

Affective Computing and Sentiment Analysis

Erik Cambria, *Nanyang Technological University*

The question is not whether intelligent machines can have any emotions, but whether machines can be intelligent without any emotions. — Marvin Minsky (1927–2016)

Emotions play an important role in successful and effective human–human communication. In fact, in many situations, emotional intelligence is more important than IQ for successful interaction.¹ There is also significant evidence that rational learning in humans is dependent on emotions.²

Affective computing and sentiment analysis, hence, are key for the advancement of AI³ and all the research fields that stem from it. Moreover, they find applications in various scenarios and companies, large and small, that include the analysis of emotions and sentiments as part of their mission. Sentiment-mining techniques can be exploited for the creation and automated upkeep of review and opinion aggregation websites, in which opinionated text and videos are continuously gathered from the Web and not restricted to just product reviews, but also to wider topics such as political issues and brand perception.

Affective computing and sentiment analysis also have great potential as a subcomponent technology for other systems. They can enhance the capabilities of customer relationship management and recommendation systems—for example, to reveal which features customers enjoy or to exclude from the recommendations items that received negative feedback. Similarly, they can be exploited for affective tutoring and affective entertainment or for troll filtering and spam detection in online social communication.

Business intelligence is also a main factor behind corporate interest in the fields of affective computing and sentiment analysis. Nowadays, companies

invest an increasing amount of money in marketing strategies and are constantly interested in both collecting and predicting the attitudes of the general public toward their products and brands. The design of automatic tools capable of mining sentiments over the Web in real time and creating condensed versions of these represents one of the most active research and development areas. The development of such systems, moreover, is not only important for commercial purposes but also for government intelligence applications able to monitor increases in hostile communications or model cyber-issue diffusion.

Several commercial and academic tools, such as those from IBM (www.ibm.com/analytics), SAS (www.sas.com/social), Oracle (www.oracle.com/social), SenticNet (www.business.sentic.net), and Luminoso (www.luminoso.com), track public viewpoints on a large scale by offering graphical summarizations of trends and opinions in the blogosphere. Nevertheless, most COTS tools are limited to a polarity evaluation or a mood classification according to a limited set of emotions. In addition, such methods rely mainly on parts of text in which emotional states are explicitly expressed and, hence, they cannot capture opinions and sentiments that are expressed implicitly. Because they are based mainly on statistical properties associated with words, in fact, many COTS tools are easily tricked by linguistic operators such as negation and disjunction.

In this article, I list common tasks of affective computing and sentiment analysis and present a general categorization for them.

Common Tasks

The Web is evolving toward an era where communities will define future products and services.⁴ In this context, public opinion is destined to gain

increasing prominence, and so are affective computing and sentiment analysis (see Figure 1).

The basic tasks of affective computing and sentiment analysis are emotion recognition^{2,5-8} and polarity detection.⁹⁻¹² Although the former focuses on extracting a set of emotion labels, the latter is usually a binary classification task with outputs such as “positive” versus “negative,” “thumbs up” versus “thumbs down,” or “like” versus “dislike.” These two tasks are highly interrelated and interdependent to the extent that some sentiment categorization models, such as the Hourglass of Emotions,¹³ treat them as a unique task by inferring the polarity associated to a sentence directly from the emotions this conveys. In many cases, in fact, emotion recognition is considered a subtask of polarity detection.

Polarity classification itself can also be viewed as a subtask of more advanced analyses. For example, it can be applied to identifying pro and con expressions that can be used in individual reviews to evaluate the pros and cons that influenced the judgements of a product and that make such judgements more trustworthy. Another instance of binary sentiment classification is agreement detection, that is, given a pair of affective inputs, deciding whether they should receive the same or differing sentiment-related labels.

