Explainable AI for Stress and Depression Detection in the Cyberspace and Beyond

Erik Cambria[1(✉)], Balázs Gulyás[1], Joyce S. Pang[1], Nigel V. Marsh[2], and Mythily Subramaniam[3]

[1] Nanyang Technological University, Singapore, Singapore
{cambria,balazs.gulyas,joycepang}@ntu.edu.sg
[2] James Cook University, Singapore, Singapore
nigel.marsh@jcu.edu.au
[3] Institute of Mental Health, Singapore, Singapore
mythily@imh.com.sg

Abstract. Stress and depression have emerged as prevalent challenges in contemporary society, deeply intertwined with the complexities of modern life. This paper delves into the multifaceted nature of these phenomena, exploring their intricate relationship with various sociocultural, technological, and environmental factors through the application of neurosymbolic AI to social media content. Through a quantitative and qualitative analysis of results, we elucidate the profound impact of technological advancements on information processing, work culture, and social dynamics, highlighting the role of digital connectivity in exacerbating stressors. Economic pressures and social isolation further compound these challenges, contributing to a pervasive sense of unease and disconnection. Environmental stressors, including climate change, add another layer of complexity, fostering existential concerns about the future. Moreover, the persistent stigma surrounding mental health perpetuates a cycle of silence and suffering, hindering access to support and resources. Addressing these issues necessitates a holistic approach, encompassing societal changes, policy interventions, and individual coping strategies.

Keywords: Stress detection · Depression Detection · Artificial Intelligence · Natural Language Processing · Affective Computing

1 Introduction

In the contemporary world, stress and depression have become pervasive issues, intricately woven into the fabric of modern life. One of the primary culprits is the rapid advancement of technology. While it has undeniably enhanced convenience and connectivity, it has also ushered in a relentless cycle of information overload. We find ourselves inundated with constant streams of data, struggling to sift through the endless notifications and updates bombarding our digital existence. This constant barrage can lead to feelings of overwhelm, anxiety, and a pervasive sense of being constantly "on".

Z. Wang and C. W. Tan (Eds.): PAKDD 2024 Workshops, LNAI 14658, pp. 108–120, 2024.
https://doi.org/10.1007/978-981-97-2650-9_9

The modern workplace culture exacerbates these pressures. Fueled by a relentless pursuit of productivity, many individuals find themselves caught in a ceaseless cycle of long hours and mounting expectations. The pressure to perform, coupled with the looming specter of job insecurity, fosters an environment ripe for stress and burnout. Economic pressures further compound these challenges. With living costs on the rise and economic uncertainty looming large, many individuals find themselves grappling with the weight of financial strain. The struggle to make ends meet adds another layer of stress, exacerbating feelings of anxiety and hopelessness. Yet, perhaps one of the most insidious aspects of the modern era is the pervasive sense of social isolation. Despite the veneer of connectivity offered by social media, many individuals find themselves grappling with profound feelings of loneliness. The digital age has paradoxically alienated us from genuine human connection, leaving us longing for meaningful relationships amidst a sea of superficial interactions.

Additionally, environmental factors such as pollution and climate change cast a looming shadow over our collective psyche, fostering a sense of existential dread about the future. Coupled with the relentless pressure to measure up to curated versions of perfection portrayed on social media, many individuals find themselves mired in a cycle of comparison and self-doubt. Compounding these challenges is the persistent stigma surrounding mental health. Despite growing awareness, discussions about mental illness are often shrouded in silence and shame. This pervasive stigma can act as a barrier to seeking help, perpetuating a cycle of suffering in silence.

In the context of computer science, early works have used linguistic analysis to detect signs of depression from text. These had two main drawbacks: they were not very accurate and they only detected presence or absence of depression. More recently, advanced AI techniques were used for a more accurate and finer-grained analysis of depression from both text and videos [2,10,11]. This was part of a larger effort to adopt state-of-the-art AI techniques for healthcare [14], with particular focus on mental health [17,18] and suicidal ideation detection [15,16]. Despite more performant, however, these new algorithms still had a crucial drawback: they were not explainable, which made them virtually useless for clinicians and mental healthcare experts.

