

A Silent Voice: A Bilingual Longitudinal Analysis of Mahsa Amini’s Movement Tweet

Sadegh Jafari
Computer Science
Iran University of Science and Technology
Tehran, Iran
sadegh_jafari@comp.iust.ac.ir

Mahsa Farsouni Nezhad
Sociology
Allameh Tabatabai University
Tehran, Iran
mahsafarsi16@gmail.com

Melika Nobakhtian
Computer Science
Khatam University (TeIAS)
Tehran, Iran
M.Noubakhtian.5740@khatam.ac.ir

Arash Chaichi Mellatshahi
Computer Science
Ghent University
Ghent, Belgium
Arash.ChaichiMellatshahi@UGent.be

Amirmohammad Kazemeini
Computer Science
Vector Institute
Ontario, Canada
amirmohammad.kazemeini@vectorinstitute.ai

Erik Cambria
College of Computing and Data Science
Nanyang Technological University
Singapore
cambria@ntu.edu.sg

Abstract—Mahsa Amini’s death was not just a tragedy; it sent shockwaves through Iranian society, igniting a firestorm of dissent that raged both online and offline. This study delves into the dramatic transformation of Twitter in the 90 days following her death. We analyze the emotional pulse of Iranian users, tracking how their sentiments evolved and the topics of discussion from both chronological and linguistic standpoints. By dissecting the changing topics of discourse and identifying the most potent words and phrases, we showcase how understanding both English and Persian tweets reflects two different yet complementary aspects based on their audiences. The progression of the movement is represented in the tweets, highlighting the necessity of considering different languages and the full timeline to capture the representation of tweets about this movement in both English and Persian.

Index Terms—Topic Modeling, Sentiment Analysis, Social Media, Twitter.

I. INTRODUCTION

Social media is crucial in disseminating information and shaping public opinion, especially during crises [1]. The tragic death of Mahsa Amini in Iran sparked widespread outrage and protests, with social media platforms like Twitter becoming a focal point for expression and solidarity. This work delves into the content and sentiment of tweets surrounding the event, particularly focusing on how Iranian users leveraged the platform despite access limitations.

Sentiment analysis, often treated as a simple polarity detection task, is a “suitcase problem” involving various natural language processing subtasks, such as microtext analysis and subjectivity detection [2].

In September 2022, the passing of Mahsa Amini while in police custody for violating Iran’s dress code ignited nationwide protests and international condemnation. The event triggered a surge of activity on social media, particularly on Twitter, where the hashtag “#Mahsa_Amini” became a global phenomenon despite the platform’s official ban within Iran.

In this paper, we present a comprehensive analysis of the Twitter discourse surrounding Mahsa Amini’s death, focusing

on the 90 days following the incident. Leveraging a dataset comprising over 1.9 million tweets in multiple languages, predominantly Persian and English, we delve into the nuanced dynamics of sentiment and topic distribution within the Twitterverse.

Building upon our prior work [3], this study extends the sentiment analysis to encompass the entire multilingual dataset, offering a more comprehensive understanding of global reactions to Mahsa Amini’s death. We further leverage topic modeling using BERTopic [4] to identify critical themes emerging from the tweets. Initially, this analysis was performed on a sample to capture the most prominent topics. Subsequently, we assigned topics to all tweets and examined their prevalence and evolution across seven-day intervals. Finally, hierarchical clustering helps us delve deeper by uncovering the relationships between these evolving thematic structures.

This article also analyzes the semantic differences between Persian and English tweets regarding the protests following the death of Mahsa Amini, based on data clustering. Additionally, by categorizing the data over time, the article compares the changes in the semantics of the tweets in representing the protests throughout different periods.

II. RELATED WORKS

A. Topic Modeling

In recent years, neural topic models have gained prominence by leveraging neural networks to enhance traditional topic modeling techniques [5]. Researchers have explored various approaches to improve topic modeling, including the incorporation of word embeddings into classical models such as Latent Dirichlet Allocation (LDA) [6]. Word embeddings, derived from pre-trained language models, offer powerful representations that enhance topic modeling performance.