Complementary to binary sentiment classification is the assignment of degrees of positivity to the detected polarity or valence to the inferred emotions. If we waive the assumption that the input under examination is opinionated and is about a single issue or item, challenging new tasks arise, such as subjectivity detection and opinion target identification.¹⁴ The capability of distinguishing whether an input is subjective or objective, in particular, can be highly

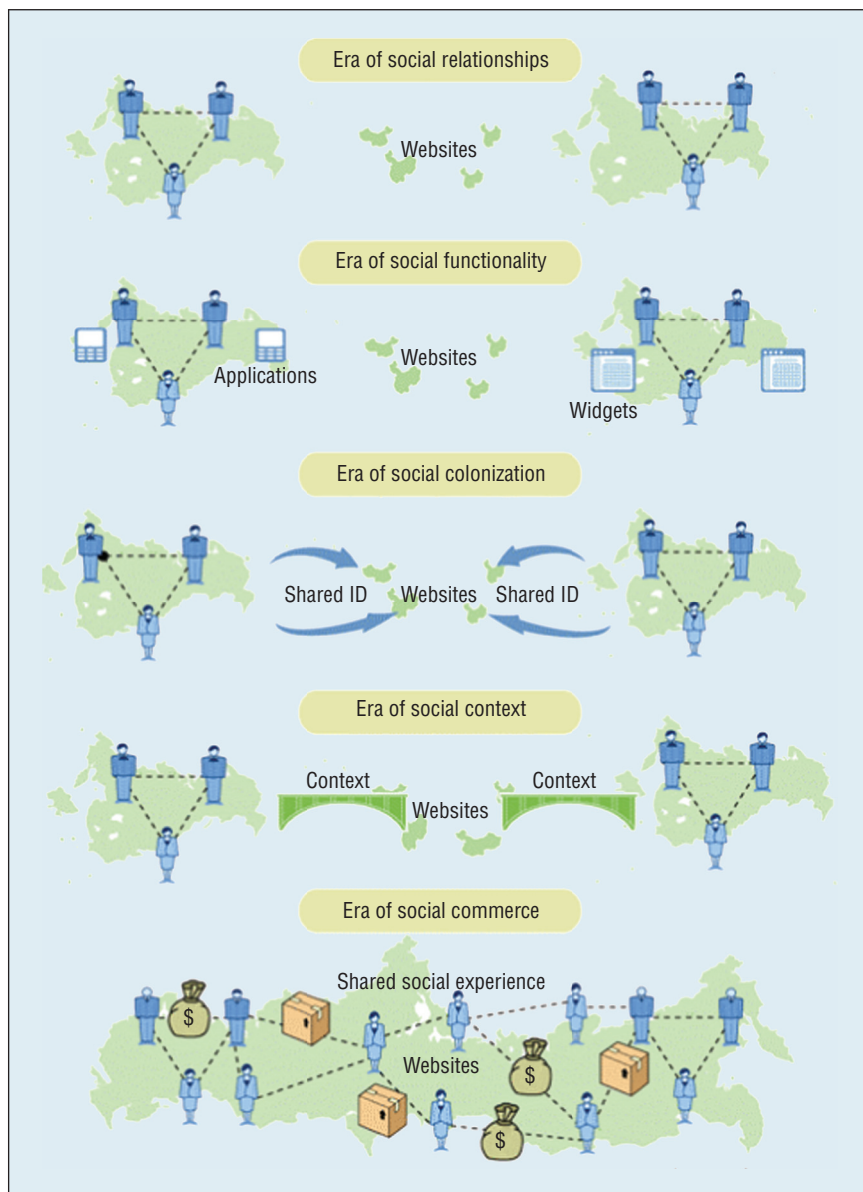


Figure 1. The five-eras vision shows that mining sentiments from the general public is becoming increasingly important for the future of the Web.⁴ According to this vision, we are gradually shifting to an era in which people’s opinions will dictate the final shape of products and services. (Source: Jeremiah Owyang; used with permission.)

beneficial for a more effective sentiment classification. Moreover, a record can have a polarity without necessarily containing an opinion—for example, a news article can be classified into good or bad news without being subjective.