In this work, we apply explainable AI (XAI) [4] to the problem of stress and depression detection from social media to provide a unique lens through which researchers and mental health professionals can observe and understand (and possibly even monitor and prevent) mental health trends in real-time across a wide and diverse population. In particular, we collected about 300,000 tweets about stress and depression and employed state-of-the-art neurosymbolic AI tools to gain a deeper, more nuanced understanding of their potential causes and contributing factors in the modern era.

The remainder of this paper is organized as follows: Sect. 2 introduces our data collection methodology; Sect. 3 describes the data analysis approach undertaken; Sect. 4 discusses results; finally, Sect. 5 offers concluding remarks and outlines future work.

2 Data Collection

We collected our stress and depression dataset from Twitter using the 10 most popular hashtags listed below. We used the new Twitter Pro API package (priced at $5,000 per month) for one month between 1st January to 1st February 2024.

- `#MentalHealth`: This hashtag is widely used to discuss various aspects of mental health, including stress, depression, anxiety, and other related conditions. It encompasses conversations about personal experiences, coping strategies, and advocacy efforts.
- `#Depression`: This hashtag specifically focuses on discussions surrounding depression, a mood disorder characterized by persistent feelings of sadness, hopelessness, and loss of interest. It is often used to share personal stories, raise awareness, and provide support to those struggling with depression.
- `#Anxiety`: Anxiety is a common mental health condition characterized by excessive worry, fear, and apprehension. The #Anxiety hashtag is used to share experiences, coping mechanisms, and resources for managing anxiety-related symptoms.
- `#Stress`: This hashtag is used to discuss the experience of stress, which refers to the body's response to perceived threats or challenges. Discussions under this hashtag include triggers of stress, coping strategies, and the impact of chronic stress on mental and physical health.
- `#SelfCare`: Self-care involves intentionally taking care of one's physical, emotional, and mental well-being. The #SelfCare hashtag is used to share tips, practices, and experiences related to self-care activities that can help alleviate stress and promote overall wellness.
- `#MentalHealthAwareness`: This hashtag is used to raise awareness about mental health issues, including stress and depression, and to promote understanding, acceptance, and support for individuals experiencing mental health challenges.
- `#EndStigma`: Stigma surrounding mental health can create barriers to seeking help and support. The #EndStigma hashtag is used to advocate for ending the discrimination and prejudice associated with mental illness, fostering a more inclusive and supportive society.
- `#MentalHealthMatters`: This hashtag emphasizes the importance of prioritizing mental health and acknowledging its significance in overall well-being. It is often used to promote conversations, initiatives, and policies aimed at addressing mental health issues such as stress and depression.
- `#Wellness`: Wellness encompasses various dimensions of health, including physical, mental, emotional, and social well-being. The #Wellness hashtag is used to share tips, resources, and practices that support holistic health and promote stress reduction and resilience.
- `#SelfLove`: Self-love involves cultivating a positive and compassionate relationship with oneself. The #SelfLove hashtag is used to promote self-acceptance, self-care, and self-compassion, which are important aspects of managing stress and improving mental health.

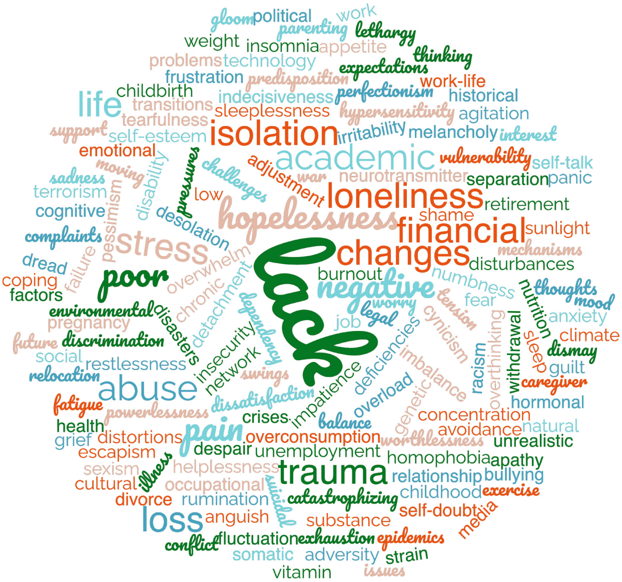

Fig. 1. Word cloud representing the top keywords in the dataset

Within one month, we collected a total of one million tweets. After preprocessing (e.g., removal of irrelevant tweets, removal of non-English tweets, removal of duplicates, removal of re-tweets, etc.), we were left with roughly one third of it. The exact distribution of tweets with respect to hashtags is illustrated in Table 1. Figure 1 proposes a visual representation of the most significant terms in the collected dataset (after stopword removal), where the size of each word is proportional to its frequency.

