# Profiling Sentiment Archetypes of Popular Twitch Emotes: A Comparative Analysis of Rule-Based Segmentation and Unsupervised Clustering Techniques

Calvina Izumi[1,*], Soeltan Abdul Ghaffar[2], Wilbert Clarence Setiawan[3]

[1]School of Management Business, Ciputra University, Surabaya, Indonesia

[2]Department of Marine Information Systems, Universitas Pendidikan Indonesia, Bandung Indonesia

[3]Faculty of Informatics Engineering, Universitas Taruma Negara, Jakarta, Indonesia

## ABSTRACT

Live-streaming platforms like Twitch have fostered unique communicative ecosystems where non-linguistic tokens, or "emotes," are central to expressing collective sentiment. However, conventional sentiment analysis tools, designed for standard text, largely fail to capture the nuanced and context-dependent meaning of these symbols. This paper addresses this gap by profiling sentiment archetypes of popular Twitch emotes through a comparative methodological framework. We contrast a theory-driven, rule-based segmentation approach, where emotes are manually assigned to predefined sentiment categories, with a data-driven, unsupervised clustering approach that groups emotes based on their contextual usage patterns. Applying this dual analysis to a dataset of Twitch chat messages, we constructed a feature set for each emote incorporating both textual context (TF-IDF) and behavioral metrics. A K-Means clustering algorithm was then used to identify emergent archetypes from the data. Our results reveal a profound divergence between the two methods, quantified by a near-zero Adjusted Rand Index (ARI) of 0.0012. This indicates that an emote's prescribed semantic meaning has virtually no correlation with its functional co-usage in practice. The clustering algorithm successfully identified coherent, function-based groups, such as a "Hype/Spam" cluster and a "Nuanced Reaction" cluster, which are not captured by the rule-based taxonomy. We conclude that emote meaning is not static but is an emergent property of community practice, defined primarily by pragmatic function rather than semantic content. This finding highlights the critical limitations of applying traditional sentiment analysis to dynamic digital cultures and underscores the necessity of data-driven, context-aware methods for understanding online communication.

**Keywords** Clustering, Computational Linguistics, Emotes, Sentiment Analysis, Twitch

Submitted 16 July 2025
Accepted 6 August 2025
Published 1 September 2025

*Corresponding author
Calvina Izumi,
calvinaizumi01@gmail.com

Additional Information and Declarations can be found on page 198



## Introduction

The rise of Twitch as a prominent live-streaming platform reflects broader transformations in digital communication, characterized by interactive participation and community engagement. Given its origins in the gaming community, Twitch has rapidly evolved to encompass diverse content, including political discourse and social issues, which highlights its critical role in contemporary media consumption. The integration of emotes within Twitch's

**How to cite this article:** C. Izumi, W. C. Setiawan, S. A. Ghaffar, "Profiling Sentiment Archetypes of Popular Twitch Emotes: A Comparative Analysis of Rule-Based Segmentation and Unsupervised Clustering Techniques r ," *J. Digit. Soc.*, vol. 1, no. 3, pp. 183-200, 2025.

communicative framework serves as a significant socio-emotional marker that facilitates nuanced interactions among users, extending beyond mere textual communication.

Twitch's function as a live-streaming platform has fundamentally altered user interactions. Its participatory nature fosters not only a sense of community but also enables emotional expressions often absent in traditional forms of online communication. Emotes—customized graphical representations—allow users to convey complex sentiments rapidly, creating a shared lexicon within streams that enhances collective audience behavior. For instance, Catá notes that the interplay of emotes among viewers reflects larger structural dynamics, paralleling issues of race and representation in the platform's expansive culture, thus showcasing the socio-political dimensions of digital interactions [1].

In analyzing Twitch's role during heightened societal tension, such as the COVID-19 pandemic, research indicates that the platform's usage surged as people sought connection and solace through shared viewing experiences [2]. Findings reveal that emotes not only facilitate communication but also serve as tools for emotional support, particularly in discussions around mental health [3], [4]. By leveraging emotes, users can signal empathy, agreement, or humor in real-time, ultimately fostering a sense of belonging, as highlighted by engagements within small, intimate stream communities [5], [6].

Furthermore, the dynamics of Twitch allow for a unique blend of personal connection and mass communication, wherein streamers often navigate their roles as entertainers, confidants, and community leaders [7], [8]. The ability of viewers to participate through chat interactions—be it through emotes or text—forms a layer of engagement that influences the nature of discussions and the perceived emotional state of the streamer [9]. Research shows that such interactions contribute to the overall viewer experience, as the responsive nature of chat aids in developing robust community ties and enhances the meaningfulness of digital interactions [7].

The challenge of conducting sentiment analysis, particularly in the context of platforms like Twitch that incorporate non-linguistic tokens such as emotes, is a significant issue in contemporary computational linguistics. Conventional sentiment analysis predominantly relies on textual data, often neglecting the emotional cues presented by non-verbal expressions like emotes. This gap is crucial as emotes significantly influence viewer interactions and community dynamics [10].

Emotes can convey nuanced emotional states that textual words might not fully encapsulate, complicating the sentiment classification process. Existing methodologies struggle to categorize these non-linguistic tokens appropriately due to their unique properties and contexts [11]. Traditional models typically utilize lexicons composed of sentiment-bearing words, which inadequately capture non-linguistic sentiments [12]. For example, while Jauk et al. highlight the importance of context in sentiment analysis, their focus on acoustic features does not directly address the incorporation of visual tokens like emotes. This indicates that the lack of systematic categorization of emotes hinders sentiment analysis's effectiveness on digital platforms like Twitch, where such elements are prevalent [11].

A significant gap in current research is the absence of a robust framework to systematically categorize and analyze sentiment conveyed through emotes.

Although sentiment analysis has advanced with new algorithms and models focusing on linguistic elements, the interplay between text and visual symbols like emotes requires further investigation [13], [14]. Previous research by Lei et al. underscores that even advanced networks that integrate sentiment resources frequently overlook non-verbal cues, thereby missing the complexities presented by multi-modal inputs commonly found in interactive platforms [12].

This lack of systematic categorization not only limits understanding user interactions but also restricts opportunities for developing more nuanced sentiment analysis tools capable of recognizing and interpreting emotional expressions depicted through emotes. The incorporation of such methods could yield richer insights into audience sentiment, enhancing our comprehension of community behavior within platforms like Twitch [15].

