

# BLM-17m: A Large-Scale Dataset for Black Lives Matter Topic Detection on Twitter

Hasan Kemik  
Information Technologies  
Carbon Consulting  
Istanbul, Turkey  
hasan.kemik@carbonconsulting.com

Nusret Özateş  
Department of Computer Science and Engineering  
Koc University  
Istanbul, Turkey  
mozates23@ku.edu.tr

Meysam Asgari-Chenaghlou  
Research And Development  
Ultimate  
Berlin, Germany  
meysam.asgarichenaghlou@ultimate.ai

Yang Li  
School of Automation  
Northwestern Polytechnical University  
Xi'an, China  
liyangu@nwpu.edu.cn

Erik Cambria  
School of Computer Science and Engineering  
Nanyang Technological University  
Singapore, Singapore  
cambria@ntu.edu.sg

**Abstract**—Protection of human rights is one of the most important problems of the modern world. In this paper, we construct a Twitter dataset that covers one of the most significant human rights contradiction in recent years which affected the whole world: the George Floyd incident. We propose a labeled dataset for topic detection that contains about 17 million tweets. These Tweets are collected from 25 May 2020 to 21 August 2020, covering about 90 days from the start of the incident. We labeled the dataset by monitoring most trending news topics from global and local newspapers and used TF-IDF and LDA as baselines. We evaluated the results of these two methods with three different  $k$  values for precision, recall and F1-score.

**Index Terms**—BlackLivesMatter, BLM, Sentiment Analysis, Natural Language Processing, AI, Social Media

## I. INTRODUCTION

The George Floyd incident has affected many people of color and defenders of liberty and democracy. People protested this horrible incident for months around the USA and all over the planet. Despite the corona virus outbreak in months before the incident, many kept their faith in democracy and liberty by protesting. The importance of this incident has motivated many around the world to start researching many aspects of it [1]–[3]. The hashtag #BlackLivesMatter or #BLM first emerged in 2013 following the acquittal of George Zimmerman in the shooting death of unarmed Black teen Trayvon Martin (Fig. 1). As the scourge of police brutality failed to subside over the years, the movement evolved and strengthened. The BLM movement sustained its momentum and activism after the 2014 shooting of Michael Brown by a white police officer. In a span of just three weeks following the grand jury's refusal to indict the officer, the #BlackLivesMatter hashtag was used over 1.7 million times on Twitter. Finally, the hashtag culminated in 2020 with the George Floyd incident [4]–[6]. To have a better understanding of the social stance for the people concerned, supporting, or even in contradiction with the George Floyd incident, we collected a dataset containing about 17 million tweets.




Fig. 1. History of the #BlackLivesMatter hashtag.

These tweets are collected by monitoring the most important hashtags related to the incident that people have used during the protests and afterwards. The data labeling involved monitoring the most important and trustworthy news outlets have published timely reports on the incident, protests and events after it. The remainder of this paper is organized as follows: Section II introduces the data collection methodology; Section III describes the data analysis methods; Section IV presents baseline results on the dataset; finally, Section V offers concluding remarks.

## II. DATA COLLECTION

We collected the BLM dataset<sup>1</sup> using various hashtags covering the majority of Tweets about this topic. Such hashtags are listed below:

- #BLM
- #BlackLivesMatter
- #JusticeForGeorgeFloyd
- #GeorgeFloyd
- #ICantBreathe
- #PoliceBrutality
- #PeacefulProtest

After data collection was completed, the tweets window was adjusted between the day of George Floyd incident is happened, 25 May 2020 [7], and 21 August 2020 [3], when George Floyd’s mural was defaced. The total number of tweets crawled is 16,782,467 (Table I). In Fig. 2, daily frequencies for Tweets are displayed in a time-series manner, showing high values on the first days of incident and a decrease on coming days but the event is still trending 90 days later.

