Advances in Consumer Research
https://acr-journal.com/
Volume-2 | Issue-4 | Sep 2025

Original Researcher Article

The Dhoni Enigma: Emotion, Composure, and the Architecture of Reputational Capital in the Age of Large Language Models

Dr Vigneshwaran SR[1], Vinodhini Ranganathan[2]

[1]Department of Decision Sciences and Information Systems Indian Institute of Management Bangalore (IIMB), India Email: vigneshwarans.r@iimb.ac.in
[2]Applied Researcher NLP Point Click Care, USA Email: vinodhini.ranganathan@pointclickcare.com

Received:10/08/2025
Revised: 20/08/2025
Accepted:12/09/2025
Published:30/09/2025

ABSTRACT
Social media platforms have become a major facet of modern information systems, providing organizations real-time insights into stakeholder sentiments that can steer strategic decisions and corporate governance. This study explores how information systems-enabled sentiment analysis can be used to track reputational capital and foster stakeholder trust. Based on stakeholder theory and reputational capital models, we applied a hybrid AI method combining large language models with clustering techniques to analyze more than 36,000 social media posts gathered during the retirement of a well-known public figure.

Our analysis revealed distinct emotional groups such as positive, neutral, and negative reflecting emotions of admiration, uncertainty, and disappointment. These insights also bring out valuable actionable information for early detection of reputational risks, emphasising transparency in corporate reporting, and improving communication with stakeholders. Academically, this research advances the information systems discipline by presenting a method to transform unstructured social media data into efficacious decision-support tools. For practitioners, it demonstrates how organizations can develop public accountability, stakeholder trust, and strengthen resilience in today's transparent digital environment.

Keyword: Large Language Model, GPT, Emotion, Clustering, NLP, Sentiment Analytics.

Introduction

1.1 Digital Emotions, Reputation, and the Information Systems Lens

In an era defined by hyperconnectivity, social media has changed how individuals, communities, and organizations interact, share information, and express their emotions. These platforms are not merely communication tools but powerful components of modern information systems (IS), shaping how reputational capital and stakeholder trust are formed, maintained, and contested in real time. Unlike traditional surveys or financial metrics of reputation, digital expressions of sentiment seize raw and immediate reactions to unfolding events. This makes them a crucial but underused resource for decision-support, corporate reporting, and governance systems.

Reputational capital has become an intangible but invaluable resource for organizations. It influences consumer loyalty, investor confidence, staff engagement, and regulatory goodwill. Stakeholder trust, often considered the foundation of sustainable competitive advantage, is increasingly mediated by digital discourse. When stakeholders make openly visible expressions of admiration, disappointment, or uncertainty, these sentiments resonate across markets and communities, affecting both the perception and value of organizational actions. Consequently, IS scholars and practitioners are bringing their attention to how digital traces of emotion can be systematically captured and transformed into actionable intelligence.

1.2 The Gap in Information Systems Research

Despite the growing significance of social media analytics, the IS literature has only partially acknowledged how unstructured emotional data can be operationalized into decision-making processes. Much of the preceding work has intensify marketing applications, such as consumer sentiment analysis or brand management, while fewer studies have grounded emotional analytics in comprehensive theoretical constructs such as stakeholder theory, legitimacy theory, and reputational capital frameworks. These viewpoints remind us that organizations are within a social contract, where stakeholders perpetually evaluate not just products and services, but also values, accountability, and transparency.

Traditional IS-enabled business intelligence systems stand out at structured data—financial statements, operational metrics, and transactional records. Yet, the scantiness in capturing the affective dimensions of stakeholder interactions that directly influence corporate reputation and trust. The gap is especially acute when it comes to high-impact events that trigger unanticiptated surges of online expression. Such events provide natural experiments for assimilation the dynamics of digital emotions and testing how IS methodologies can detect, classify, and interpret them.

This study addresses this gap by demonstrating how large language models (LLMs) and clustering techniques can analyze stakeholder emotions at scale and integrate them into decision-support systems for reputational monitoring. In doing so, we position emotional analytics not simply as a technical exercise, but as a theory-anchored IS contribution that enhances corporate reporting, strengthens transparency, and builds resilience in organizational governance.

1.3 Case Context: The Retirement of a Public Figure

To illustrate our approach, we analyze the retirement announcement of Mahendra Singh Dhoni, a former captain of the Indian cricket team and a global sports icon. Dhoni's persona extended far beyond cricket. Renowned for his leadership, composure under pressure, and strategic acumen, he became a figure of public trust, shaping brand endorsements and stakeholder perceptions across multiple industries. His career highlights—including leading India to the 2007 ICC T20 World Cup and the 2011 ODI World Cup victories—cemented his reputation as a national hero and international role model.

On 15 August 2020, Dhoni announced his retirement from international cricket via a succinct Instagram post: "From 1929 hrs consider me as retired." This understated message triggered an extraordinary wave of emotional responses across digital platforms, with over 36,000 tweets posted within two weeks. These reactions included admiration, nostalgia, disappointment, and uncertainty, offering a rich corpus of unstructured data. We can use this event to naturally examine how widespread emotional responses impact reputational capital.

Although framed within the context of sports, this case offers significant insights into how high-profile events in business, politics, or entertainment develop and spread across the digital sphere. Organizations, like individuals, are under the gaze public opinion, and their reputational capital is influenced by analogous dynamics of admiration, critique, and scrutiny. By exerting on Dhoni's retirement as an empirical case, we exemplify a reorientable methodology that organizations can tweak for real-time reputation management, crisis communication, and stakeholder engagement.

1.4 Linking Case to Theory and IS Contribution

This study is anchored in stakeholder theory (Freeman, 1984), which argues that an organization's success depends on earning the trust of various stakeholders. It also draws on legitimacy theory, which suggests that organizations must align with societal values to maintain their public standing. Finally, we use reputational capital frameworks, which view reputation as a valuable asset that's key to long-term survival.

Applying these ideas to social media data lets us interpret digital emotions as more than just "fan reactions." Instead, we see them as signals of a company's trust, legitimacy, and reputation. The retirement of a public figure like Dhoni shows how sudden events can create ripple effects in reputation, with emotions clustering around themes of admiration, uncertainty, and disappointment. By turning these clusters into actionable insights, we can show how information systems methodologies can take a theoretical idea and apply it to real-world governance and accountability.

1.5 Methodological Overview

We employ a hybrid AI approach that integrates large language models for zero-shot emotion detection with semantic embedding (Word2Vec) and Agglomerative Clustering. This methodology enables us to move beyond binary sentiment analysis and capture nuanced emotional landscapes. By analyzing 36,000+ tweets, we identify three dominant clusters: Positive and Motivational, Neutral and Ambivalent, and Negative and Distressed each reflecting distinct dimensions of reputational perception.

While the technical details of this methodology are elaborated in Section 3 (Methodology), its placement here underscores how IS-enabled tools can scale emotional analysis into structured intelligence. Such approaches have significant implications for corporate reporting on intangible assets, proactive reputational risk management, and the design of transparent communication strategies.

1.6 Contributions to Research and Practice

This study makes two primary contributions.

For research, it advances IS scholarship by:

Demonstrating how LLM-enabled emotion analytics can operationalize theoretical constructs of reputational capital and stakeholder trust.

Bridging the gap between unstructured emotional data and structured decision-support systems.

Expanding the IS literature on intangible asset measurement through the integration of AI and social media analytics.

For practice, it provides organizations with:

A scalable framework for real-time reputational monitoring using digital emotions.

Tools for identifying reputational risks before they escalate into crises.

Evidence that transparent communication informed by stakeholder emotions can strengthen trust and accountability.

1.7 Structure of the Paper

The remainder of the paper is structured as follows. Section 2 reviews the literature on reputational capital, stakeholder trust, and the role of IS in advancing emotion analytics. Section 3 outlines the research methodology, detailing the data collection process and the hybrid AI techniques employed. Section 4 provides an in-depth analysis of the emotional responses to Dhoni's retirement, situating them within the broader digital discourse. Section 5 presents the results, including the identification and distribution of distinct emotional clusters. Section 6 discusses the managerial implications of these findings, emphasizing how organizations can harness emotion analytics to strengthen stakeholder trust, enhance reputational monitoring, and design effective communication strategies. Section 7 highlights the study's limitations, while Section 8 concludes with a synthesis of key contributions. Finally, Section 9 suggests avenues for future research, positioning our framework as transferable to reputational events across industries and contexts.