Table 1. Distribution of collected tweets with respect to hashtags.

Hashtag	Start Date	End Date	Tweet Count
#MentalHealth	01-01-2024	01-02-2024	98,034
#Depression	01-01-2024	01-02-2024	82,736
#Anxiety	01-01-2024	01-02-2024	70,304
#Stress	01-01-2024	01-02-2024	31,036
#SelfCare	01-01-2024	01-02-2024	6,802
#MentalHealthAwareness	01-01-2024	01-02-2024	5,124
#EndStigma	01-01-2024	01-02-2024	4,293
#MentalHealthMatters	01-01-2024	01-02-2024	2,640
#Wellness	01-01-2024	01-02-2024	2,479
#SelfLove	01-01-2024	01-02-2024	1,935

Total: *305,383*

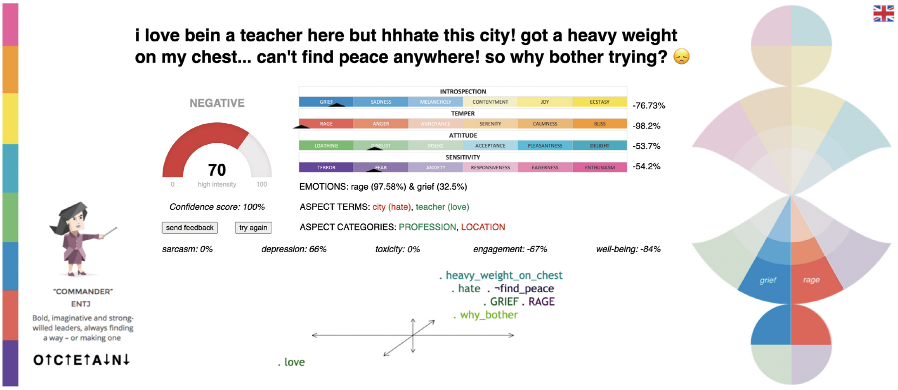

Fig. 2. Sentic API user interface sample.

3 Data Analysis

In order to gain insights from the collected data, we leverage sentiment analysis, a natural language processing (NLP) field in which computational methods are used to extract emotions from text. Different AI techniques have been leveraged to improve both accuracy and interpretability of sentiment analysis algorithms, including symbolic AI, subsymbolic AI, and neurosymbolic AI [5]. Besides traditional algorithms [20] focusing on English text, multilingual [21,26] and multimodal [24] sentiment analysis have also attracted increasing attention recently. Typical applications of sentiment analysis include social network analysis [8], finance [28], and healthcare [3]. In this work, we use Sentic APIs[1], a suite of application programming interfaces available in 80 languages, which employ neurosymbolic AI to perform various sentiment analysis tasks in a fully interpretable manner (Fig. 2). A short description of each API and its usage within this work is provided in the next 12 subsections.

3.1 Concept Parsing

This API provides access to Sentic Parser [6], a knowledge-specific concept parser based on SenticNet [7], which leverages both inflectional and derivational morphology for the efficient extraction and generalization of affective multiword expressions from 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 meaningful multiword expressions. We use the API for extracting words and multiword expressions from text in order to better understand what are the key concepts related to stress and depression. As shown in Fig. 2, for example, some of the concepts extracted are `hate`, `why_bother`, and `heavy_weight_on_chest`.

[1] https://sentic.net/api.

