

# Bridging Culture and Finance: A Multimodal Analysis of Memecoins in the Web3 Ecosystem

Hou-Wan Long  
houwanlong@link.cuhk.edu.hk  
The Chinese University of Hong Kong  
Shatin, Hong Kong

Nga-Man Wong  
s1145167@s.edu.hk  
Education University of Hong Kong  
Tai Po, Hong Kong

Wei Cai  
weicaics@uw.edu  
University of Washington  
Tacoma, Washington, USA

## Abstract

Memecoins, driven by social media engagement and cultural narratives, have rapidly grown within the Web3 ecosystem. Unlike traditional cryptocurrencies, they are shaped by humor, memes, and community sentiment. This paper introduces the Coin-Meme dataset, an open-source collection of visual, textual, community, and financial data from the Pump.fun platform on the Solana blockchain. We also propose a multimodal framework to analyze memecoins, uncovering patterns in cultural themes, community interaction, and financial behavior. Through clustering, sentiment analysis, and word cloud visualizations, we identify distinct thematic groups centered on humor, animals, and political satire. Additionally, we provide financial insights by analyzing metrics such as Market Entry Time and Market Capitalization, offering a comprehensive view of memecoins as both cultural artifacts and financial instruments within Web3. The Coin-Meme dataset is publicly available at <https://github.com/hwlongCUHK/Coin-Meme.git>.

## CCS Concepts

• Information systems → World Wide Web; • Human-centered computing → Empirical studies in collaborative and social computing.

## Keywords

Memecoin, Web3, Multimodal Analysis, Blockchain

### ACM Reference Format:

Hou-Wan Long, Nga-Man Wong, and Wei Cai. 2025. Bridging Culture and Finance: A Multimodal Analysis of Memecoins in the Web3 Ecosystem. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3701716.3715561>

## 1 Introduction

The memecoin phenomenon represents a rapidly expanding segment within the Web3 ecosystem, where internet culture and decentralized finance intersect. Unlike Bitcoin [6], which was purposefully designed as a decentralized digital currency, memecoins are primarily fueled by social media engagement, humor, and cultural narratives, often driven by community sentiment rather than intrinsic utility [7]. Memecoins such as Dogecoin, Shiba Inu, and Pepe Coin epitomize this trend, with Dogecoin reaching an \$80

billion Market Capitalization in 2021, propelled by viral trends and celebrity endorsements. Platforms like Pump.fun have further democratized the creation of memecoins, enabling users to launch tokens without requiring technical expertise, while also empowering decentralized community-driven governance. By September 2024, Pump.fun had facilitated over 1 million token launches, generating \$100 million in revenue and dominating daily token creation on the Solana blockchain.

Unlike memes, which are primarily cultural and humorous, memecoins introduce a financial dimension that connects internet culture to finance. Despite their significant impact on both cultural and financial spheres, memecoins have been understudied. Most research focuses on their financial aspects, such as market dynamics and price fluctuations, while overlooking the cultural and community-driven forces behind their creation and growth [2, 3, 7]. Meanwhile, research on memes only examines their cultural attributes, since memes were initially viewed as non-commercial cultural artifacts, without any inherent financial dimension [5]. This paper aims to bridge both gaps by analyzing memecoins through their cultural and financial dimensions, offering a comprehensive view of their role in the Web3 ecosystem.

This paper makes the following contributions: **(i) Open-Source Multimodal Dataset:** We introduce Coin-Meme, the first open-source dataset for studying memecoins in the Web3 space, combining visual (logos), textual (descriptions), community (user comments), and financial (Market Capitalization, Market Entry Time) data to explore the intersection of culture, community, and decentralized finance. **(ii) Multimodal Analytical Framework:** We propose a framework to analyze memecoins across their visual, textual, community, and financial dimensions, providing a holistic view of how cultural narratives, community engagement, and Web3 technologies influence memecoin adoption and market behavior. **(iii) Cultural and Community Insights:** Applying our framework to Coin-Meme, we uncover patterns in memecoin adoption, community sentiment, and thematic trends, offering insights into how decentralized communities and culture shape the financial dynamics of memecoins in Web3.