Related Work

2.1 Advancements in LLM-based Sentiment and Emotion Analysis

Over the last two decades, sentiment and emotion analysis has transitioned from a niche research area in computational linguistics to a mainstream field at the intersection of artificial intelligence (AI), psychology, and management sciences. Early methods, grounded in lexicons and rule-based heuristics, offered the first attempts to quantify sentiment from text. While these approaches were computationally inexpensive and interpretable, they were also brittle—incapable of capturing polysemy, context-dependence, or the subtle interplay between emotion and culture. For instance, phrases like "Dhoni's silence spoke volumes" would often be misclassified by lexicon-based tools, since the word "silence" could be mapped to neutrality or negativity without considering its situational connotation of respect or admiration.

The advent of machine learning marked the next major stage. Bag-of-words (BoW) and n-gram–based models, often coupled with algorithms such as Naïve Bayes, Support Vector Machines (SVM), or Random Forests, improved predictive accuracy by identifying statistical correlations in labeled datasets (Pang & Lee, 2008). Yet, these methods were still shallow, struggling with sarcasm, irony, or evolving slang, all of which are particularly prevalent in social media platforms like Twitter. The sports domain illustrates this well: when a fan tweets "Only Dhoni could fail so brilliantly," traditional classifiers often misinterpret it as negative, whereas human readers can decode the admiration embedded in the apparent criticism.

The deep learning revolution in natural language processing (NLP), particularly with Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM) networks, began to fix these issues (Mehta et al., 2022; Rao et al., 2025, Venkatesh and Gokul, 2025). These architectures were built to model how words depend on each other in a sequence, helping systems understand context across sentences. However, the Transformer architecture (Vaswani et al., 2017) truly changed the game. By using a "self-attention" mechanism, Transformers get going on capturing subtle and long-range contextual relationships with incredible efficiency.

The surge of Large Language Models (LLMs)—like BERT (Devlin et al., 2019), the GPT series (Radford et al., 2018; Brown et al., 2020), and their abundant successors—is the result of this progress. These models, pre-trained on text which are tremendous, are now able to perceive nuanced meanings and real-world intelligence. This allows them to work well across different areas without needing much fine-tuning. Most importantly, LLMs introduced zero-shot and few-shot learning, which has been a game-changer for sentiment and emotion analysis. This is crucial for fast-moving situations, like a sudden celebrity retirement or a corporate scandal, where getting a lot of labeled data is simply not an option.

Recent studies highlight this shift. For example, Liu et al. (2024) created a family of specialized models called EmoLLMs. These models, trained explicitly for affective computing, surprisingly surpassed the performance of the much larger GPT-4 in several emotion classification benchmarks. This shows that a model's specific focus can be just as effective as its sheer size. Similarly, Siddiqui et al. (2024) found that a fine-tuned GPT-2 variant could do better than bigger, more general models like GPT-3.5 at multi-label emotion classification on Twitter. This challenges the popular idea that "bigger is always better" and suggests that smaller, specialized architectures might be more practical for industry work where computing costs are a major concern.

The drive for efficiency and clarity also led to smaller models like EmoBERTTiny (Stigall et al., 2024). Despite its size, this model was just as accurate as many larger models, proving that not every emotion analysis task needs a massive, resource-heavy system. For companies that need to monitor social media in real-time, these lightweight models are very appealing because they lower deployment costs while still providing a high level of analytical detail.

Beyond just being accurate, researchers are now much more focused on explainability. As Borg (2025) argues, LLMs are moving past being "black boxes." For example, a model could do more than just label a tweet about Dhoni's retirement as "nostalgic." It could also explain why, saying, "The text expresses longing for the past by mentioning Dhoni's leadership era." This kind of transparency makes the analysis far more useful for managers, journalists, or policymakers who need to justify decisions based on AI insights.

Research is also quickly expanding beyond just text to multimodal sentiment and emotion recognition. Chiorrini et al. (2022) showed that combining text with visuals—like memes, GIFs, or photos—significantly improves how well a model can detect emotion. This is especially relevant in sports, where fan expression often goes beyond words. Yazhou et al. (2023) looked into multilingual sentiment classification, finding that LLMs can understand culturally specific emotional expressions, which is key for analyzing discourse that uses different languages and slang, like the conversation around Dhoni's retirement.

While the use of large language models (LLMs) in sports-specific applications shows their great potential, it also highlights their limitations. For example, a study by Yang et al. (2025) found that while LLMs and Vision-Language Models (VLMs) were getting better at answering complex sports questions, they still struggled with domain-specific metaphors and cultural context. Terms like "finisher" or "captain cool," which have a lot of emotional meaning in Indian cricket, might be misinterpreted by a model without specific cricket knowledge.

Despite these advancements, some serious methodological and ethical challenges remain. A major concern is algorithmic bias. Since LLMs are trained on enormous amounts of often unfiltered internet data, they can pick up and amplify societal biases related to race, gender, and culture. For example, a model might misclassify emotions in text that uses slang from a marginalized community or reinforce negative stereotypes about a fan group. As a result, researchers are pushing for more Fairness, Accountability, and Transparency (FAT) in these systems to prevent unfair or skewed results (Borg, 2025).

The black box problem still exists. Even with some LLMs that can offer explanations, their internal decision-making processes are often opaque, making it hard to interpret them for critical applications like media reporting or policy-making, where the "why" is as important as the "what."

Another problem comes from relying on zero-shot learning. While efficient, these methods often work without any real-world, verified data, making their results questionable in specific cultural or domain-rich contexts. For instance, Indian cricket fans often use code-switching, sarcasm, or specific idioms that simply don't fit into the general training data of global LLMs.

The temporal drift problem adds another layer of complexity. Social media language changes quickly; a model fine-tuned on 2020 data might not understand the latest slang or memes from 2025. These issues like bias, fairness, and the ever-changing language—show that a consistent, rigorous method and regular updates are essential if emotion analysis is going to remain useful over time.

These points, when you look at them together, really show the maturity and the fragility of emotion analysis that relies on LLMs. On the one hand, these models let researchers get insights from millions of data points in real time—something that was impossible just a decade ago. But on the other hand, it's a reminder that this massive scale shouldn't be a trade-off for ethical integrity, contextual awareness, and a clear, transparent methodology.

In our study, we build on these advancements while acknowledging their challenges. We use GPT-4 for zero-shot emotion detection, combine it with Word2Vec embeddings for better semantic understanding, and apply Agglomerative Clustering to find emerging emotional patterns. This hybrid approach uses the power of LLMs while also addressing issues of interpretability, bias, and domain-specific context. Our work thus stands at the forefront of LLM-enabled emotion analytics in the information systems field.

2.2 Theoretical Foundations in IS and Management

To position our study within the IS discipline, we anchor it in a multi-theoretical framework that captures both the reputational and communicative dynamics of digital emotion analysis. Specifically, we draw on Reputational Capital Theory, Stakeholder Theory, Signaling Theory, and Media Richness/Social Presence Theory. Together, these perspectives provide a robust conceptual lens through which the emotional discourse surrounding Mahendra Singh Dhoni's retirement can be examined.

2.2.1 Reputational Capital Theory

Reputational capital is defined as an organization's accumulated stock of perceptual and relational assets (Fombrun & Shanley, 1990). Unlike tangible resources, reputation is socially constructed and evolves through stakeholder evaluations. A positive reputation enhances legitimacy, customer loyalty, and resilience in times of crisis; a negative reputation, conversely, erodes trust and damages long-term value (Deephouse, 2000).

Applying this lens, public figures such as Mahendra Singh Dhoni embody a form of reputational capital that extends beyond personal achievement to the brands and institutions with which they are associated. For example, Dhoni's image as "Captain Cool" translated into a reputational asset for corporate sponsors and the Indian cricketing establishment. Analyzing emotional responses to his retirement provides methodological insights into how organizations can track fluctuations in reputational capital in real time using digital emotion analytics. In this way, emotion analysis moves beyond abstract categorization and becomes a practical tool for evaluating how reputational capital is created, sustained, or threatened in digital ecosystems.

