Cognitive Computation (2022) 14:1–4
https://doi.org/10.1007/s12559-021-09972-0

Guest Editorial: A Decade of Sentic Computing

Erik Cambria[1] · Amir Hussain[2]

Published online: 4 January 2022
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Introduction

The opportunity to capture the opinions of the general public has raised growing interest both within the scientific community, leading to many exciting open challenges, and in the business world due to the remarkable range of benefits envisaged, including from marketing, business intelligence, and financial prediction. Mining opinions and sentiments from natural language, however, is an extremely difficult task as it involves a deep understanding of most of the explicit and implicit, regular and irregular, syntactical and semantic rules appropriate of a language.

Existing approaches to sentiment analysis mainly rely on parts of text in which opinions are explicitly expressed such as polarity terms, affect words, and their co-occurrence frequencies. However, opinions and sentiments are often conveyed implicitly through latent semantics, which make purely syntactical approaches ineffective. Concept-level approaches, instead, leverage knowledge representation and reasoning techniques to accomplish semantic text analysis. This helps the system grasp the conceptual and affective information associated with natural language opinions. By relying on large semantic knowledge bases, such approaches step away from blindly using keywords and word co-occurrence counts and instead rely on the implicit features associated with natural language concepts. Superior to purely syntactical techniques, concept-based approaches can detect subtly expressed sentiments. Concept-based approaches, in fact, can analyze multi-word expressions that do not explicitly convey emotion, but are related to concepts that do so.

Sentic computing [1] is a multidisciplinary approach to natural language processing and understanding at the crossroads between affective computing, semiotics, and commonsense reasoning, and exploits both computer and human sciences to better interpret and process social information on the Web. In sentic computing, whose term derives from the Latin 'sentire' (root of words such as sentiment and sentience) and 'sensus' (as in common sense), the analysis of natural language is based on commonsense computing [2], which enable the analysis of text not only at the document, page, or paragraph level but also at the sentence, clause, and concept level.

In the past ten years, sentic computing positioned itself as a horizontal technology that served as a back-end to many different applications in the areas of e-business, e-commerce, e-health, e-governance, e-security, e-learning, e-tourism, e-mobility, e-entertainment, and more. Some examples of such applications include financial forecasting [3], social media monitoring [4], marketing [5], and multimodal sentiment analysis [6].

In this light, this Cognitive Computation Special Issue focused on the introduction, presentation, and discussion of novel approaches that further develop and apply sentic computing models (such as the Hourglass of Emotions [7] and Sentic Patterns [8]), resources (such as AffectiveSpace [9] and SenticNet [10]), algorithms (such as Sentic LDA [11] and Sentic LSTM [12]), and applications (such as Sentic PROMs [13] and Sentic Album [14]).

Contents of the Special Issue

This special issue represents the inaugural issue for the Special Section of Cognitive Computation on Sentic Computing[1]. We received over 40 valid paper submissions. After several rounds of rigorous reviews and revisions, we decided to publish 26 of them in this special issue.

The first article, which was independently reviewed, is entitled "Ten Years of Sentic Computing" and opens the special issue with an overview of sentic computing models, resources, algorithms and applications developed between

✉ Erik Cambria
cambria@ntu.edu.sg

1 School of Computer Science and Engineering, Nanyang Technological University, Singapore, Singapore

2 School of Computing, Edinburgh Napier University, Edinburgh, United Kingdom

1 https://sentic.net/scs.pdf

2010 and 2020. Additionally, Susanto et al. explain sentic computing's key shifts and key tasks. Finally, they provide insights on future directions of the research field.

The second article, entitled "A Decade of Sentic Computing: Topic Modeling and Bibliometric Analysis", also provides an overview of the development of sentic computing in the last ten years by combining bibliometric analysis and structural topic modeling. In particular, the work examines various aspects of sentic computing literature, including tendency of annual article count, top journals, countries, institutions, authors, scientific collaborations between major contributors, as well as the major topics and their tendencies.

Next, "An Ensemble Method for Radicalization and Hate Speech Detection Online Empowered by Sentic Computing" proposes two novel feature extraction methods that use AffectiveSpace and SenticNet. In addition, this paper presents a machine learning framework using an ensemble of different features to improve the overall classification performance. Authors perform a thorough evaluation of the proposed features across five different datasets that cover radicalization and hate speech detection tasks.