Hashtag	Start Date	End Date	Tweet Count
blm	2020-05-25	2020-08-21	7,735,737
blacklivesmatter	2020-05-25	2020-08-21	1,341,533
justiceforgeorgefloyd	2020-05-25	2020-08-21	1,054,728
georgefloyd	2020-05-25	2020-08-21	4,545,748
icantbreathe	2020-05-25	2020-08-21	251,494
policebrutality	2020-05-25	2020-08-21	1,822,548
peacefulprotest	2020-05-25	2020-08-21	30,679

Total: 16,782,467

TABLE I  
CRAWLED TWEET NUMBER AND DATE RANGE BASED ON LABELS.

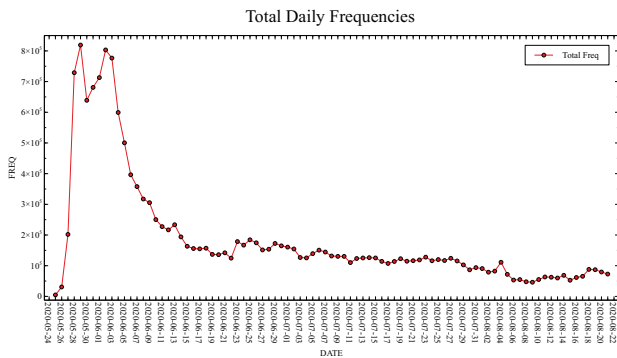


Fig. 2. Total daily tweet numbers.

Data labeling is accomplished using following steps:

- Local and international newspapers are examined in order to correctly determine the days of the events.
- Websites are also examined in the same way as the newspapers.
- After noting down important dates and articles found on newspapers and websites, keywords are extracted from each of them by hand.

<sup>1</sup>freely available for download at <https://github.com/senticnet/BLM>

Fig. 3 represents the top keywords extracted from the dataset on daily basis. We considered these keywords after lemmatization and stop-word removal. We kept the preprocessing minimal not to alter the dataset and its main characteristics. Some of the extracted keywords are listed below.




Fig. 3. Frequency analysis of top keywords.

- abetting
- accident
- activist
- american
- anger
- arrest
- black
- blacklivesmatter
- breathe
- burial
- chaos
- governor
- ground
- guard
- head
- icantbreathe
- injure
- justice
- killed
- manslaughter
- murder
- charged
- civil
- clash
- confederate
- crisis
- crowd
- curfew
- death
- fire
- floyd
- george
- officer
- peace
- police
- protest
- riot
- shooting
- symbol
- trump
- washington
- whitehouse

In Fig. 4, most frequent weekly words are visualized after lemmatization and stopword removal. This weekly view shows trending different keywords on different weeks. We labeled 5k sentences from the collected tweets and preprocessed them by removing urls, user mentions, twitter picture links and emojis. We labeled the tweets based on their sentiment and relatedness to BLM. We extracted keywords of the positive and negative tweets. For sentiment analysis, we use four labels: “1” for the positive tweets, “2” for the negative tweets, “3” for the Neutral tweets, and “4” for the “No data” tweets (for the tweets that only contain hashtags or mentions). The dataset contains 568 positive tweets, 1972 negative tweets, 1363 neutral tweets, and 648 “No Data” tweets. For the relatedness to BLM, we used 3 labels: “0” for not related, “1” for Related, and “4” for No data. This dataset contains 648 “No data” tweets, 560 “Not related” tweets, and 3343 “Related” tweets.