2.2.2 Stakeholder Theory

Stakeholder theory (Freeman, 1984) argues that organizational success depends on managing the expectations of diverse stakeholder groups. In today's digital economy, stakeholders include not only shareholders, employees, and regulators, but also digitally networked publics who actively shape reputational narratives through platforms like Twitter. These publics cannot be considered passive observers; rather, they are co-producers of meaning, capable of amplifying or contesting reputational signals through likes, retweets, memes, and hashtags.

Monitoring and responding to collective emotions allows organizations to strengthen trust, pre-empt reputational crises, and sustain long-term stakeholder engagement. Dhoni's retirement illustrates this vividly: while fans celebrated his legacy with gratitude and nostalgia, others expressed disappointment or uncertainty about Indian cricket's future. For brands and sports institutions, failure to monitor such emotional currents risks alienating segments of their stakeholder base. Thus, stakeholder theory provides a normative rationale for why digital emotion analytics is indispensable in managing reputation and legitimacy in real time.

2.2.3 Signaling Theory

Signaling theory explains how one party conveys credible information to another under conditions of asymmetry (Spence, 1973). In our context, Dhoni's retirement announcement functioned as a powerful signal—one that was interpreted by millions of stakeholders in diverse ways. For fans, it signaled the end of an era; for the Indian cricket board, it signaled the need for leadership transition; for corporate sponsors, it signaled a shift in brand associations.

The interpretive multiplicity of such signals highlights the importance of capturing their emotional reception. Companies that misread or ignore such signals risk reputational dissonance, while those that respond with sensitivity can reinforce stakeholder trust and legitimacy. By grounding our study in signaling theory, we emphasize that tweets, posts, and memes are not mere noise but signals whose emotional interpretation carries material consequences for reputational capital.

2.2.4 Media Richness and Social Presence Theory

While the above theories emphasize what is being signaled and to whom, Media Richness Theory (Daft & Lengel, 1986) and Social Presence Theory (Short et al., 1976) draw attention to how emotions are transmitted in the digital environment. Media Richness Theory posits that communication channels differ in their capacity to reduce ambiguity and transmit nuanced information, with face-to-face interaction being the richest. On first glance, platforms like Twitter might appear "lean" media, given their character limits and absence of non-verbal cues. However, in practice, Twitter and Instagram have evolved into rich platforms by enabling multimodal forms of emotional expression—hashtags, emojis, GIFs, memes, and short videos that amplify and intensify emotional communication.

This richness directly intersects with Social Presence Theory, which emphasizes a medium's ability to create a sense of being "with another." On social media, fans often feel a heightened sense of intimacy and presence with a public figure like Dhoni, despite the physical distance. The immediate, viral, and participatory nature of platforms creates what scholars term "networked presence": the ability of individuals to feel collectively connected through shared emotional expression (Jeo & Sun, 2020). For example, Dhoni's simple Instagram post announcing his retirement did not merely deliver information; it generated a profound emotional contagion across millions of users who simultaneously expressed nostalgia, pride, grief, and admiration.

By integrating Media Richness and Social Presence theories into our framework, we underscore that the emotional signals we analyze are not isolated linguistic outputs but are shaped by the affordances of the medium itself. A hashtag like #ThankYouDhoni becomes a carrier of collective emotion precisely because Twitter affords rapid amplification, searchable aggregation, and participatory meaning-making. This highlights that digital emotion analytics must account not only for the semantics of content but also for the medium through which content acquires emotional intensity.

2.2.5 Integrative Perspective

Taken together, these theoretical lenses situate our study at the intersection of IS, management, and communication research. Reputational Capital Theory highlights the stakes, Stakeholder Theory underscores the audiences, Signaling Theory emphasizes the interpretive dynamics, and Media Richness/Social Presence Theory foregrounds the communicative environment. By bridging these perspectives, we ensure that our methodological innovations are not only technically sound but also conceptually grounded in the broader socio-technical context in which reputational capital is produced, contested, and transformed.

2.3 Applications of Sentiment and Emotion Analysis in Sports Contexts

The sports industry offers a particularly fertile ground for sentiment and emotion analysis due to its emotionally charged and highly visible nature. Sporting events and athlete-related announcements routinely generate passionate, polarized, and highly expressive digital discourse. Fans express joy, pride, admiration, disappointment, or anger across platforms such as Twitter, Instagram, and YouTube, creating vast real-time datasets. These features make sports a natural laboratory for testing, refining, and applying advanced sentiment analysis methods.

2.3.1 Event-Based Sentiment Analysis

Several studies have investigated emotional reactions to sporting events. Qian et al. (2024) applied aspect-based sentiment analysis to college football stadium experiences, identifying how emotional dimensions such as excitement, frustration, or satisfaction influenced fan ratings and overall satisfaction. Yonghwan and Yuhei (2021) showed that spectator emotions directly impact psychological vigor, linking game outcomes to affective states that shape both immediate and long-term engagement. These works highlight the causal influence of emotions on broader consumer behavior, illustrating how micro-level emotional reactions aggregate into macro-level patterns of consumption, loyalty, and fan identity.

2.3.2 Media-Sentiment Dynamics

Beyond live events, research has shown that the interaction between media narratives and public sentiment is deeply reciprocal. Wang and Sant (2022) explored how media coverage of athlete protests both shaped and was shaped by public sentiment. Their big data analysis revealed a feedback loop where media amplified public emotions while public reactions influenced subsequent media framing. This demonstrates the co-evolutionary relationship between media discourse and collective opinion—an insight highly relevant not only to sports journalism but also to corporate communications, where reputation is continuously negotiated between organizational messages and public feedback.

2.3.3 Self-Construal and Emotional Expression

Emotions in sports contexts are also mediated by cultural and cognitive orientations. Jang et al. (2021) examined how self-construal influences sports consumption, finding that independent self-construal was linked to admiration for skill and achievement, while interdependent self-construal was tied to moral excellence and collective pride. Such findings underscore that emotion in sports is not uniform but shaped by cultural background, identity, and interpretive frames. For global brands that sponsor athletes or sporting events, this indicates the importance of interpreting sentiment contextually, rather than assuming emotional universality.

2.3.4 The Gap in Sports-Oriented IS Research

Despite the richness of these insights, much of the existing research in sports sentiment remains descriptive rather than operational. Studies often identify what emotions are present but stop short of systematically linking these emotions to IS-enabled decision-support tools or organizational strategies. This creates a gap in translating emotional insights into actionable outcomes such as reputational monitoring, stakeholder trust-building, or corporate reporting. Our study addresses this shortfall by bridging emotion analytics with IS applications, demonstrating how sports-related emotions can inform broader organizational strategies for communication, crisis management, and stakeholder engagement.

2.3.5 Transcending the Sports Context: Toward a Generalizable Framework

While our empirical case study is rooted in the sports domain, the emotional dynamics observed are not unique to this context. Rather, they are highly relevant to other high-stakes reputational events in business, politics, and entertainment. The retirement of Mahendra Singh Dhoni, a figure whose personal brand transcended his profession, is analogous to a CEO's retirement, a major corporate merger, or a political figure's public departure. In each scenario, a single event functions as a critical incident that unifies and triggers a wide range of stakeholder emotions.

For instance, a controversial merger could elicit a similar emotional landscape, with employees expressing disappointment and uncertainty, while investors may feel excitement or optimism. Likewise, a successful product launch could generate joy and satisfaction among consumers, while competitors and analysts might adopt more neutral or reflective tones. These parallels demonstrate that emotion-driven reputational dynamics cut across domains.

The value of our methodology, therefore, lies not in its specific application to a sports icon, but in its scalability and transferability. By categorizing and clustering emotions in real time, organizations in any sector can better anticipate risks, reinforce trust, and capitalize on positive narratives (Rathore et al., 2021; Wanniarachchi et al., 2023). This positions our research as a foundational tool for proactive reputational management and crisis communication across industries, extending the relevance of sports-based emotion analysis to the broader information systems discipline.

2.4 Research Gap and Our Contribution

Despite robust progress in LLM-based sentiment analysis and sports applications, significant gaps remain. First, much of the literature focuses on binary sentiment polarity (positive vs negative), overlooking the granularity and multidimensionality of emotions. Second, studies often prioritize model accuracy over interpretability and practical deployment in organizational settings. Third, the link between emotion analytics and business-critical processes—such as corporate reporting, reputational risk management, and stakeholder trust—is underdeveloped.