The article "DomainSenticNet: An Ontology and a Methodology Enabling Domain-Aware Sentic Computing" proposes an extension of the OntoSenticNet ontology, named DomainSenticNet, and contributes an unsupervised methodology to support the development of domain-aware Sentic applications. Authors developed an unsupervised methodology that, for each concept in OntoSenticNet, mines semantically related concepts from WordNet and Probase knowledge bases and computes domain distributional information from the entire collection of Kickstarter domain-specific crowdfunding campaigns.

In "An Effective Sarcasm Detection Approach Based on Sentimental Context and Individual Expression Habits", the authors focus on the problem of detecting sarcastic content in social media. A dual-channel convolutional neural network is proposed that analyzes not only the semantics of the target text, but also its sentimental context. In addition, SenticNet is used to add commonsense to the LSTM model. The attention mechanism is then applied to take the user's expression habits into account. A series of experiments were carried out on several public datasets, the results of which show that the proposed approach can significantly improve the performance of sarcasm detection tasks.

The paper "CAT-BiGRU: Convolution and Attention with Bi-Directional Gated Recurrent Unit for Self-Deprecating Sarcasm Detection" also focuses on the problem of detecting sarcastic content. Authors propose a novel Convolution and Attention with Bi-directional Gated Recurrent Unit (CAT-BiGRU) model, which consists of an input, embedding, convolutional, BiGRU, and two attention layers. The efficacy of the proposed model is evaluated on two SenticNet-based sentic computing resources: Amazon word embedding and AffectiveSpace. The experimental results are significantly better than many neural network-based baselines and state-of-the-art methods.

"A Multitask Framework to Detect Depression, Sentiment and Multi-label Emotion from Suicide Notes" proposes a deep multitask framework that features a knowledge module that uses SenticNet's IsaCore and AffectiveSpace vector spaces to infuse external knowledge specific features into the learning process. The system models depression detection and sentiment classification simultaneously. Evaluation results show that all proposed multitask models perform better than their single-task variants.

The work "Sentic Computing for Aspect-Based Opinion Summarization Using Multi-Head Attention with Feature Pooled Pointer Generator Network" proposes a cognitive aspect-based opinion summarizer, Feature Pooled Pointer Generator Network (FP2GN), which selectively attends to thematic and contextual cues to generate sentiment-aware review summaries. In particular, FP2GN identifies aspect terms in review text using sentic computing (SenticNet and concept frequency-inverse opinion frequency) and statistical feature engineering.

In "Stock Price Prediction Incorporating Market Style Clustering", the authors investigate how to characterize market styles to improve stock prediction performance under varying market styles. Experiments are conducted with five years of real Hong Kong Stock Exchange data that includes both stock prices and corresponding news. Two famous sentiment dictionaries (i.e., SenticNet and the Loughran-McDonald financial sentiment dictionary 2018) are employed to analyze the news sentiments.

The article "Hybrid Deep Learning Models for Thai Sentiment Analysis" proposes a framework for sentiment analysis in Thai along with the Thai-SenticNet corpus. The framework employs different types of features, namely, word embedding, part-of-speech, sentic features, and all combinations of these features. Three datasets in Thai were used in this work: ThaiTales, ThaiEconTwitter, and Wisesight datasets. The experimental results show that combining all three features and fusing deep learning algorithms were able to improve overall performance.

Next, "Ordered Weighted Averaging for Emotion-Driven Polarity Detection" introduces a fuzzy framework for computing user mood based on SenticNet and sentic patterns, which are used to guide an ordered weighted averaging operator. This operator allows the aggregation to be computed in such a way as to provide an understanding of why some positive or negative aspects are considered to a greater or lesser extent. The performance and advantages of this proposal are illustrated in depth via a variety of scenarios applied to real data.

In "Multitask Learning for Complaint Identification and Sentiment Analysis", the authors propose a deep multitask framework that features a knowledge element that uses AffectiveSpace to infuse commonsense knowledge specific features into the learning process. The framework models complaint identification (the primary task) and sentiment classification (supplementary task) simultaneously. The proposed multitask system outperforms the single-task systems indicating a strong correlation between sentiment analysis and complaint classification tasks, thus benefiting from each other when learned concurrently.

The work by Weichselbraun et al. entitled "Automatic Expansion of Domain-Specific Affective Models for Web Intelligence Applications" combines common and commonsense knowledge graphs with language models and affective reasoning, improving coverage and consistency as well as supporting domain-specific interpretations of emotions. An extensive evaluation compares the performance of different expansion techniques: (i) a quantitative evaluation based on the revisited Hourglass of Emotions model to assess performance on complex models that cover multiple affective categories, using manually compiled gold standard data, and (ii) a qualitative evaluation of a domain-specific affective model for television program brands.