Beyond LDA-like models, recent advancements have focused on embedding-based topic modeling techniques. These methods build directly on word embeddings, demonstrating the potential of this approach. For instance, the Continuous Topic



Fig. 1: Mahsa Amini’s legacy has sparked a new movement for Iranian women.

Model (CTM) [7] leverages pre-trained language models, anticipating that improvements in language models will lead to better topic representations [8].

BERTopic builds upon the clustering embeddings approach (The default approach is HDBSCAN, and we used it). It extends this technique by incorporating a class-based variant of Term Frequency-Inverse Document Frequency (TF-IDF) to create coherent and interpretable topic representations. By combining BERT embeddings with class-based TF-IDF, BERTopic achieves dense clusters while retaining essential keywords in topic descriptions.

B. Sentiment Analysis

Sentiment analysis has undergone a significant transformation with the emergence of deep learning, particularly transformer-based architectures [9], [10]. Traditional machine learning models like Support Vector Machines (SVMs) [11] and Naive Bayes [12] are often outperformed by models like Bidirectional Encoder Representations from Transformers (BERT) [13] and RoBERTa [14]. These models leverage the power of large-scale pre-training on massive text corpora, allowing them to capture complex linguistic relationships and achieve superior performance.

Furthermore, domain-specific pre-trained models offer additional advantages. CardiffNLP’s Twitter-XLM-RoBERTa [15] model exemplifies this approach, achieving state-of-the-art results in sentiment analysis on social media text by specifically accounting for the nuances and slang commonly found in online conversations.

C. Protest Representation on Social Media

As [16] have demonstrated, linguistic and geographical differences shape the representation patterns of protest news. This article explores these variations by analyzing news coverage of various protests worldwide in English and Spanish on social media. [17] compared the media coverage of protests following the death of Michael Brown in Ferguson, United States, with that of the protests in Ayotzinapa, Mexico, where justice was sought for 43 missing students. Their analysis focuses on differences in framing, narrative devices, and news sources in domestic versus international protest coverage, as well as variations in how news related to these protests was shared on social media.

As discussed by [16] and [17], the representation of protests varies based on factors like language and audience location. In this article, we aim to examine the differences in the representation of the Mahsa movement protests based on language. To achieve this, we have analyzed both English-language and Persian-language tweets, utilizing data from two distinct periods for comparison.

III. DATA

Utilizing the “sncrape”¹ Python library, we systematically crawled Twitter data encompassing tweets associated with the hashtag “#mahsa_amin” and Persian equivalent and related hashtags spanning from September 21, 2022, to December 19, 2022. Our dataset comprises approximately 2 million tweets from different languages.

To ensure data uniformity and facilitate subsequent analysis, we conducted preprocessing steps aimed at enhancing

¹<https://github.com/JustAnotherArchivist/sncrape>

the quality and consistency of the dataset. Specifically, we removed usernames, hashtags, and URLs from the raw tweet data. Furthermore, leveraging the insights garnered from the "langdetect"² Python library, we conducted an analysis of language distribution within the preprocessed dataset. Table 1 presents a comprehensive overview of the distribution of tweets across the five most prevalent languages identified in our dataset.

TABLE I: Number of tweets in different languages

Language	Number of tweets
Persian	1,445,537
English	317,046
Arabic	54,106
Urdu	28,880
German	13,919

IV. METHODOLOGY

To analyze our tweets dataset, we applied topic modeling and sentiment analysis. In the following subsections, we explain each technique and present the results in detail.