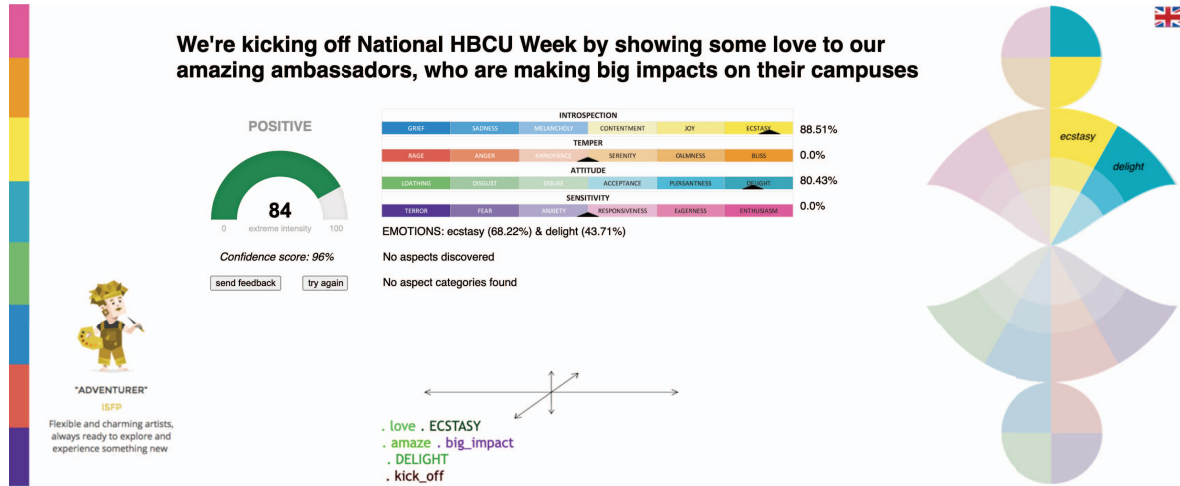


Fig. 7. Sentic API user interface sample.

### III. DATA ANALYSIS

To further enrich the manually labeled data, we use automatic tools for sentiment analysis. In recent years, sentiment analysis has become increasingly popular for social data analysis [8]. Different AI techniques have been leveraged to improve both accuracy and interpretability of sentiment analysis algorithms, including symbolic AI [9], [10], subsymbolic AI [11], [12], and neurosymbolic AI [13], [14]. Besides traditional algorithms [15] focusing on English text, multilingual [16] and multimodal [17] sentiment analysis have also attracted increasing attention recently. Typical applications of sentiment analysis include social network analysis [18], finance [19], and healthcare [20]. In particular, we use Sentic APIs<sup>2</sup>, a suite of application programming interfaces, which employ neurosymbolic AI to perform various sentiment analysis tasks in a fully interpretable manner (Fig. 7). A short description of each API and its usage within this work is provided in the next 12 subsections.

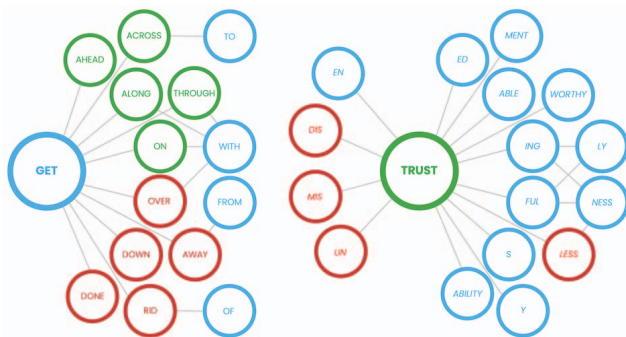


Fig. 8. Sentic Parser graph sample.

<sup>2</sup>freely available at <https://sentic.net/api>

#### A. Concept Parsing

This API provides access to Sentic Parser [21], a knowledge-specific concept parser based on SenticNet [14], which leverages both inflectional and derivational morphology for the efficient extraction and generalization of affective multiword expressions from English text. In particular, Sentic Parser is a hybrid semantic parser that uses an ensemble of constituency and dependency parsing and a mix of stemming and lemmatization to extract ‘semantic atoms’ like `pain_killer`, `go_bananas`, or `get_along_with`, which would carry different meaning and polarity if broken down into single words (Fig. 8). We use the API for extracting words and multiword expressions from tweets in order to better understand what are the key concepts related to BLM as shown in Fig. 4.

#### B. Subjectivity Detection

Subjectivity detection is an important NLP task that aims to filter out ‘factual’ content from data, i.e., objective text that does not contain any opinion. This API leverages a knowledge-sharing-based multitask learning framework powered by a neural tensor network, which consists of a bilinear tensor layer that links different entity vectors [22]. We used the API to identify BLM text as either objective (unopinionated) or subjective (opinionated) but also to handle neutrality, that is, text that is opinionated but neither positive nor negative (ambivalent stance towards the opinion target).