Next, "Towards Sentiment-Aware Multi-Modal Dialogue Policy Learning" presents a multi-intent-based dialogue policy by utilizing a unified hierarchical deep reinforcement learning framework. The system is trained to serve multiple intents of the user for a particular domain by leveraging from abstractions exhibited at different hierarchical levels developed to execute diverse tasks at different time-steps. Multimodal information elicitation strategy is employed to identify user's preference for different entities in the dialogue system.

"Design and Deployment of an Image Polarity Detector with Visual Attention" presents a novel hardware-friendly detector of image polarity, enhanced with the ability of saliency detection. The approach stems from a hardware-oriented design process, which trades off prediction accuracy and computational resources. The final solution combines lightweight deep-learning architectures and post-training quantization. Experimental results on standard benchmarks confirmed that the design strategy can infer automatically the salient parts and the polarity of an image with high accuracy.

In "Affective Concept-Based Encoding of Patient Narratives via Sentic Computing and Neural Networks", the authors propose biomedical embeddings built using online patient narratives combined with label propagation from a widely used comprehensive biomedical vocabulary. Based on the dependence on biomedical definitions, drug reaction sample selection is also suggested. These feature settings are then used to build and recognize affective semantics and sentics. Finally, a semisupervised LSTM-BiLSTM model for biomedical sentiment analysis is constructed.

The article "Context Aware Sentiment Link Prediction in Heterogeneous Social Network" employs Sentic LSTM to incorporate various explicit context knowledge and implicit multi-dimensional user interaction association knowledge into representation learning and sentiment link prediction. Experimental results on a real-world dataset show that the proposed approach is effective and feasible for detecting unobserved sentiment links from online social networks and outperforms the state-of-the-art baselines in sentiment link prediction tasks.

Next, "Mood of the Planet: Challenging Visions of Big Data in the Arts" describes an interactive physical-digital sculpture, which emits colored light and sound as a form of visualization and sonification of the changing, live emotions expressed by people all around the Earth. It is the product of several disciplines, including the arts, computer science, linguistics and psychology. Feedback from the audience proved the Mood of the Planet to provide a more accurate, personal and tangible experience about the data-emotions dichotomy.

"Emotionally Informed Hate Speech Detection: A Multi-target Perspective" explores the ability of hate speech detection models to capture common properties from topic-generic datasets and transfer this knowledge to recognize specific manifestations of hate speech. Authors experiment with the development of models to detect both topics (racism, xenophobia, sexism, misogyny) and hate speech targets, going beyond standard binary classification. Finally, authors study the impact of affective knowledge encoded in sentic computing resources (SenticNet) and in semantically structured hate lexicons (HurtLex) in determining specific manifestations of hate speech.

Following, "To BAN or Not to BAN: Bayesian Attention Networks for Reliable Hate Speech Detection" proposes a Bayesian method using Monte Carlo dropout within the attention layers of the transformer models to provide well-calibrated reliability estimates. Authors evaluate and visualize the results of the proposed approach on hate speech detection problems in several languages. Additionally, the authors test whether affective dimensions can enhance the information extracted by the BERT model in hate speech classification.

In "Does Twitter Affect Stock Market Decisions? Financial Sentiment Analysis During Pandemics: A Comparative Study of the H1N1 and the COVID-19 Periods", the authors analyze Twitter data and important worldwide financial indices to answer the following question: How does the polarity generated by Twitter posts influence the behavior of financial indices during pandemics? This study is based on the financial sentiment analysis of influential Twitter accounts and its relationship with the behavior of important financial indices. To carry out this analysis, authors used fundamental and technical financial analysis combined with a lexicon-based approach on financial Twitter accounts.

The article “Words, Tweets, and Reviews: Leveraging Affective Knowledge Between Multiple Domains” proposes a method for transferring affective knowledge between words, tweets, and movie reviews using two representation techniques: Word2Vec static embeddings and BERT contextualized embeddings. Their experimental results show that affective knowledge can be successfully transferred between the three domains, that contextualized embeddings tend to outperform their static counterparts, and that better transfer learning results are obtained when the source domain has longer textual units than the target domain.