Our study makes three contributions to addressing this gap:

Methodological Innovation: We demonstrate the use of GPT-4 for zero-shot emotion detection, enabling nuanced classification of stakeholder emotions without requiring pre-labeled data. This scalability is essential for analyzing sudden, high-volume events.

Holistic Emotional Mapping: By integrating GPT-4 outputs with Word2Vec embeddings and Agglomerative Clustering, we identify semantically coherent emotional clusters that go beyond simple sentiment polarity.

Actionable IS Contribution: We translate these clusters into strategic business intelligence, offering tools for real-time reputational monitoring, proactive communication design, and transparent stakeholder engagement. This bridges the divide between technical NLP methods and their practical application in IS and management contexts.

3Methodology

This study adopts a novel hybrid AI methodology to systematically examine public emotional reactions to Mahendra Singh Dhoni's retirement announcement. Unlike traditional sentiment analysis, which often reduces complex reactions into broad categories such as positive, neutral, or negative, our approach leverages the power of Large Language Models (LLMs), semantic embeddings, and unsupervised clustering to uncover richer, more actionable emotional landscapes. The methodology is designed to be scalable, replicable, and adaptable for analyzing other corporate or celebrity-driven events that carry significant reputational implications.

By integrating GPT-4 for zero-shot emotion detection, Word2Vec embeddings for semantic representation, and Agglomerative Hierarchical Clustering for emotion grouping, we construct a framework that goes beyond surface-level sentiment analysis. The approach ensures that nuanced reactions such as nostalgia, admiration, disappointment, or distress are preserved and translated into meaningful insights for reputation management, corporate governance, and strategic communication design.

3.1 Data Collection and Preprocessing

3.1.1 Data Source

The dataset was collected from Twitter (now X), a widely recognized platform for capturing spontaneous, real-time stakeholder sentiment. Specifically, we focused on tweets related to Mahendra Singh Dhoni's retirement announcement on August 15, 2020, a landmark moment in Indian cricket that triggered massive online discourse.

A total of 67,874 tweets were gathered between 15 and 30 August 2020 using relevant hashtags (e.g., #Dhoni, #ThankYouDhoni, #CaptainCool), keywords, and user mentions associated with Dhoni's retirement. Metadata fields such as tweet ID, creation time, source, user description, location, follower count, and retweet status were initially extracted. However, for the purposes of emotion analysis, the text field was retained as the primary input, since it carries the richest emotional signal.

3.1.2 Preprocessing Pipeline

Raw social media data is typically noisy and inconsistent. To ensure data quality and reliability, a rigorous multi-step preprocessing pipeline was applied:

Noise Removal: Deleted URLs, hashtags, user mentions, emojis, numbers, and non-alphabetic symbols.

Deduplication: Removed retweets and near-duplicate tweets to avoid redundancy.

Case Normalization: Converted all text to lowercase to ensure uniformity.

Stopword Filtering: Eliminated high-frequency but semantically unimportant words (e.g., "the," "is," "at").

Irrelevant Mentions: Excluded tweets about Suresh Raina, who coincidentally retired at the same time, to maintain Dhoni-centric focus.

Minimum Context Filtering: Discarded tweets with fewer than three meaningful words to ensure sufficient emotional context.

This process refined the dataset to 36,000 high-quality tweets, representing a cleaner, more representative emotional corpus for subsequent analysis.

3.1.3 Data Assurance and Quality Control

To enhance validity and reliability, multiple assurance mechanisms were applied:

Cross-verification of hashtags and keyword queries to confirm comprehensive data coverage.

Contextual validation by manually inspecting random samples to confirm Dhoni-related relevance.

Exclusion of anomalies, including bot-generated posts and promotional spam.

Together, these measures ensure that the dataset is both contextually relevant and analytically reliable.

3.2Emotion Detection using GPT-4

The main part of this study is the use of GPT-4 in a zero-shot learning framework to detect emotions. Unlike traditional models that need to be trained on labeled data, zero-shot classification allows GPT-4 to figure out emotional categories without any prior training for that specific task. This is a huge help for analyzing real-time events, where getting a labeled dataset is difficult or even impossible.

3.2.1 Model Configuration

Model: GPT-4 (OpenAI API).

Temperature: 0.5 (balancing creativity with determinism).

Max Tokens: 10 (to restrict responses to concise emotion labels).

Prompt Engineering: Each tweet was transformed into a structured prompt asking the model to classify its underlying emotion.

3.2.2 Rationale

We went with GPT-4 for this because it's just so good at understanding context, cultural nuances, and those subtle emotional cues. Traditional methods, which just look for a set list of words, often miss things like sarcasm or figurative language. This is a big problem on Indian social media, where the language is often casual and full of colloquialisms. But GPT-4 is different; it can recognize all those complexities and even figure out blended sentiments, which gives us a much more human-like emotional categorization.

3.2.3 Output

Each tweet we analyzed was tagged with a primary emotional label, such as admiration, nostalgia, disappointment, sadness, or reflection. This diverse emotional vocabulary gave us a solid foundation to start with, which we could then use to group the tweets into broader, more understandable themes.

3.3 Semantic Representation with Word2Vec

To turn our emotional labels into something a computer could read, we used Word2Vec embeddings. This process converted each emotion into a dense vector—essentially a numerical representation based on how words tend to appear together in large amounts of text.

3.3.1 Why Word2Vec?

Captures semantic proximity (e.g., "joy" and "happiness" appear closer in vector space).

Enables contextual grouping of related emotions, reducing noise in clustering.

Provides interpretable vector distances for visualization and cluster analysis.

This embedding process allowed us to retain the subtle shades of meaning between emotions while enabling computational grouping.

3.4 Emotion Clustering

3.4.1 Determining Optimal Clusters

To uncover higher-order emotional structures, clustering was performed. The Elbow Method and Silhouette Score were used to determine the optimal cluster number, with results initially suggesting 7 clusters.

3.4.2 Agglomerative Hierarchical Clustering

We applied Agglomerative Hierarchical Clustering with cosine similarity as the distance metric. This approach does not assume a fixed number of clusters but instead builds a hierarchy, making it ideal for exploring natural groupings in emotions.

3.4.3 Refinement

Initial clustering produced granular but overlapping groups.

By leveraging Word2Vec embeddings, clusters were refined to capture semantic coherence (e.g., grouping joy, pride, happiness together).

Manual review ensured interpretability and alignment with the study's goals.

3.4.4 Final Thematic Clusters

Ultimately, emotions were synthesized into three primary clusters:

Positive & Motivational – admiration, pride, nostalgia, gratitude.

Neutral & Ambivalent – reflection, uncertainty, acceptance.

Negative & Distressed – sadness, disappointment, anger.

This thematic reduction preserved nuance while making results actionable for reputational analysis.

3.5 Visualization with t-SNE

Finally, to visualize the emotional landscape, we employed t-distributed Stochastic Neighbor Embedding (t-SNE) for dimensionality reduction. This enabled us to project high-dimensional emotion embeddings into a two-dimensional space for intuitive interpretation.

The resulting plots revealed distinct separations between positive and negative clusters, with neutral emotions forming transitional bridges. Such visualizations provide managers and analysts with an accessible snapshot of public sentiment dynamics, reinforcing the decision-support value of the framework.

Experiment

To conduct a comprehensive analysis of the emotional landscape surrounding Mahendra Singh Dhoni's retirement, it was necessary to systematically prepare, process, and interpret the dataset before subjecting it to advanced emotion detection methods. The analysis stage not only involved the technical pipeline of preprocessing, model-driven detection, and clustering, but also the interpretive work of validating the results, comparing them with alternative techniques, and linking findings to the broader construct of reputational capital.

4.1 Text Preprocessing

Before analyzing emotional signals, the dataset had to be cleaned to ensure that the content was contextually relevant and semantically interpretable. Twitter data is inherently noisy, often containing spam, abbreviations, emojis, hashtags, and other artifacts that can distort results. Hence, our preprocessing module was designed to enhance both accuracy and reliability.