Typically, affective computing and sentiment analysis are performed over

on-topic documents (for example, on the result of a topic-based search engine). However, several studies suggested that managing these tasks jointly can be beneficial for overall performance. For example, off-topic passages of a document could contain irrelevant affective information and misleading results for the global

sentiment polarity about the main topic. Also, a document can contain material on multiple topics that might interest the user. In this case, it is therefore necessary to identify the topics and separate the opinions associated with each of them.

Another important task of affective computing and sentiment analysis is multimodal fusion. With increasing amounts of webcams installed in end-user devices such as smart phones, touchpads, or netbooks, an increasing amount of affective information is being posted to social online services in an audio or audiovisual format rather than a purely textual basis. For a rough impression on the extent, consider that, on average, two days of video material are uploaded to YouTube per minute. Besides speech-to-text recognition, this allows for additional exploitation of acoustic information, facial expression and body movement analysis, or even the “mood” of the background music or color filters.

Multimodal fusion means to integrate all single modalities into a combined single representation. Two types of fusion techniques—feature level and decision level—have been used to improve reliability in emotion recognition from multimodal information.¹⁵ Stephan Raaijmakers and colleagues fused acoustic and linguistic information,¹⁶ but linguistic information is based on the transcript of the spoken content rather than on automatic speech recognition output. Louis-Philippe Morency and colleagues combined acoustic, textual, and video features to assess opinion polarity in 47 YouTube videos.¹⁷ They demonstrated a significant improvement in a leave-one-video-out evaluation using hidden Markov models for classification. They identified polarized words, smile, gaze, pauses, and voice pitch as relevant features. However, they based the textual analysis

only on the manual transcript of spoken words.

Soujanya Poria and colleagues proposed a novel methodology for multimodal sentiment analysis that comprises harvesting sentiments from Web videos by demonstrating a model that uses audio, visual, and textual modalities as sources of information.¹⁸ They used both feature- and decision-level fusion methods to merge affective information extracted from multiple modalities, achieving an accuracy of nearly 80 percent.

General Categorization

Existing approaches to affective computing and sentiment analysis fall into three main categories: knowledge-based techniques, statistical methods, and hybrid approaches.

Knowledge-based techniques are popular because of their accessibility and economy. Text is classified into affect categories on the basis of the presence of fairly unambiguous affect words, such as “happy,” “sad,” “afraid,” and “bored.” Popular sources of affect words or multiword expressions include the Affective Lexicon,¹⁹ linguistic annotation scheme,²⁰ WordNet-Affect,²¹ SentiWordNet,²² SenticNet,²³ and other probabilistic knowledge bases trained from linguistic corpora.^{24–26}

The major weakness of knowledge-based approaches is poor recognition of affect when linguistic rules are involved. For example, although a knowledge base can correctly classify the sentence “Today was a happy day” as being happy, it is likely to fail on a sentence like “Today wasn’t a happy day at all.” To this end, more sophisticated knowledge-based approaches exploit linguistics rules to distinguish how each specific knowledge base entry is used in text.²⁷

The validity of knowledge-based approaches, moreover, depends heav-

ily on the depth and breadth of the employed resources. Without a comprehensive knowledge base that encompasses human knowledge, in fact, it is not easy for a sentiment-mining system to grasp the semantics associated with natural language or human behavior.

Another limitation of knowledge-based approaches lies in the typicality of their knowledge representation, which is usually strictly defined and does not allow handling different concept nuances, because the inference of semantic and affective features associated with concepts is bounded by the fixed, flat representation.

Statistical methods, such as support vector machines and deep learning, have been popular for affect classification of texts, and researchers have used them on projects such as a movie review classifier²⁸ and many others.^{29–32}

By feeding a machine learning algorithm a large training corpus of affectively annotated texts, it is possible for the system to not only learn the affective valence of affect keywords (as in the keyword-spotting approach) but also to consider the valence of other arbitrary keywords (like lexical affinity) and word co-occurrence frequencies.