In the work “Distant Supervised Construction and Evaluation of a Novel Dataset of Emotion-Tagged Social Media Comments in Spanish”, a dataset of social media comments in Spanish is compiled, selected, filtered, and presented. A sample of the dataset is reclassified by a group of psychologists and validated using the Fleiss Kappa interrater agreement measure. Error analysis is performed by using the Sentic Computing tool BabelSenticNet. Results indicate that the agreement between the human raters and the automatically acquired tag is moderate, similar to other manually tagged datasets, with the advantages that the presented dataset contains several hundreds of thousands of tagged comments and it does not require extensive manual tagging.

“Incremental Word Vectors for Time-Evolving Sentiment Lexicon Induction” is about automatically inducing continuously updated sentiment lexicons from Twitter streams by training incremental word sentiment classifiers from time-evolving distributional word vectors. Authors experiment with various sketching techniques for efficiently building incremental word context matrices and study how the lexicon adapts to drastic changes in the sentiment pattern.

In “Automatically Building Financial Sentiment Lexicons While Accounting for Negation”, the authors construct new lexicons by leveraging 200,000 messages from StockTwits. They evaluate the constructed financial sentiment lexicons in two different sentiment classification tasks (unsupervised and supervised). In addition, the created financial sentiment lexicons are compared with each other and with other existing sentiment lexicons.

Finally, the work by Wawer et al. entitled “Single and Cross-Disorder Detection for Autism and Schizophrenia” explores the limits of automated detection of autism spectrum disorder and schizophrenia. The article tests the effectiveness of several baseline approaches, such as bag of words and dictionary-based vectors, followed by a machine learning model. Authors employed two more refined Sentic text representations using affective features and concept-level analysis on texts. The best breed of automated methods outperformed human raters (psychiatrists).

Acknowledgements The guest editors would like to thank the Journal Publishing Editor, Danielle Lupo, for facilitating the timely publication of this special issue. Hussain would also like to acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC) - Grants Ref. EP/M026981/1, EP/T021063/1, EP/T024917/1, and the Scottish Government Chief Scientist Office (Grant Ref. COV/NAP/20/07).

References

1. Cambria E, Hussain A. Sentic Computing: A Common-Sense-Based Framework for Concept-Level Sentiment Analysis. Cham, Switzerland: Springer; 2015.
2. Cambria E, Hussain A, Havasi C, Eckl C. Common sense computing: From the society of mind to digital intuition and beyond. In: Biometric ID Management and Multimodal Communication, vol. 5707. Lecture Notes in Computer Science. Berlin Heidelberg: Springer; 2009. p. 252–9.
3. Picasso A, Merello S, Ma Y, Oneto L, Cambria E. Technical analysis and sentiment embeddings for market trend prediction. Expert Syst Appl. 2019;135:60–70.
4. Chandra JK, Cambria E, Nanetti A. One belt, one road, one sentiment? A hybrid approach to gauging public opinions on the New Silk Road initiative. In: ICDM, 2020;7–14.
5. Cambria E, Grassi M, Hussain A, Havasi C. Sentic computing for social media marketing. Multimedia Tools and Applications. 2012b;59(2):557–77.
6. Stappen L, Baird A, Cambria E, Schuller BW. Sentiment analysis and topic recognition in video transcriptions. IEEE Intell Syst. 2021;36:2.
7. Susanto Y, Livingstone A, Ng BC, Cambria E. The hourglass model revisited. IEEE Intell Syst. 2020;35(5):96–102.
8. Poria S, Cambria E, Winterstein G, Huang GB. Sentic patterns: Dependency-based rules for concept-level sentiment analysis. Knowl-Based Syst. 2014;69:45–63.
9. Cambria E, Fu J, Bisio F, Poria S. AffectiveSpace 2: Enabling affective intuition for concept-level sentiment analysis. In: AAAI, 2015;508–514.
10. Cambria E, Li Y, Xing F, Poria S, Kwok K. SenticNet 6: Ensemble application of symbolic and subsymbolic AI for sentiment analysis. In: CIKM, 2020;105–114.
11. Poria S, Chaturvedi I, Cambria E, Bisio F. Sentic LDA: Improving on LDA with semantic similarity for aspect-based sentiment analysis. In: IJCNN 2016;4465–4473
12. Ma Y, Peng H, Cambria E. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive LSTM. In: AAAI, 2018;5876–5883.
13. 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. 2012a;39(12):10533–43.
14. Cambria E, Hussain A. Sentic album: Content-, concept-, and context-based online personal photo management system. Cogn Comput. 2012;4(4):477–96.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.