The preprocessing steps are summarized below, with justification for each:

Handling Missing Values: Tweets containing NaN or incomplete values were removed. Retaining incomplete entries would not only bias the dataset but also compromise semantic embedding quality. In total, approximately 2.3% of the raw data was discarded at this stage.

Removing URLs, Mentions, and Hashtags: Links to videos or articles (e.g., "watch Dhoni's farewell match highlights") and mentions of users (@BCCI, @ChennaiIPL) were stripped. While these elements are useful for network analysis, they contribute little to emotional semantics. If left unremoved, they would inflate clusters around irrelevant tokens such as "http," "pic.twitter," or "@username."

Excluding Tweets about Suresh Raina: Since Raina announced his retirement on the same day as Dhoni, many users tweeted about both simultaneously. To preserve a singular focus on Dhoni, any tweets explicitly mentioning Raina were removed. Pilot testing revealed that leaving these tweets in created false emotional signals—particularly confusion and sadness—that were directed toward Raina rather than Dhoni. This exclusion step increased the internal validity of our findings.

Lowercasing and Removing Non-Alphabetic Characters: To maintain uniformity, all tweets were converted to lowercase and stripped of numbers, punctuation, and non-alphabetic symbols. Without this, "Dhoni," "dhoni," and "DHONI!!!" would be treated as separate tokens, fragmenting semantic patterns.

Emoji and Symbol Removal: Emojis were excluded from the dataset. Although emojis are powerful conveyors of emotion, we opted for removal because LLM-based analysis in our pipeline was optimized for textual rather than visual-symbolic input. A shortcoming here is the potential loss of subtle affective nuances—such as the difference between and —but this decision ensured linguistic consistency.

Duplicate Removal: Retweets and near-identical tweets were eliminated to avoid redundancy. This step was important because Dhoni's retirement triggered mass retweeting of tribute messages, which could have created artificially skewed emotional distributions.

Short Text Filtering: Tweets consisting of only one or two words (e.g., "Legend," "Miss you") were removed. Although emotionally expressive, these tweets lacked sufficient context for meaningful LLM classification. This accounted for roughly 7% of the data.

4.2Emotion Analysis using Large Language Models

Following the preprocessing phase, the cleaned dataset comprising 36,000 tweets was utilized for emotion detection using the GPT-4 model. The emotion detection process was facilitated through an API designed to interact with the OpenAI Chat Completion service. The core function find_emotion was implemented to extract emotional sentiments from each tweet. The code snippet is shown in Figure 1.

In this function, the following parameters were set for the API call:

Model: The gpt-4 model was employed to leverage its advanced natural language processing capabilities for nuanced emotion detection.

Prompt Structure: Each tweet was formatted as a prompt to the model, instructing it to identify the emotion conveyed.

Max Tokens: Set to 10, this parameter limited the response length to ensure concise emotion classification.

Temperature: Configured at 0.5, this setting controlled the randomness of the model's responses, balancing between creativity and accuracy.

```
def find_emotion(text):

    prompt = f"Identify the emotion conveyed in the following
tweet:\n'{text}'\nEmotion:"

    try:
        response = openai.ChatCompletion.create(
            model="gpt-4",
            messages=[{"role": "system", "content": "You are an emotion
detection assistant."},
                      {"role": "user", "content": prompt}],
            max_tokens=10,
            temperature=0.5,
        )
        emotion = response['choices'][0]['message']['content'].strip()
        return emotion
    except Exception as e:
        print(f"Error detecting emotion for tweet: {text}\nError: {e}")
        return "Error"
```

Figure 1: Code Snippet of GPT-4

The system was designed to handle potential exceptions gracefully, allowing for continuous processing of the dataset despite any isolated failures in emotion detection. Through this methodology, a diverse range of emotions was captured for each tweet, providing a comprehensive overview of public sentiment regarding Dhoni's retirement. We obtained various emotions from tweets using large language models, capturing a wide range of emotional expressions present in the data. The results are visually represented in the accompanying Figure 2, showcasing the diverse emotional expressions found within the data.

```
unique_first_words = set(df1['first_word'].dropna())

print(unique_first_words)
```

```
{'betrayal', 'Responding', 'Bias.', 'Empathy', 'shyness', 'Discomfort.', 'Anticipation', 'Sarcasm', 'regret', 'skepticism', 'sa
tisfaction', 'Devotion', 'power', 'Aggression', 'kindness', 'hopeful.', 'Wonder.', 'happiness', 'happiness.', 'insecurity', 'ha
te.', 'respect.', 'compassion', 'surprise', 'Joy.', 'excitement.', 'incomplete', 'patriotism', 'hopeful', 'Optimism', 'optimis
m', 'Happy', 'Shame', 'Pride.', 'Patriotism', 'affection?', 'confidence', 'Contemplation', 'Disgust', 'Competition', 'Understan
ding', 'busy', 'joy', 'Satisfaction', 'independence', 'Contempt', 'sentiment', 'Apology', 'There', 'hate', 'Disrespect', 'confu
sion.', "Can't", 'remorse', 'pride.', 'Neutral\\*', 'suspicion', 'curiosity', 'cannot', 'attachment', 'inappropriate', 'indiffe
rent', 'Sorry.', 'gratitude', 'sadness.', 'Headache', 'good', 'Admiration', 'Jealousy', 'uncertainty\\*', 'It', 'annoyance', 'S
orry,', 'Confusion', 'satisfaction)', 'Craving', 'Irritation', 'This', 'legend', 'Hate.', 'inconclusive', 'loyalty', 'affectio
n', 'joyful', 'Longing', 'excitement', 'Hatred', 'agreement', 'Empathy.', 'Confidence.', 'apathy', 'Excited', 'calmness', 'Sa
d.', 'pride', 'Impatience', 'hunger', 'Prejudice', 'Disappointment', 'No', 'bitterness', 'sentimentality', 'fear', 'disagreemen
t', 'Inspiration.', 'Cannot', 'Indecision', 'Happiness', 'complexity', 'melancholy', 'love', 'Compassion', 'euphoria', 'Love.',
'Surprise.', 'Honor.', 'hope', 'affection.', 'humor', 'confusion', 'sentiment.', 'Love', 'Greed.', 'Confused', 'Caution', 'help
ness', 'Courage', 'concern', 'Kindness', 'hopeful?', 'Inspiration', 'stress', 'desire', 'satisfaction.', 'Patience', 'frustrati
on', 'Anxiety', 'sadness', 'Humor', 'angst', 'neutral', 'Danger', 'Surprise', 'jealousy', 'patience', 'Defensive', 'sympathy',
'irritation', 'Joy', 'Loyalty', 'boredom', 'Longing.', 'N/A', 'Hate', 'embarrassment', 'appreciation', 'trust', 'loyalty?', 're
spect', 'anticipation', 'Support', 'neutrality', 'interest', 'shame', 'Bitterness', 'Respect', 'excited', 'Grief.', 'serenity',
'injustice', 'Envy', 'Difficult', 'uncertainty', 'Determination', 'Worry.', 'Confidence', 'Excitement', 'duty', 'Pride', 'amuse
ment', 'relief', 'Excited.', 'Hopeful.', 'Respect.', 'Fear', 'laughter', 'Sadness.', 'envy', 'disappointment', 'exhaustion', 'S
hame.', 'Fear.', 'gratitude.', 'guilt', 'Serene', "can't", 'Awe.', 'caution', 'anxiety', 'Guilt', 'Anger', 'Gratitude', 'longin
g', 'disgust', 'urgency', 'Numbness.', 'wisdom', 'lust', 'Unclear', 'calm', 'sad', 'Neutral', 'determination', 'anger', 'conten
tment', 'Sadness', 'tranquility', 'empathy', 'Anger.', 'nostalgia', 'happy', 'obsessiveness', 'sarcasm', 'impatience', 'Gratitu
de.', 'Awe', 'neutral.', 'curiosity.', 'romance', 'confused', 'None.', 'peace', 'confident', 'The', 'unsure', 'Caution.', 'pers
istence', 'admiration', 'Frustration', 'Neutral.', 'Curiosity', 'dissatisfaction', "I'm", 'heartbroken', 'Neutral?', 'desperati
on', 'Nostalgia', 'Satisfaction.', 'rage', 'responsibility', 'Power.', 'indifference', 'creativity', 'Inconclusive', 'corruptio
n', 'Indifference', 'worry', 'Detected', 'aggression', 'hatred', 'Inspired', 'greed', 'inspiration', 'ignorance', 'Affection.'}
```

Figure 2: Emotional Expressions

4.3Clustering the identified emotions

After getting various emotions, we systematically analyzed emotional expressions using advanced machine learning clustering techniques to better understand their relationships and inherent structure. Our approach involved several key steps, which we detail below:

Step 1: Visualizing the Optimal Number of Clusters (Elbow Method)

The first step in our clustering process involved determining the optimal number of clusters. We utilized the Elbow Method by calculating the silhouette score for various cluster numbers. The silhouette score measures how well a data point is assigned to its cluster compared to neighboring clusters, providing a metric for cluster quality. We plotted the silhouette scores against different numbers of clusters to visualize the trend. As shown in Figure 3, the silhouette score increased initially and then started to plateau, with the "elbow point" occurring at 7 clusters. This peak or plateau in the silhouette score indicated the optimal cluster count for our dataset, confirming that this was the ideal number of clusters to proceed with for further analysis.