#### C. Polarity Classification

Once opinionated text is extracted using the *Subjectivity Detection API*, the *Polarity Classification API* further categorizes such text as either positive or negative. This is one of the most important APIs we use to understand the stance of tweeters towards BLM. It leverages an explainable fine-grained multiclass sentiment analysis method [23], which involves a multi-level modular structure designed to mimic natural language understanding processes, e.g., ambivalence handling process, sentiment strength handling process, etc.

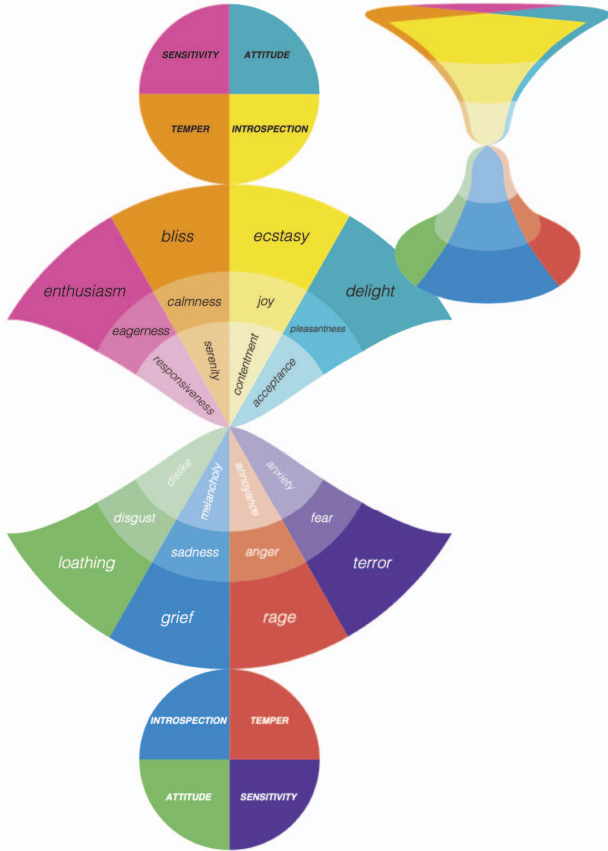


Fig. 9. The Hourglass of Emotions.

#### D. Intensity Ranking

For a finer-grained analysis, we further process the text classified by the *Polarity Classification API* using the *Intensity Ranking API* to infer its degree of negativity (floating-point number between -1 and 0) or positivity (floating-point number between 0 and 1). In particular, the API leverages a stacked ensemble method for predicting sentiment intensity by combining the outputs obtained from several deep learning and classical feature-based models using a multi-layer perceptron network [24].

#### E. Emotion Recognition

This API employs the Hourglass of Emotions [25], a biologically-inspired and psychologically-motivated emotion categorization model, that represents affective states both through labels and through four independent but concomitant affective dimensions, which can potentially describe the full range of emotional experiences that are rooted in any of us (Fig. 9). We use the API to go beyond polarity and intensity by examining what are the specific emotions elicited by BLM in both supporters and opposers of the movement.

#### F. Aspect Extraction

This API uses a meta-based self-training method that leverages both symbolic representations and subsymbolic learning for extracting aspects from text. In particular, a teacher model is trained to generate in-domain knowledge (e.g., unlabeled data selection and pseudo-label generation), where the generated pseudo-labels are used by a student model for supervised learning. Then, a meta-weighter is jointly trained with the student model to provide each instance with sub-task-specific weights to coordinate their convergence rates, balancing class labels, and alleviating noise impacts introduced from self-training [26]. We use the API to better understand the BLM phenomenon in terms of subtopics or opinion targets. Instead of simply identifying a polarity associated with the whole tweet, the *Aspect Extraction API* deconstructs input text into a series of specific aspects or features to then associate a polarity to each of them. This is particularly useful to process antithetic tweets, e.g., tweets where users list pros and cons of BLM.