3.2 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 [23]. We use the API to classify stress and depression-related tweets as either objective (unopinionated) or subjective (opinionated) but also to handle neutrality, that is, a tweet that is opinionated but neither positive nor negative (ambivalent stance towards the opinion target). All labels come with a confidence score based on how much SenticNet concepts contributed to the classification output. As depicted in Fig. 2, the confidence score of the proposed example is 100%. Finally, the *Subjectivity Detection module* is also responsible for identifying the language of the input, as indicated in the top-right corner of the UI.

3.3 Polarity Classification

Once an opinionated tweet is detected using the *Subjectivity Detection API*, the *Polarity Classification API* further categorizes this tweet as either positive or negative. This is one of the most important APIs we use to understand the stance of tweeters towards stress and depression. It leverages an explainable fine-grained multiclass sentiment analysis method [27], which involves a multi-level modular structure designed to mimic natural language understanding processes, e.g., ambivalence handling process, sentiment strength handling process, etc. As illustrated in Fig. 2, for example, the extracted polarity is NEGATIVE.

3.4 Intensity Ranking

For a finer-grained analysis, we further process stress and depression classified by the *Polarity Classification API* using the *Intensity Ranking API* to infer their degree of negativity (floating-point number between -100 and 0) or positivity (floating-point number between 0 and 100). 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 [1]. As shown in Fig. 2, the extracted polarity of the proposed example is 70 (high intensity).

3.5 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. We use the API to go beyond polarity and intensity by examining what are the specific emotions elicited by stress and depression in both their ardent supporters and vocal opposers. As depicted in Fig. 2, for example, the emotion spectrum of the input is visualized in terms of the Hourglass Model's affective dimensions, namely: -76.73% Introspection, -98.2% Temper, -53.7% Attitude, and -54.2% Sensitivity. From these, the API also extracts the two top resulting emotion labels, rage and grief, with an intensity of 97.58% and 32.5%, respectively.

3.6 Aspect Extraction

This API uses a meta-based self-training method that leverages both symbolic representations and subsymbolic learning for extracting aspects from text. A teacher model is trained to generate in-domain knowledge, where the generated pseudo-labels are used by a student model for supervised learning [13]. We use the API to better understand stress and depression in terms of subtopics. 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 containing both positive and negative opinion targets. As illustrated in Fig. 2, the aspect terms extracted from the proposed example are `city` and `teacher`, which belong to the aspect categories `LOCATION` and `PROFESSION`, respectively. The UI also displays the affective concepts most relevant to each aspect term (in brackets) which are also colored according to their respective polarities (green for positive and red for negative).

3.7 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. 3). 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 [30]. We use the API to study the different personalities and personas involved in stress and depression discussions and, hence, better understand the possible drivers of such discussions. As shown in Fig. 2, for example, the MBTI personality extracted is ENTJ (Extraverted, iNtuitive, Thinking, and Judging) and the OCEAN personality traits extracted are O↑C↑E↑A↓N↓, i.e., high Openness, high Conscientiousness, high Extraversion, low Agreeableness, and low Neuroticism.

3.8 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 [29]. We use the API to understand how much stress and depression are subject to satire and critique but also

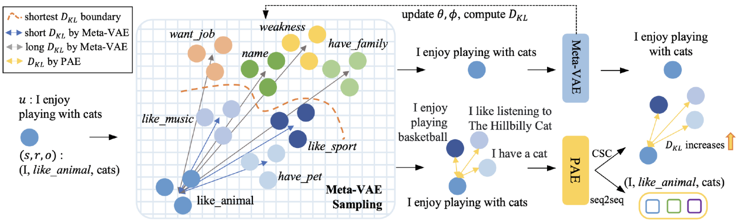

Fig. 3. Personality prediction visualization sample

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. The sarcasm score goes from zero (no sarcasm detected) to 100 (extremely sarcastic content). As depicted in Fig. 2, no sarcasm was detected in the proposed example.

3.9 Depression Categorization

This API employs a novel encoder combining hierarchical attention mechanisms and feed-forward neural networks. To support psycholinguistic studies, the model leverages metaphorical concept mappings as input. Thus, it not only detects depression, but also identifies features of such users' tweets and associated metaphor concept mappings [12]. We use it to discover common causes of depression but also to study different reactions to it by different users. The depression score ranges from zero (no depression detected) to 100 (severe depression). As illustrated in Fig. 2, for example, the depression score is 66%.