The paper is structured as follows: Section 2 introduces Pump.fun and reviews related meme analysis. Section 3 describes the Coin-Meme dataset. Section 4 presents the analysis framework. Section 5 discusses the results, including clustering, sentiment, and financial metrics. Section 6 concludes the paper.

## 2 Pump.fun Memecoins and Related Work

Pump.fun is a decentralized platform on the Solana blockchain that simplifies memecoin creation through an intuitive interface. Users can create memecoins by specifying attributes like name, symbol,



This work is licensed under a Creative Commons Attribution 4.0 International License. *WWW Companion '25, April 28-May 2, 2025, Sydney, NSW, Australia*  
© 2025 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1331-6/25/04  
<https://doi.org/10.1145/3701716.3715561>

and supply without technical expertise. Once created, the coin attracts community engagement through comments and feedback, which can enhance visibility or slow growth. If its Market Capitalization reaches 69,000\$, the coin can be listed on Raydium, a decentralized exchange (DEX), allowing for liquidity and broader exposure within the Solana DeFi ecosystem.

The rise of Web2 platforms has made memes a key element of internet culture, shaping public opinion, trends, and humor. The study of memes through multimodal analysis has gained attention, enabling researchers to understand how meaning is constructed and shared across digital platforms. For example, Kiela et al. [1] proposed a multimodal framework for detecting hate speech in memes, combining visual and textual embeddings for classification. Suryawanshi et al. [8] developed the MultiOFF dataset for detecting offensive content in memes by combining image and text modalities, while Xu et al. [9] introduced the MET-Meme dataset, which incorporates metaphorical features for sentiment analysis. Additionally, Ling et al. analyzes visual elements that distinguish viral image memes from non-viral ones and develops a machine learning model to predict meme virality [4].

However, the intersection of memes and cryptocurrency within the Web3 ecosystem has been largely underexplored. In this study, we build on existing multimodal research to analyze memecoins using a framework that integrates visual, textual, community, and financial data, providing deeper insights into their role within Web3.

### 3 Dataset Description

The Coin-Meme dataset consists of four modalities: Textual Descriptions, Visual Content, Community Interactions, and Financial Data. Textual Descriptions are the narratives associated with each memecoin. Visual Content includes the logo image of memecoin. Community Interactions encompass user comments. Financial Data consists of market-related metrics, including Market Capitalization and the Market Entry Time, which quantifies the speed at which a token gains liquidity and market presence on decentralized exchanges. The dataset includes all 3,751 memecoins that were created between January 2024 and November 2024 and subsequently moved to Raydium, providing a comprehensive overview of the memecoin ecosystem.

#### 3.1 Data Acquisition

Data acquisition began by extracting token metadata from Dune.com using Pump.fun's program address<sup>1</sup> including token addresses, creation times, names, and creators. Token-specific URLs were generated by appending token addresses to the base URL<sup>2</sup>. A web scraping pipeline built with Selenium and Pandas automated the extraction of textual descriptions, token images, user comments, and financial data.

#### 3.2 Data Cleansing

Following data acquisition, the dataset was cleaned to ensure accuracy and consistency. Text data (token descriptions and user comments) were preprocessed by converting text to lowercase,

removing non-alphabetic characters (e.g., URLs), eliminating stopwords, and applying tokenization and lemmatization. Entries with missing descriptions or comments were excluded. Images were resized to 224x224 pixels to match ResNet50 input size, non-RGB images were converted to RGB, and corrupted or missing images were replaced with zero vectors. The time difference between token launch and being moved to Raydium was calculated as Market Entry Time.

## 4 Multimodal Analysis Framework

The multimodal analysis framework presented in this study integrates visual, textual, community interaction and financial data to analyze the cultural and financial dynamics of memecoins.

