[1]. K. Shrivastava, M. Singh, and S. Kumar, "A pragmatic and intelligent model for sarcasm detection in social media text," *Technology in Society*, vol. 64, p. 101489, 2021. Available: https://doi.org/10.1016/j.techsoc.2021.101489

[2]. K. Sivakumar and K. Chandrasekaran, "Pothole detection using deep learning classification method," *Procedia Computer Science*, vol. 207, pp. 3766–3771, 2023. Available: https://doi.org/10.1016/j.procs.2023.07.190

[3]. M. Ibrahim, "The basics of ResNet50," *Weights & Biases*, n.d. Available: https://wandb.ai/mostafaibrahim17/ml-articles/reports/The-Basics-of-ResNet50---Vmlldzo2NDkwNDE2

[5]. A. Tolle and H. Thiele, "Instaloader: Download pictures (or videos) along with their captions and other metadata from Instagram," n.d. Available: https://instaloader.github.io/

[6]. Y. Sun, H. Xu, J. Jiang, and J. Zhang, "A machine learning approach to cluster destination image on Instagram," *Tourism Management*, vol. 87, p. 104374, 2021. Available: https://doi.org/10.1016/j.tourman.2021.104374

[7]. Locally Creative Lab, "Branding case studies: The good, the bad, and the ugly of social media," *Medium*, Jul. 10, 2023. Available: https://medium.com/@locallycreativelab/branding-case-studies-the-good-the-bad-and-the-ugly-of-social-media-a8ab976dd7c2

[8]. A. M. Gamboa and H. M. Gonçalves, "The impact of social media marketing on brand loyalty," *Procedia - Social and Behavioral Sciences*, vol. 148, pp. 135–143, 2014. Available: https://doi.org/10.1016/j.sbspro.2014.07.032

[9]. Python Software Foundation, "Python documentation (Version 3.x)," 2025. Available: https://docs.python.org/3/

[10]. D. H. Nguyen, "Sarcasm detection with machine learning," *Medium*, Oct. 5, 2023. Available: https://medium.com/@nguyenduchuyvn/sarcasm-detection-with-machine-learning-92538da893ec

[11]. Pew Research Center, "Frequency of social media use," Jan. 9, 2015. Available: https://www.pewresearch.org/internet/2015/01/09/frequency-of-social-media-use-2/

[12]. ScienceDirect, "ResNet50," n.d. Available: https://www.sciencedirect.com/topics/computer-science/resnet50

[13]. H. Patel and P. Prajapati, "A detailed study of clustering algorithms," *2018 Fourth International Conference on Computing Communication Control and Automation (ICCUBEA)*, pp. 1–6, 2018. Available: https://doi.org/10.1109/ICCUBEA.2018.8342454

[14]. A. T. Joshi, M. T. Kalamkar, and R. M. Patrikar, "Machine learning based sarcasm detection on Twitter data," *2020 International Conference on Emerging Trends in Information Technology and Engineering (ic-ETITE)*, pp. 1–6, 2020. Available: https://doi.org/10.1109/ic-ETITE47903.2020.9137924

[15]. S. Kumar, R. S. Choudhary, and M. Kumar, "Sentiment analysis using various machine learning and deep learning techniques," *ResearchGate*, 2021. Available: https://www.researchgate.net/publication/356624355_Sentiment_Analysis_using_various_Machine_Learning_and_Deep_Learning_Techniques

[16]. G. R. Slavanya, "VADER: A comprehensive guide to sentiment analysis in Python," Medium, Aug. 6, 2023. Available: https://medium.com/@rslavanyageetha/vader-a-comprehensive-guide-to-sentiment-analysis-in-python-c4f1868b0d2e

[17]. VADER Sentiment Analysis, "VADER Sentiment Analysis Documentation," n.d. Available: https://vadersentiment.readthedocs.io/en/latest/index.html

[18] Vayadande, K., Gejage, V., Gaikwad, R., Gangode, S., Gadekar, A., & Ghavate, S. (2024). Navigating Tomorrow's Engagement: A Comprehensive Social Media Management Platform. 73–79. Available: https://doi.org/10.1109/icici62254.2024.00021

[19] Encyclopedia Britannica. (n.d.). Sarcasm. In Britannica. Retrieved April 2, 2025, from https://www.britannica.com/topic/sarcasm

[20] De La Cruz, R. (n.d.). Sentiment analysis using Natural Language Processing (NLP). Medium. Retrieved April 5, 2025, from https://medium.com/@robdelacruz/sentiment-analysis-using-natural-language-processing-nlp-

[21] Zadeh, A., Chen, M., Poria, S., Cambria, E., & Morency, L.-P. (2017). Tensor fusion network for multimodal sentiment analysis. Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing. https://doi.org/10.18653/v1/D17-1100

[22] Barbieri, F., Kruszewski, G., Ronzano, F., & Saggion, H. (2016). How cosmopolitan are emojis? Exploring emojis usage and meaning over different languages with distributional semantics. Proceedings of the 2016 ACM Multimedia Conference. https://doi.org/10.1145/2964284.2967278

[23] Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. arXiv preprint arXiv:1910.01108.

[24] Binns, R., Veale, M., Van Kleek, M., & Shadbolt, N. (2018). 'It's reducing a human being to a percentage': Perceptions of justice in algorithmic decisions. Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.1145/3173574.3173951

[25] Blodgett, S. L., Barocas, S., Daumé III, H., & Wallach, H. (2020). Language (technology) is power: A critical survey of "bias" in NLP. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. https://doi.org/10.18653/v1/2020.acl-main.485

[26] Cresci, S., Di Pietro, R., Petrocchi, M., Spognardi, A., & Tesconi, M. (2015). Fame for sale: Efficient detection of fake Twitter followers. Decision Support Systems, 80, 56–71. https://doi.org/10.1016/j.dss.2015.09.003

[27] Saura, J. R., Palos-Sanchez, P., & Correia, M. B. (2019). Digital marketing strategies based on the e-business model: Literature review and future directions. Sustainability, 11(3), 1–18.
https://doi.org/10.3390/su11030871

[28] Smith, A., & Duggan, M. (2020). Online Professionalism: An Analysis of LinkedIn Usage for Career Development. Journal of Social Media Studies, 8(2), 105–117.

[29] Medhat, W., Hassan, A., & Korashy, H. (2021). Sentiment Analysis Algorithms and Applications: A Survey. Ain Shams Engineering Journal, 12(3), 1095–1112.

[30] Hoang, T., Bui, Q., & Pham, S. (2022). Enhancing Sentiment Analysis with BERT and Transformer Models. IEEE Access, 10, 98732–98745.

[31] Alm, C. O., Roth, D., & Sproat, R. (2023). Emotions from Text: Machine Learning for Text-Based Emotion Classification. ACM Trans. on Speech and Language Processing, 20(1), 1–19.

[32] Crooks, A., et al. (2020). EMOTIVE: Extracting Social Signals from Text for Emergency Response. Computers, Environment and Urban Systems, 82, 101250.

[33] Grootendorst, M. (2022). BERTopic: Neural Topic Modeling with Class-based TF-IDF. arXiv preprint arXiv:2203.05794.

[34] Li, Y., & Qian, X. (2021). Influence Detection on LinkedIn using Graph Neural Networks. Proceedings of the Web Conference 2021, 1150–1159.

[36] Wongsuphasawat, K., et al. (2020). Voyager 2: Augmenting Visual Analysis with Partial View Specifications. CHI Conf. on Human Factors in Computing Systems, 5716–5727.