3.10 Toxicity Spotting

This API is based on a multichannel convolutional bidirectional gated recurrent unit for detecting toxic comments in a multilabel environment [19]. 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'. The toxicity score goes from zero (no toxicity detected) to 100 (highly toxic content). As shown in Fig. 2, there was no toxicity detected in the proposed example.

3.11 Engagement Measurement

Measuring engagement is important to understand which specific topics or events are more impactful for both stress and depression. 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 [9]. The engagement score ranges from -100 (high disengagement) to 100 (high engagement). As depicted in Fig. 2, for example, the engagement score is -67%.

3.12 Well-Being Assessment

This API leverages a mix of lexicons, embeddings, and pretrained language models for stress detection from social media texts [22]. 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. The well-being score ranges from -100 (high stress) to 100 (high well-being). As illustrated in Fig. 2, the well-being score is -84% in the proposed example.

4 Results

In this section, we discuss the most important insights gained through the use of Sentic APIs on the collected dataset. The *Concept Parsing API* enabled us to discover what are the current hot topics related to stress and depression, e.g., `anxiety`, `chronic_stress`, `burnout`, `sadness`, `hopelessness`, `despair`, `worry`, `isolation`, `exhaustion`, and `social_withdrawal`.

Through the *Subjectivity Detection API*, we realized that the vast majority of stress and depression tweets were opinionated. The unopinionated tweets were mostly promotional and advertising posts. This was further validated by the results of the *Intensity Ranking API*, which were high for both negative and positive spectrum.

By processing subjective text using the *Polarity Classification API*, we obtained 59% negative tweets and 41% positive tweets (at least for the time window of our analysis). For the former group, the most common MBTI personality type was ENTJ and the predominant emotion was anxiety. The latter group (the tweeters promoting self-help), instead, was characterized by an INFJ personality trait and a predominant emotion of acceptance.

The *Sarcasm Identification API* has flagged a subtle presence of sarcasm within the context of this particular topic, albeit at a low level. Unlike many other topics discussed on social media, stress and depression do not seem to lend themselves well to sarcasm, probably because of the lack of ambiguity and straightforwardness gravitating around the subject.

As expected, the *Depression Categorization API* observed high levels of depression. Some individuals used Twitter as a means of reaching out for support and validation from their online communities. Some others shared their thoughts and experiences to externalize their internal struggles and potentially receive empathy and support from others. Finally, some users used Twitter as an outlet to vent their frustration, anger, or despair.

The *Toxicity Spotting API* also did not pick up much toxic content. One significant factor is the nature of stress and depression, which inherently lends itself to more neutral or consensus-based discussions, minimizing the potential for conflict or toxicity. Most individuals approached the topic with openness, curiosity, and a willingness to listen to differing perspectives and, hence, fostered an environment conducive to rather constructive dialogues without personal attacks nor hostility.

The *Engagement Measurement API* exhibited high levels of disengagement, most likely caused by emotional exhaustion (characterized by feelings of emptiness, apathy, and detachment) but also anhedonia (the inability to experience pleasure) and social withdrawal.

The *Well-being Assessment API* detected high levels of stress involving emotional intensity, uncertainty, conflict, pressure, personal vulnerability, information overload, lack of control, and negative social dynamics.

Finally, some very useful insights came from the *Aspect Extraction API*, which helped us individuate key causes of stress and depression. We list the 10 most frequent ones below along with a short elucubration on why such aspects emerged from the over 300,000 tweets as the most prominent.