A. Topic Modeling


We have almost 2 million tweets, which we divided into 4 groups. We first filter the data to keep only Persian and English tweets, then split the remaining data into Persian and English tweets, and we also split it based on the date into two equal periods: one from the start of the event to the middle of the date range, and one from the middle to the end of the date range. We avoided pre-imposing time interval boundaries based on sociological events to ensure an unbiased, objective analysis of the data over 2-time intervals. We selected 160,000 samples from each group with the greatest length due to memory limitations (requiring more than 12 GB of RAM) to train the clustering models. Thus, we trained 4 clustering models. We used BERTopic tools to cluster the data, utilizing the "BAAI/bge-m3" [18] sentence transformer (because this model is one of the best sentence transformers for multi-language corpora) to extract the features for each sample. This clustering is hierarchical, and we only analyze the 2 most important clusters for each clustering model.

In this study, we focused on analyzing the top two topics identified by the model. While it is possible to delve into a broader range of topics, the selection of the top two was driven by practical constraints. Given the complexity of the discourse surrounding the protests, a detailed and thorough analysis of additional topics would require significantly more time and resources, particularly in terms of expert validation from social media analysts. By concentrating on the most prominent topics, we aimed to provide a meaningful yet manageable exploration of the discourse, ensuring that the insights remain focused and relevant without compromising the depth of the analysis.


²<https://github.com/Mimino666/langdetect>

B. Sentiment Analysis


We used the "cardiffnlp/twitter-xlm-roberta-base-sentiment" [15] model to assign tags to each tweet for sentiment analysis. After that, you can see the distribution of sentiment tags for each topic in each of the 4 groups explained in subsection 4.1 in Figure 2.




(a) Percentage Distribution of Sentiments in Two Topics (First half of protest in English tweets)



(b) Percentage Distribution of Sentiments in Two Topics (First half of protest in Persian tweets)



(c) Percentage Distribution of Sentiments in Two Topics (Second half of protest in English tweets)



(d) Percentage Distribution of Sentiments in Two Topics (Second half of protest in Persian tweets)

Fig. 2: Sentiment analysis for each 4 groups based on their topics.

V. RESULTS AND ANALYSIS

The main focus of English tweets in Topic 1 is on having the voices of protesters heard by international communities and describing their form and extent of repression. For example, the selected tweets of Topic 1 English attempt to describe the situation in Iran, internet shutdowns, repression, and arrests by the government to the global community. On the other hand, the Farsi tweets in Topic 1 emphasize the solidarity of protesters within Iran and highlight the reflection of Iran's protests in international communities to give hope and motivation to internal protesters. Unlike the English tweets, which focus on receiving international help and sanctioning important government figures, the Farsi tweets talk about the need for internal solidarity. For instance, one tweet mentions that the national football team of Iran should support the protesters (Topic 1-Farsi).

In the Farsi tweets of Topic 2, there is talk of the resistance, protests, and bravery of women and girls, with a major focus on women. These women, who have faced years of discrimination and double oppression, have now taken to the streets. In contrast, none of the English tweets on the topics focus in this manner on women and girls and their significant role in the recent protests. The Farsi tweets speak of the necessity of continuous struggle for justice and freedom and organizing for it, and they are more radical, while the English tweets emphasize human rights violations in Iran and their reflection. For example, in Topic 1-English tweets, there is talk of the brutal repression of protesters, whereas in the Farsi tweets of Topics 1 and 2, the emphasis is on resistance, hope for the future, and continuing the protests.

The English tweets of Topics 3 and 4, which have been collected from the second half of the protest period, speak of intensified repression and the execution of protesters, focusing on specific requests from the international community. For instance, declaring the Islamic Revolutionary Guard Corps as a terrorist organization and eliminating the bank accounts of senior officials of the Islamic Republic. These tweets also talk about the repression in Baluchistan and Kurdistan and call for the widespread echoing of their voices.

In contrast, the Farsi tweets of Topics 3 and 4, collected from the second half of the protest period, focus on resistance and unity, speaking of resistance in universities and steadfastness in Kurdistan. These tweets also extensively talk about severe internet shutdowns in Iran and the increasingly violent nature of the protests.