The primary objective of this research is to profile the sentiment archetypes of popular emotes used on the Twitch live-streaming platform. To achieve this, we employ a comparative analytical approach, directly contrasting two distinct methodologies for categorization. The first is a rule-based segmentation, which relies on a handcrafted taxonomy to assign emotes to predefined archetypes based on established community understanding. This is compared against a data-driven, unsupervised clustering method that groups emotes based on their contextual usage patterns, allowing for the discovery of emergent, community-specific archetypes without prior assumptions.

This study makes two primary contributions to the fields of computational linguistics and digital media studies. First, it provides a comprehensive and replicable methodological pipeline that combines exploratory data analysis (EDA), a human-curated rule-based taxonomy, and unsupervised clustering to create a holistic framework for analyzing non-linguistic digital tokens. Second, and more significantly, this paper offers crucial empirical insight into the divergence between theory-driven sentiment categories and organic, usage-based community patterns. By quantifying the misalignment between these two approaches, we illuminate the dynamic and context-dependent nature of meaning-making in online communities.

## Literature Review

### Sentiment Analysis in Social Media

Traditional sentiment analysis methods have predominantly focused on text-based data, employing lexicon-based approaches, models like VADER (Valence Aware Dictionary and sEntiment Reasoner), and deep learning techniques to assess users' emotions through language. However, such methods face considerable limitations when confronted with non-linguistic tokens such as emojis and emotes, which carry significant emotional weight but are often overlooked in conventional frameworks.

Lexicon-based methods, including VADER, primarily rely on predefined dictionaries of words associated with specific sentiments, which may not encompass the full breadth of emotional expression conveyed through emojis or emotes. For instance, while textual sentiment classes typically categorize opinions as positive, negative, or neutral, emojis and emotes introduce additional nuances that these lexicons fail to capture effectively. This oversight is particularly impactful on platforms like Twitch, where visual expressions significantly influence user interactions and overall sentiment [2]. Traditional frameworks thus miss valuable contextual elements that emojis and emotes provide, such as irony, humor, or sarcasm, which can drastically alter the intended sentiment of a message [16].

Moreover, deep learning models, while powerful in analyzing large datasets for pattern recognition, often do not integrate non-verbal tokens into their training sets. As a result, these models may misinterpret or fail to account for the significance of emotes in conveying sentiment, leading to skewed or incomplete analyses of the data [17]. The inability to incorporate these non-verbal elements limits the efficacy of sentiment assessment, particularly within communities where such expressions are commonplace, like Twitch, where engagement and emotional resonance are critical [18].

Research indicates that the emotional context established by emotes is crucial for interpreting community dynamics and sentiments accurately. For example, engagement on platforms like Twitch can help explain emotional states that may not be expressed solely through words, reflecting the complexity of user interactions and the significance of visual symbols in communication [19]. Therefore, a dedicated approach focusing specifically on identifying and categorizing the sentiment conveyed through emojis and emotes is essential. Such methodologies could enhance the understanding of online interactions across platforms, enabling deeper insights into community behavior and user sentiment [20].

### Emotes as Cultural Artifacts

The study of emotes and emojis within the digital communication landscape reveals their multifaceted function as cultural artifacts that significantly enhance emotional expression. Platforms such as Twitch, Twitter, and Discord utilize these non-verbal symbols to convey sentiment, facilitating emotional contagion and collective affect among users. Understanding the semantic roles of these visual tokens is crucial in recognizing their impact on interpersonal and community dynamics in virtual environments.

Previous research indicates that the semantics of emojis and emotes diverge from traditional text-based communication forms. For instance, emotes on Twitch are not merely decorative; they serve specific communicative purposes, such as signaling reactions to events, enhancing humor, or fostering a sense of community cohesion [21]. This dynamic is particularly illustrative of "emotional contagion," where collective affect spreads within chat rooms and streams. The presence of specific emotes can amplify viewers' emotional responses, creating an environment where emotions are collectively experienced, which suggests that digital symbols can evoke social emotional charges and foster connections among users [2].

Furthermore, as noted by Dolin et al., emotes can function as linguistic neologisms, with varying meanings that evolve over time within specific communities [22]. This fluidity presents challenges for sentiment analysis tools, which typically rely on fixed lexicons and may not accommodate the evolving significance of emotes. During significant global events, such as the COVID-19 pandemic, Twitch witnessed substantial user engagement as individuals flocked to express feelings and solidarity through emotes, indicating robust emotional responses that traditional text analysis struggles to encapsulate effectively [2].

Moreover, the collective dynamics fostered by emotes on live streams indicate a shift in social norms regarding emotional expression. Such platforms capitalize on the immediacy of interaction, where viewers can deploy emotes to mirror or heighten their emotional states and those of the streamer, leading to a shared sense of belonging and identity among participants [23]. This collective engagement is visualized through the vibrant and diverse use of emotes, which represent communal experiences and sentiments unique to each streaming community.

The literature underscores a critical need for developing frameworks that address the specific emotional semantics of emotes and emojis in digital communication. The challenges of contextual understanding and sentiment categorization of these non-linguistic tokens signal a gap in existing sentiment analysis methodologies, pushing for interdisciplinary approaches that incorporate emotional intelligence and cultural studies to enhance comprehension in rapidly evolving digital interactions [24].

## Rule-based vs. Data-driven Methods

The discourse surrounding sentiment analysis in social media often contrasts rule-based lexicon methods, such as VADER and LIWC, with data-driven techniques like unsupervised clustering. Each methodology has its strengths and limitations, particularly concerning the representation and interpretation of complex user emotions and behaviors.

Rule-based sentiment analysis utilizes predefined dictionaries, such as VADER (Valence Aware Dictionary and sEntiment Reasoner) and LIWC (Linguistic Inquiry and Word Count), to assess the emotional content in text, relying heavily on lexical and grammatical rules. These lexicons provide clear sentiment scores by matching words or phrases with associated emotional weights, enabling efficient and interpretable sentiment analysis. However, their effectiveness diminishes when addressing non-linguistic tokens such as emojis or emotes, which convey sentiment nuances beyond simple word associations. For instance, while VADER is effective in social media contexts, it may not distinguish the varied emotional implications of similar emojis used in different contexts, thus leading to potential misinterpretations.