- `Relationship issues`: Problems within intimate relationships or family conflicts can impact mental health and contribute to depressive symptoms.
- `Financial problems`: Financial stress, such as debt, unemployment, or financial instability, can lead to feelings of hopelessness.
- `Social isolation`: Lack of social support and feelings of loneliness can cause depression, as social connections are essential for emotional well-being.
- `Work-life balance`: Difficulty balancing work responsibilities with personal life and self-care can lead to chronic stress and impact mental well-being.
- `Academic pressure`: Students experience stress and depression due to academic demands, performance pressure, or difficulty coping with coursework.
- `Discrimination`: Experiencing discrimination based on race, ethnicity, gender identity, sexual orientation, or other factors can lead to chronic stress.
- `Chronic pain`: Living with chronic health conditions or experiencing persistent pain can be emotionally draining and exacerbate feelings of depression.
- `Trauma`: Past trauma, including physical, emotional, or sexual abuse, can have long-lasting effects on mental health and increase the risk of depression.
- `Media exposure`: Overexposure to negative news, social media comparison, or unrealistic portrayals of success can contribute to feelings of inadequacy.
- `Environmental factors`: Environmental stressors such as pollution, noise, or overcrowding can contribute to chronic stress and impact mental health.

5 Conclusion

Stress and depression have emerged as prevalent challenges in contemporary society, deeply intertwined with the complexities of modern life. This paper delves into the multifaceted nature of these phenomena, exploring their intricate relationship with various socio-cultural, technological, and environmental factors through the application of neurosymbolic AI to social media content. Through a quantitative and qualitative analysis of results, we elucidate the profound impact of technological advancements on information processing, work culture, and social dynamics, highlighting the role of digital connectivity in exacerbating stressors. Economic pressures and social isolation further compound these challenges, contributing to a pervasive sense of unease and disconnection.

Environmental stressors, including climate change, add another layer of complexity, fostering existential concerns about the future. Moreover, the persistent stigma surrounding mental health perpetuates a cycle of silence and suffering, hindering access to support and resources. Addressing these issues necessitates a holistic approach, encompassing societal changes, policy interventions, and individual coping strategies. By fostering greater awareness, empathy, and collective action, we can strive towards a more resilient and compassionate society, better equipped to navigate the complexities of the modern era.

Acknowledgments. This research/project is supported by the Ministry of Education, Singapore under its MOE Academic Research Fund Tier 2 (STEM RIE2025 Award MOE-T2EP20123-0005).

References

1. Akhtar, M.S., Ekbal, A., Cambria, E.: How intense are you? Predicting intensities of emotions and sentiments using stacked ensemble. IEEE Comput. Intell. Mag. **15**(1), 64–75 (2020)
2. Ansari, L., Ji, S., Chen, Q., Cambria, E.: Ensemble hybrid learning methods for automated depression detection. IEEE Trans. Comput. Soc. Syst. **10**(1), 211–219 (2023)
3. Cambria, E., Benson, T., Eckl, C., Hussain, A.: Sentic PROMs: application of sentic computing to the development of a novel unified framework for measuring health-care quality. Expert Syst. Appl. **39**(12), 10533–10543 (2012)
4. Cambria, E., Malandri, L., Mercorio, F., Mezzanzanica, M., Nobani, N.: A survey on XAI and natural language explanations. Inf. Process. Manage. **60**, 103111 (2023)
5. Cambria, E., Mao, R., Chen, M., Wang, Z., Ho, S.B.: Seven pillars for the future of artificial intelligence. IEEE Intell. Syst. **38**(6), 62–69 (2023)
6. Cambria, E., Mao, R., Han, S., Liu, Q.: Sentic parser: a graph-based approach to concept extraction for sentiment analysis. In: Proceedings of ICDM Workshops, pp. 413–420 (2022)
7. Cambria, E., Zhang, X., Mao, R., Chen, M., Kwok, K.: SenticNet 8: fusing emotion AI and commonsense AI for interpretable, trustworthy, and explainable affective computing. In: Proceedings of the 26th International Conference on Human-computer Interaction (2024)
8. Cavallari, S., Cambria, E., Cai, H., Chang, K., Zheng, V.: Embedding both finite and infinite communities on graph. IEEE Comput. Intell. Mag. **14**(3), 39–50 (2019)
9. Chaturvedi, I., Thapa, K., Cavallari, S., Cambria, E., Welsch, R.E.: Predicting video engagement using heterogeneous DeepWalk. Neurocomputing **465**, 228–237 (2021)
10. Chen, Q., Chaturvedi, I., Ji, S., Cambria, E.: Sequential fusion of facial appearance and dynamics for depression recognition. Pattern Recogn. Lett. **150**, 115–121 (2021)