However, statistical methods are generally semantically weak—that is, lexical or co-occurrence elements in a statistical model have little predictive value individually. As a result, statistical text classifiers work with acceptable accuracy only when given a sufficiently large text input. So, although these methods might be able to affectively classify a user’s text on the page or paragraph level, they do not work well on smaller text units such as sentences or clauses.

Hybrid approaches to affective computing and sentiment analysis, finally, exploit both knowledge-based

techniques and statistical methods to perform tasks such as emotion recognition and polarity detection from text or multimodal data. Sentic computing,³³ for example, exploits an ensemble of knowledge-driven linguistic patterns and statistical methods to infer polarity from text (see Figure 2). Yunqing Xia and colleagues used SenticNet and a Bayesian model for contextual concept polarity disambiguation.³⁴ Mauro Dragoni and colleagues proposed a fuzzy framework that merges WordNet, ConceptNet, and SenticNet to extract key concepts from a sentence.³⁵ iFeel is a system that lets users create their own sentiment analysis framework by combining SenticNet, SentiWordNet, and other sentiment analysis methods.³⁶ Jose Chenlo and David Losada used SenticNet to extract bag-of-concepts and polarity features for subjectivity detection and other sentiment analysis tasks.³⁷ Jay Kuan-Chieh Chung and colleagues used SenticNet concepts as seeds and proposed a method of random walk in ConceptNet to retrieve more concepts along with polarity scores.³⁸ Other works have proposed the joint use of knowledge bases and machine learning for Twitter sentiment analysis,³⁹ short text message classification,⁴⁰ and frame-based opinion mining.⁴¹

The passage from a read-only to a read-write Web made users more enthusiastic about sharing their emotions and opinions through social networks, online communities, blogs, wikis, and other online collaborative media. In recent years, this collective intelligence has spread to many different areas of the Web, in particular to fields related to everyday life, such as commerce, tourism, education, and health.

Despite significant progress, affective computing and sentiment analysis




Figure 2. Sentic computing's hybrid framework for polarity detection. Text is first deconstructed into concepts: if these are available in SenticNet, sentic patterns are triggered; if no match is found, machine learning is applied.

are still finding their own voice as new interdisciplinary fields. Engineers and computer scientists use machine learning techniques for automatic affect classification from video, voice, text, and physiology. Psychologists use their long tradition of emotion research with their own discourse, models, and methods. Affective computing and sentiment analysis are research fields that are inextricably bound to affective sciences that attempt to understand human emotions. Simply put, the development of affect-sensitive systems cannot be divorced from the century-long psychological research on emotion.

Hybrid approaches aim to better grasp the conceptual rules that govern sentiment and the clues that can convey these concepts from realization to verbalization in the human mind. In recent years, such approaches are gradually setting affective computing and sentiment analysis as interdisciplinary fields between mere natural language processing and natural

language understanding by gradually shifting from syntax-based techniques to more and more semantics-aware frameworks,⁴² which consider both conceptual knowledge and sentence structure (see Figure 3).

So far, sentiment-mining approaches from text or speech have been based mainly on the bag-of-words model because, at first glance, the most basic unit of linguistic structure appears to be the word. Single-word expressions, however, are just a subset of concepts, multiword expressions that carry specific semantics, and sentics (that is, the denotative and connotative information commonly associated with objects, actions, events, and people). Sentics, in particular, specifies the affective information associated with real-world entities, which is key for emotion recognition and polarity detection, the basic tasks of affective computing and sentiment analysis.