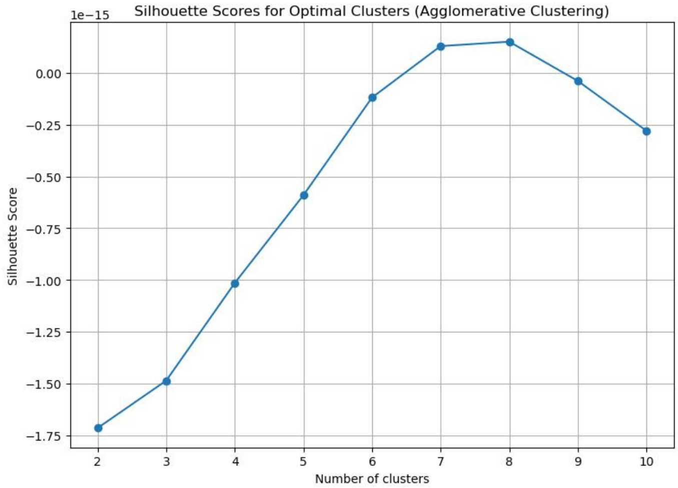

Figure 3: Elbow Method

Step 2: Clustering with Agglomerative Clustering and t-SNE Visualization

Following the determination of the optimal number of clusters, we applied Agglomerative Clustering, a hierarchical clustering technique, with an increased number of clusters set to 7. This adjustment allowed for capturing more granular emotional distinctions. The clustering algorithm grouped data points by iteratively merging clusters based on similarity, resulting in a detailed hierarchy of emotional groupings. As shown in Figure 4, the resulting emotional clusters revealing distinct groupings of emotions based on their semantic similarities. Each cluster represents a unique emotional state, allowing us to observe how various emotions were related to one another in the context of the tweets surrounding Dhoni's retirement.

Cluster 1: dissatisfaction, headache, devotion, attachment, apathy, apology, sorry, injustice, melancholy, excited, appreciation, agreement, happiness, calm, hunger, persistence, peace, sympathy, joyful, indifferent, happy, duty, good, busy, difficult, worry, responsibility, sentiment, cannot, desire, regret, envy, unsure, contentment, fear, na, heartbroken, confidence, irritation, legend, understanding, disgust, loyalty, concern, optimism, sadness, indecision, suspicion, it, love, caution, empathy, indifference, boredom, corruption, annoyance, complexity, contempt, awe, nostalgia, exhaustion, this, independence, hope, cant, affection, surprise, no, responding, hopeful, numbness, betrayal, impatience, sad, admiration, disrespect, creativity, disagreement, patriotism, guilt, craving, confused, inappropriate, sarcasm, shame, hatred, rage, inspired, discomfort, ignorance, neutrality, urgency, defensive, insecurity, excitement, tranquility, bitterness, amusement, euphoria, inconclusive
Cluster 2: relief, courage, remorse, contemplation, calmness, skepticism, jealousy, angst, satisfaction, desperation, anxiety, joy, respect
Cluster 3: stress, serene, uncertainty, inspiration, pride, competition, embarrassment, none, compassion, anger
Cluster 4: wisdom, interest, bias, curiosity, unclear, im, romance, grief, hate, detected
Cluster 5: confident, support, gratitude, greed, trust, patience, obsessiveness, aggression, prejudice, danger
Cluster 6: incomplete, helpness, there, sentimentality, humor, anticipation, power, disappointment, longing, neutral
Cluster 7: confusion, laughter, shyness, determination, wonder, serenity, lust, kindness, honor, frustration

Figure 4: Emotion Clusters

For better visualization of the high-dimensional data, we employed t-SNE (t-distributed Stochastic Neighbor Embedding), a dimensionality reduction technique that effectively preserves relative distances between data points. Using t-SNE, we projected the data into a 2D space, enabling us to observe the distribution and relationships among the clusters. The visualization revealed several distinct clusters, each represented with unique markers for clarity as shown in Figure 5. However, certain clusters exhibited overlaps, reflecting the inherent complexity and interrelated nature of emotions.

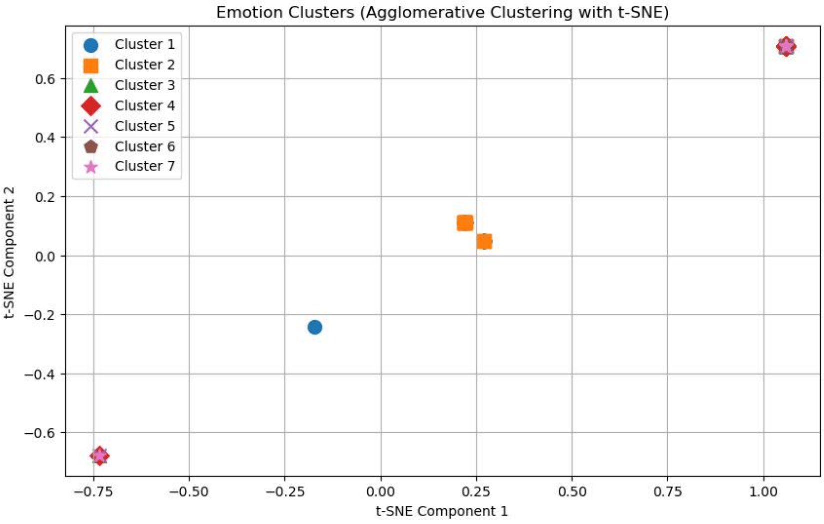

Figure 5: Agglomerative Clustering with t-SNE

Step 3: Refining the Clusters with Pre-Trained Semantic Embeddings

Our initial clustering gave us a solid feel for the emotional landscape, but it wasn't perfect. In the observation there are lot of overlaps between some emotional groups, that showcased that the model was not capturing the subtle differences. To solve this, the approach was modified in this study. The Word2Vec embeddings was introduced, that brought out a much better handle on the relationships between words. Think of it like a smart thesaurus; Word2Vec analyzes how often words appear together and then groups words with similar meanings closer to one another. This allowed us to refine our emotional clusters, making sure we were grouping emotions more accurately based on their real context.

We started by gathering a set of emotion-related words, such as "joy," "anger," "fear," and "hopeful," and used Word2Vec to generate embeddings for them. This allow researchers to go beyond coherent clustering and plunge into the much deeper understanding of language built into the Word2Vec model.

After that, Agglomerative Clustering on the new embeddings was used. The result, as shown in Figure 6, revealed much more nuanced emotional relationships. The words with similar meanings together was grouped, so "joy," "happiness," and "excitement" were clustered because of their strong semantic similarity. This step was crucial for getting a truly accurate picture of the emotions in the data.