VI. CONCLUSION

Based on the obtained results, a significant difference can be observed between the representation of the protests over Mahsa Amini's death in English and Persian. In English representations, users emphasize international support, describing human rights violations, and appeals to international organizations. In contrast, the Persian representations of the protests emphasize internal solidarity, resistance, courage, and the struggle for freedom and justice.

From a chronological standpoint, the difference highlights a progression in measures both from the repression by the government and from the response of society. The violence has escalated from milder measures to more severe repressive actions, such as executions. Society's calls for awareness have become more tangible and specific. The resistance strategies have also evolved from discussing potential tools and strategies to more urgent and practical tips for protests. While the initial time interval tweets had a tone of desperation and urgent pleas, the second interval shifted to a determined tone, signifying a solution-oriented shift in the movement.

VII. LIMITATIONS AND FUTURE WORK

Our study, while comprehensive in many aspects, is subject to several limitations that should be acknowledged:

- 1) **Temporal Scope of Data:** The dataset comprises nearly 2 million tweets related to Mahsa Amini and associated hashtags, collected over three months. A longer data collection period, such as one year, would likely yield more comprehensive insights into the evolution and dynamics of the movement.
- 2) **Bias in Data Collection:** The dataset is potentially biased as it predominantly includes tweets supportive of the movement. We lack data representing the opposing perspectives, which could provide a more balanced and nuanced understanding of public opinion. Including tweets from dissenting voices would allow for a more thorough analysis of the discourse surrounding the movement.
- 3) **Focus on Textual Data:** Our analysis is limited to the textual content of the tweets. We did not incorporate user metadata, which could enable a user-based analysis and offer deeper insights into the demographics and influence patterns within the movement. The inclusion of user data, such as user location, follower count, and network connections, would enhance the granularity of the analysis.
- 4) **Twitter Data Collection Limitations:** Due to Twitter's data access restrictions, we were unable to gather a more extensive dataset. These limitations hindered our ability to perform a more detailed and segmented analysis.
- 5) **Temporal Chunking of Data:** The data was only split into two temporal chunks for analysis. A more segmented temporal analysis would help in detecting finer-grained changes and trends over time. However, time constraints prevented us from conducting a more detailed temporal segmentation and analysis.
- 6) **Personality Detection in User Analysis:** While our study focuses primarily on textual content, an additional

dimension worth exploring in future work is the personality traits of users who tweet about the movement. Utilizing approaches such as the one proposed by [19] could allow us to analyze users' personality traits based on their tweets. This could provide deeper insights into how personality influences engagement with the movement and the nature of the content shared. Such an analysis could further enhance our understanding of user behavior patterns within the social media discourse.

- 7) **Enhancing Multilingual Sentiment Analysis:** Future work could also involve leveraging resources like BabelSenticNet [20] to improve the accuracy and depth of sentiment analysis across languages. Given the multilingual nature of the dataset, incorporating a commonsense reasoning framework, such as BabelSenticNet, can help better align sentiment detection with cultural nuances and conceptual understanding in different languages. This would allow a more nuanced analysis of sentiment trends across Persian, English, and other languages within the Twitter dataset.

Addressing these limitations in future research would involve extending the data collection period, incorporating a more diverse set of tweets including those opposing the movement, utilizing user metadata for more comprehensive analyses, overcoming data access limitations, conducting more detailed temporal analyses, and incorporating personality detection models. This would lead to a richer, more accurate understanding of the movement and its impacts.