The best way forward for these two fields, hence, is the ensemble application

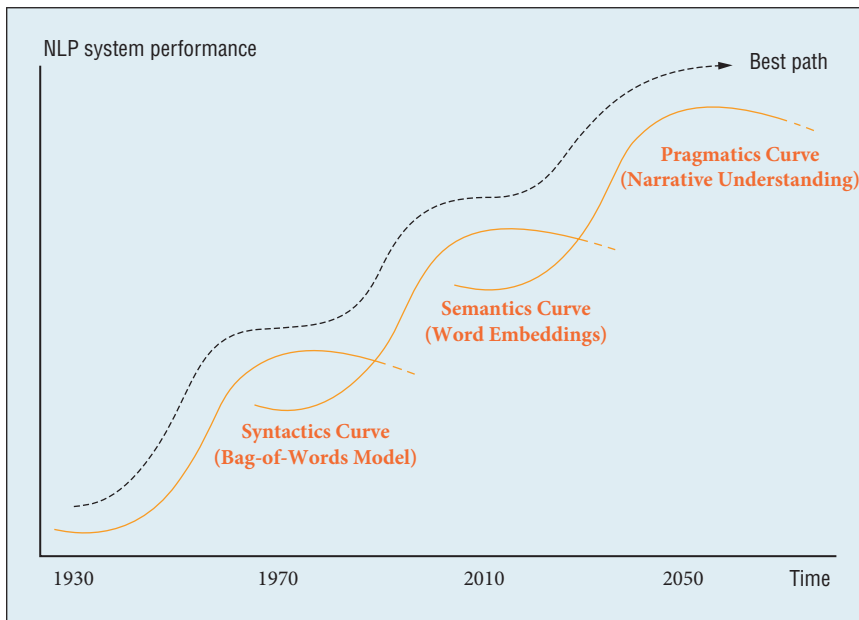


Figure 3. Jumping natural language processing (NLP) curves. Borrowed from the field of business management and marketing prediction, this paradigm reinterprets the evolution of NLP research as the intersection of three overlapping curves, which will eventually lead NLP research to evolve into natural language understanding.

of semantic knowledge and machine learning, in which different approaches can cover for each other's flaws. In particular, the combined application of linguistics and knowledge bases will allow sentiments to flow from concept to concept on the basis of the input sentence's dependency relations, while machine learning acts as backup for missing concepts and unknown linguistic patterns.

Next-generation sentiment-mining systems need broader and deeper common and commonsense knowledge bases, together with more brain-inspired and psychologically motivated reasoning methods, to better understand natural language opinions and, hence, more efficiently bridge the gap between (unstructured) multimodal information and (structured) machine-processable data. Looking ahead, blending scientific theories of emotion with the practical engineering goals of analyzing sentiments in natural language and human behavior will pave the way for the development of more bioinspired approaches to the

design of intelligent sentiment-mining systems that can handle semantic knowledge, make analogies, learn new affective knowledge, and detect, perceive, and "feel" emotions. ■


References

1. M. Pantic et al., "Affective Multimodal Human-Computer Interaction," *Proc. ACM Int'l Conf. Multimedia*, 2005, pp. 669–676.
2. R. Picard, *Affective Computing*, MIT Press, 1997.
3. M. Minsky, *The Emotion Machine: Commonsense Thinking, Artificial Intelligence, and the Future of the Human Mind*, Simon & Schuster, 2006.
4. J. Owyang, "The Future of the Social Web: In Five Eras," blog, 27 Apr. 2009; www.web-strategist.com/blog/2009/04/27.
5. R. Calvo and S. D'Mello, "Affect Detection: An Interdisciplinary Review of Models, Methods, and Their Applications," *IEEE Trans. Affective Computing*, vol. 1, no. 1, 2010, pp. 18–37.