#### G. Personality Prediction

This API uses a novel hard negative sampling strategy for zero-shot personality trait prediction from text using both OCEAN and MBTI models (Fig. 10). In particular, the API leverages an interpretable variational autoencoder sampler, to pair clauses under different relations as positive and hard negative samples, and a contrastive structured constraint, to disperse the paired samples in a semantic vector space [27]. We use the API to study the different personalities and personas involved in BLM discussions and, hence, better understand the possible drivers of such discussions.

#### H. Sarcasm Identification

This API combines commonsense knowledge and semantic similarity detection methods to better detect and process sarcasm in text. It also employs a contrastive learning approach with triplet loss to optimize the spatial distribution of sarcastic and non-sarcastic sample features [28]. We use the API to understand how much the BLM movement is subject to satire and critique but also to increase the accuracy and reliability of the *Polarity Classification API*. As sarcasm often involves expressing a sentiment that is opposite to the intended emotion, in fact, it may lead to polarity misclassification and, hence, generate wrong insights and conclusions.




Fig. 10. Personality prediction visualization sample.

### I. Depression Categorization

This API employs ensemble hybrid learning methods for automated depression categorization. In particular, the API combines symbolic AI (lexicon-based models) with subsymbolic AI (attention-based deep neural networks) to enhance the overall performance and robustness of depression detection [29]. We use it to study different reactions to BLM events by different users, especially those who are negatively affected by them.

### J. Toxicity Spotting

Given the controversy associated with BLM, it is important to measure the different types and intensities of toxicity associated with some of the tweets. This API is based on a multichannel convolutional bidirectional gated recurrent unit for detecting toxic comments in a multilabel environment [30]. In particular, the API extracts local features with many filters and different kernel sizes to model input words with long term dependency and then integrates multiple channels with a fully connected layer, normalization layer, and an output layer with a sigmoid activation function for predicting multilabel categories such as ‘obscene’, ‘threat’, or ‘hate’ (Fig. 11).

### K. Engagement Measurement

Measuring engagement is important to understand which specific events or topics are more impactful for the BLM movement (Fig. 1). This API employs a graph-embedding model that fuses heterogeneous data and metadata for the classification of engagement levels. In particular, the API leverages hybrid fusion methods for combining different types of data in a heterogeneous network by using semantic meta paths to constrain the embeddings [31].

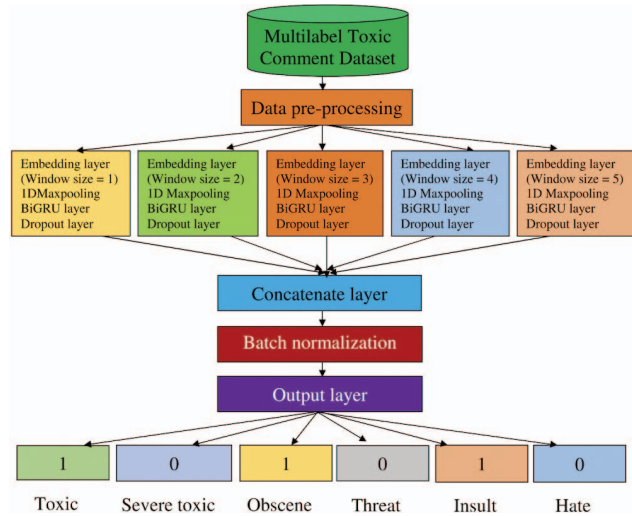


Fig. 11. Toxicity spotting framework.

### L. Well-being Assessment

Besides levels of toxicity and engagement, another important dimension for understanding BLM tweeters is their level of stress. This API leverages a mix of lexicons, embeddings, and pretrained language models for stress detection from social media texts [32]. In particular, the API employs a transformer-based model via transfer learning to capture the nuances of natural language expressions that convey stress in both explicit and implicit manners.