### 4.1 Feature Representation

To enable a multimodal representation, textual, visual and financial features were combined into a unified feature space: **Textual Features:** Latent Dirichlet Allocation (LDA) was applied to tokenize and model topics within the textual descriptions. Each memecoin was represented by a vector of topic distributions, capturing its thematic essence. **Visual Features:** High-dimensional image vectors extracted using ResNet50 encapsulated visual patterns, including color schemes, shapes, and design motifs. **Financial Features:** The time difference between the creation and the time used to move the token to Raydium was calculated. Additionally, the Market Capitalization (mcap) of each token was included as a feature, capturing its market value and investment potential. Both quantities were scaled to reflect their relative impact on token adoption and liquidity.

Textual features, visual features and the financial features were concatenated to form a single high-dimensional representation for each token, effectively capturing its narrative, visual, and financial dynamics.

### 4.2 Clustering

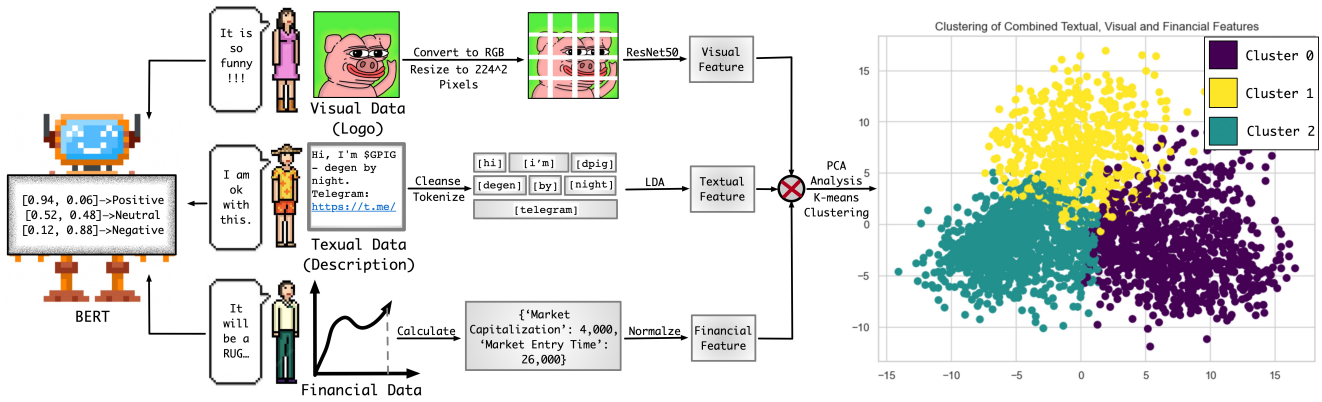
K-Means clustering was used to group memecoins based on their multimodal features due to its efficiency, simplicity, and ability to handle large datasets. The optimal number of clusters was determined using the silhouette score, which evaluates cluster cohesion and separation. K-Means is particularly effective when the number of clusters is predefined, as it minimizes within-cluster variance, making results easy to interpret. To enhance interpretability and reduce dimensionality, Principal Component Analysis (PCA) was applied before clustering.

### 4.3 Analysis

After clustering, Word Cloud Analysis was conducted on the memecoin descriptions to visualize the frequency of words. Prominent words appear larger in the cloud, helping to identify the dominant themes and keywords associated with each group. This technique uncovers the central narratives and cultural motifs within the memecoin ecosystem. For instance, certain clusters emphasized playful, irreverent themes, while others reflected lighthearted, animal-centric content or politically charged, satirical discourse.

<sup>1</sup>6EF8rrecthR5Dkzon8Nwu78hRvfCKubJ14M5uBEwF6P

<sup>2</sup><https://pump.fun/>



**Figure 1: Multimodal clustering of memecoins based on combined textual, visual, and financial feature, with representative descriptions, logos and comments illustrating sentiment categories**

Sentiment Analysis was performed on user comments for each cluster using the sentiment-analysis pipeline from the Transformers library with BERT as the model. The model outputs probabilities for positive and negative sentiments, classifying comments as neutral when both probabilities are close. Comments were categorized into positive, neutral, and negative, revealing sentiment variations across thematic groups and offering insights into emotional tones and cultural dynamics in the memecoin space. Engagement metrics like comments per user were also calculated. Sentiment variability was calculated by calculating the standard deviation of sentiment scores for each coin. Scores were derived from the average of the highest probabilities (positive or negative) across comments. Variability in sentiment helps gauge genuine engagement versus potential bot activity or automated manipulation.