Cluster 3: relief, remorse, angst, desperation, dissatisfaction, headache, apathy, apology, sorry, injustice, hunger, worry, regret, fear, heartbroken, irritation, disgust, concern, sadness, indecision, suspicion, caution, indifference, boredom, corruption, annoyance, contempt, exhaustion, surprise, numbness, betrayal, impatience, sad, disrespect, disagreement, guilt, shame, hatred, discomfort, ignorance, urgency, defensive, insecurity, bitterness, confusion, shyness, frustration, power, obsessiveness, aggression, prejudice, interest, romance, hate, detected, serene, competition, compassion, anger
Cluster 1: courage, contemplation, calmness, jealousy, skepticism, satisfaction, devotion, melancholy, appreciation, happiness, calm, persistence, sympathy, joyful, sentiment, desire, envy, contentment, confidence, legend, loyalty, optimism, love, empathy, awe, nostalgia, affection, admiration, patriotism, craving, sarcasm, rage, inspired, excitement, tranquility, amusement, euphoria, laughter, determination, wonder, serenity, lust, kindness, honor, there, respect, humor, disappointment, support, gratitude, trust, patience, danger, bias, im, grief, stress, uncertainty, inspiration, none, anxiety
Cluster 2: attachment, excited, agreement, peace, indifferent, happy, duty, good, busy, difficult, responsibility, cannot, unsure, na, understanding, it, complexity, this, independence, hope, cant, no, responding, hopeful, creativity, confused, inappropriate, neutrality, inconclusive, incomplete, helpness, sentimentality, anticipation, longing, confident, greed, wisdom, curiosity, unclear, pride, embarrassment, joy

Figure 6: Refined Emotion Clusters

To visualize these refined clusters, later the t-SNE was used. This t-SNE visualization clearly portrayed distinct clusters (Figure 7), with more defined separations between positive and negative emotions, as well as nuanced combinations of complex emotions.

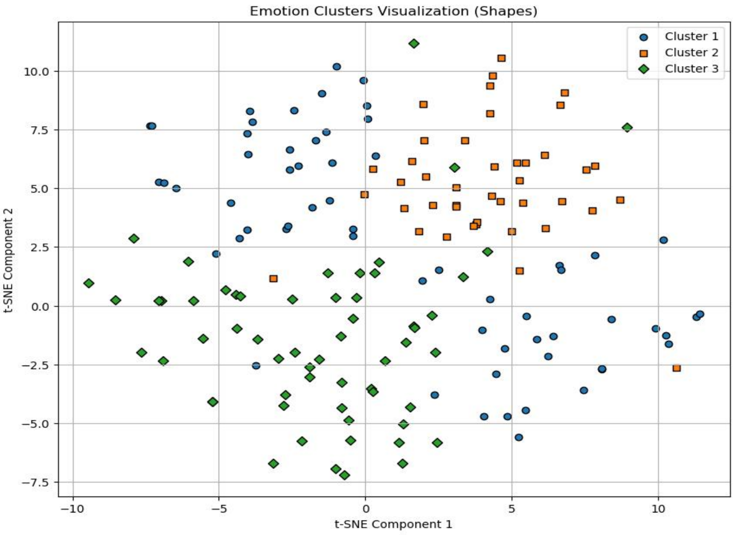

Figure 7: Refined Agglomerative Clustering with t-SNE

This approach was a success; this approach produced a graph that was much better to get deeper meanings and also to observe the relationships between emotions. By combining Word2Vec embeddings with Agglomerative Clustering and then visualizing the results with t-SNE, we were able to group similar emotions more accurately. This led to a more meaningful categorization. After that, we did a manual review and identified three main emotional themes: Positive & Motivational, Neutral & Ambivalent, and Negative & Distressed. Our refined method gave us a much clearer understanding of public sentiment, which was especially important for an event like Dhoni's retirement, where emotions were so varied and subtly overlapped.

Results

The emotion analysis of tweets during the Mahendra Singh Dhoni's retirement revealed three distinct emotional clusters, each carrying significant implications for brand management, stakeholder engagement, and reputational risk.

Cluster 1: Positive & Motivational Emotions: This cluster consisted of uplifting, energizing, and optimistic emotions, such as courage, happiness, joy, optimism, love, confidence, and inspiration. These emotions reflect states of motivation, well-being, and a positive outlook on life. Fans expressed their appreciation for Dhoni's contributions to cricket and their optimism for the future of Indian cricket, celebrating his legacy and the impact he made.

From the perspective of reputational capital, this cluster strengthens dimensions of trust, legitimacy, and loyalty. Dhoni's positive emotional resonance among stakeholders reinforces his symbolic value as a leader, which sponsors and organizations can leverage to build sustained goodwill. In theoretical terms, these emotions enhance reputational capital by accumulating intangible assets of admiration and credibility, which can be mobilized for brand continuity, leadership positioning, and stakeholder alignment.

Cluster 2: Neutral & Ambivalent Emotions: The second cluster captured emotions characterized by uncertainty, indecisiveness, or detachment. Words like indifferent, unsure, neutrality, helplessness, and inconclusive indicated emotional states that were more neutral or lacked strong positive or negative charge. This cluster reflects the ambivalence and emotional uncertainty felt by some fans in response to Dhoni's retirement, as they grappled with the challenges of adjusting to a future without him on the field. In reputational capital theory, such ambivalence represents a zone of fragility and uncertainty. It does not erode reputational capital outright but signals a vulnerable equilibrium where stakeholders neither fully endorse nor reject a brand or figure. This "neutral" space can easily shift toward positivity if managed through reassurance and transparency, or drift into negativity if ignored. Thus, organizations and sponsors must treat neutrality not as absence of sentiment but as a latent risk domain within reputational capital that requires proactive engagement.

Cluster 3: Negative & Distressed Emotions: This third cluster captured a range of negative and painful emotions, from sadness and remorse to anger and frustration. These feelings showcased us the discomfort, dissatisfaction, and emotional turmoil some fans have come across when Dhoni retired. Their reactions, including disappointment, frustration, and even grief, prove just how big an impact his departure had on the Indian cricket community. From a reputational standpoint, these emotions represent a clear risk. If it was left unaddressed, such strong negative feelings can completely takes away at public trust, create negative narratives, and weaken the goodwill a public figure or organization has built over time. The good news is, by acknowledging these feelings and also by communicating with empathy, a potential crisis can be turned into a chance to build renewed credibility.

5.1 Theoretical Synthesis

Taken together, the three clusters illustrate the dynamic nature of reputational capital. Positive emotions strengthen and accumulate capital (trust, admiration, loyalty), neutral emotions signal fragility and the need for stewardship, while negative emotions highlight risks that can deplete the accumulated stock of legitimacy and goodwill. This confirms the conceptualization of reputational capital as a strategic intangible asset that is constantly negotiated through stakeholder emotions, particularly during high-impact events such as the retirement of a public figure.

Managerial Implications

This study provides a practical blueprint for organizations to leverage AI for strategic decision-making and ethical practice.

Real-Time Reputational Risk Management: Managers can use a similar hybrid AI system to monitor public sentiment during critical events. By identifying emerging negative emotions (e.g., frustration, anger) in real time, a company can deploy a proactive crisis management strategy before public discourse escalates. For instance, a brand could issue an empathetic statement acknowledging public sentiment to avoid reputational damage.

Data-Driven Communication Strategy: The nuanced emotional clusters (Positive, Neutral, Negative) provide a more sophisticated guide than simple sentiment. Managers can tailor their communication to specific audiences. For example, a sponsor could create a marketing campaign that resonates with the "Positive and Motivational" cluster by celebrating Dhoni's legacy, while simultaneously preparing a different, more empathetic message for the "Negative and Distressed" group to acknowledge their feelings of loss.

Enhancing Stakeholder Trust and Transparency: The ability to understand and respond to the emotional landscape of the public demonstrates an organization's commitment to listening to its stakeholders. This transparency, particularly in an era of heightened public scrutiny, can build significant trust and serve as a form of non-financial reporting on intangible assets like reputational capital.

Limitations

While this study offers valuable insights, it is important to acknowledge its limitations.

Based on our findings, we see some clear and exciting directions for future research.

First off, it would be really interesting to do a long-term study. This would let us track how emotions shift over an extended period. Organizations could then see if a public reaction to a specific event eventually turns into a lasting change in public perception, which is a huge deal. On a related note, we could broaden this methodology to other social media platforms and languages to get a much more complete, holistic view of what global public sentiment looks like.

Another crucial step is to connect our work to the business world. Researchers could integrate our AI-driven sentiment analysis with financial market data to put a number on how public emotional responses impact things like a company's stock price or market value. Doing this would really prove the critical role reputational capital plays in a business.