### M. API Results

Through Sentic APIs, we realized that most users used Twitter to share news articles, updates, and announcements related to BLM, including information about protests, rallies, legal developments, and social justice initiatives. Many also used Twitter to express their support for the BLM movement, share hashtags like #BlackLivesMatter, and show solidarity with the cause. Some tweets contained calls to action, encouraging followers to donate to BLM-related organizations, sign petitions, participate in protests, or engage in other forms of activism. Some users shared personal stories, experiences, or encounters related to racism and discrimination, illustrating the ongoing need for BLM’s goals. Finally, Twitter also served as a platform for artists, poets, and writers to share their creative work related to BLM, including visual art, poetry, and essays. In the figures below, we propose some visualizations of Sentic API results.




Fig. 12. Distribution of depression found in the tweets.




Fig. 13. Distribution of engagement found in the tweets.

Score Comparison for Sentiment Classification							
Method Name	Accuracy	Precision		Recall		F-1 Score	
		-	Macro	Weighted	Macro	Weighted	Macro
Linear SVC	0.62	0.61	0.62	0.53	0.62	0.54	0.61
KMeans	0.47	0.34	0.41	0.36	0.47	0.33	0.42
Logistic Regression	0.63	0.62	0.63	0.50	0.63	0.49	0.59

TABLE II  
SCORE COMPARISON FOR SENTIMENT CLASSIFICATION.




Fig. 14. Distribution of toxicity found in the tweets.

Method	Keyword Count	Precision	Recall	F1-score
TF-IDF	Top 10	0.0438	0.0250	0.0318
LDA	Top 10	0.1314	0.0750	0.0955
TF-IDF	Top 25	0.0949	0.0217	0.0650
LDA	Top 25	0.2044	0.0467	0.1400
TF-IDF	Top 50	0.2920	0.0333	0.1000
LDA	Top 50	0.3066	0.0350	0.1050

TABLE III  
SCORE COMPARISON FOR KEYWORD EXTRACTION.

Based on Fig. 13, we can say that the tweets that includes advertisement goes down to -100 range, while the tweets that are actually related to the topic lies right to the 0 point. Based on the similarities of distribution in the Fig. 12 and Fig. 14, we can also see that the majority of the data that we collected mostly convey sadness, rather than aggressiveness, towards the topic.

#### IV. BASELINES

TF-IDF and LDA extraction methods on daily basis is used to provide baselines for this dataset. Top 10, 25 and 50 extracted keywords is used to check if they exist in the ground truth labels or not. This step can also be considered as a quality check because it ensures that with minimum effort and easier to use methods such as TF-IDF and LDA, results keeps getting better on various top-k analysis. We used standard metrics as main evaluation anchors are precision, recall and f-measures (as specified in equations 1, 2 and 3).

$$Precision = \frac{|\{extracted\_keywords\} \cap \{ground\_truth\}|}{|\{ground\_truth\}|} \quad (1)$$

$$Recall = \frac{|\{extracted\_keywords\} \cap \{ground\_truth\}|}{|\{extracted\_keywords\}|} \quad (2)$$

$$F1 - score = \frac{2 * Precision * Recall}{(Precision + Recall)} \quad (3)$$

Results are presented in Table III for various top-k precision, recall and F1-score. These metrics present an overall improvement between the two baselines, TF-IDF and LDA. According to this table, best obtained results for top-10, 25 and 50 are 0.1314, 0.2044 and 0.3066 for precision. Finally, we have also conducted another experiment by using the 17M tweets and a TF-IDF Vectorizer to classify the sentiment scores of the provided tweets. We transformed the tweets from the dataset to vectors by using the trained TF-IDF Vectorizer. By using these vectors, we trained three different sentence classification methods as a baseline. All methods were trained using sklearn with default values.

#### V. CONCLUSION

In this work, we have collected a dataset of about 17 million tweets about the George Floyd incident. We have used several machine learning and AI techniques to analyze the collected data and produce several insights. AI holds immense potential to advance and protect human rights on a global scale. We hope this work can foster future investigations on similar topics, such as identifying and combating human rights abuses, enhancing access to justice, and promoting accountability.