Moreover, rule-based lexicons often struggle with the evolving nature of language in digital communication. As new slang and expressions emerge—particularly within platforms like Twitch and Discord—these static lexicons may become outdated, necessitating continuous updates that are resource-intensive. This limitation is critical as digital expressions become progressively more fluid and context-dependent, yielding the necessity for a more dynamic analytical approach.

In contrast, data-driven approaches, particularly those employing unsupervised clustering techniques, provide an alternative methodology for discovering latent user behaviors. By analyzing large datasets without predefined labels, these methods can uncover inherent patterns and structures that may not be evident through rule-based systems. For example, unsupervised algorithms such as clustering can effectively identify distinct user groups based on shared behavioral traits, enabling insights into collective sentiment dynamics within communities.

While unsupervised methods have shown promise in detecting user behavior patterns, they also face challenges, particularly related to their interpretability.

Unlike rule-based methods that provide straightforward sentiment categorizations, the results of unsupervised clustering can be complex and less transparent, making it difficult for analysts to assign meaningful interpretations to the identified clusters.

Furthermore, the blend of supervised and unsupervised methods can yield better results. By incorporating human-annotated datasets to train models, hybrid approaches can enhance the identification of sentiment while still benefiting from the adaptability of data-driven techniques. Such models could significantly improve the accuracy of behavior identification while maintaining the capacity to adapt to evolving user expressions.

### Research Gap

The existing literature on sentiment analysis and user behavior in digital communication platforms reveals a notable gap in comparative studies that explore the alignment between rule-based archetypes—like those employed in sentiment lexicons—and cluster-based community dynamics. Although significant advancements have been made in understanding user interaction through both rule-based lexicons, such as VADER and LIWC, and data-driven methods like unsupervised clustering, there is a scarcity of studies that systematically compare these methodologies.

Rule-based sentiment analysis methods utilize lexicons to classify sentiments based on predefined emotional weights assigned to specific words or phrases. Tools like VADER have been particularly effective in assessing sentiment in social media contexts, where language is often informal and rich with emoticons and abbreviations. However, these lexicons may struggle to capture the evolving semantic meanings that arise with new forms of communication, such as emotes or memes, which convey complex emotions and cultural contexts. Morina and Bernstein demonstrated that cultural artifacts like image memes significantly influence community engagement and sentiment, showing that traditional methods often lack the adaptability required to keep pace with emergent discourse [25].

Unsupervised clustering approaches aim to discover latent patterns in user behavior without relying on predefined labels. These methods can reveal complex user communities characterized by shared emotional experiences and interests. However, since they rely on statistical properties, they often lack the granular emotional insights that rule-based systems can provide. The discrepancies between sentiments categorized by lexicons and the emotional expressions captured within clustered communities pose crucial questions around the alignment between these two approaches.

Given the distinct strengths and weaknesses of rule-based and data-driven methods, comparative analyses are essential to understanding how these methodologies can complement each other in interpreting user sentiment. This investigation is critical for building a more comprehensive framework for sentiment analysis that recognizes the multifaceted nature of user emotions, especially in communities that heavily utilize non-linguistic tokens such as emotes. Previous work has examined cultural variances in customer emotions within e-commerce and social platforms, highlighting a need to apply similar comparative methodologies to the analysis of online communities and digital interactions [26] .

Furthermore, studies focused on emotional design and cultural response suggest that enhancing user engagement requires an understanding not only of individual sentiments but also of collective emotional behaviors within these digital spaces [27], [28]. Bridging the gap between rule-based lexicons and data-driven clustering could foster new insights into how user sentiments disseminate and evolve across platforms, such as Twitch, where community dynamics play a critical role in communication practices.

# Method

This study employs a comparative computational framework to profile the sentiment archetypes of Twitch emotes, contrasting a predefined, rule-based taxonomy with an emergent, data-driven clustering approach. This dual methodology is designed to reveal the potential divergence between the intended, dictionary-like meaning of an emote and its actual, community-driven usage patterns. The entire pipeline, from data preprocessing to comparative evaluation, was implemented in Python 3.9 using a suite of scientific computing libraries, including pandas for data manipulation, scikit-learn for machine learning tasks, and VADER for lexical sentiment analysis. The methodology is divided into four main stages: (1) data preprocessing and feature extraction; (2) exploratory data analysis; (3) rule-based archetype assignment; and (4) unsupervised clustering and comparative evaluation.

## Dataset and Preprocessing

The analysis began with a dataset of Twitch chat messages provided in a single CSV file, representing a corpus of raw, unstructured text typical of live-streaming environments. To ensure the pipeline's robustness and reusability, it was engineered to automatically detect the relevant columns for message content, timestamp, user, and channel by searching for a list of common candidate names (e.g., "message," "text," "chat"). This step obviates the need for manual dataset configuration.

The core preprocessing challenge involved the accurate identification and extraction of emotes. Recognizing that a static list would fail to capture the fluid and ever-evolving lexicon of individual streaming communities, we adopted a hybrid strategy. A comprehensive list of known emotes was first generated by combining a curated seed list of popular global emotes (e.g., Kappa, PogChamp) with a dynamic discovery process. This process tokenized all messages and identified potential new emotes based on a set of orthographic heuristics designed to mimic their common structures, such as CamelCase (peepoHappy) or all-caps (KEKW). To reduce noise from random capitalization, this heuristic was coupled with frequency constraints, including any token appearing with a minimum frequency of 5 (min_freq=5) and the top 400 most common tokens overall (top_cap=400). This ensures that only tokens with established usage are classified as emotes. For each message, a list of its contained emotes was extracted and stored. Concurrently, a cleaned "context" version of each message was created by programmatically stripping it of all identified emotes and standard Unicode emojis. This crucial step isolates the surrounding textual environment, preventing the emotes themselves from confounding the VADER sentiment analysis, which is calibrated for natural language rather than ideograms.