6. Z. Zeng et al., "A Survey of Affect Recognition Methods: Audio, Visual, and Spontaneous Expressions," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 31, no. 1, 2009, pp. 39–58.
7. B. Schuller et al., "Recognising Realistic Emotions and Affect in Speech: State of the Art and Lessons Learnt from the First Challenge," *Speech Communication*, vol. 53, nos. 9/10, 2011, pp. 1062–1087.
8. H. Gunes and B. Schuller, "Categorical and Dimensional Affect Analysis in Continuous Input: Current Trends and Future Directions," *Image and Vision Computing*, vol. 31, no. 2, 2012, pp. 120–136.
9. B. Pang and L. Lee, "Opinion Mining and Sentiment Analysis," *Foundations and Trends in Information Retrieval*, vol. 2, nos. 1–2, 2008, pp. 1–135.
10. B. Liu, *Sentiment Analysis and Opinion Mining*, Morgan and Claypool, 2012.
11. T. Wilson, J. Wiebe, and P. Hoffmann, "Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis," *Proc. Conf. Human Language Technology and Empirical Methods in Natural Language Processing*, 2005, pp. 347–354.
12. E. Cambria et al., "New Avenues in Opinion Mining and Sentiment Analysis," *IEEE Intelligent Systems*, vol. 28, no. 2, 2013, pp. 15–21.
13. E. Cambria, A. Livingstone, and A. Hussain, "The Hourglass of Emotions," *Cognitive Behavioral Systems*, LNCS 7403, 2012, pp. 144–157.
14. E. Cambria et al., "The CLSA Model: A Novel Framework for Concept-Level Sentiment Analysis," *Computational Linguistics and Intelligent Text Processing*, LNCS 9042, 2015, pp. 3–22.
15. A. Konar and A. Chakraborty, *Emotion Recognition: A Pattern Analysis Approach*, Wiley & Sons, 2015.
16. S. Raaijmakers, K. Truong, and T. Wilson, "Multimodal Subjectivity