By providing tools for early detection, evidence collection, and advocacy, AI has the potential to significantly contribute to the protection and promotion of human rights worldwide. However, it is crucial to navigate the ethical and privacy implications of AI in this context to ensure that these technologies are used responsibly and in alignment with human rights principles.

#### ACKNOWLEDGMENTS

This research is supported by the Agency for Science, Technology and Research (A\*STAR) under its AME Programmatic Funding Scheme (Project #A18A2b0046).

#### REFERENCES

- [1] T. B. Oriola and W. A. Knight, "Covid-19, george floyd and human security," *African Security*, vol. 13, no. 2, pp. 111–115, 2020.
- [2] B. P. Dreyer, M. Trent, A. T. Anderson, G. L. Askew, R. Boyd, T. R. Coker, T. Coyne-Beasley, E. Fuentes-Afflick, T. Johnson, F. Mendoza *et al.*, "The death of george floyd: Bending the arc of history toward justice for generations of children," *Pediatrics*, vol. 146, no. 3, 2020.



- [3] T. A. Press, "George floyd mural in minneapolis defaced with black paint," Aug 2020. [Online]. Available: <https://abcnews.go.com/US/wireStory/george-floyd-mural-minneapolis-defaced-black-paint-72504250>
- [4] J. Clement, "Twitter global mda 2020," Jul 2020. [Online]. Available: <https://www.statista.com/statistics/970920/monetizable-daily-active-twitter-users-worldwide/>
- [5] C. Castillo, M. Mendoza, and B. Poblete, "Information credibility on twitter," in *Proceedings of the 20th international conference on World wide web*, 2011, pp. 675–684.
- [6] S. Petrovic, M. Osborne, R. McCreddie, C. Macdonald, and I. Ounis, "Can twitter replace newswire for breaking news?" 2013.
- [7] Reuters, "Timeline: Key events in the month since george floyd's death," Jun 2020. [Online]. Available: <https://www.reuters.com/article/us-minneapolis-police-usa-onemonth-timel-idUSKBN23W1NR>
- [8] L. Oneto, F. Bisio, E. Cambria, and D. Anguita, "Statistical learning theory and ELM for big social data analysis," *IEEE Computational Intelligence Magazine*, vol. 11, no. 3, pp. 45–55, 2016.
- [9] E. Cambria, A. Hussain, C. Havasi, and C. Eckl, "Common sense computing: From the society of mind to digital intuition and beyond," in *Biometric ID Management and Multimodal Communication*, ser. Lecture Notes in Computer Science. Berlin Heidelberg: Springer, 2009, vol. 5707, pp. 252–259.
- [10] F. Xing, F. Pallucchini, and E. Cambria, "Cognitive-inspired domain adaptation of sentiment lexicons," *Information Processing and Management*, vol. 56, no. 3, pp. 554–564, 2019.
- [11] E. Cambria, T. Mazzocco, A. Hussain, and C. Eckl, "Sentic medoids: Organizing affective common sense knowledge in a multi-dimensional vector space," ser. Lecture Notes in Computer Science. Berlin Heidelberg: Springer-Verlag, 2011, vol. 6677, pp. 601–610.
- [12] E. Cambria, B. Schuller, B. Liu, H. Wang, and C. Havasi, "Statistical approaches to concept-level sentiment analysis," *IEEE Intelligent Systems*, vol. 28, no. 3, pp. 6–9, 2013.
- [13] A. Valdivia, V. Luzón, E. Cambria, and F. Herrera, "Consensus vote models for detecting and filtering neutrality in sentiment analysis," *Information Fusion*, vol. 44, pp. 126–135, 2018.
- [14] E. Cambria, Q. Liu, S. Decherchi, F. Xing, and K. Kwok, "SenticNet 7: A commonsense-based neurosymbolic AI framework for explainable sentiment analysis," in *LREC*, 2022, pp. 3829–3839.