## Exploratory Data Analysis (EDA)

To gain foundational insights into the dataset's characteristics, a series of exploratory analyses were conducted before formal modeling. First, the frequency distribution of all discovered emotes was calculated and visualized. This step is critical for understanding the power-law distribution typical of linguistic elements, where a small number of emotes account for a vast majority of usage, informing which emotes are most central to the platform's culture. Second, the VADER (Valence Aware Dictionary and sEntiment Reasoner) sentiment analysis tool was applied to the cleaned textual context of each message to generate compound, positive, and negative sentiment scores. The distribution of these compound scores was plotted in a histogram to assess the overall sentiment polarity of the chat environment. This analysis served as a direct test of the hypothesis that conventional sentiment lexicons, when applied to the short, meme-dense, and often sarcastic nature of Twitch chat, would yield a distribution heavily skewed toward a neutral score of 0.0. Finally, a co-occurrence heatmap was generated for the top 20 most frequent emotes (topn=20). This visualization reveals which emotes tend to appear together in the same messages, providing a preliminary look at potential "emote syntax" and functional relationships, such as emotes used for amplification (LUL followed by OMEGALUL) or to form common affective phrases.

## Rule-Based Sentiment Archetype Profiling

The first primary analytical approach involved assigning each unique emote to a sentiment archetype based on a predefined, human-curated taxonomy. This expert-driven method is grounded in established community understanding of emote meanings. The rulebook mapped dozens of common emotes to descriptive archetypes (e.g., "Laughter/Sarcasm," "Hype/Excitement," "Sadness/Empathy") and an associated polarity score (+1 for positive, -1 for negative, 0 for neutral or highly context-dependent). This approach acknowledges that the meaning of emotes is culturally specific and cannot be readily inferred by general-purpose sentiment dictionaries. To handle emotes discovered dynamically that were not in the rulebook, a data-informed fallback rule was implemented. For these unknown emotes, the emote was assigned to a "Positive," "Negative," or "Neutral" archetype based on the average VADER compound score of all message contexts in which it appeared, using thresholds of +0.3 and -0.3. This hybrid technique combines deep domain knowledge with an automated, data-driven method for scalability. The final output of this stage was a comprehensive table detailing each emote's frequency, average contextual sentiment, and its assigned rule-based archetype, representing a "ground truth" of intended meaning.

## Unsupervised Clustering and Comparative Evaluation

The second, contrasting approach sought to derive sentiment archetypes empirically from usage data using unsupervised machine learning, operating under the premise that an emote's meaning is defined by the contexts in which it is used. To achieve this, a sophisticated feature matrix was constructed where each row represented a unique emote. This matrix integrated two distinct types of information: textual context and behavioral metrics. Textual features were created by aggregating all context strings for a given emote into a single document and then vectorizing these documents using a Term Frequency-Inverse Document Frequency (TF-IDF) model, capped at a maximum of 1,500 features (max_features=1500) to maintain computational feasibility. To ensure robustness against datasets with sparse textual context, the vectorizer was programmed with a fallback from a word-based model to a character n-gram model (with n-grams ranging from 3 to 5 characters).

These textual vectors were then augmented with a set of nine scaled numeric features summarizing the emote's pragmatic usage patterns. These included its average contextual VADER scores, the average length of messages it appears in, and the average prevalence of linguistic cues for excitement (exclamation points), inquiry (question marks), and emphasis (uppercase letters). This feature engineering strategy creates a holistic representation of each emote, capturing not only what is said around it but how it is said.

Using this combined feature matrix, the K-Means clustering algorithm was employed to group emotes into distinct, data-driven clusters. The optimal number of clusters, k, was determined by iterating through k values from 3 to 12 and selecting the value that maximized the mean silhouette score. This score was calculated using the cosine distance metric, a choice particularly well-suited for high-dimensional and sparse text data as it measures the angle between vectors rather than their Euclidean distance. For reproducibility, the K-Means algorithm was configured with a fixed random state (random_state=42) and was run 20 times with different centroid initializations (n_init=20) for each value of k, ensuring the stability of the final clusters.

Finally, the outputs of the two approaches were compared both quantitatively and qualitatively. The quantitative alignment was measured using the Adjusted Rand Index (ARI), a metric that calculates the similarity between two data clusterings (in this case, the rule-based archetypes and the K-Means labels), corrected for chance. A low ARI score would signify a strong divergence between prescribed meaning and actual use. For qualitative analysis, the clusters were visualized by projecting the high-dimensional feature matrix into a 2D space using Principal Component Analysis (PCA) and generating a scatter plot. Furthermore, word clouds were created from the aggregated contexts of each cluster to provide an interpretable summary of its dominant linguistic themes. This dual analysis provides a rich framework for comparing prescribed sentiment categories with emergent community usage patterns.

## Result and Discussion

### Exploratory Data Analysis (EDA) Results

Before conducting the primary comparative analysis, an exploratory data analysis was performed to understand the fundamental characteristics of the dataset. This stage involved examining the frequency of emotes, the distribution of sentiment in the surrounding text, and the co-occurrence patterns between different emotes. The findings from this EDA were crucial for validating the study's core hypotheses and shaping the subsequent modeling steps.