These metrics offer valuable insights into how different themes resonate with users and contribute to the cultural and financial dynamics of the memecoin ecosystem.

### 5 Analysis Results on Coin-Meme Dataset

This section presents the results of applying the proposed multimodal analysis framework on the Coin-Meme dataset. The clustering process identified three distinct memecoin groups as shown in Figure 1: Cluster 0 (1,127 memecoins), Cluster 1 (1,535 memecoins), and Cluster 2 (1,089 memecoins).

#### 5.1 Word Cloud Analysis Results

Cluster 0 focuses on humor, parody, and internet niche culture. The word cloud (Figure 2) highlights terms like "meme" and "pepe," reflecting the irreverent and playful nature of this group. Bold visuals, such as caricatures and fantastical creatures, cater to an audience that values creativity and humor. This cluster exemplifies the humorous and internet culture-driven side of memecoin creation.

Cluster 1 revolves around animal themes and whimsical designs. The word cloud (Figure 2) reveals keywords like "cat," "dog," and "cutest," reflecting the lighthearted and relatable nature of the tokens. Descriptions often feature humorous and endearing animal scenarios, while visuals showcase animals like tigers, red pandas, and popular meme animals like Grumpy Cat. The use of words such



**Figure 2: Word Cloud and Representative Images for Each Cluster**

as "friend" and "fun" underscores the cluster's appeal to users who enjoy playful, emotional content with a whimsical brand identity.

Cluster 2 is characterized by political satire and references to iconic cultural figures, particularly Donald Trump. Descriptions frequently include terms like "Trump," while visuals often feature exaggerated caricatures and hyperbolic portrayals of authority. The word cloud (Figure 2) highlights terms such as "trump," "crypto," and "world," emphasizing the cluster's political and engagement-driven nature. This cluster's provocative humor fosters active, often polarizing discussions, making it one of the most controversial and engaging groups in the memecoin space.

Despite the distinct themes of each cluster, several common terms appear across all three, reflecting the shared strategies used by memecoin creators to capture attention. "world" is often used in

Cluster	Positive to Negative Comment Ratio	Comments per User	Mean Sentiment Variability	Mean Market Entry Time (s)	95% Quantile of MCap (\$)
0	1.23	2.97	0.3761	25,897	33,557.5
1	1.36	3.74	0.3754	7,129	24,970.5
2	1.64	2.78	0.3394	4,728	29,508.5

**Table 1: Sentiment, comment and financial statistics per cluster**

superlative phrases like "the best in the world," designed to evoke a sense of uniqueness and global appeal. "community" emphasizes the importance of user engagement and collective identity, highlighting the social aspect of memecoins. "meme" is central to the memecoin concept, embodying the cultural and viral nature that drives their popularity. These shared terms demonstrate the common goal of leveraging internet culture and community-building to attract and retain users.

## 5.2 Sentiment Analysis Results

Cluster 0, focused on humor, parody, and internet niche culture, has the highest positive-to-negative comment ratio (1.23), reflecting the playful and irreverent nature of its content. This suggests a more passive, niche audience with moderate engagement, as indicated by a comments per user rate of 2.97. Cluster 1, centered on animal themes and whimsical designs, demonstrates a positive-to-negative comment ratio slightly higher than cluster 0, highlighting its broader, non-confrontational appeal. The community engagement is notably higher, with 3.74 comments per user, suggesting that its playful, lighthearted nature encourages frequent interactions. Cluster 2, with its focus on political satire and controversial figures, stands out with the highest positive-to-negative comment ratio (1.64), reflecting strong resonance with a passionate, polarized audience. However, its lower comments per user (2.78) suggest less frequent engagement.