Finally, it's vital that future work explores the ethical issues of using AI to analyze public emotions. We need to make sure these technologies are used responsibly to benefit stakeholders, not to manipulate them.

Conclusion

The emotional responses to Mahendra Singh Dhoni's retirement underscore the significant influence of public figures on collective sentiment, stakeholder engagement, and brand perception. This study affirms that Dhoni is not merely a name but an enduring emotional touchstone for stakeholders. Using GPT-4 for zero-shot emotion detection on 36,000 tweets and clustering via Word2Vec embeddings, this study identified three emotional clusters: Positive and Motivational (admiration, nostalgia), Neutral and Ambivalent (uncertainty, reflection), and Negative and Distressed (disappointment, sadness). These findings extend the abstract's argument by demonstrating how major public events, even outside financial markets, shape digital discourse with tangible implications for corporate governance, stakeholder trust, and brand strategy.

From an organization perspective, AI-powered sentiment analysis has emerged as an essential tool for managing intangible assets such as brand reputation and public trust. Foreseeing emotional trends enables organizations to mitigate reputational risks, design responsible communication strategies, and foster ethical brand partnerships. For entities in the public eye, such insights are essential to ensure transparency, accountability, and alignment with societal expectations.

Future Research

Building on these findings, future research could explore several avenues:

Longitudinal Analysis: Future studies could track sentiment shifts over a longer period to understand how public emotions evolve from an initial event-based reaction to a long-term change in public perception.

Cross-Platform and Multi-Lingual Analysis: Expanding the methodology to other social media platforms and languages would provide a more holistic view of global public sentiment.

Integration with Financial Metrics: Researchers could integrate AI-driven sentiment analysis with financial market data to quantify the financial impact of public emotional responses on a company's stock price or market valuation. This would further validate the business-critical role of reputational capital.

Ethical AI in Sentiment Analysis: Future work could explore the ethical implications of using AI to analyze and operationalize public emotions, ensuring that such technologies are used responsibly to benefit stakeholders without manipulating them.

References

1. Borg, M. (2025). Sentiment Analysis for the Masses – How LLMs Changed the Game. IEEE Software. Doi: 10.1109/MS.2024.3478474
2. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language Models Are Few-Shot Learners. In Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS '20), 159, 1877–1901.
3. doi.org/10.48550/arXiv.2005.14165.

4. Chiorrini, C. Diamantini, A. Mircoli, D. Potena, E. Storti (2022). Emotion alberto: Emotion recognition of italian social media texts through bert. 26th International Conference on Pattern Recognition (ICPR), 1706–1711. doi:10.1109/ICPR56361.2022.9956403.
5. Daft, R.L. & Lengel, R.H. (1986). Organizational Information Requirements, Media Richness and Structural Design. Management Science, 32, 554-571.
6. dx.doi.org/10.1287/mnsc.32.5.554
7. Deephouse, D.L. (2000). Media Reputation as a Strategic Resource: An Integration of Mass Communication and Resource-Based Theories. Journal of Management, 26, 1091-1112.
8. doi.org/10.1177/014920630002600602
9. Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. North American Chapter of the Association for Computational Linguistics, 1, 4171–4186. doi: 10.18653/v1/N19-1423
10. Fombrun, C. J., & Shanley, M. (1990). What's in a Name? Reputation Building and Corporate Strategy. Academy of Management Journal, 33, 233-258.
11. doi.org/10.5465/256324
12. Freeman, R.E. (1984). Strategic Management: A Stakeholder Approach. Pitman, Boston.
13. Jang., W.E, Wu., L, & Wen., T.J. (2021). Understanding the effects of different types of meaningful sports consumption on sports consumers' emotions, motivations, and behavioral intentions. Sport Management Review. doi: 10.1016/j.smr.2020.07.002
14. Jeo., N, & Sun., H. P. (2020). Value of Friendship: Instrumental and Sentimental Motivations for Corporate Executives' Networking Behavior, Seoul Journal of Business, 26 (2),50-87. doi:10.35152/snusjb.2020.26.2.002
15. Liu, Z., Yang, K., Cie, Q., Zhang., T, & Ananiadou., S. (2024). "EmoLLMs: A series of emotional large language models and annotation tools for comprehensive affective analysis", Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 33, 5487–5496. doi:10.1145/3637528.3671552.
16. Mehta, M., De, S. J., & Chattopadhyay, M. (2022). Elucidating the role of emotion in privacy-concerns: A text-convolutional neural network (text-cnn)-based tweets analysis of contact tracing apps. Australasian Journal of Information Systems, 26. doi: doi.org/10.3127/ajis.v26i0.3687
17. Pang, B., & Lee, L. (2008). Opinion Mining and Sentiment Analysis. Foundations and Trends® in Information Retrieval, 2, 1-135. Doi:10.1561/1500000011
18. Radford, A., & Narasimhan, K. (2018). Improving Language Understanding by Generative Pre-Training.
19. Rao V S et al. (2025). Rethinking Risk Management: The Role of AI and Big Data in Financial Forecasting. Advances in Consumer Research, 2 (1), 178-185.
20. Rathore, A. K., Maurya, D., & Srivastava, A. K. (2021). Do policymakers use social media for policy design? A Twitter analytics approach. Australasian Journal of Information Systems, 25. doi: 10.3127/ajis.v25i0.2965
21. Qian, T. Y., Li, W., Gong, H., Seifried, C., & Xu, C. (2024). Experience is all you need: a large language model application of fine-tuned GPT-3.5 and RoBERTa for aspect-based sentiment analysis of college football stadium reviews, Sport Management Review, 1–25. Doi: 10.1080/14413523.2024.2386467
22. Short, J., Williams, E. & Christie, B. (1976). The Social Psychology of Telecommunications. John Wiley and Sons Ltd., Hoboken.
23. Siddiqui, M. H. F., Inkpen, D., & Gelbukh, A. (2024). Instruction Tuning of LLMs for Multi label Emotion Classification in social Media content. Proceedings of the Canadian Conference on Artificial Intelligence. https://caiac.pubpub.org/pub/lezimqvm/release/1.
24. Spence, M. (1973). Job Market Signaling. Quarterly Journal of Economics, 87, 355- 374. https://doi.org/10.2307/1882010
25. Stigall, W., Khan, M. a. a. H., Attota, D., Nweke, F., & Pei, Y. (2024). Large Language Models Performance Comparison of Emotion and Sentiment Classification, 2024 ACM Southeast Conference. Doi: 10.1145/3603287.3651183.
26. Vaswani, A., Shazeer, N.M., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., & Polosukhin, I. (2017). Attention is All you Need. Neural Information Processing Systems. Doi: doi.org/10.48550/arXiv.1706.03762
27. Venkatesh, N.A,., & Gokul, C. (2025). Chat Gpt And Super Intelligence: Revolutionizing Human Computer Interaction For The Future Of Work In The Ai Era. Advances in Consumer Research, 2 (4), 2315-2322.
28. Wanniarachchi, V. U., Scogings, C., Susnjak, T., & Mathrani, A. (2023). Hate speech patterns in social media: A methodological framework and fat stigma investigation incorporating sentiment analysis, topic modelling and discourse analysis. Australasian Journal of Information Systems, 27. Doi: 10.3127/ajis.v27i0.3929
29. Wang, W., & Sant, S. L. (2022). A big data analysis of social media coverage of athlete protests, Sport Management Review, 26(2), 224–245. Doi: 10.1080/14413523.2022.2051393
30. Yang., Z, Xia., H, Li., J, Chen., Z, Zhu., Z, & Shen., W. (2025). Sports Intelligence: Assessing the Sports Understanding Capabilities of Language Models through Question Answering from Text to Video.", arXiv preprint.
31. arXiv:2406.14877. Doi: 10.48550/arXiv.2406.14877
32. Yazhou., Z, Wang., M, Tiwari., P, Li., Q, Wang., B & Qin., J. (2023). DialogueLLM: Context and Emotion Knowledge-Tuned LLaMA Models for Emotion Recognition in Conversations, arXiv preprint. arXiv:2310.11374.
33. Yonghwan., C & Yuhei., I. (2021). Spectator emotions in predicting psychological vigor: emotional meta experience and affect valuation

perspectives. Sport Management Review, 24:4, 594-619. DOI: 10.1080/14413523.2021.1880759.