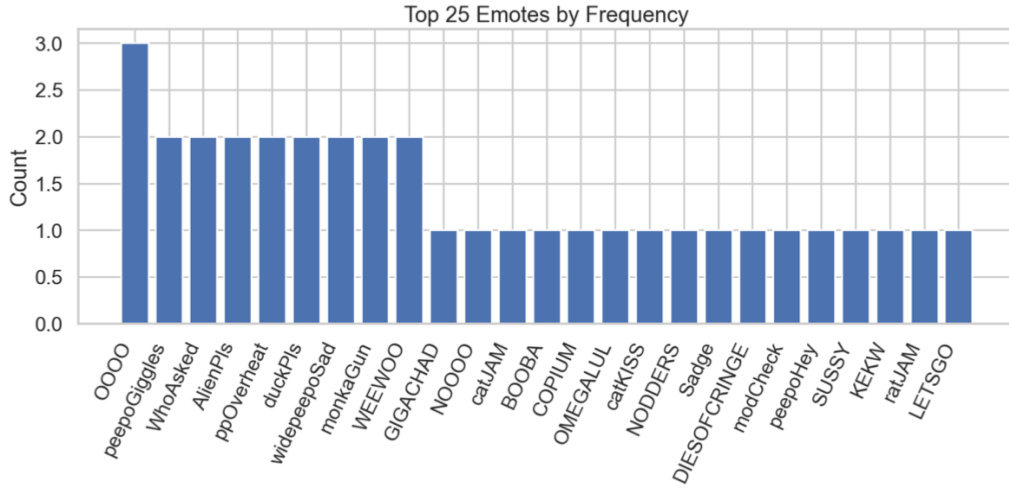


**Figure 1** **Emote Frequency Distribution**

The initial analysis of emote frequencies revealed a distribution typical of linguistic data, characterized by a long tail where a few items dominate the usage. As shown in figure 1, the emote OOOO is the most frequent in this dataset, followed by a group of emotes with similar counts, including peepoGiggles, WhoAsked, AlienPls, and ppOverheat. This skewed distribution, often referred to as a power-law or Zipfian distribution, is significant because it indicates that a relatively small number of emotes form the core vocabulary of this specific chat log. Understanding which emotes are most prevalent is essential for interpreting their cultural importance and ensuring that both the rule-based and clustering analyses are grounded in the most relevant data.

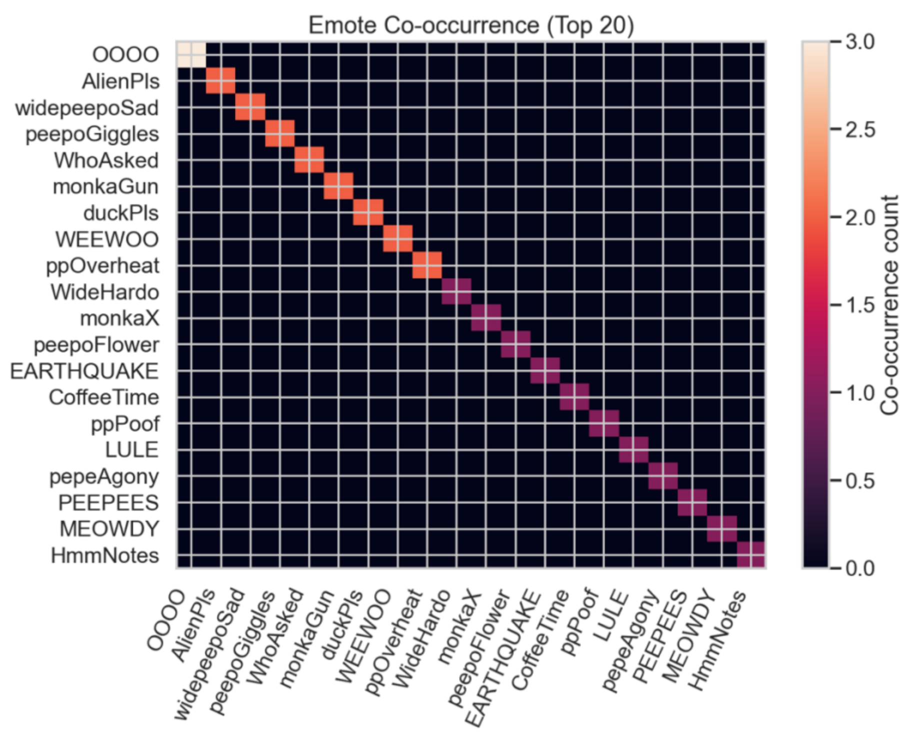


**Figure 2** **Emote Co-occurrence Patterns**

The co-occurrence heatmap was generated to investigate which of the top 20 emotes tend to appear together in the same messages, offering insights into potential "emote phrases" or functional pairings. The heatmap shows a strong diagonal line, which is expected, as it represents each emote co-occurring with itself (i.e., its presence in a message). The off-diagonal cells, which represent the co-occurrence between two different emotes, are uniformly dark, indicating a co-occurrence count of zero or near-zero across the board. For this particular dataset, this result suggests that the top emotes are most often used in isolation rather than in specific, frequently repeated pairs. While some emotes on Twitch are known to be used in sequence (e.g., monkaS followed by monkaW), this pattern was not prevalent among the most common emotes within this specific corpus. This highlights the importance of analyzing individual emote contexts, as complex, multi-emote phrases appear to be rare in this sample.

## Exploratory Findings and VADER's Limitations

The analysis of VADER sentiment scores on the textual context surrounding these emotes was particularly revealing and served as a crucial validation of our core problem statement. The resulting sentiment histogram in figure 3 showed a sharp, narrow peak centered precisely at a compound score of 0.0. This indicates that the vast majority of messages were classified as affectively neutral, a finding that is facially inconsistent with the highly emotional and reactive nature of live-stream chat. This outcome empirically demonstrates that conventional sentiment analysis lexicons, which are trained on more formal text corpora like product reviews or news articles, are ill-equipped to handle the non-standard grammar, irony, sarcasm, and meme-based lexicon of Twitch. Their failure validates the need for domain-specific, context-aware methods to decipher affective meaning in this environment.

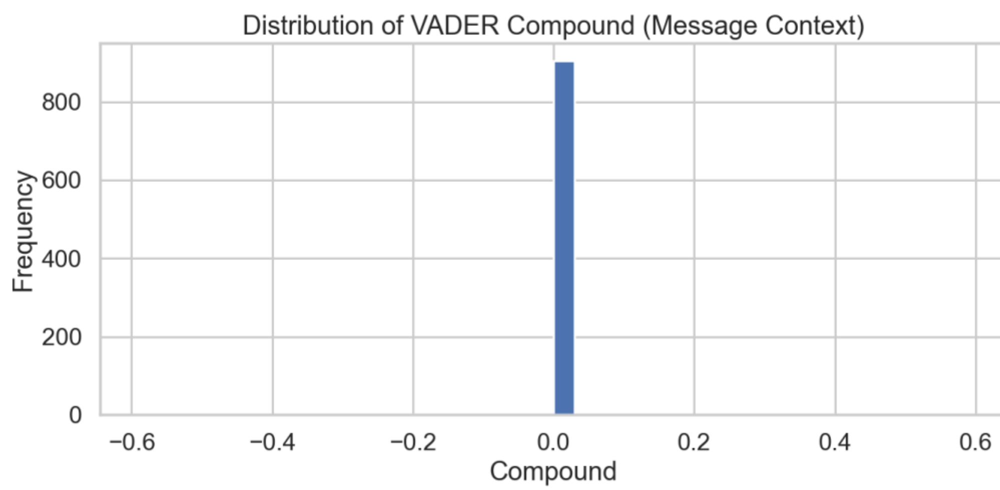


**Figure 3** **Distribution of VADER Compound**

## Rule-Based Archetype Profiling

The rule-based assignment process successfully categorized all 433 unique emotes into 14 distinct archetypes, ranging from "Hype/Excitement" to "Sadness/Empathy." As evidenced by the data in top_emotes_per_rule.csv, this method produced coherent and intuitively correct groupings based on established community knowledge. For instance, the "Anxiety/Tension" archetype correctly grouped the semantically related trio of monkaS, monkaW, and monkaOMEGA, while the "Applause" archetype logically contained peepoClap and peepoCheer. This approach excels at capturing the canonical or intended meaning of well-established emotes, reflecting a shared cultural dictionary among users.