REFERENCES

- [1] E. Cambria, R. Mao, M. Chen, Z. Wang, and S.-B. Ho, "Seven pillars for the future of artificial intelligence," *IEEE Intelligent Systems*, vol. 38, no. 6, pp. 62–69, 2023.
- [2] E. Cambria, *Understanding Natural Language Understanding*. Springer, ISBN 978-3-031-73973-6, 2024.
- [3] S. Jafari, A. Fathi, A. Hajizadegan, A. Kazemeini, and S. Eetemadi, "A longitudinal study about gradual changes in the Iranian online public sphere pre and post of 'mahsa moment': Focusing on Twitter," in *Proceedings of the 4th Workshop on Computational Approaches to Historical Language Change*, N. Tahmasebi, S. Montariol, H. Dubossarsky, A. Kutuzov, S. Hengchen, D. Alfter, F. Periti, and P. Cassotti, Eds. Singapore: Association for Computational Linguistics, Dec. 2023, pp. 46–52. [Online]. Available: <https://aclanthology.org/2023.lchange-1.5>
- [4] M. Grootendorst, "Bertopic: Neural topic modeling with a class-based tf-idf procedure," 2022.
- [5] Q. Liu, S. Han, Y. Li, E. Cambria, and K. Kwok, "PrimeNet: A framework for commonsense knowledge representation and reasoning based on conceptual primitives," *Cognitive Computation*, 2024.
- [6] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent dirichlet allocation," *Journal of machine Learning research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [7] R. Alghamdi and K. Alfalqi, "A survey of topic modeling in text mining," *Int. J. Adv. Comput. Sci. Appl.(IJACSA)*, vol. 6, no. 1, 2015.
- [8] S. Chang, R. Wang, P. Ren, and H. Huang, "Enhanced short text modeling: Leveraging large language models for topic refinement," *arXiv preprint arXiv:2403.17706*, 2024.
- [9] I. Chaturvedi, Y.-S. Ong, I. Tsang, R. Welsch, and E. Cambria, "Learning word dependencies in text by means of a deep recurrent belief network," *Knowledge-Based Systems*, vol. 108, pp. 144–154, 2016.
- [10] E. Cambria, X. Zhang, R. Mao, M. Chen, and K. Kwok, "SenticNet 8: Fusing emotion AI and commonsense AI for interpretable, trustworthy, and explainable affective computing," in *Proceedings of the International Conference on Human-Computer Interaction (HCI)*, Washington DC, USA, 2024.
- [11] C. Cortes and V. Vapnik, "Support-vector networks," *Machine learning*, vol. 20, pp. 273–297, 1995.
- [12] T. Bayes, "Naive bayes classifier," *Article Sources and Contributors*, pp. 1–9, 1968.
- [13] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," 2019.
- [14] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," 2019.
- [15] F. Barbieri, L. Espinosa Anke, and J. Camacho-Collados, "XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond," in *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, Jun. 2022, pp. 258–266. [Online]. Available: <https://aclanthology.org/2022.lrec-1.27>
- [16] S. Harlow, D. K. Brown, R. Salaverria, and V. Garcia-Perdomo, "Is the whole world watching? building a typology of protest coverage on social media from around the world," *Journalism Studies*, vol. 21, no. 11, pp. 1590–1608, 2020.
- [17] D. K. Kilgo, S. Harlow, V. Garcia-Perdomo, and R. Salaverria, "From# ferguson to# ayotzinapa: Analyzing differences in domestic and foreign protest news shared on social media," in *Social Media News and Its Impact*. Routledge, 2021, pp. 122–146.
- [18] J. Chen, S. Xiao, P. Zhang, K. Luo, D. Lian, and Z. Liu, "Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation," 2024.
- [19] A. Kazemeini, S. S. Roy, R. E. Mercer, and E. Cambria, "Interpretable representation learning for personality detection," in *2021 International conference on data mining workshops (ICDMW)*. IEEE, 2021, pp. 158–165.
- [20] D. Vilares, H. Peng, R. Satapathy, and E. Cambria, "BabelSenticNet: A commonsense reasoning framework for multilingual sentiment analysis," in *IEEE SSCI*, 2018, pp. 1292–1298.