- [15] H. T. Nguyen, P. H. Duong, and E. Cambria, "Learning short-text semantic similarity with word embeddings and external knowledge sources," *Knowledge-Based Systems*, vol. 182, no. 104842, 2019.
- [16] 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.
- [17] E. Cambria, N. Howard, J. Hsu, and A. Hussain, "Sentic blending: Scalable multimodal fusion for continuous interpretation of semantics and sentics," in *IEEE SSCI*, Singapore, 2013, pp. 108–117.
- [18] S. Cavallari, E. Cambria, H. Cai, K. Chang, and V. Zheng, "Embedding both finite and infinite communities on graph," *IEEE Computational Intelligence Magazine*, vol. 14, no. 3, pp. 39–50, 2019.
- [19] F. Xing, E. Cambria, and R. Welsch, "Intelligent asset allocation via market sentiment views," *IEEE Computational Intelligence Magazine*, vol. 13, no. 4, pp. 25–34, 2018.
- [20] E. Cambria, T. Benson, C. Eckl, and A. Hussain, "Sentic PROMs: Application of sentic computing to the development of a novel unified framework for measuring health-care quality," *Expert Systems with Applications*, vol. 39, no. 12, pp. 10 533–10 543, 2012.
- [21] E. Cambria, R. Mao, S. Han, and Q. Liu, "Sentic parser: A graph-based approach to concept extraction for sentiment analysis," in *Proceedings of ICDM Workshops*, 2022, pp. 413–420.
- [22] R. Satapathy, S. Rajesh Pardeshi, and E. Cambria, "Polarity and subjectivity detection with multitask learning and bert embedding," *Future Internet*, vol. 14, no. 7, p. 191, 2022.
- [23] Z. Wang, Z. Hu, S.-B. Ho, E. Cambria, and A.-H. Tan, "Mimusa—mimicking human language understanding for fine-grained multi-class sentiment analysis," *Neural Computing and Applications*, vol. 35, no. 21, pp. 15 907–15 921, 2023.
- [24] M. S. Akhtar, A. Ekbal, and E. Cambria, "How intense are you? predicting intensities of emotions and sentiments using stacked ensemble," *IEEE Computational Intelligence Magazine*, vol. 15, no. 1, pp. 64–75, 2020.
- [25] Y. Susanto, A. Livingstone, B. C. Ng, and E. Cambria, "The hourglass model revisited," *IEEE Intelligent Systems*, vol. 35, no. 5, pp. 96–102, 2020.
- [26] K. He, R. Mao, T. Gong, C. Li, and E. Cambria, "Meta-based self-training and re-weighting for aspect-based sentiment analysis," *IEEE Transactions on Affective Computing*, vol. 15, 2024.
- [27] L. Zhu, W. Li, R. Mao, V. Pandelea, and E. Cambria, "PAED: Zero-shot persona attribute extraction in dialogues," in *ACL*, 2023, pp. 9771–9787.
- [28] T. Yue, R. Mao, H. Wang, Z. Hu, and E. Cambria, "Knowlenet: Knowledge fusion network for multimodal sarcasm detection," *Information Fusion*, vol. 100, p. 101921, 2023.
- [29] L. Ansari, S. Ji, Q. Chen, and E. Cambria, "Ensemble hybrid learning methods for automated depression detection," *IEEE Transactions on Computational Social Systems*, vol. 10, no. 1, pp. 211–219, 2023.
- [30] J. A. Kumar, S. Abirami, T. E. Trueman, and E. Cambria, "Comment toxicity detection via a multichannel convolutional bidirectional gated recurrent unit," *Neurocomputing*, vol. 441, pp. 272–278, 2021.
- [31] I. Chaturvedi, K. Thapa, S. Cavallari, E. Cambria, and R. E. Welsch, "Predicting video engagement using heterogeneous deepwalk," *Neurocomputing*, vol. 465, pp. 228–237, 2021.
- [32] A. Rastogi, Q. Liu, and E. Cambria, "Stress detection from social media articles: New dataset benchmark and analytical study," in *IJCNN*, 2022.