However, the severe limitations of this manual approach became apparent in its reliance on a fallback rule for emotes not explicitly defined. A staggering 381 out of 433 emotes (88%) were not present in the predefined rulebook. Due to the aforementioned failure of VADER to detect sentiment, these were almost universally labeled as "Neutral/Contextual." This created an overwhelmingly large and functionally useless neutral category, masking the rich diversity of the emote lexicon. This finding highlights a critical scalability problem: manually curating taxonomies for a rapidly evolving and expanding set of cultural symbols is an intractable task. The result is a system that is accurate for a small core of emotes but fails to classify the vast "long tail" of the emote ecosystem.

## Unsupervised Clustering Outcomes

In stark contrast to the rule-based method, the K-Means clustering algorithm partitioned the emotes into three distinct clusters based on their holistic usage profiles. The optimal number of clusters was determined to be k=3, which achieved a silhouette score of 0.444. While not a perfect score, this value suggests that the clusters are moderately dense and reasonably well-separated, indicating that the algorithm found genuine structure within the data. The clusters varied significantly in size, with Cluster 0 containing 267 emotes, Cluster 1 containing 128, and Cluster 2 being a very small, specific group of 8 emotes. Qualitative analysis of the top emotes within each cluster allows for a rich functional interpretation that moves beyond simple sentiment.

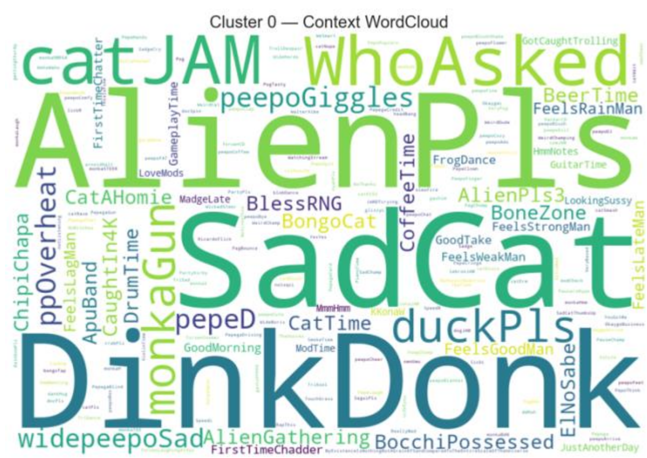


**Figure 4** **WordCloud of Cluster 0**

Cluster 0: The "Pepe" Emote Vernacular & Nuanced Reactions. This largest cluster represents the core of modern Twitch emote culture, dominated by a wide variety of Pepe the Frog variants and other general-purpose reaction emotes, as shown in wordcloud in figure 4. It includes emotes for complex emotional states like sadness (widepeepoSad, PepeHands), anxiety (monkaS), contemplation, and niche reactions. Its sheer size and diversity suggest it functions as a broad, foundational lexicon. These are not merely "neutral" emotes; they are tools for expressing a wide range of nuanced, often self-referential community sentiments that defy simple positive/negative classification. This cluster captures the subcultural dialect of the platform.

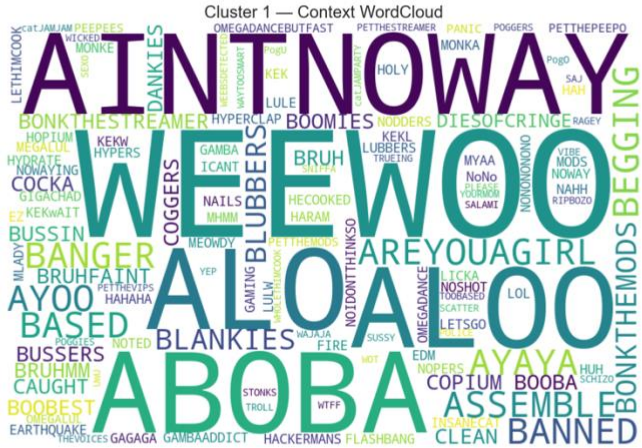


Figure 5 WordCloud of Cluster 1

Cluster 1: Hype, Spam, and Performative Excitement. This cluster is clearly defined by its pragmatic function: expressing high-arousal emotions like excitement, hype, and active engagement, as shown in wordcloud in figure 5. It contains classic hype emotes like HYPERS, POGGERS, and PogU, as well as emotes often used in rapid succession or "spam" during climactic moments (WEEWOO, AYAYA). This group represents a functional category of emotes used to collectively perform and amplify energy during key events in a stream, such as a difficult part of a speedrun or a winning play in a competitive game.


Figure 6 WordCloud of Cluster 2

Cluster 2: The "OOOO" Anomaly & Auditory Mimicry. This small but highly coherent cluster is almost entirely composed of emotes that are variations of elongated vowel sounds, such as OOOO and AAAA, as shown in wordcloud in figure 6. This grouping is a testament to the power of data-driven methods, as it reflects a unique co-usage pattern that a human-curated taxonomy would likely overlook. These emotes are likely used in similar contexts to express surprise, drawn-out reactions, or as a form of onomatopoeia—a collective vocal mimicry of a streamer's own reaction. Its emergence demonstrates the algorithm's ability to identify niche, pattern-based relationships that are purely functional.

### Comparative Analysis: The Divergence of Intention and Usage

The central goal of this study was to compare the rule-based archetypes with the emergent, data-driven clusters. Figure 7 visualizes the cross-tabulation between these two categorization schemes. Each cell shows the number of emotes that belong to a specific rule-based archetype (y-axis) and were assigned to a particular cluster (x-axis).