11. Chiong, R., Budhi, G., Cambria, E.: Detecting signs of depression using social media texts through an ensemble of ensemble classifiers. IEEE Trans. Affect. Comput. **15** (2024)
12. Han, S., Mao, R., Cambria, E.: Hierarchical attention network for explainable depression detection on twitter aided by metaphor concept mappings. In: Proceedings of COLING, pp. 94–104 (2022)
13. He, K., Mao, R., Gong, T., Li, C., Cambria, E.: Meta-based self-training and re-weighting for aspect-based sentiment analysis. IEEE Trans. Affect. Comput. **15** (2024)
14. He, K., et al.: A survey of large language models for healthcare: From data, technology, and applications to accountability and ethics. arXiv preprint arXiv:2310.05694 (2024)
15. Ji, S., Li, X., Huang, Z., Cambria, E.: Suicidal ideation and mental disorder detection with attentive relation networks. Neural Comput. Appl. **34**, 10309–10319 (2022)
16. Ji, S., Pan, S., Li, X., Cambria, E., Long, G., Huang, Z.: Suicidal ideation detection: a review of machine learning methods and applications. IEEE Trans. Comput. Soc. Syst. **8**(1), 214–226 (2021)
17. Ji, S., Zhang, T., Yang, K., Ananiadou, S., Cambria, E.: Rethinking large language models in mental health applications. arXiv preprint arXiv:2311.11267 (2024)
18. Ji, S., Zhang, T., Yang, K., Ananiadou, S., Cambria, E., Tiedemann, J.: Domain-specific continued pretraining of language models for capturing long context in mental health. arXiv preprint arXiv:2304.10447 (2024)
19. Kumar, J.A., Abirami, S., Trueman, T.E., Cambria, E.: Comment toxicity detection via a multichannel convolutional bidirectional gated recurrent unit. Neurocomputing **441**, 272–278 (2021)
20. Nguyen, H.T., Duong, P.H., Cambria, E.: Learning short-text semantic similarity with word embeddings and external knowledge sources. Knowl. Based Syst. **182**, 104842 (2019)
21. Peng, H., Ma, Y., Poria, S., Li, Y., Cambria, E.: Phonetic-enriched text representation for Chinese sentiment analysis with reinforcement learning. Inf. Fusion **70**, 88–99 (2021)
22. Rastogi, A., Liu, Q., Cambria, E.: Stress detection from social media articles: new dataset benchmark and analytical study. In: IJCNN (2022)
23. Satapathy, R., Rajesh Pardeshi, S., Cambria, E.: Polarity and subjectivity detection with multitask learning and BERT embedding. Future Internet **14**(7), 191 (2022)
24. Stappen, L., Baird, A., Cambria, E., Schuller, B.: Sentiment analysis and topic recognition in video transcriptions. IEEE Intell. Syst. **36**(2), 88–95 (2021)
25. Susanto, Y., Livingstone, A., Ng, B.C., Cambria, E.: The hourglass model revisited. IEEE Intell. Syst. **35**(5), 96–102 (2020)
26. Vilares, D., Peng, H., Satapathy, R., Cambria, E.: BabelSenticNet: a commonsense reasoning framework for multilingual sentiment analysis. In: IEEE SSCI, pp. 1292–1298 (2018)
27. Wang, Z., Hu, Z., Ho, S.B., Cambria, E., Tan, A.H.: MiMuSA–mimicking human language understanding for fine-grained multi-class sentiment analysis. Neural Comput. Appl. **35**(21), 15907–15921 (2023)
28. Xing, F., Cambria, E., Welsch, R.: Intelligent asset allocation via market sentiment views. IEEE Comput. Intell. Mag. **13**(4), 25–34 (2018)

29. Yue, T., Mao, R., Wang, H., Hu, Z., Cambria, E.: KnowleNet: knowledge fusion network for multimodal sarcasm detection. Inf. Fusion **100**, 101921 (2023)
30. Zhu, L., Li, W., Mao, R., Pandelea, V., Cambria, E.: PAED: zero-shot persona attribute extraction in dialogues. In: ACL, pp. 9771–9787 (2023)