Knowledge-Based Systems 238 (2022) 107920

Contents lists available at ScienceDirect

Knowledge-Based Systems

journal homepage: www.elsevier.com/locate/knosys

Editorial

Guest Editorial: Explainable artificial intelligence for sentiment analysis

1. Abstract

Social media analytics have proven valuable in numerous research areas as a pragmatic tool for public opinion mining and analysis [1]. Sentiment analysis addresses the dynamics of complex socio-affective applications that permeate intelligence and decision making in the sentient and solution-savvy Social Web [2]. Having started as simple polarity detection, contemporary sentiment analysis has advanced to a more nuanced analysis of affect and emotion sensing [3]. Detecting fine-grained sentiment in natural language, however, is tricky even for humans, making its automated detection very complicated. Moreover, online opinions can be put forth in the form of text reviews or ratings, for a product as a whole, or each of its individual aspects [4]. Multiple and lengthy reviews, usage of casual dialect with microtext (wordplay, neologism and slang), use of figurative language (sarcasm, irony), multilingual content (code-mixed and code-switched) and opinion spamming add challenges to the task of extracting opinions.

Recently memes, GIFs, typo-graphic (artistic way of text representation), info-graphic (text embedded along with an image) visual content and edited videos also dominate social feeds. Consequently, the intra-modal modeling and inter-modal interactions between the textual, visual and acoustic components add to the linguistic challenges [5]. Therefore, conceptualization and development of multi-faceted sentiment analysis models to adequately capture observed opinion-sensitive information are imperative. Artificial-intelligence driven models, especially deep learning models, have achieved state-of-the-art results for various natural language processing tasks including sentiment analysis. We get highly accurate predictions using these in conjunction with large datasets, but with little understanding of the internal features and representations of the data that a model uses to classify into sentiment categories. Most techniques do not disclose how and why decisions are taken. In other words, these black-box algorithms lack transparency and explainability. Explainable artificial intelligence (XAI) is an emerging field in machine learning that aims to address how artificial-intelligence systems make decisions in contexts such as dialogue systems [6], financial forecasting [7], healthcare [8], and more. It refers to artificial-intelligence methods and techniques that produce human-comprehensible solutions by leveraging the ensemble application of symbolic and subsymbolic AI [9]. XAI solutions will enable enhanced prediction accuracy with decision understanding and traceability of actions taken. XAI aims to improve human understanding, determine the justifiability of decisions made by the machine, introduce trust and reduce bias.

2. Contents of the special issue

This special issue focused on emerging techniques and trendy applications of XAI in fields such as sentiment analysis, intention recognition, commonsense reasoning, narrative understanding, speech emotion recognition, and empathetic dialogue systems. We received over 60 valid paper submissions. After several rounds of rigorous reviews and revisions, we decided to include 10 of them in this special issue.

The first article is entitled "A Constrained Optimization Approach for Cross-Domain Emotion Distribution Learning" [10] and proposes a constrained optimization approach based on NMTF for cross-domain emotion distribution learning. In this model, the relationship between document clusters and emotion labels is not always one-to-one. A novel content-based constraint is also endowed based on the hypothesis that documents belonging to the same clusters must have similar content. Authors solve the optimization problem by leveraging an alternately iterative algorithm and show the proof of convergence. Experiments on 12 real-world cross-domain emotion distribution learning tasks validate the effectiveness of the proposed method.

In "AMFF: A New Attention-based Multi-Feature Fusion Method for Intention Recognition" [11], the authors propose AMFF, an attention-based multi-feature fusion method for intention recognition. In particular, they enrich short text features by fusing features extracted from different word representations. In order to measure important features, they utilize attention mechanisms to assign weights for the fusion features. Experimental results on different datasets demonstrate that the proposed AMFF model outperforms traditional machine learning models and typical deep learning models on short text classification.

The article "A Commonsense Reasoning Framework for Explanatory Emotion Attribution, Generation and Re-classification" [12] presents DEGARI (Dynamic Emotion Generator And ReclassIfier), an explainable system for emotion attribution and recommendation. This system relies on a recently introduced commonsense reasoning framework, the TCL logic, which is based on a human-like procedure for the automatic generation of novel concepts in a Description Logics knowledge base. Starting from an ontological formalization of emotions based on the Plutchik model, known as ArsEmotica, the system exploits the logic TCL to automatically generate novel commonsense semantic representations of compound emotions.