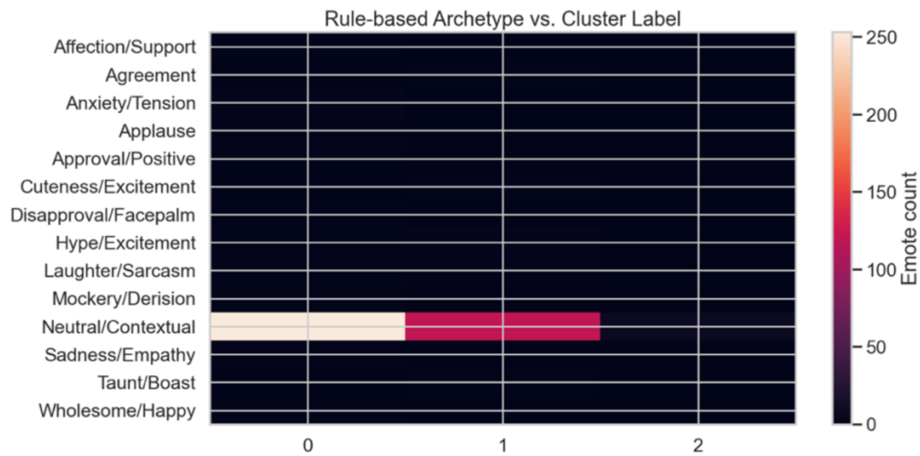


**Figure 7** **Heatmap of Rule-based vs Cluster Label**

The visualization reveals a near-total lack of correspondence between the two methods. The vast majority of emotes, regardless of their rule-based label (e.g., "Affection/Support," "Anxiety/Tension"), fall into the massive "Neutral/Contextual" category, which is then split across Clusters 0 and 1. There is no clear pattern where a specific rule-based archetype maps cleanly to a specific cluster. This visual evidence strongly supports the quantitative finding of a near-zero Adjusted Rand Index (ARI), confirming that the two labeling systems are fundamentally different. The rule-based method attempts to classify emotes by their semantic meaning, while the clustering method successfully groups them by their pragmatic function and co-usage patterns. The stark contrast in this heatmap is the primary finding of this paper, demonstrating that how emotes are defined is very different from how they are actually used.

The central and most compelling finding of this study is the stark lack of alignment between the rule-based archetypes and the data-driven clusters. The quantitative evaluation yielded an Adjusted Rand Index (ARI) of 0.0012. An ARI value close to 0.0 indicates that the correspondence between the two labeling systems is no better than random chance, signifying a near-total disagreement between the two methodologies. This profound divergence is visualized in the cross-tabulation table, which shows that the vast "Neutral/Contextual" rule-based category is scattered across all three clusters, while emotes from a single rule-based archetype like "Hype/Excitement" are found in both Cluster 0 (Pog) and Cluster 1 (PogU, HYPERS).

This result forces a critical discussion that forms the core of this paper's contribution: the rule-based method categorizes emotes by their intended or semantic meaning, while the clustering method groups them by their functional and pragmatic co-usage. An emote's dictionary definition is not the same as its job in a sentence. For example, while Pog (rule: Hype) and PepeLaugh (rule:

Laughter/Sarcasm) have different semantic meanings, the community may frequently use them in similar pragmatic contexts (e.g., reacting to an unexpected or absurd event), causing them to be clustered based on this shared function. The clustering reveals that the pragmatic role of an emote in a specific conversational context is a more powerful organizing principle than its static definition.

This implies that emote meaning is not fixed but is emergent, fluid, and heavily dependent on the shared conversational patterns of the community. The two methodologies are not simply "right" or "wrong"; rather, they measure two different dimensions of meaning. The rule-based approach captures the slow-moving, culturally agreed-upon symbol, while the clustering approach captures the fast-moving, context-dependent usage. The significant gap between these two dimensions is where the most interesting sociolinguistic phenomena reside, revealing a dynamic process of meaning-making in digital spaces.

## Conclusion

This study demonstrates that Twitch emotes operate within a complex system of meaning that cannot be adequately captured by static, rule-based sentiment taxonomies alone. The core contribution of this paper is the empirical validation of a significant divergence between an emote's prescribed semantic meaning and its emergent, pragmatic function within the community. The near-zero alignment between our handcrafted archetypes and the data-driven clusters (ARI = 0.0012) provides strong evidence that an emote's role is defined less by a dictionary-style definition and more by the contextual patterns of its use. Unsupervised clustering revealed coherent functional groupings—such as a "Hype/Spam" cluster and a "Nuanced Reaction" cluster—that would be invisible to a purely semantic analysis, proving the value of data-driven methods in uncovering how communities actually use these symbolic tokens. Ultimately, this research underscores the limitations of applying traditional sentiment analysis frameworks to dynamic and subcultural digital environments like Twitch. The failure of VADER and the intractability of manual classification highlight the need for more sophisticated, hybrid approaches that can account for context, co-usage, and the fluid nature of online communication. Future work should aim to scale this analysis with larger, multi-streamer datasets and incorporate advanced contextual embeddings from models like BERT to build a more robust understanding of this rich and evolving digital vernacular. Such efforts will be crucial for developing analytical tools that can keep pace with the rapid innovation of online culture.

## Declarations

### Author Contributions

Conceptualization: C.I.; Methodology: W.C.S.; Software: W.C.S.; Validation: S.A.G.; Formal Analysis: C.I.; Investigation: W.C.S.; Resources: S.A.G.; Data Curation: S.A.G.; Writing Original Draft Preparation: C.I.; Writing Review and Editing: S.A.G.; Visualization: C.I.; All authors have read and agreed to the published version of the manuscript.

### Data Availability Statement

The data presented in this study are available on request from the corresponding author.

## Funding

The authors received no financial support for the research, authorship, and/or publication of this article.

## Institutional Review Board Statement

Not applicable.

## Informed Consent Statement

Not applicable.

## Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References

[1] A. S. Catá, "Convergence of Rhetoric, Labour, and Play in the Construction of Inactive Discourses on Twitch," *Digit. Cult. Soc.*, 2019, doi: 10.14361/dcs-2019-0209.