- Analysis of Multiparty Conversation,” *Proc. Conf. Empirical Methods in Natural Language Processing*, 2008, pp. 466–474.
17. L.-P. Morency, R. Mihalcea, and P. Doshi, “Towards Multimodal Sentiment Analysis: Harvesting Opinions from the Web,” *Proc. Int’l Conf. Multimodal Interfaces*, 2011, pp. 169–176.
 18. S. Poria et al., “Fusing Audio, Visual and Textual Clues for Sentiment Analysis from Multimodal Content,” *Neurocomputing*, vol. 174, part A, 2016, pp. 50–59.
 19. A. Ortony, G. Clore, and A. Collins, *The Cognitive Structure of Emotions*, Cambridge Univ. Press, 1988.
 20. J. Wiebe, T. Wilson, and C. Cardie, “Annotating Expressions of Opinions and Emotions in Language,” *Language Resources and Evaluation*, vol. 39, no. 2, 2005, pp. 165–210.
 21. C. Strapparava and A. Valitutti, “WordNet-Affect: An Affective Extension of WordNet,” *Proc. 4th Int’l Conf. Language Resources and Evaluation*, 2004, pp. 1083–1086.
 22. A. Esuli and F. Sebastiani, “SentiWordNet: A Publicly Available Lexical Resource for Opinion Mining,” *Proc. 5th Int’l Conf. Language Resources and Evaluation*, 2006, pp. 417–422.
 23. E. Cambria, D. Olsher, and D. Rajagopal, “SenticNet 3: A Common and Common-Sense Knowledge Base for Cognition-Driven Sentiment Analysis,” *Proc. 28th AAAI Conf. Artificial Intelligence*, 2014, pp. 1515–1521.
 24. R. Stevenson, J. Mikels, and T. James, “Characterization of the Affective Norms for English Words by Discrete Emotional Categories,” *Behavior Research Methods*, vol. 39, 2007, pp. 1020–1024.
 25. S. Somasundaran, J. Wiebe, and J. Ruppenhofer, “Discourse Level Opinion Interpretation,” *Proc. Int’l Conf. Computational Linguistics*, 2008, pp. 801–808.
 26. D. Rao and D. Ravichandran, “Semi-Supervised Polarity Lexicon Induction,” *Proc. 12th Conf. European Chapter of the Assoc. for Computational Linguistics*, 2009, pp. 675–682.
 27. S. Poria et al., “Sentiment Data Flow Analysis by Means of Dynamic Linguistic Patterns,” *IEEE Computational Intelligence*, vol. 10, no. 4, 2015, pp. 26–36.
 28. B. Pang, L. Lee, and S. Vaithyanathan, “Thumbs Up? Sentiment Classification Using Machine Learning Techniques,” *Proc. Empirical Methods on Natural Language Processing*, 2002, pp. 79–86.
 29. M. Hu and B. Liu, “Mining and Summarizing Customer Reviews,” *Proc. 10th ACM SIGKDD Int’l Conf. Knowledge Discovery and Data Mining*, 2004., pp. 168–177.
 30. X. Glorot, A. Bordes, and Y. Bengio, “Domain Adaptation for Large-Scale Sentiment Classification: A Deep Learning Approach,” *Proc. 28th Int’l Conf. Machine Learning*, 2011, pp. 97–100.
 31. R. Socher et al., “Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank,” *Proc. Empirical Methods on Natural Language Processing*, 2013, pp. 1642–1654.
 32. R. Lau, Y. Xia, and Y. Ye, “A Probabilistic Generative Model for Mining Cybercriminal Networks from Online Social Media,” *IEEE Computational Intelligence*, vol. 9, no. 1, 2014, pp. 31–43.
 33. E. Cambria and A. Hussain, *Sentic Computing: A Common-Sense-Based Framework for Concept-Level Sentiment Analysis*, Springer, 2015.
 34. Y. Xia et al., “Word Polarity Disambiguation Using Bayesian Model and Opinion-Level Features,” *Cognitive Computation*, vol. 7, no. 3, 2015, pp. 369–380.
 35. M. Dragoni, A.G. Tettamanzi, and C. da Costa Pereira, “A Fuzzy System for Concept-Level Sentiment Analysis,” *Semantic Web Evaluation Challenge*, Springer, 2014, pp. 21–27.
 36. M. Araújo et al., “iFeel: A System that Compares and Combines Sentiment Analysis Methods,” *Proc. 23rd Int’l Conf. World Wide Web*, 2014, pp. 75–78.
 37. J.M. Chenlo and D.E. Losada, “An Empirical Study of Sentence Features for Subjectivity and Polarity Classification,” *Information Sciences*, vol. 280, 2014, pp. 275–288.
 38. J.K.-C. Chung, C.-E. Wu, and R.T.-H. Tsai, “Improve Polarity Detection of Online Reviews with Bag-of-Sentimental-Concepts,” *Proc. 11th European Semantic Web Conf.*, 2014; http://2014.eswc-conferences.org/sites/default/files/eswc2014-challenges_cl_submission_33.pdf.
 39. F. Bravo-Marquez, M. Mendoza, and B. Poblete, “Meta-Level Sentiment Models for Big Social Data Analysis,” *Knowledge-Based Systems*, vol. 69, 2014, pp. 86–99.
 40. G. Gezici et al., “SU-Sentilab: A Classification System for Sentiment Analysis in Twitter,” *Proc. Int’l Workshop Semantic Evaluation*, 2013, pp. 471–477.
 41. D.R. Recupero et al., “Sentilo: Frame-Based Sentiment Analysis,” *Cognitive Computation*, vol. 7, no. 2, 2014, pp. 211–225.
 42. E. Cambria and B. White, “Jumping NLP Curves: A Review of Natural Language Processing Research,” *IEEE Computational Intelligence*, vol. 9, no. 2, 2014, pp. 48–57.

Erik Cambria is an assistant professor in the School of Computer Engineering at Nanyang Technological University. He is also affiliated with the Rolls-Royce@NTU Corporate Lab, A*STAR SIMTech, and MIT Synthetic Intelligence Lab. Contact him at cambria@ntu.edu.sg.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.