The Mean Sentiment Variability values as shown in Table 1 suggest similar levels of bot activity in clusters 0 and 1, with cluster 2 showing slightly lower variability. This indicates that while all three clusters experience some level of automated influence, cluster 2, with its politically charged content, seems to have less fluctuation in sentiment, possibly reflecting a more controlled or targeted bot activity compared to the other two.

## 5.3 Financial Implication

We observe distinct financial patterns across clusters, shedding light on the economic behavior of memecoins. The Mean Market Entry Time and 95% Quantile Market Capitalization were selected to capture key aspects of market adoption, liquidity, and financial stability. The Mean Market Entry Time reflects how quickly a token gains traction on decentralized exchanges, while the 95% Quantile Market Capitalization reveals the market cap of top-performing tokens, highlighting their scale and potential value.

Cluster 0, with a Mean Market Entry Time of 25,897 seconds and a 95% Quantile Market Capitalization of \$33,557.5, shows slower liquidity but higher market cap potential, indicating value due to niche appeal. Cluster 1, with a faster Mean Market Entry Time of 7,129 seconds and a market cap of \$24,970.5, demonstrates quicker adoption and broader appeal but lower market cap. Cluster 2, with the fastest Mean Market Entry Time of 4,728 seconds and a market

cap of \$29,508.5, shows rapid entry but potentially lower long-term stability, reflecting the polarizing nature of its content. These metrics provide valuable insights into liquidity, market adoption, and financial stability within each cluster.

## 6 Conclusion

This study introduces the Coin-Meme dataset and a novel multi-modal analytical framework to explore the dynamics of memecoins within the Web3 ecosystem. By integrating visual, textual, community interaction, and financial data, we offer a comprehensive understanding of how memecoins evolve through cultural narratives, community engagement, and market forces in decentralized environments. Our analysis highlights the diversity of memecoin themes, from humorous and playful to politically charged, and uncovers key patterns in user sentiment, community interaction, and potential bot-driven behaviors. Financial implications are also provided. These insights underscore the importance of studying memecoins as both cultural and financial phenomena, bridging internet culture and DeFi. This research not only sheds light on their role within the cryptocurrency market but also offers implications for the future of Web3 communities and governance. Future work could further investigate the impact of social media platforms, sentiment manipulation, and the broader role of memecoins in the evolving Web3 landscape, particularly in relation to DeFi and Web3 governance mechanisms.

## 7 Acknowledgments

This research is supported by Department of Statistics and New Asia College of The Chinese University of Hong Kong.

## References

- [1] Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes. In *Advances in Neural Information Processing Systems*.
- [2] David Krause. 2024. Beyond the Hype: A Meme Coin Reality Check for Retail Investors. Available at SSRN 4891841 (2024).
- [3] Tao Li, Donghua Shin, Chuyi Sun, and Baolian Wang. 2022. The dark side of decentralized finance: Evidence from meme tokens. Available at SSRN (2022).
- [4] Chen Ling, Ihab AbuHilal, Jeremy Blackburn, Emiliano De Cristofaro, Savvas Zannettou, and Gianluca Stringhini. 2021. Dissecting the meme magic: Understanding indicators of virality in image memes. *Proceedings of the ACM on human-computer interaction* 5, CSCW1 (2021), 1–24.
- [5] Maria D Molina. 2020. What makes an internet meme a meme? Six essential characteristics. *Handbook of visual communication* (2020), 380–394.
- [6] Satoshi Nakamoto. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. *Manuscript published online* (2008). <https://bitcoin.org/bitcoin.pdf>
- [7] Adrian Stencil. 2023. What is a meme coin? Dogecoin to the moon! (2023).
- [8] Shardul Suryawanshi, Bharathi Raja Chakravarthi, Mihael Arcan, and Paul Buiteelaar. 2020. Multimodal meme dataset (MultiOFF) for identifying offensive content in image and text. In *Proceedings of the second workshop on trolling, aggression and cyberbullying*. 32–41.
- [9] Bo Xu, Tingting Li, Junzhe Zheng, Mehdi Naseriparsa, Zhehuan Zhao, Hongfei Lin, and Feng Xia. 2022. Met-meme: A multimodal meme dataset rich in metaphors. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 2887–2899.