[2] A. Leith and E. Gheen, "Twitch in the Time of Quarantine: The Role of Engagement in Needs Fulfillment.," *Psychol. Pop. Media*, 2022, doi: 10.1037/ppm0000372.

[3] J. Uttarapong, N. LaMastra, R. Gandhi, Y.-H. Lee, C. W. Yuan, and D. Y. Wohn, "Twitch Users' Motivations and Practices During Community Mental Health Discussions," *Proc. Acm Hum.-Comput. Interact.*, 2022, doi: 10.1145/3492824.

[4] N. LaMastra, J. Uttarapong, R. Gandhi, C. Cook, and D. Y. Wohn, "How a Live Streamer's Choice in Played Game Affects Mental Health Conversations," 2020, doi: 10.1145/3383668.3419894.

[5] J. T. Sheng and S. Kairam, "From Virtual Strangers to IRL Friends," *Proc. Acm Hum.-Comput. Interact.*, 2020, doi: 10.1145/3415165.

[6] C. Flores-Saviaga, J. Hammer, J. P. Flores, J. Seering, S. Reeves, and S. Savage, "Audience and Streamer Participation at Scale on Twitch," 2019, doi: 10.1145/3342220.3344926.

[7] R. Gandhi, C. Cook, N. LaMastra, J. Uttarapong, and D. Y. Wohn, "An Exploration of Mental Health Discussions in Live Streaming Gaming Communities," *Front. Psychol.*, 2021, doi: 10.3389/fpsyg.2021.575653.

[8] J. Li, X. Gui, Y. Kou, and Y. Li, "Live Streaming as Co-Performance," *Proc. Acm Hum.-Comput. Interact.*, 2019, doi: 10.1145/3359166.

[9] K. Ask, H. S. Spilker, and M. H. Hansen, "The Politics of User-Platform Relationships: Co-Scripting Live-Streaming on Twitch.tv," *First Monday*, 2019, doi: 10.5210/fm.v24i7.9648.

[10] Z. Teng, D. T. Vo, and Y. Zhang, "Context-Sensitive Lexicon Features for Neural Sentiment Analysis," 2016, doi: 10.18653/v1/d16-1169.

[11] I. Jauk, J. Lorenzo-Trueba, J. Yamagishi, and A. Bonafonte, "Expressive Speech Synthesis Using Sentiment Embeddings," 2018, doi: 10.21437/interspeech.2018-2467.

[12] Z. Lei, Y. Yang, M. Yang, and Y. Liu, "A Multi-Sentiment-Resource Enhanced Attention Network for Sentiment Classification," 2018, doi: 10.18653/v1/p18-2120.

[13] J. Barnes, E. Velldal, and L. Øvrelid, "Improving Sentiment Analysis With

Multi-Task Learning of Negation," *Nat. Lang. Eng.*, 2020, doi: 10.1017/s1351324920000510.

[14] P. Ke, H. Ji, S. Liu, X. Zhu, and M. Huang, "SentiLARE: Sentiment-Aware Language Representation Learning With Linguistic Knowledge," 2020, doi: 10.18653/v1/2020.emnlp-main.567.

[15] K. R. Mabokela, T. Çelik, and M. Raborife, "Multilingual Sentiment Analysis for Under-Resourced Languages: A Systematic Review of the Landscape," *Ieee Access*, 2023, doi: 10.1109/access.2022.3224136.

[16] T. Wulf, F. M. Schneider, and S. Beckert, "Watching Players: An Exploration of Media Enjoyment on <i>Twitch</I>," *Games Cult.*, 2018, doi: 10.1177/1555412018788161.

[17] D. Spińczyk, K. Nabrdalik, and K. Rojewska, "Computer Aided Sentiment Analysis of Anorexia Nervosa Patients' Vocabulary," *Biomed. Eng. Online*, 2018, doi: 10.1186/s12938-018-0451-2.

[18] S. W. Chae and S. H. Lee, "Sharing Emotion While Spectating Video Game Play: Exploring Twitch Users' Emotional Change After the Outbreak of the COVID-19 Pandemic," *Comput. Hum. Behav.*, 2022, doi: 10.1016/j.chb.2022.107211.

[19] S. L. Evans and S. M. Llano, "Tryhard With a Vengeance: Meaning Making and Boundary Keeping on Twitch," *J. Electron. Gaming Esports*, 2023, doi: 10.1123/jege.2022-0042.

[20] N. Pröllochs, D. Bär, and S. Feuerriegel, "Emotions Explain Differences in the Diffusion of True vs. False Social Media Rumors," *Sci. Rep.*, 2021, doi: 10.1038/s41598-021-01813-2.

[21] K. Kobs *et al.*, "Emote-Controlled," *Acm Trans. Soc. Comput.*, 2020, doi: 10.1145/3365523.

[22] П. И. Долин, L. d'Hauthuille, and A. Vattani, "FeelsGoodMan: Inferring Semantics of Twitch Neologisms," 2021, doi: 10.48550/arxiv.2108.08411.

[23] T. M. Erle, K. Schmid, S. H. Goslar, and J. Martin, "Emojis as Social Information in Digital Communication.," *Emotion*, 2022, doi: 10.1037/emo0000992.

[24] G. Bhuvaneswari and S. Vijayakumar, "Emotional Intelligence and Values in Digital World Through Emoticons Among Indian Students and Faculty," *Int. J. Asian Educ.*, 2021, doi: 10.46966/ijae.v2i2.142.

[25] D. Morina and M. S. Bernstein, "A Web-Scale Analysis of the Community Origins of Image Memes," *Proc. Acm Hum.-Comput. Interact.*, 2022, doi: 10.1145/3512921.

[26] V. Truong, "Cultural Differences in Customer Emotions," 2024, doi: 10.21203/rs.3.rs-4747823/v1.

[27] X. Feng, L. Yu, J. Wang, W. Kong, and W. Chen, "Research on the Mechanism of Emotional Design in Chinese Cultural and Creative Products," 2022, doi: 10.21203/rs.3.rs-1413692/v1.

[28] M. Mestiri, M. Khadhar, A. Fergombé, and A. Huftier, "Fostering Social Interaction Variability in the Metaverse: A Case Study of the Museum of l'Avenois in Fourmies," 2024, doi: 10.20944/preprints202412.2494.v1.