Next, the authors of "Attention Uncovers Task-Relevant Semantics in Emotional Narrative Understanding" [13] investigate

https://doi.org/10.1016/j.knosys.2021.107920

the capability of an attention mechanism to 'attend to' semantically meaningful words. Using a dataset of naturalistic emotional narratives, they first build a Window-Based Attention (WBA) consisting of a hierarchical, two-level long short-term memory (LSTM) with softmax attention. Their model outperforms state-of-the-art models at predicting emotional valence, and even surpassing average human performance. Experimental results using six different pre-trained word embeddings suggest that deep neural network models which achieve human-level performance may learn to place greater attention weights on words that humans find semantically meaningful to the task at hand.

The work entitled "Empathetic Response Generation through Graph-based Multi-hop Reasoning on Emotional Causality" [14] considers emotional causality, namely, what feelings the user expresses (i.e., emotion) and why the user has such feelings (i.e., cause). Most existing works, in fact, merely focus on what the emotion is and ignore how the emotion is evoked, thus weakening the capacity of the model to understand the emotional experience of the user for generating empathetic responses. The authors, instead, propose a novel graph-based model with multi-hop reasoning to model the emotional causality of the empathetic conversation. Finally, they demonstrate the effectiveness of their model on EmpatheticDialogues in comparison with several competitive models.

Following, the authors of "Knowledge enabled BERT for Aspect based Sentiment Analysis" [15] propose a knowledge-enabled language representation model BERT for aspect-based sentiment analysis. Specifically, their proposal leverages the additional information from a sentiment knowledge graph by injecting sentiment domain knowledge into the language representation model, which obtains the embedding vectors of entities in the sentiment knowledge graph and words in the text in a consistent vector space. In addition, the model is capable of achieving better performance with a small amount of training data by incorporating external domain knowledge into the language representation model to compensate for the limited training data.

"Sentiment Analysis Using Novel and Interpretable Architectures of Hidden Markov Models" [16] introduces novel, interpretable HMM-based methods for recognizing sentiments in text and examines their performance under various architectures, training methods, orders and ensembles. The introduced models possess interpretability, they can indicate the sentimental parts of a sentence and illustrate the way that the overall sentiment evolves from the start to the end of it. A concrete experimental study is conducted and the results show that the introduced HMMs methods and the training approaches are quite competitive with machine learning methods and that they outperform traditional HMMs. Furthermore, the designed HMMs methods possess great interpretability and can be an efficient approach for sentiment analysis.

In "ADOPS: Aspect Discovery OPinion Summarisation Methodology based on deep learning and subgroup discovery for generating explainable opinion summaries" [17], the authors present the Aspect Discovery for OPinion Summarisation (ADOPS) methodology, which is aimed at generating explainable and structured opinion summaries. ADOPS is built upon aspect-based sentiment analysis methods based on deep learning and Subgroup Discovery techniques. The resultant opinion summaries are presented as interesting rules, which summarize in explainable terms for humans the state of the opinion about the aspects of a specific entity. Authors annotate and release a new dataset of opinions about a single entity on the restaurant review domain for assessing the ADOPS methodology. The results show that ADOPS is able to generate interesting rules with high values of support and confidence, which provide explainable and insightful knowledge about the state of the opinion of a certain entity.

The paper "Sentiment Lossless Summarization" [18] introduces a sentiment compensation mechanism into document summarization and proposes a graph-based extractive summarization approach named Sentiment Lossless Summarization (SLS). SLS first creates a graph representation for a document to obtain the importance score (i.e., literal indicator) of each sentence. Second, sentiment dictionaries are leveraged to analyze the sentence sentiments. Third, during each summarization iteration, the sentences with the lowest scores are iteratively removed, and the sentiment compensation weights of the remaining sentences are updated. With the help of sentiment compensation during the summarization process, sentiment consistencies between candidate summaries and the original documents are maintained.

Finally, "Combining Cross-modal Knowledge Transfer and Semi-supervised Learning for Speech Emotion Recognition" [19] proposes a new architecture that combines cross-modal knowledge transfer from visual to audio modality into a semi-supervised learning method with consistency regularization. Authors posit that introducing visual emotional knowledge by the cross-modal transfer method can increase the diversity and accuracy of pseudo-labels and improve the robustness of the model. To combine knowledge from cross-modal transfer and semi-supervised learning, authors design two fusion algorithms, i.e., weighted fusion and consistent and random. Experiments on two datasets show that the proposed method can effectively use additional unlabeled audio-visual data to outperform state-of-the-art results.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank the Managing Editor-in-Chief, Prof Jie Lu, and Knowledge-Based Systems journal managers for their valuable support during the preparation of this special issue. In addition, we are deeply grateful to the reviewers for their contributions to the decision process. This work is supported by the Agency for Science, Technology and Research (A*STAR), Singapore under its AME Programmatic Funding Scheme (Project #A18A2b0046).

References

[1] Erik Cambria, Haixun Wang, Bebo White, Guest editorial: Big social data analysis, Knowl.-Based Syst. 69 (2014) 1–2.
[2] Marco Grassi, Erik Cambria, Amir Hussain, Francesco Piazza, Sentic web: A new paradigm for managing social media affective information, Cogn. Comput. 3 (3) (2011) 480–489.
[3] Zhaoxia Wang, Sengbeng Ho, Erik Cambria, A review of emotion sensing: Categorization models and algorithms, Multimedia Tools Appl. 79 (2020) 35553–35582.
[4] J. Ashok Kumar, Tina Esther Trueman, Erik Cambria, A convolutional stacked bidirectional LSTM with a multiplicative attention mechanism for aspect category and sentiment detection, Cogn. Comput. (2021).
[5] Luka Stappen, Alice Baird, Erik Cambria, Björn Schuller, Sentiment analysis and topic recognition in video transcriptions, IEEE Intell. Syst. 36 (2) (2021).
[6] Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, Minlie Huang, Augmenting end-to-end dialogue systems with commonsense knowledge, in: AAAI, 2018, pp. 4970–4977.
[7] Andrea Picasso, Simone Merello, Yukun Ma, Luca Oneto, Erik Cambria, Technical analysis and sentiment embeddings for market trend prediction, Expert Syst. Appl. 135 (2019) 60–70.
[8] Erik Cambria, Amir Hussain, Tariq Durrani, Catherine Havasi, Chris Eckl, James Munro, Sentic computing for patient centered applications, in: IEEE ICSP, 2010, pp. 1279–1282.

[9] Erik Cambria, Yang Li, Frank Xing, Soujanya Poria, Kenneth Kwok, SenticNet 6: Ensemble application of symbolic and subsymbolic AI for sentiment analysis, in: CIKM, 2020, pp. 105–114.
[10] Xiaorui Qin, Yufu Chen, Yanghui Rao, Haoran Xie, Man Leung Wong, Fu Lee Wang, A constrained optimization approach for cross-domain emotion distribution learning, Knowl.-Based Syst. 227 (2021) 107160.
[11] Cong Liu, Xiaolong Xu, AMFF: A New attention-based multi-feature fusion method for intention recognition, Knowl.-Based Syst. 233 (2021) 107525.
[12] Antonio Lieto, Gianluca Pozzato, Stefano Zoia, Viviana Patti, Rossana Damiano, A commonsense reasoning framework for explanatory emotion attribution, generation and re-classification, Knowl.-Based Syst. (2021).
[13] Thanh-Son Nguyen, Zhengxuan Wu, Desmond C. Ong, Attention uncovers task-relevant semantics in emotional narrative understanding, Knowl.-Based Syst. 226 (2021) 107162.
[14] Jiashuo Wang, Wenjie Li, Peiqin Lin, Feiteng Mu, Empathetic response generation through graph-based multi-hop reasoning on emotional causality, Knowl.-Based Syst. 233 (2021) 107547.
[15] Anping Zhao, Yu Yu, Knowledge-enabled BERT for aspect-based sentiment analysis, Knowl.-Based Syst. 227 (2021) 107220.
[16] Isidoros Perikos, Spyridon Kardakis, Ioannis Hatzilygeroudis, Sentiment analysis using novel and interpretable architectures of hidden Markov models, Knowl.-Based Syst. 229 (2021) 107332.
[17] Miguel López, Eugenio Martínez-Cámara, M. Victoria Luzón, Francisco Herrera, ADOPS: Aspect Discovery OPinion Summarisation Methodology based on deep learning and subgroup discovery for generating explainable opinion summaries, Knowl.-Based Syst. 231 (2021) 107455.
[18] Xiaodong Li, Pangjing Wu, Chenxin Zou, Haoran Xie, Fu Lee Wang, Sentiment lossless summarization, Knowl.-Based Syst. 227 (2021) 107170.
[19] Sheng Zhang, Min Chen, Jincai Chen, Yuan-Fang Li, Yiling Wu, Minglei Li, Chuanbo Zhu, Combining cross-modal knowledge transfer and semi-supervised learning for speech emotion recognition, Knowl.-Based Syst. 229 (2021) 107340.

Erik Cambria *
Nanyang Technological University, Singapore
E-mail address: cambria@ntu.edu.sg.

Akshi Kumar
Delhi Technological University, India

Mahmoud Al-Ayyoub
Jordan University of Science and Technology, Jordan

Newton Howard
Oxford University, UK

5 December 2021
Available online 11 December 2021

* Corresponding editor.